Phenomenology and the Cognitive Sciences
https://doi.org/10.1007/s11097-024-10025-8

Phenomenal transparency and the boundary of cognition

Julian Hauser[1] · **Hadeel Naeem**[2]

Accepted: 15 August 2024

Abstract
Phenomenal transparency was once widely believed to be necessary for cognitive extension. Recently, this claim has come under attack, with a new consensus coalescing around the idea that transparency is neither necessary for internal nor extended cognitive processes. We take these recent critiques as an opportunity to refine the concept of transparency relevant for cognitive extension. In particular, we highlight that transparency concerns an agent's employment of a resource – and that employment is compatible with an agent consciously apprehending (or attending to) a resource. This means it is possible for an object to be transparent and opaque to an agent, even at a single moment time. Once we understand transparency in this way, the detractors' claims lose their bite, and existing arguments for transparency's necessity for cognitive extension return to apply with full force.

Keywords Phenomenal transparency · Extended cognition · Extended mind · Cognitive resources

1 Introduction

Early proponents of the hypothesis of the extended cognition argued that *phenomenal transparency*[1] is necessary for cognitive extension (Clark, 2004, 2008; Thompson & Stapleton, 2009; Wheeler, 2005, 2019).[2] These authors argued that an agent can only

[1] Phenomenal transparency is distinct from other types of transparency discussed in the literature, in particular procedural transparency and informational transparency. In addition, a concept of transparency is also employed in the literature on AI systems, where it is used to talk about the degree to which the inner workings of such systems are accessible (Andrada et al., 2022). In this paper, we focus squarely on *phenomenal* transparency, which is also sometimes – including by us in this paper – simply called transparency.

[2] Others have argued that transparency importantly contributes to cognitive extension, without quite making the step of arguing that it is necessary (Heersmink, 2013, 2015, 2016). At times, Clark can also be interpreted as belonging in this camp, as when he writes that '*typical* extended mind scenarios rely upon fluid unreflective use' (2015, p. 3373, emphasis ours).

✉ Hadeel Naeem
hadeel.naeem@khk.rwth-aachen.de

[1] LOGOS, Universitat de Barcelona, Barcelona, Spain

[2] RWTH Aachen, The Käte Hamburger Kolleg, Aachen, Germany

Published online: 29 August 2024

extend their cognitive system to some resource if she manipulates the resource without consciously apprehending it as an object. Today, this view is less popular, with a slew of authors arguing that transparency isn't necessary for cognitive extension. We push back against this and argue that, once we understand transparency properly – as not mutually exclusive with the conscious apprehension of the resource – the detractors' arguments are left without bite. This means that the existing arguments showing that transparency is necessary for cognitive extension can move forward unimpeded.

Generally put, the hypothesis of extended cognition affirms that in certain cases, cognitive processes may involve resources external to the body (Clark & Chalmers, 1998).[3] Clark and Chalmer's original theory – now widely called the first wave of extended cognition (Sutton, 2010) – was based on the *parity principle*, the idea that if we consider some body-internal resource a part of the cognitive system because it has a certain functional profile, then a body-external resource with the same functional profile is also part of the cognitive system. If Otto's notebook fulfils the same function as his internal biological memory, then the notebook is a part of his cognitive system.

The second wave of thinking about extended cognition moved away from the requirement of parity and instead emphasised complementarity between internal and external resources (Menary, 2008). On this view, a body-external resource is part of the cognitive system if it gives rise to extended cognitive skills. Third-wave accounts focus on the links between predictive processing and cognitive extension (Clark, 2015, 2022; Facchin, 2022; Kirchhoff & Kiverstein, 2019). Note that this talk of three waves oversimplifies the debate and that some authors combine features attributed to different waves. Both Clark and Facchin, for instance, combine their predictive processing accounts with the parity principle.

In this paper, we do not argue for any of these accounts; we assume that cognitive extension is possible and discuss one of the conditions potentially required for it. Where we do have a substantial commitment is to the view that extended cognition is, fundamentally, about extended *processes*. This view – now standard in the literature (Clark, 2022; Menary, 2012; Pritchard, 2018) – looks at cognitive processes and investigates whether these extend into the environment. It contrasts with accounts according to which the logically primary question is about the objects to which our cognition extends.

In all accounts of cognitive extension, the question has come up whether transparency is a necessary requirement for extension. The concept of transparency, borrowed from the phenomenological literature (Heidegger, 1990; Merleau-Ponty, 2002), describes how agents, when effortlessly and fluidly employing a resource, tend not to focus on the resource, but rather on the task at hand.[4]

[3] A related hypothesis of *extended mind* also exists (Clark & Chalmers, 1998). It claims that not only cognitive states and processes may extend, but that mental states (most commonly, beliefs) extend too. While we think that much of our discussion applies to both notions, we only explicitly discuss – and put forward claims about – extended cognition.

[4] To avoid being unnecessarily wordy, we speak of the *conscious apprehension* of a resource (or even just of *being conscious of* a resource) to mean that an agent consciously apprehends the resource *as an object*. This means, roughly speaking, that the agent is conscious of the resource as an independent object with (at least potentially) determinate properties.

The classical example is that of a master carpenter (Heidegger, 1990): when employing her hammer, the carpenter doesn't focus her attention on the hammer, but rather on the task at hand. The hammer becomes "transparent" and the agent "sees through" it to the task she is trying to accomplish. The object at the centre of her consciousness is not the hammer but the table she is building.

Some authors have argued that the transparent employment of a resource shifts the subject-world boundary (Thompson & Stapleton, 2009; Wheeler, 2019), allowing the agent to extend outwards to the resource. The expert carpenter isn't conscious of the hammer as a tool but rather experiences it as part of her body. The object she is conscious of is the table she is building with the hammer. To her, the hammer is not a problem to be solved, but is rather a bodily resource with which to solve problems (Clark, 2004, 2008). When transparently employing a resource, the resource gets thus located on the subject side of the subject-world boundary, and this motivates the claim that transparency is a necessary condition on cognitive extension.

As mentioned, this view has recently fallen out of favour (Andrada, 2020, 2021; Facchin, 2022; Smart et al., 2022). Some authors argue that even in the case of body-internal cognitive resources, it's sometimes the case that the resource is consciously apprehended (Andrada, 2020, 2021; Facchin, 2022). Thus, following the parity principle, we can't demand transparency for body-external cognitive resources. Others argue that even in paradigm cases of cognitive extension, transparency doesn't obtain (Facchin, 2022). And finally, authors influenced by second-wave thinking on cognitive extension have argued that even highly skilful actions can come without transparency – implying that when such skills are extended, they need not come with transparency either (Andrada, 2021).

While these arguments are ostensibly quite diverse, they all repose on a similar approach (or so we'll argue in this paper). They look at cases of resource use and show how the agent consciously apprehends the resource in question. For instance, an author might argue that it's possible for an agent to consciously apprehend their own brain states when these are visualised on a screen by means of a so-called neurofeedback device. Further, the thinking goes, as there is conscious apprehension of the resource, there therefore cannot be any transparency. Thus, neurofeedback devices show that even in the case of our internal cognitive machinery, transparency may be absent.

We think that the move from conscious apprehension of a resource to a lack of transparent employment of the resource is too quick. We can see this once we notice that the sort of conscious apprehension often discussed in the literature isn't at all about how the resource in question is *employed*. While it's true that, say, a neurofeedback device – a device that can, for instance, visualise our brain states on a monitor – allows us to consciously apprehend our own brain states, this doesn't impact the agent's employment of her own brain. If the brain had been employed transparently before, then there's no reason to think that simply looking at one's own brain states causes transparent employment to go away. What is crucial for transparency isn't whether an agent consciously apprehends a resource but rather whether the agent's employment of the resource relies on this conscious apprehension. It's possible to consciously apprehend a resource *while* transparently employing it.

In the following pages, we first (in Sect. 2) turn our attention to introducing transparency and how it relates to cognitive extension. In Sect. 3, we discuss arguments

formulated against the claim that transparency is necessary for cognitive extension and show how they all depend on the claim that conscious apprehension is mutually exclusive with transparency. Thereafter, in Sect. 4, we spell out our own understanding of transparency. Section 5 shows how this understanding means that the counter-arguments fail. Finally, Sect. 6 highlights some our account's more general implications for the literature on extended cognition and concludes the paper.

2 Transparency and cognitive extension

In the present section, we look at the original formulation of the concept of phenomenal transparency and how it was later adapted and employed in the literature on extended cognition.

Heidegger (1990; see also Clark, 2004) points out that when an agent engages in a skilled activity – such as when a master carpenter employs a hammer – she will no longer focus her attention on the tool but rather on the task she is trying to achieve with it. A skilled carpenter isn't focusing on the hammer, but rather on sinking the nails into the wood and making the table.

Merleau-Ponty (2002) gave us the second core example of the literature: a blind person navigating the world with her white cane. Her focus, he says, isn't on the tactile vibrations that the cane elicits in her hand but rather on the world at the tip of it. She will, for instance, encounter the pavement as her cane bumps into its edge.

If the *resource* isn't at the focus of attention during its transparent employment – what is? In short, when using a resource transparently, we focus on the task for which it is employed. This idea has often been described by saying that when a resource is used transparently, we see *through* it. A blind person navigating the world with a cane doesn't focus on the vibrations transmitted to her hand by the cane, but rather on the world they perceive through it. A skilled carpenter employing a hammer is focusing on joining pieces of wood rather than on how she needs to employ the hammer to do so. It is in this sense that the resource becomes transparent.

Note that we focus on the transparent *employment* or *use* of a resource while some others, such as Clark (2008), primarily discuss transparent *resources* or *equipment*. We think that the debate concerns a relational property of the agent and the object, namely whether the object is transparent to the agent (and thus, it is never about whether the object is, as such, transparent). Moreover, we are interested in how transparency affects cognitive extension, and cognitive extension necessarily involves resources *used by* agents. Thus, at least for our purposes, the interesting phenomenon concerns an agent employing a resource such that the resource becomes transparent to the agent (in the manner discussed previously). This is what we mean when we say "transparent employment" or "employing transparently".

Some of extended cognition's early proponents think that employing a resource transparently is necessary for cognitive extension. Clark (2004, 2008) juxtaposes transparent and opaque technologies and shows how the transparent use of resources underlies the kinds of cases in which he thinks cognition extends. Similarly, Wheeler

(2019) says that when an external resource is transparent, "a necessary condition [is] met for its constitutive incorporation into the user's mental machinery" (p. 862). Finally, Thompson and Stapleton (2009) propose a *transparency constraint* according to which an external resource must be employed transparently if it is to become part of an agent's cognitive system.

It's the link between transparency and the constitution of the subject-world interface that drove authors to claim that transparency is necessary for cognitive extension (Clark, 2004; Thompson & Stapleton, 2009; Wheeler, 2019). The skilled carpenter's attention is focused on the table; that's the object of which she is conscious. Likewise, the blind person perceives the world at the tip of her cane; this is the world from which she conceives herself to be distinct. In contrast, the hammer and the cane are part of the subject – or "body-as-subject" as Thompson and Stapleton (2009) write – which allows them to perceive and act in these ways. Transparency is therefore essential in establishing the boundary between that which allows an agent to perceive and act on the world – the cognitive machinery – and that which is located outside the boundary, to be acted on and perceived.

Consider tactile-visual sensory substitution (TVSS) devices (see, for instance, Clark, 2004; Kiverstein & Farina, 2011; Palermos, 2014). Such devices allow blind people to perceive the locations of objects in space in a way akin to how sighted people perceive the world using their visual systems. A device is attached to some part of the body (for instance the forehead) and hooked up to a video camera. An array of (electro- or vibrotactile) activators in the device are triggered based on information captured by the camera. The resulting tactile stimulus can then be used by the agent to gain spatial information about their environments. For instance, a tactile sensation at the centre of the array might indicate an object located in front of the agent.

When an agent is new to such a device, she needs to focus on the tactile sensations and use these to infer information about, say, the distance to a chair she is facing. In this case, the agent is using the resource opaquely; she is focusing on the tactile stimuli and is using her knowledge of the device's functionality (in particular, her knowledge of how certain stimuli correlate with object location) to achieve her cognitive task (that is, inferring the location of the chair).

With time, agents begin to use TVSS devices effortlessly and automatically. They no longer focus on their tactile experiences to infer the spatial organisation of their surroundings. When they enter a novel room, they no longer attend to the stimuli on their foreheads to infer that there's a chair blocking their path; rather, they simply *experience* the chair being located in their path. When they wonder about the locations of objects in their surroundings, their focus is on the content of their "visual" experience rather than the tactile sensations. In fact, they no longer consciously experience the tactile sensations or the fact that they're employing a TVSS device – they are using the device transparently and perceive their surroundings "through" it.

A recent account by Grush and Springle (2019) can clarify just how transparency helps constitute the subject-world interface. Let's re-examine what happens when the user of a TVSS device encounters a chair: Light is reflected off the chair

and travels to the lens of the TVSS device, where it gets converted into digital signals. These are processed and sent to actuators that press onto the skin, where the pressure is detected by touch receptors. The resulting signals are transduced, sent along nerves, and finally processed in the brain. This is the causal chain with which the agent is connected to the world and which enables her to gain information about it.

Grush and Springle focus on explaining why agents perceive objects that are more (the chair) or less (the TVSS device) far along the causal chain. They argue that perception depends on internal models, which map from internal causes (say, certain neural signals) to external causes.[5] When an agent lacks experience with a type of causal chain, her internal model may map neural signals only to causes quite close to her. Hence, a novice's internal model might map a given neural signal to the fact that the TVSS device is pressing against her skin in a particular way. Such a user will then, in a second step, need to infer what worldly cause (the chair) might be behind those tactile sensations. Here, the agent attributes properties to an object (the TVSS device) and then uses her knowledge of those properties to make further inferences. In contrast, an experienced user of a TVSS device may map neural signals directly to causes further out in the environment. A given neural signal may let her know that she has encountered a chair. She doesn't need to consciously *infer* that she's encountered the chair.

Sophisticated internal models map directly to more distant causes, skipping intermediary causal links. Skipped causal links are transparent to the agent. The agent *perceives* the object and properties to which the internal model maps. The further out in a causal chain the object is, the further out an agent will experience an interface between herself and the objective world. Because objects at intermediary steps are transparent, they aren't consciously apprehended as objects and are rather experienced – if they are experienced at all – as a part of ourselves. A novice user of a TVSS device is conscious of the device whereas the expert user is conscious of the chair in front of her.

3 Arguments against transparency in extension

Recently, a number of authors (Andrada, 2020, 2021; Facchin, 2022; Smart et al., 2022) have argued against the view that transparency is necessary for cognitive extension. Their arguments generally fall into one of the following categories: (1) claims that internal processes may be fail to be transparent and that, following the parity principle, we can't demand transparency for external processes either, (2) claims that there are paradigm cases of extension in which transparency is absent, and (3) claims that the incorporation of skill (as demanded by second wave views of cognitive extension) doesn't depend on transparency. All of these claims, we

[5] Grush and Springle (2019) discuss perception *and* action. To fully understand the true force of their argument, it is necessary to consider them jointly. However, for the point we are making in this paper, it suffices to look at perception alone.

argue, repose on the unwarranted assumption that conscious apprehension of some resource is mutually exclusive with the transparent employment of the resource.[6]

The parity principle is the starting point of some of the most forceful arguments against the claim that transparency is necessary for cognitive extension. This principle, first introduced by Clark and Chalmers (1998), captures the idea that if some body-internal process is cognitive because it possesses, certain properties, then an extended process exhibiting those very same properties should also be considered cognitive. From this principle it then follows that when a given functional property *isn't* required for an internal process to count as cognitive, then it cannot be required of an external process either.

Here's an example of such an argument taken from Facchin (2022), who is inspired by earlier discsussions by Clark (2015; see also Bagherzadeh et al., 2020). As mentioned, neurofeedback devices allow users to observe their own evolving brain states (for instance on a monitor). Facchin submits that this is a case in which our own body-internal (and clearly cognitive) processes are opaque to us – after all, we see them on a monitor right in front of us. If there are such opaque (internal) cognitive processes, then opaque (extended) cognitive processes are also possible. And thus, transparency cannot be necessary for cognitive extension.

Similarly, Andrada (2021) asks us to imagine an agent, careful Inga, who engages in various activities to improve her brain's memory faculty (for instance, "taking pills of Ginkgo Biloba, practic[ing] mnemonic techniques" [p. 4699]). By doing so, Inga is taking this faculty as an object of her cognition, as something that is opaque to her. But doing that doesn't change the functional profile of said faculty – she still uses it in the same sort of way to remember things – and therefore the faculty remains a memory faculty. But that faculty is clearly cognitive, thus showing that transparency isn't required for an internal process to count as cognitive. Thus, if another agent – careful Otto – employs their notebook in a careful way that involves attending to the notebook, this conscious apprehension of the notebook is compatible with Otto extending to the notebook.

The second kind of argument concerns cases in which some resource is (a) clearly a part of an agents cognitive system, but is (b) intransparent and (c) body-external.

[6] As a reviewer pointed out, Smart, Andrada, and Clowes (2022) argue that non-human animals present another challenge to the claim that transparency is necessary for cognitive extension. Our arguments in this paper do not address this challenge, but we think there are independent reasons to think the challenge fails. First, we do not think that there are reasons, as such, to think it's "difficult to see how we might apply the notion of transparency to something like a spider – an entity whose phenomenology (if it has any) is inaccessible to us" (p. 334). We agree that insisting on the importance of phenomenology might make it difficult to determine whether non-human animals have extended cognitive processes, but this epistemic difficulty alone cannot support claims regarding the determinants of our cognitive boundaries. Moreover, just because our claims concern phenomenology, it doesn't follow that we lack any means to study non-human animals. For instance, Grush and Springle's (2019) account (discussed in Sect. 2) seems well-poised to investigate transparency in non-human animals. Their information processing account argues that transparency has to do with the predictive powers of internal models and as such can be applied to human beings and non-human animals alike. Much more needs to be said – especially with regard to non-human animals' ability to encounter objects – but this will have to wait for future discussions.

Facchin (2022) discusses a number of cases that we think are of this sort. First, Facchin (2022) mentions a case originally developed by Dennett (1981) in which an agent's brain is transplanted into a vat and hooked up to the rest of the body with a wireless communication system. We are then asked to imagine the agent looking at their own brain. Clearly, that brain still powers the agent's cognition (a), but just as clearly, the brain isn't within the agent's body (c) and the agent is attending to it (b).[7]

Facchin also discusses a host of other – what he calls *paradigmatic* – examples of extended cognition in which transparency is absent. For instance, in Tetris, players use extended cognitive processes to rotate the falling zoids to figure out how they can best slot them in at the bottom (Facchin, 2022; see also Kirsh & Maglio, 1994). That is, rather than figuring out how the zoids look when oriented in certain ways using their brain-internal processes, they rotate the zoids on the screen. At first glance, given the fact that players may be unaware of rotating the zoids to this end, it could seem like this is a case of transparent employment. But not so, claims Facchin: players still attend to the zoids, the zoids and their properties figure as objects in their perception, and are thus opaque to them. Thus, here we have a paradigmatic case of cognitive extension – but no transparency. The conclusion is obvious: transparency cannot be necessary for cognitive extension.

Some authors – following the general trend of second wave extended cognition – question the very idea that parity with internal processes is somehow essential to cognitive extension. This isn't the place to get into this debate, but it's worth noting here that, as some (Andrada, 2021; Smart et al., 2022) have argued, the case for parity might be more challenging to make for transparency than for functional properties. As originally conceived, the parity principle applied to functional properties as it is based on the functionalist conception of cognitive and mental properties as functional kinds. If the relevant properties are individuated by functional characteristics, then there is, at least *prima facie*, some merit to the idea that a given functional kind should be called cognitive (or mental) no matter whether it is body-internal or extended. However, as these authors point out, it's far from clear that there's a dependable relation between phenomenal properties and functional properties, and thus, *even if functional parity is a goal*, it seems that this doesn't straightforwardly engender any requirements for phenomenal parity. We therefore needn't assume that the phenomenology of extended cognition needs to "mimic our experience when" (Andrada, 2020, p. 8) relying on internal processes.

In line with her more general misgivings with first wave extended cognition thinking, Andrada (2021) also presents a positive argument to show why complementarity-based accounts of extended cognition do not entail a requirement of

[7] Facchin seems to think that this argument is one in which is shown that internal processes can be transparent (thus of the parity kind disscussed above), but we think this classification is false. It's quite clear that the brain isn't in the body anymore (it's in the vat). It's part of what powers the cognitive system (and thus, might be considered "internal" to the cognitive system), but as such it derives it's power from being a case of extended cognition in which we're sure, for whatever reasons, that the resource in question is part of the agent's cognitive machinery.

transparency. Second-wave views focus on how and when agents incorporate skills and thus, what matters for extension isn't parity but whether the conditions for the incorporation of an extended skill are fulfilled. To show that the incorporation of extended skills doesn't depend on transparency, Andrada shows how there are a variety of skills in which agents attend to the relevant resource. Even expert athletes require a "self-awareness [that] involves *focally attending* to cues provided by movements in order to make precise adjustments during performance" (2021, p. 4704, *emphasis original*). If expert skill doesn't come with transparency, then, if the incorporation of skills is what drives cognitive extension, transparency cannot be necessary for it.

While the above arguments cover a wide gamut of ideas, they all repose on the same idea: that the conscious apprehension of a resource entails that the resource in question isn't employed transparently. Smart, Andrada, and Clowes (2022) write that for transparency to obtain, the resource "should disappear from the conscious apprehension of the user, such that the user is no longer aware of the resource as an independent object". This implies that when such conscious apprehension doesn't happen, the agent fails to transparently employ the resource. Similarly Facchin (2022) defines transparency such that "[a resource] is phenomenally transparent to [a subject] only if, when [the subject] uses [the resource], [the resource] is not an intentional object [the subject] is thematically directed at". And Andrada says that her view "make[s] room for some forms of conscious thought and attention within extended cognitive routines", implying that such conscious thought is sufficient to end transparency.[8]

The idea that conscious apprehension of an object precludes its transparent employment is crucial to the above arguments. All of the arguments discussed are based on the identification of cases (of internal cognitive processes, of external cognitive processes, and of skills) in which the agent consciously apprehends the resource in question. Neurofeedback devices imply it's possible to consciously apprehend one's internal cognitive processes, Tetris shows that even paradigm cases of extended cognition may come with conscious apprehension (of the zoids), and professional athletes' skills show that skills, likewise, can come with the conscious apprehension of the relevant resources. This, these authors submit, is sufficient to show that there is no transparency in these cases and that, therefore, transparency cannot be required for cognitive extension. However, if there is, as we'll show in the next section, still a real sense in which transparency obtains in these cases *despite the conscious apprehension of the resource*, these arguments fail to go through.

[8] This view of transparency is also shared by authors who do not explicitly argue against the view that transparency is necessary for cognitive extension. Farina and Lavazza (2022), for instance, write, "[t]ransparency – a property acquired by incorporated objects – indicates that the object has been 'appropriated', which means that it has faded away from the user's awareness." They then use this understanding of transparency to argue that the concept cannot be used to adjudicate between extended and embedded views of cognition.

4 Transparency and cognitive extension, revisited

As we've shown, existing arguments against transparency's necessity for cognitive extension repose on the idea that an agent who consciously encounters some resource cannot transparently employ it. The present section discusses some of the details of transparency to show why this – seemingly innocuous – assumption is mistaken. Transparency and opacity are revealed not to be mutually exclusive once we understand that transparency – at least as it is relevant for cognitive extension – characterises an agent's *employment* of a resource. Such employment may be transparent while the resource in question also figures as an object in consciousness (viz. Wheeler, 2019).

To show how transparency is about the employment of a resource (and not just any encounter of it as an object), it is helpful to consider how we employ external resources to realise certain contents. Note, first, that extended cognition means that the vehicles that realise some given cognitive or mental content are extended. Thus, rather than, say, Otto's brain realising the belief that MoMA is at a certain location, that very belief might be realised by an extended system involving Otto's brain, body, and notebook. Cognitive extension (at least as we're considering it in this paper) is therefore about the extension of contentful states. What we want to highlight with two examples below is that there are two distinct ways in which we employ (body-external) resources to bring about cognitive or mental contents.

Consider a biologist who is highly skilled at using her microscope. Her use of it is automatic, effortlessly, and seamless. When she wants to look at a bacterium, she reaches for her microscope, clips in the slide, looks through the lens, adjusts the knobs to bring the slide into focus, and looks at the specimen. In all this, she doesn't need to think about the functionality of the microscope, how to fasten the slide, how to turn the knobs, how to move the slide to centre the specimen, and so forth. For instance, when focusing, she may turn the knob a bit, get visual feedback through the lens, and use this to fine-tune the focus without this back-and-forth being consciously apprehended at all. The vehicle properties of the microscope aren't consciously apprehended by the expert biologist: she intends to identify the specimen, and she doesn't need to think about any of the steps required to achieve that goal.

Opaque use of a microscope differs. A novice biologist may have the same goal as the expert – to look at a specimen – but her path to achieving that goal is different. She needs to not only form the intention to look at the specimen, but needs to think about the intermediate steps required to get her there. She must think about how to use the various bits of the microscope so that the specimen ends up in focus and at the centre of the field of vision. Unlike in the transparent case, the agent needs to attend to, and consciously manipulate, the physical properties of the microscope to realise the contents that interest her. For instance, when focusing on the slide, the novice needs to think about which knobs to manipulate and how to manipulate them, and she needs to pay attention to the changes in sharpness of the image to gauge whether she has overshot the point of focus or not. This illustrates how in opaque employment, an agent cannot achieve her end goal without first achieving some intermediary task that requires attending to the physical properties of the resource.

Note that the above example is meant to illustrate a difference between the transparent and opaque employment of resources. To do this, it doesn't need to be a case of extended cognition (because, even if our argument succeeds, transparency is at most a necessary – but not sufficient – condition for extended cognition). However, given that this paper is about transparency as it relates to extended cognition, we'll now turn to a case that illustrates the difference between opaque and transparent employment as it applies to a (hopefully) clear-cut case of extended cognition.

Let's return to the TVSS devices introduced earlier. As we've said, an expert user of such a device simply perceives their spatial environment. She tokens certain cognitive states (about, say, the locations of objects) with the help of her device without having to consciously manipulate any of the device's physical properties. The novice user, in contrast, who is still trying to figure out how to employ the device, needs to think about a variety of the device's properties. She has to consider the tactile sensations on her skin, think about the device's functionality, and *infer* what is located in front of her. Her employment is opaque insofar as she must consciously apprehend the resource's vehicle properties in order to manipulate it so that she can realise the desired contents.

We now understand how the transparent employment of a resource (to cognitive ends) is linked to a lack of conscious manipulation (and, by implication, conscious apprehension) of the resource's vehicle properties. Transparency is, in short, about a specific kind of *employment* of a resource. With this understanding, we are now in a position to turn to the core of our argument and show how transparency and conscious apprehension aren't mutually exclusive.

Imagine a cognitive scientist who is employing a TVSS device to navigate her life. She has done so for a long time, does so automatically and fluently, and would describe her experience of using the device as enabling her simply to "see" things. She wants to study how cognition works when employing such a device, but lacking suitable test subjects, she has hooked her own TVSS device to a analysis software. She is now studying (with the help of her TVSS device) this software's output on a monitor in front of her. Is she transparently or opaquely employing her TVSS device?

In a sense, her TVSS device is opaque to her. After all, she sees its states evolve on a screen in front of her, she is attending to these, and they are very much visible and present at the centre of her conscious experience. Note, however, how her *employment* of her TVSS device seems to be entirely independent of what she's seeing on the monitor. Nothing of relevance changed in how she employs her own TVSS device when she sat down and started studying the monitor. If she had been employing the TVSS device transparently before she started observing it, then it seems she must be employing it transparently now, too. The scientist therefore transparently employs her TVSS device while also (opaquely) studying it.

For those who aren't convinced that our scientist employs the TVSS device transparently, let's suppose that the monitor shows someone else's TVSS device. In this case, the cognitive scientist is clearly employing her own TVSS device transparently (after all, she doesn't at all consciously apprehend her own TVSS device). Moreover, there doesn't seem to be a relevant difference between how she employs her device here, and the previous case, in which the device was her own. In both cases,

her employment doesn't require any conscious apprehension of the device's physical properties. If you're *still* unconvinced, imagine our scientist looking at a TVSS device she *thinks* belongs to another person, but which is in fact her own. Phenomenologically speaking, nothing might distinguish this case from the one in which she is looking at another person's TVSS device. Clearly, when looking at someone else's TVSS device, the scientist is employing her own device transparently. And if it's possible that looking at one's own TVSS device and looking at another person's TVSS device gives rise to the same phenomenal states, then – given that transparency is a phenomenological notion – it is possible to look at one's own TVSS device without thereby losing transparency.

The kind of transparency relevant to cognitive extension is about the *employment* of a resource, and such employment may co-exist with the agent consciously apprehending the resource. As we've seen above, transparent employment means a resource is employed without the resource's physical properties needing to be consciously manipulated. Note how the above cognitive scientist is, even when studying her own device, still employing her TVSS device to "see through" it. She doesn't manipulate any of its vehicle properties to form the various beliefs she has about it. In fact, she gains those beliefs *because of* her effortless and automatic employment of the device. Merely being conscious of (or attending to) the relevant resource therefore doesn't spell the end for transparency.

5 Replying to others

Armed with this improved understanding of transparency – and in particular, the knowledge that transparently employing a resource is compatible with consciously apprehending it – we are now in a position to respond to the various arguments against transparency's necessity for cognitive extension. Addressing these arguments will also present us with an opportunity to work out some of the more complex aspects of our proposal.

Before we discuss those argument's failure, some important preliminary remarks are in order regarding what is entailed by the discovery of a case of opaque or transparent employment of a resource *if* we're right that transparency is necessary for cognitive extension. What is entailed depends on what theory of cognitive extension one subscribes to – something that is missed in the literature and has led to mistakes as well as allegations of *reductio* (see Smart et al., 2022).

Call an account of cognitive extension *dispositional* when it focuses on the extension of dispositional cognitive states. Clark and Chalmer's (1998) original account serves as a useful prototype. They focus on the extension of dispositional beliefs (such as Otto's belief regarding the location of the MoMA) and argue that such a belief may be (partially) realised in body-external objects (such as Otto's notebook). If we subscribe to such an account, what matters first and foremost isn't how some agent is *actually* employing a resource, but rather how they are disposed to do so. Otto has an extended belief not because he is right now employing the notebook in a certain way, but because he is so disposed. Were he interested in the location of the

MoMA, he would pick up his notebook, look up the address, and generally engage in behaviours characteristic of the possession of a dispositional belief.

Here, transparency is necessary for cognitive extension in this way: it's necessary in the sense that the user must be *disposed* to use the device transparently. If the kind of disposition at play here is a probabilistic disposition,[9] exemplification of the disposition doesn't entail that the agent never uses the device opaquely. Just as someone may be disposed to cycle to work even if they take the bus on isolated occasions, it may be possible to be disposed to use a device transparently without always doing so. Thus, on dispositional accounts, if transparency is necessary for extension, it is a modal condition: what matters isn't whether the agent actually employs the resource transparently, but whether she is disposed to do so.

There is another, *dynamical*, way of understanding cognitive extension. On this view, agents extend when they employ resources in a certain way and then shrink down when they stop employing them in this way. Wheeler (2019; see also Clark, 2007), for instance, seems to subscribe to such a view according to which agents dynamically expand and contract. According to this view of extension, if an agent stops transparently employing a resource, they shrink and employ the process without extending into it.

Both these accounts provide space to argue that transparency is necessary for cognitive extension. It's not true that "if one subscribes to the view that transparency is necessary for cognitive extension, then a shift in transparency will reconfigure the borders of the mental/cognitive machine" (Smart et al., 2022). Such a reconfiguration is possible, but not necessary. In what follows, we remain neutral regarding the choice between these two views.

With that out of the way, let's now look at parity-style arguments against transparency's necessity for cognitive extension. As you may recall, Andrada (2021) and Facchin (2022) argue that there are instances in which internal cognitive resources are employed opaquely (and that, consequently, transparency cannot be a condition on cognitive extension). However, when we look at the cases of Inga (who uses various strategies to keep her memory faculty sharp) and neurofeedback devices (which allow a user to observe their own brain states), we can see these aren't, in fact, instances in which agents opaquely *employ* a resource for cognition. In both cases, the agents employ their resources transparently, that is, their manipulation of the resource to achieve their cognitive ends doesn't require them to consciously manipulate any of the resource's vehicle properties. Inga, when remembering something, does so just as the rest of us, without thinking about how her brain works, the chemicals it requires, and so forth. Her employment of her memory faculty doesn't change just because she additionally relates to it in a way that makes it an object of her cognition. Similarly, neurofeedback devices give an agent a view on their cognitive system, but they do not as such change how the agent employs their cognitive system.

[9] We can contrast probabilistic dispositions with surefire dispositions (see Choi & Fara, 2018). An object exemplifies a probabilistic disposition if it *tends* to manifest the disposition in circumstances of manifestation. For instance, a mug is fragile if it's likely to break when dropped *even if it doesn't always break in such situations*. In contrast, an object exemplifies a surefire disposition if it always manifests the disposition in circumstances of manifestation. Arguably, for a person to be loyal, they may never betray others.

However, there exists another way in which an agent may opaquely relate to one of their internal cognitive resources. What if, for instance, an agent used the information from their neurofeedback device to manipulate their own brain states. As discussed in the literature, neurofeedback information can be used by agents to manipulate their brain's alpha waves (Bagherzadeh et al., 2020; see also Clark, 2015), and this in turn can have beneficial effects on agent's ability to control their attention. If our agent did *that*, she would be manipulating her brain's physical properties, the alpha waves, to bring about a change in her cognitive states. Here, we seemingly have an example in which a resource isn't just consciously apprehended, but where said apprehension is used *to employ the resource*.

This presents a challenge to our argument, as we can no longer simply claim that opacity and transparency are not mutually exclusive because, sometimes, conscious apprehension of a resource isn't about manipulating the resource's vehicle properties. It would seem that in this case, we have to agree that given the resource's opaque *employment*, we do not extend to it. But once we agree to this, we are then led to the absurd conclusion that the brain is no longer a part of the agent's cognitive system.

To understand how to address this challenge, we need to take seriously the idea that extended cognition is about extended *processes* (Clark, 2022; Menary, 2012; Pritchard, 2018). Once we do that, we can see that a given resource may help sustain a variety of an agent's cognitive processes and that, as we'll show now, these can – even at a single moment in time – be of the transparent and opaque kind.

In the case of the alpha wave-manipulating agent, there is one process (or, more likely, a multitude of processes) that undergird typical cognition. These are processes in which the agent is employing their brain transparently, and these processes have been taking place before – and will continue after – she started manipulating her alpha waves. These processes don't depend on the agent consciously apprehending and manipulating the brain's vehicle properties. However, once the agent starts manipulating their alpha waves, a new process arises, and that process depends on the conscious manipulation of vehicle properties.

Arguments based on cases of opaque extended cognitive processes can be answered analogously. For instance, the Dennettian case of looking at one's own brain in a vat is analogous to the above case in which a neurofeedback device is used to merely observe one's own brain. The only difference is that when an agent uses a neurofeedback device, their brain is still in the body whereas in Dennett's case, the brain quite clearly isn't. Thus, the agent who is looking at their own brain in a vat transparently employs their brain for various cognitive ends just as they did before their brain got removed from their body.

The case of Tetris is more complex. Here, the player pays attention to the zoids in order to better place them, which means that the agent employs the zoids opaquely to succeed at their task. The zoids don't just figure as objects in consciousness, but the player is using her knowledge of the zoids to play the game. In this, the case is akin to the previously discussed case where an agent manipulates her alpha waves.

And, similar to that case, we suggest that a resource (a computing device running Tetris) is here supporting multiple processes, only some of which are characterised by transparency. To place the zoids efficiently, players need to know how the zoids look in various orientations so that they can find complementary gaps at the bottom

of the screen. Players employ the computing device to do two things: first, they must consider how zoids fit at the bottom of the screen given various orientations, and second, they need to slot the zoids into place. Kirsh and Maglio call the first an epistemic and the second a pragmatic action (Kirsh & Maglio, 1994).

We submit that to achieve her first task, a player transparently employs the device to rotate the zoids, and this doesn't require consciously apprehending the zoids and their properties. Of course, the player will consciously apprehend the rotated zoids, but, we believe, her rotation of these zoids doesn't depend on that apprehension. Note, for instance, that a player simply presses the appropriate button to rotate a zoid 90 degrees to the left – she needn't consider the shape of the zoid at all. Moreover, we think that the players' unawareness regarding the goals of their epistemic actions provides further evidence that no conscious manipulation of physical properties is required to rotate the zoids. The second task – slotting the zoids into place – does seem to require consciously attending to the zoids. Thus, the confusion surrounding this case is due to the fact that a single external resource is involved in multiple processes – and is only employed transparently in some of them.[10]

Similar considerations apply to arguments based on the idea that incorporated skills need not come with transparency. Consider an athlete, say, an expert cricket player, attempting to hit a ball with her bat. She may very well be attending to the ball, and Andrada takes this as evidence that attention can be a part even of expert skill (and therefore can be a part of extended cognitive skills). However, note how such an expert attends to different properties than a novice. A novice might attend to how the bat is oriented in their hands, their stance, how they have to move their arms to move the bat, and a whole host of other things. The expert, in contrast, might attend to the opponent and try to estimate, given her knowledge of the opponent's idiosyncrasies, what move would be appropriate. She might attend to the ball looking to identify the kind of technique employed by the opponent and then react so that she maximises her chances of outplaying the opponent. However, the expert can hit the ball in just the right way without needing to think about how to hold the bat or move her arms.

We see the kinds of skills looked at by Andrada and others as involving a variety of processes, and expert skills are likely such that many of these processes involve the transparent employment of resources. In these processes, the agent does not think about the resource in question and can rely entirely on automatic and seamless manipulation. Other processes involved in the exercising of the skill might still require conscious apprehension – and, thus, so does the overall exercise of the skill. However – and this is the crucial point – as long as some processes are transparent in the fashion described above, it's possible (but not necessary) that these processes are extended cognitive processes.[11]

[10] Clark (2022) makes a similar point when he discusses how players shuffle their tiles in Scrabble to increase their chances of winning the game.

[11] Note that we do not hereby endorse the view – critised by Andrada (2021) and Montero (2010) – according to which expert skill must be transparent. Quite the opposite, we think our account makes space for expert skills that involve some attentional processing. Moreover, our claims do not concern the nature of expert skill – we only claim that any skill, whether expert or not, is transparent (and, hence, potentially extended) to the degree that it involves processes that do not depend on the conscious manipulation of vehicle properties.

Overall, our treatment of cases such as the alpha wave-manipulating agent, the Tetris player, and the expert cricketer all repose on the idea that we must distinguish between the various processes involved in a cognitive task. Those processes that do not depend on the conscious manipulation of vehicle properties (and are, hence, transparent), are in the running to be extended cognitive processes. Those processes that do involve the conscious manipulation of vehicle properties aren't transparent and cannot extend.

Existing arguments against transparency's necessity for cognitive extension fail when applied to our account because they all depend – as we have shown above – on cases where only *some* of the involved processes are non-transparent. To counter our argument, our opponents would need to identify clear cases of cognitive extension where all of the relevant processes rely on the conscious manipulation of vehicle properties. We think our opponents will be hard-pressed to identify such cases. Think, for instance, of the Tetris player rotating zoids: if this involved the player thinking about the zoids, realising that she could figure out how they look by rotating them on screen, and consciously manipulating the game to rotate them, we would no longer be driven to consider this a paradigm case of extension.

To conclude, the various arguments against transparency's necessity for cognitive extension fail because they repose on the idea that opacity and transparency are mutually exclusive. As we've shown, there are at least two ways in which this isn't the case: First, sometimes a resource is consciously apprehended without the apprehended properties being employed to manipulate the resource for cognitive ends. Here, transparent employment takes place in just the same way as it does when an agent doesn't consciously apprehend the resource. Second, even in cases where the apprehended properties are used to manipulate the resource to cognitive ends, this doesn't mean that the resource isn't also transparently employed. It's possible that various cognitive processes run through a given resource, only some of which may involve transparent employment. Hence, it's far from "incoherent" (Facchin, 2022) to talk about an agent both transparently and opaquely relating to some resource – in fact, understanding how transparency and opacity aren't mutually exclusive is crucial to giving a proper account of (extended) cognition.

6 Closing remarks

In this paper, we have focused on showing why various arguments against transparency's necessity for cognitive extension fail. We concluded that argument in the previous section. However, before we end this paper, we want to outline how our account might impact thinking about cognitive extension more generally.

First, and as already mentioned, our claims imply that we can only properly make sense of transparency and its importance to cognitive extension when we subscribe to a process-based account of cognitive extension (see Menary, 2012). For such an account, the question of whether certain *objects* are a part of the cognitive machinery isn't primary – rather, we start by investigating whether certain *processes* reaching through those objects can be properly called cognitive.

Second, when we subscribe to such a process-based view, transparency becomes, in a certain sense, a graded property. A variety of processes might run

through a given resource, some of which have the agent consciously encountering it as an object and others not. The higher the ratio of transparent to opaque processes, the more we might want to say that the agent employs the resource transparently. This view is in tension with some authors' proposals, such as Facchin's (2022), which posit that transparency is an all-or-nothing affair, but it lends supports and theoretical sustenance to accounts (Farina & Lavazza, 2022; Smart et al., 2022) that see transparency as graded.

Such a graded approach to transparency can then lead to a graded account of cognitive extension. As a given resource may be involved in a variety of processes, only some of which might be cognitive, it's no longer guaranteed that we can give a yes-or-no answer to the question regarding some resource's inclusion in the cognitive machinery (or, maybe more precisely, it no longer follows from a resource's inclusion in the cognitive machinery that *all* processes running through it are cognitive). Resources can be more or less strongly involved in the cognitive machinery.

Arguably, our notion of transparency is, in certain ways, quite unlike many discussed in the literature. After all, we are saying that it's possible to consciously apprehend a resource while transparently employing it, and, maybe even more controversially, that it's possible to transparently employ a resource while opaquely employing it. Have we stretched the concept of transparency too thin and turned it into something else entirely? We do not think so. When we say that transparency is necessary for cognitive extension, we are saying that a specific kind of phenomenological criterion needs to be fulfilled for extension: it must be the case that the extended cognitive process doesn't rely on the conscious manipulation of the resource's vehicle properties. In other words, the (extended) process must take place without the agent consciously manipulating (and thus, *a fortiori*, consciously apprehending) the resource's physical properties. Note that this isn't the claim that an agent must automatically or seamlessly interact with a resource to extend her cognition to it – rather, it's the claim that a certain *phenomenology* is required. Because the criterion is phenomenological, we believe that it constitutes a demand for transparency.

It has not been our goal in this paper to provide an argument for transparency's necessary for cognitive extension – we have only aimed at invalidating certain counterarguments. What we do want to note here is that existing arguments for transparency's necessity for cognitive extension are *compatible* with our account of transparency. A lack of transparency institutes a subject-world boundary so that the opaque object is cognised by the agent rather than a part of the machinery with which the agent does the cognising. The only complication we introduce is that different cognitive processes can set the subject-world boundary at different places and that, just because *some* process sets the subject-world boundary so that a resource comes to be outside of it, we cannot conclude that *none* of the agent's cognitive processes extend to the resource.

Our view of transparency and cognitive extension does imply that we agree with our opponents' claims to some extent: we both think that the conscious apprehension of a resource doesn't preclude the possibility of cognitive extension. However, we think that when our opponents consequently abandon transparency, they are throwing the baby out with the bathwater. As we've shown throughout the paper, transparency still has a role to play in setting the subject-world boundary. For each process

involved in some cognitive task, for this process to be extended, it must be the case that it doesn't rely on the conscious manipulation of vehicle properties. This implies, as we have said above, that when looking at the overall cognitive task (comprising various processes), we arrive at graded notions of transparency and extension. However, that doesn't in any way invalidate the claim that transparency is necessary – if none of the processes involved are transparent, then there cannot be any cognitive extension.

Things have gotten complex – but such is cognition. What is important for our purposes is that no matter the complexity, one constant remains: for a cognitive process to extend to a resource, the agent must employ the resource transparently. It doesn't follow that this resource cannot be an object of cognition (as when seeing one's TVSS device in the mirror) or that there cannot be other processes which employ the object opaquely (as when playing Tetris). Transparency is necessary – no less and no more.

Acknowledgements The authors want to thank the reviewers for their insightful comments.

Authors' contributions Both authors (JH and HN) contributed to the conception, design, and writing of the manuscript, and they both read and approved the final draft.

Funding Open Access funding enabled and organized by Projekt DEAL. Julian Hauser's work is funded by the Swiss National Research Foundation (project number: P500PH_202829/1). Hadeel Naeem's research is a part of her fellowship at the Käte Hamburger Kolleg: Cultures of Research, funded by the Germany Ministry of Education and Research (project number: BMBF 01UK2104).

Data availability Not applicable.

Declarations

Ethics approval and consent to participate Not applicable.

Consent for publication Not applicable.

Competing interests The authors have no relevant financial or non-financial interests to disclose.

References

Andrada, G. (2020). Transparency and the phenomenology of extended cognition. *Límite: Interdisciplinary Journal of Philosophy & Psychology, 15*(20), 1–17.

Andrada, G. (2021). Mind the notebook. *Synthese, 198*, 4689–4708. https://doi.org/10.1007/s11229-019-02365-9

Andrada, G., Clowes, R. W., & Smart, P. R. (2022). Varieties of transparency: Exploring agency within AI systems. *AI & SOCIETY*. https://doi.org/10.1007/s00146-021-01326-6

Bagherzadeh, Y., Baldauf, D., Pantazis, D., & Desimone, R. (2020). Alpha synchrony and the neurofeedback control of spatial attention. *Neuron, 105*, 577-587.e5. https://doi.org/10.1016/j.neuron.2019.11.001

Choi, S., & Fara, M. (2018). Dispositions. In Edward N. Zalta (Ed.), *The stanford encyclopedia of philosophy* (Fall 2018). Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/fall2018/entries/dispositions/. Accessed 24 Mar 2024

Clark, A. (2004). *Natural-born cyborgs: Minds, technologies, and the future of human intelligence*. Oxford University Press.

Clark, A. (2007). Soft selves and ecological control. In *Distributed cognition and the will: Individual volition and social context* (pp. 101–122). MIT Press.

Clark, A. (2008). *Supersizing the mind: Embodiment, action, and cognitive extension*. Oxford University Press.

Clark, A. (2015). What 'extended me' knows. *Synthese, 192*, 3757–3775. https://doi.org/10.1007/s11229-015-0719-z

Clark, A. (2022). Extending the predictive mind. *Australasian Journal of Philosophy*, 1–12. https://doi.org/10.1080/00048402.2022.2122523

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

Dennett, D. (1981). *Brainstorms: Philosophical essays on mind and psychology*. Bradford Books. https://search.ebscohost.com/login.aspx?direct=true&scope=site&db=nlebk&db=nlabk&AN=48930. Accessed 24 Mar 2024

Facchin, M. (2022). Phenomenal transparency, cognitive extension, and predictive processing. *Phenomenology and the Cognitive Sciences*. https://doi.org/10.1007/s11097-022-09831-9

Farina, M., & Lavazza, A. (2022). Incorporation, transparency and cognitive extension: Why the distinction between embedded and extended might be more important to ethics than to metaphysics. *Philosophy & Technology, 35*,. https://doi.org/10.1007/s13347-022-00508-4

Grush, R., & Springle, A. (2019). Agency, perception, space and subjectivity. *Phenomenology and the Cognitive Sciences, 18*, 799–818. https://doi.org/10.1007/s11097-018-9582-y

Heersmink, R. (2013). Embodied tools, cognitive tools and brain-computer interfaces. *Neuroethics, 6*, 207–219. https://doi.org/10.1007/s12152-011-9136-2

Heersmink, R. (2015). Dimensions of integration in embedded and extended cognitive systems. *Phenomenology and the Cognitive Sciences, 13*(3), 577–598.

Heersmink, R. (2016). The cognitive integration of scientific instruments: Information, situated cognition, and scientific practice. *Phenomenology and the Cognitive Sciences, 15*, 517–537. https://doi.org/10.1007/s11097-015-9432-0

Heidegger, M. (1990). *Being and time*. Basil Blackwell.

Kirchhoff, M. D., & Kiverstein, J. (2019). *Extended consciousness and predictive processing a third wave view*. Routledge.

Kirsh, D., & Maglio, P. (1994). On distinguishing epistemic from pragmatic action. *Cognitive Science, 18*, 513–549.

Kiverstein, J., & Farina, M. (2011). Do sensory substitution extend the conscious mind? In Fabio Paglieri (Ed.), *Consciousness in interaction: The role of the natural and social context in shaping consciousness"*. John Benjamins. John Benjamins.

Menary, R. (2008). *Cognitive integration: Mind and cognition unbounded*. Palgrave Macmillan.

Menary, R. (2012). Cognitive practices and cognitive character. *Philosophical Explorations, 15*, 147–164. https://doi.org/10.1080/13869795.2012.677851

Merleau-Ponty, M. (2002). *Phenomenology of perception*. Routledge.

Montero, B. (2010). Does bodily awareness interfere with highly skilled movement? *Inquiry, 53*, 105–122. https://doi.org/10.1080/00201741003612138

Palermos, S. O. (2014). Loops, constitution and cognitive extension. *Cognitive Systems Research, 27*, 25–41.

Pritchard, D. (2018). Extended virtue epistemology. *Inquiry, 61*, 632–647. https://doi.org/10.1080/0020174x.2017.1355842

Smart, P. R., Andrada, G., & Clowes, R. W. (2022). Phenomenal transparency and the extended mind. *Synthese, 200*,. https://doi.org/10.1007/s11229-022-03824-6

Sutton, J. (2010). Exograms and interdisciplinarity: History, the extended mind, and the civilizing process. In R. Menary (Ed.), *The extended mind* (pp. 189–225). MIT Press.

Thompson, E., & Stapleton, M. (2009). Making sense of sense-making: Reflections on enactive and extended mind theories. *Topoi, 28*, 23–30. https://doi.org/10.1007/s11245-008-9043-2

Wheeler, M. (2005). *Reconstructing the cognitive world: The next step*. MIT Press. https://search.ebscohost.com/login.aspx?direct=true&scope=site&db=nlebk&db=nlabk&AN=138503. Accessed 24 Mar 2024

Wheeler, M. (2019). The reappearing tool: Transparency, smart technology, and the extended mind. *AI & SOCIETY, 34*, 857–866. https://doi.org/10.1007/s00146-018-0824-x