



The extended mind argument against phenomenal intentionality

Cody Turner^{1,2} 

Accepted: 3 August 2021 / Published online: 18 August 2021

This is a U.S. government work and not under copyright protection in the U.S.; foreign copyright protection may apply 2021

Abstract

This paper offers a novel argument against the phenomenal intentionality thesis (or PIT for short). The argument, which I'll call the extended mind argument against phenomenal intentionality, is centered around two claims: the first asserts that some source intentional states extend into the environment, while the second maintains that no conscious states extend into the environment. If these two claims are correct, then PIT is false, for PIT implies that the extension of source intentionality is predicated upon the extension of phenomenal consciousness. The argument is important because it undermines an increasingly prominent account of the nature of intentionality. PIT has entered the philosophical mainstream and is now a serious contender to naturalistic views of intentionality like the tracking theory and the functional role theory (Loar 1987, 2003; Searle 1990; Strawson 1994; Horgan and Tienson 2002; Pitt 2004; Farkas 2008; Kriegel 2013; Montague 2016; Bordini 2017; Forrest 2017; Mendelovici 2018). The extended mind argument against PIT challenges the popular sentiment that consciousness grounds intentionality.

Keywords Extended mind · Phenomenal intentionality · Extended consciousness · Propositional attitudes

This paper offers a novel argument against the phenomenal intentionality thesis. *The phenomenal intentionality thesis* (or “PIT”) claims that source intentionality is grounded in phenomenal consciousness, where by “source intentionality” I mean states that are the source of intentionality, being intrinsically intentional. The argument, which I call *the extended mind argument against phenomenal intentionality*, is centered around two claims: the first asserts that some source intentional mental states extend into the environment, while the

✉ Cody Turner
cody.turner@uconn.edu

¹ Philosophy Department, The University of Connecticut, Storrs, USA

² Willimantic, USA

second maintains that no conscious mental states extend into the environment. If these two claims are correct, then PIT is false, for PIT implies that the extension of source intentionality is predicated upon the extension of phenomenal consciousness. Put differently, I submit that the following three propositions, when properly understood, constitute an inconsistent triad: (1) the extended mind thesis is true, (2) the extended consciousness thesis is false, and (3) PIT is true.¹ To avoid a contradiction it must be the case that at least one of these propositions is false. The argument presented here motivates (1) and (2) in an effort to refute (3).

The argument is important because it undermines an increasingly prominent account of the nature of intentionality. PIT has entered the philosophical mainstream and is now a serious contender to naturalistic views of intentionality like the tracking theory and the functional role theory (Bordini, 2017; Farkas, 2008; Forrest, 2017; Horgan & Tienson, 2002; Kriegel, 2013; Loar, 1987, 2003; Mendelovici, 2018; Montague, 2016; Pitt, 2004; Searle, 1990; Strawson, 1994). The extended mind argument against PIT challenges the popular sentiment that intentionality is grounded in consciousness. Notably, this paper aims not to prove with certainty that the extended mind argument against PIT is sound but rather to show that each premise of the argument is highly plausible given certain philosophical assumptions about the nature of consciousness and extended cognition.

The general structure of the paper is as follows. Section 1 introduces PIT and briefly contrasts the view with naturalistic theories of intentionality. Section 2 presents the extended mind argument against PIT and describes Clark and Chalmers' (1998) extended mind thesis, clarifying that the thesis does not tacitly assume the falsity of PIT by presupposing the functional role theory of intentionality. Section 3 draws on an argument from Clark (2008), which I call *the self-stimulating loop argument*, to motivate the idea that the extended mind thesis is not restricted to dispositional states with derived intentionality but includes within its scope some occurrent propositional attitudes with source intentionality (i.e. some source intentional states are extended). Section 4 then draws on an argument from Chalmers (2018), which I call *the direct access argument*, to substantiate the claim that the extended consciousness thesis is false (i.e. no conscious states are extended). The conjunction of these two sub-arguments functions to undermine PIT by showing that some extended mental states possess source intentionality but lack phenomenal consciousness. Section 5 concludes.

¹ The extended consciousness thesis in this context denotes the nonmodal claim that some *existing* conscious states *are* extended, not the modal claim that conscious mental states *can be* extended. The extended mind thesis is also understood to be a nonmodal claim and is taken to encompass both dispositional states with derived intentionality and occurrent propositional attitudes with source intentionality.

1 The phenomenal intentionality thesis

The phenomenal intentionality thesis (PIT) is an increasingly popular view in the philosophy of mind about the nature of intentionality. Before unpacking the view, it will be instructive to distinguish between source intentionality and derived intentionality and introduce the two primary theories of intentionality that serve as rivals to PIT.

There are many different theories about the origin of intentionality. Nearly everyone agrees that, regardless of its origin, intentionality can be ‘passed around’ (Kriegel, 2013), so to speak, to other things which did not previously have it.² For example, both linguistic signs and physical road signs exhibit ‘aboutness’, or have a representational nature, but they do not do so intrinsically; rather, the aboutness of both types of signs derives from the aboutness of the intentional mental states of the humans responsible for fixing their respective meanings. Reflection upon the intentionality of signs suggests that one can draw a distinction between two basic kinds of intentionality: (a) *source intentionality* (or ‘original intentionality’ as it is sometimes called), and (b) *derived intentionality*. ‘Source intentionality’ refers to those things that are intrinsically intentional (and so that serve as the *source* of intentionality), whereas ‘derived intentionality’ refers to those things that have intentionality in virtue of some other thing that has intentionality. Providing an account of derived intentionality is an essential task for any complete theory of intentionality.³ However, the main goal of such a theory is to explain the nature of source intentionality.

During the twentieth century, the popular philosophical project concerning source intentionality was to ‘naturalize’ it by providing a reductionistic explanation of the phenomenon in terms of properties and processes that are fully comprehensible by natural science (Dretske, 1981, 1995; Fodor, 1987, 1990; Millikan, 1984; Neander, 1996; Papineau, 1984; Rupert, 1999; Stampe, 1977).⁴ The two most popular naturalistic theories of this sort are *tracking theories of intentionality* and *functional role theories of intentionality*. Tracking theories conceptualize source intentionality in terms of tracking relations, where ‘tracking’ is a matter of brain states “detecting, carrying information about, or otherwise corresponding with external items in the environment” (Mendelovici & Bourget, 2014, p. 326). Functional roles theories, by contrast, define source intentionality in terms of the functional roles that brain states play.

The *phenomenal intentionality thesis* (PIT) is a relatively new theory of intentionality that is increasingly regarded as a promising alternative to tracking theories and functional role theories. While the *idea* of phenomenal intentionality derives from the work of Brian Loar at the end of the 1980s (Loar, 1987), the *term* ‘phenomenal

² There are exceptions. Angela Mendelovici (2018), for instance, denies the existence of derived intentionality.

³ Some prominent accounts of derived intentionality include the views known as potentialism (Searle 1992), inferentialism (Loar 2003), interpretivism (Kriegel 2011), and self-ascriptivism (Mendelovici 2018).

⁴ The naturalistic project seeks to specify the relationship between intentionality and natural properties not only in terms of reduction, but also in terms of supervenience.

intentionality' did not officially enter the philosophical lexicon until the early 2000s (Horgan & Tienson, 2002; Loar, 2003). PIT can be understood as the conjunction of the following two claims:

- (1) There is a kind of intentionality, called 'phenomenal intentionality', grounded in a type of phenomenal character (i.e. a type of conscious experience).
- (2) Phenomenal intentionality is the only kind of source intentionality.

There are many reasons to endorse PIT, a significant one: the view arguably avoids *the problem of content determinacy* that afflicts naturalistic theories of intentionality. Source intentional states possess determinate content, which is to say, the intentional objects of such states are represented in a fine-grained, unambiguous manner. Kriegel defines the concept of 'determinate content' in the following manner:

"By 'determinate content', I simply mean content which is as fine-grained as one's intentional contents appear pre-theoretically to be. For example, pre-theoretically it seems that one's thoughts are fine-grained enough to be about rabbits rather than undetected rabbit parts, about Phosphorus rather than Hesperus, about triangles rather than closed trilateral figures, and so on. If a kind of intentional state is *not* this fine-grained, I say that its content is indeterminate" (Kriegel, 2013, p. 10).

Functional role theories and tracking theories of intentionality both appear unable to account for the fact that source intentional states bear determinate content. Bourget and Mendelovici (2014) formulate the problem in an epistemic fashion via Quine's (1960) famous 'rabbit/undetached rabbit parts' example. They point out that complete knowledge of all of the relevant tracking relations or functional role relations does not translate into knowledge of determinate content: "A Martian looking down on Earth and having complete knowledge of all Earthly physical facts could not tell whether we are representing rabbits or undetected rabbit parts. Thus, it appears that a physical-functional theory of intentionality will predict that one's concept RABBIT is indeterminate between the two contents" (Bourget & Mendelovici, 2014). The idea, to paraphrase, is that naturalistic theories of intentionality confront an epistemic problem of content determinacy that arguably has ontological implications. Natural properties are seemingly unable to secure determinate mental content if knowledge of all relevant naturalistic facts does not convert into knowledge of content determinacy. This is where PIT enters into the picture. Graham et al. (2007), Horgan and Graham (2012), and Searle (1990) all contend that phenomenal consciousness is the only thing capable of securing determinate mental content. This is known as the *content determinacy argument for PIT*, which is just one out of many existing arguments for the view.⁵

⁵ Horgan and Tienson (2002) present a phenomenological argument supporting PIT, according to which it is introspectively obvious that some conscious experiences are source intentional and source intentional in virtue of being phenomenal. Two additional arguments for the view come from the work of Charles Siewert and Brian Loar, respectively. Siewert (1998) avers that PIT is true in virtue of the fact that conscious experiences are assessable for accuracy, whereas Loar (2003) connects internalism about mental content with PIT, arguing that the latter theory is true on the basis of the former.

2 The extended mind argument against phenomenal intentionality

I will now present the extended mind argument against PIT. As is indicated by its name, the argument appeals to conceptual tools from the *extended mind literature* (Chalmers, 2018; Clark, 2008; Clark & Chalmers, 1998; Colombetti & Roberts, 2015; Farkas, 2012; Muller, 2012; Rowlands, 2009; Vold, 2015), and in particular, is motivated by a perceived asymmetry according to which some source intentional mental states extend into the environment (i.e. the extended mind thesis is true), but no conscious mental states extend into the environment (i.e. the extended consciousness thesis is false). Call this the *bipartite extension intuition (BEI)*. If BEI is true, then PIT is false, for BEI entails that there are cases of extended cognition that are source-intentional but non-phenomenal while PIT claims that source-intentional states are necessarily phenomenal. The argument has the following structure:

2.1 The extended mind argument against phenomenal intentionality

P1 Some source intentional states are extended (i.e. the extended mind thesis is true).

P2 No conscious states are extended (i.e. the extended consciousness thesis is false).

P3 P1 and P2 are mutually inconsistent if PIT is true.

C Therefore, PIT is false.

P3 is true by definition. If phenomenal intentionality is the only kind of source intentionality (i.e. PIT is true), then the extension of source intentional states implies the extension of phenomenally conscious states. That is, the combination of PIT and P1 entails the extended consciousness thesis. However, P2 rejects the extended consciousness thesis, meaning that P1 and P2 are mutually inconsistent, assuming the veracity of PIT. The proponent of PIT must, therefore, either uphold P1 and reject P2 (by affirming that some conscious states are extended), uphold P2 and reject P1 (by affirming that no source intentional states are extended), or reject both P1 and P2. The proponent cannot support both premises without violating the fundamental tenet of PIT because P1 and P2 jointly imply the existence of an extended source intentional state that lacks phenomenal consciousness, and therefore, phenomenal intentionality. So if it is held fixed that both P1 and P2 are true and mutually consistent with one another, then it must be the case that PIT is false. Put another way, the argument asserts that the conjunction of PIT and the extended mind thesis leads to a contradiction because this conjunction entails the extended consciousness thesis, but the extended consciousness thesis is false.

The strength of the extended mind argument against PIT turns on the plausibility of P1 and P2, the combination of which constitutes BEI. Before turning towards these premises, it is vital first to address an objection that one might raise to the argument *as a whole*. Namely,

Objection 1 (O1) The extended mind thesis tacitly assumes the truth of the functional role theory of intentionality, and thus, tacitly assumes that PIT is false. Therefore, one cannot invoke the extended mind thesis to argue against PIT without begging the question.

O1 suggests that I beg the question right out of the gate simply by casting my argument against PIT within the context of the extended mind framework. One cannot fully appreciate the reasoning behind O1 without first having a grasp of the extended mind thesis. The extended mind thesis derives from Andy Clark and David Chalmers (1998). They contend that cognitive processes can under certain conditions transcend the boundaries of the skull and seep out into the external world such that extra-cranial entities partly constitute them.⁶ Clark & Chalmers' original argument for the extended mind is motivated primarily by what they call *the parity principle*, which establishes the possibility of extended cognitive processes based on functional equivalency considerations: "If, as we confront some task, a part of the world functions as a process which, *were it done in the head*, we would have no hesitation in recognizing as part of the cognitive process, then that part of the world is (so we claim) part of the cognitive process" (Clark & Chalmers, 1998, p. 8). Clark and Chalmers apply the parity principle to the now-famous case of Otto, a fictional Alzheimer's patient who carries around a notebook as a substitute for his impaired biological memory. Given that Otto's notebook plays the same functional role in his cognitive economy that biological memory would otherwise play, Clark and Chalmers argue that the notebook and the writings therein should be conceptualized as a part of the realization base for some of Otto's dispositional belief states (specifically his belief about the location of the Museum of Modern Art).

It is because the extended mind thesis presupposes the parity principle that O1 may appear compelling. The parity principle seems to assume a functionalist picture of mentality, according to which X counts as mental if and only if X plays the requisite functional role in the overall cognitive system.⁷ What is true about mentality, so the idea goes, is also true about intentionality. Therefore, the parity principle, and by extension, the extended mind thesis, is committed to the functional role theory of intentionality.

O1 is flawed because it falsely assumes that what is true about the nature of mentality must also be true about the nature of source intentionality. It does not follow

⁶ Clark and Chalmers label the extended mind thesis as a version of what they call 'active externalism'. The view is 'active' because the extra-cranial entities in question are proximal (as opposed to distal), meaning that these entities can causally influence the cognitive subject, and the cognitive subject is in turn able to influence them causally. Active externalism contrasts with passive externalism, including Putnam's natural kind externalism (Putnam 1975) and Burge's social externalism (Burge 1979).

⁷ Colombetti and Roberts make this point in the context of considering how the extended mind theorist conceptualizes the 'mark of the mental': "The simplest mark thesis that could be extracted from this [parity] principle is a broadly functionalist one: x is a mental state iff x occupies the right causal role" (Colombetti & Roberts 2015: 360).

from the fact that the extended mind theorist is committed to a functionalist theory of mentality to the conclusion that she must also embrace a functionalist theory of source intentionality. This is because the question of what constitutes the mark of the mental is distinct from the question of what constitutes source intentionality. Some philosophers, to be sure, would deny the conceptual independence of these questions precisely because they identify source intentionality as the mark of the mental. But the philosophers who make this identification are not advocates of the extended mind thesis. Quite to the contrary, Adams & Aizawa (2008) equate source intentionality with the mark of the mental as a way to argue against the extended mind thesis.⁸

The extended mind theorist can thus uphold a functional role theory of mentality without thereby being committed to a functional role theory of intentionality. Functionalism about mentality is compatible with PIT, which is to say, mentality may be defined in functional terms while source intentionality is defined in phenomenal terms. O1 can thus be discarded. I will now examine P1 and P2 (i.e. BEI) of the extended mind argument against PIT, starting with P1, the premise that some source intentional states are examples of extended cognition.

3 Premise 1: Some source intentional states are extended

Upon first glance, P1 may even strike proponents of the traditional extended mind thesis as untenable, for they might maintain that only *dispositional* mental states with derived intentionality can be extended. Dispositional mental states are states that an agent is disposed to instantiate given certain conditions, whereas *occurrent* mental states refer to states that an agent is currently instantiating. A mental state is occurrent if and only if a subject is actively entertaining the state. Dispositional mental states (e.g. my standing belief that Washington, D.C. is the United States capital) ostensibly inherit their intentionality from their occurrent counterparts, and therefore, only have derived intentionality.

Clark and Chalmers' original presentation of the extended mind thesis argues for extended *dispositional* mental states, not extended *occurrent* mental states. Namely, Otto's *dispositional belief* concerning the location of the Museum of Modern Art extends into the environment in virtue of being constituted by a written proposition located in his notebook. Otto's *occurrent belief* about the museum's location, however, which he instantiates after consulting his notebook, is presumably a completely internal affair that takes place within the confines of his head. What one needs to support P1 is an independent argument for why some *occurrent* mental states are

⁸ Interestingly, Adams & Aizawa (2008) argue that the extended mind thesis is false because (a) source intentionality constitutes the mark of the mental and (b) source intentional states cannot be extended. In this paper I contend that (b) is false (i.e. source intentional states can be extended) in an effort to motivate the extended mind argument against PIT.

extended. In what follows, I draw attention to a prominent case promulgated by Clark (2008) that lends credence to the idea that the extended mind thesis encompasses some occurrent states.

Clark's (2008) argument for extended occurrent states, like the original extended mind argument for extended dispositional states (1998), assumes a broadly functionalist picture of mentality. The novelty of Clark's newer argument derives from his use of the functional concept of a *self-stimulating loop*. A system X enacts a self-stimulating loop when X produces outputs that it then recycles back into inputs. Clark illustrates the concept via the example of a turbo-driven engine, which uses its own emissions as a self-generating boost. His central contention is that occurrent cognitive processes count as extended when these processes become transiently coupled with external entities to produce a self-stimulating feedback loop. Clark begins to make this case by drawing attention to bodily gestures, arguing that gestures are not merely *expressions* of thought but are often *constitutive components* of occurrent thought processes which function as both systemic outputs and self-stimulated inputs in a subject's cognitive system: "the act of gesturing is part and parcel of a coupled neural-bodily unfolding that is itself usefully seen as an organismically extended process of thought" (Clark, 2008, p. 144).⁹ The concern for Clark's appeal to bodily gesture is that such an appeal at best establishes that cognition is *embodied*, not *extended*. To be extended, cognitive processes must be partly constituted by entities in the outer environment, not simply by extra-cranial body parts. Clark recognizes this, of course, and proceeds to explain that occurrent cognitive processes can form self-stimulating loops with external entities as well. In offering this explanation, he focuses on processes that are aided by the act of writing:

This kind of cognitively pregnant unfolding need not stop at the boundary of the biological organism. Something very similar may, as frequently remarked, occur when we are busy writing and thinking at the same time. It is not always that fully formed thoughts get committed to paper. Rather, the paper provides a medium in which, this time via some kind of coupled neural-scribbling-reading unfolding, we are enabled to explore ways of thinking that might otherwise be unavailable to us (Clark, 2008, p. 144).

The mental acts of calculation and planning are good examples of the kind of occurrent cognitive processes that Clark seems to have in mind: processes enhanced through the medium of writing by forming self-stimulating feedback loops with the independent artifacts constitutive of writing (e.g. pen and paper). The marks that I make on a piece of paper when thinking through a math problem, for example, are external outputs of cognition which are then redistributed as stimulating inputs during my occurrent act of calculation. So the written marks are not merely material

⁹ To justify this assertion, Clark draws on empirical research conducted by McNeil (2005) on the use of gesture in spontaneous speech and Goldin-Meadow (2003) on how gesturing functions to reduce the neural cognitive load for memory tasks.

manifestations of thought but are active drivers of my mathematical thinking operation. Clark's basic stance, to repeat, is that occurrent cognitive processes extend into the environment when deeply and reciprocally interwoven with external entities in a feedback loop of this sort.

My intention in this section is to transform Clark's discussion of self-stimulating loops into a formal argument for P1. The key sub-premise of this argument, which I call *the self-stimulating loop argument*, can be represented as follows:

SP1 The occurrent states involved in self-stimulating feedback loops are extended cognitive systems.

Notice that SP1 does not on its own entail the veracity of P1. To reach the conclusion that some source intentional states are extended, something like the following additional sub-premise needs to be made explicit: *All occurrent states possess source intentionality*. This additional sub-premise appears to be on solid footing, for as previously mentioned, it is commonly assumed in the literature that the intentionality of dispositional mental states derives from the source intentionality of their occurrent counterparts. There are various philosophers of mind, however, who reject the assumption that all occurrent states are source intentional. For example, David Woodruff Smith (2011) argues that there are occurrent states of 'pure' consciousness achievable via meditation that are entirely devoid of intentionality.¹⁰ I propose that the key to strengthening the self-stimulating loop argument for P1 is to restrict the scope of the argument to the class of occurrent mental states known as propositional attitudes. *Propositional attitudes* are types of mental states wherein a subject bears a cognitive relation to a proposition. Examples of such attitudes include (but are not limited to) states of belief, desire, understanding, and imagination. The assumption that all occurrent propositional attitudes (or "OPAs") have source intentionality, combined with the claim that some OPAs extend into the environment via self-stimulating loops, ostensibly entails the veracity of P1. When conceptualized in this way, the self-stimulating loop argument has the following structure:

3.1 The self-stimulating loop argument for P1

SP1 The occurrent states involved in self-stimulating feedback loops are extended cognitive systems.

SP2 Some occurrent states involved in self-stimulating feedback loops are propositional attitudes.

¹⁰ Smith describes these alleged states of pure (non-intentional) consciousness as follows: "In sustained meditation the sense of object disappears, i.e. the intentional content or object-specifying sense disappears, and the sense of self or subject disappears. That is to say, the usual intentional structure of subject object-or subject-act-content-object-dissolves in such a meditative state" (Smith, 2011, p. 481).

SC1 Therefore, some occurrent propositional attitudes are extended.

SP3 All occurrent propositional attitudes possess source intentionality.

P1 Therefore, some source intentional states are extended.

Consider first SP2, which I take to be well substantiated. The occurrent states of <planning> and <calculating> that Clark focuses on when presenting the self-stimulating loop argument appear to be propositional attitudes given that these states are both linguistically represented by verbs which embed ‘that-clauses.’ Moreover, even if the states of <planning> and <calculating> are not propositional attitudes, it is easy to imagine how one might expand Clark’s argument for extended occurrent states so that it includes within its purview standard propositional attitudes such as beliefs, judgments, desires, and fears. Colombetti and Roberts (2015), for instance, illustrate how Clark’s argument can be understood to include occurrent *judgments* within its scope. They consider the fictional case of Eve, a teenage girl who writes in her diary when she is upset with her parents: “The case of Eve, as we imagine it, is one where, as she writes that her parents do not listen to her, do not appreciate her, and so on, she is engaged in unfolding and articulating a specific evaluative judgment, which the act of writing down helps to clarify and structure. This act also feeds back into Eve, influencing her overall evaluative perspective as she continues to be engaged in the activity” (Colombetti & Roberts, 2015, p. 1258). The case of Eve is supposed to be an example of how a standard OPA can meet Clark’s ‘self-stimulating loop’ criteria for mental extension. Specifically, Eve’s act of writing in her notebook serves as a self-stimulating feedback loop in her occurrent act of judging her parents.¹¹

Turn now to SP3, which also strikes me as highly credible. SP3 gains support from the fact that OPAs are paradigmatic examples of source intentional states in the literature, which is to say, all theories of source intentionality (including PIT, tracking theories, and functional role theories) appear to either explicitly or implicitly recognize SP3 to be true. While there is controversy in the philosophy of mind over whether all mental states have intentionality, and even over whether all occurrent states have source intentionality, it is uncontroversial to declare that all OPAs have source intentionality. This being said, the objector might purport that SP3 does not necessarily follow from the fact that OPAs are paradigmatic examples of source intentional states. Just because a genus X is a paradigmatic example of that which has property P does not mean that all species of X instantiate P. For instance, birds are paradigmatic examples of flying animals, but it is not the case that all species of birds can fly (e.g. ostriches). Analogously, the objector might grant that OPAs are paradigmatic examples of source intentionality, but

¹¹ Colombetti and Roberts argue for the extension of occurrent judgments only as a part of a larger project of arguing for extended occurrent emotions.

deny that *all* OPAs are source-intentional. The most natural way for the proponent of PIT to maintain such a position is to allege that *some OPAs lack phenomenal consciousness*.¹² If OPAs include non-conscious states, then to assume that all OPAs are source intentional is to presuppose that PIT is false, for PIT maintains that source intentionality is phenomenally conscious intentionality. Thus, one might attempt to resist SP3 on the grounds that it begs the question against PIT:

Objection 2 (O2) The self-stimulating loop argument fails because SP3 begs the question against PIT.

The advocate of O2 is correct to suggest that a species X does not necessarily have property P just because the genus to which X belongs is a paradigmatic example of that which instantiates P. However, the fact that OPAs are paradigmatic examples of source intentional states lends a significant degree of credence to SP3, and the burden of proof is surely on those who wish to reject this sub-premise. The proponent of PIT might try to disavow SP3 by affirming that some OPAs are non-conscious but I deny that said proponent has the conceptual resources to make this affirmation.

First, PIT is arguably wedded to the idea that all OPAs are phenomenally conscious in virtue of being motivated by what Kriegel (2013) calls *inseparatism*, the view that “paradigmatic sensory states in fact exhibit intentionality, which is moreover grounded by their phenomenality, and that paradigmatic cognitive states in fact boast a phenomenality, which moreover grounds their intentionality” (Kriegel, 2013, p. 5). Inseparatism contrasts with a more traditional view of mind in analytic philosophy that Horgan and John (2002) label ‘separatism’, which holds that sensory mental states are phenomenally conscious but non-intentional, and cognitive mental states are intentional but non-phenomenal. Since PIT is partly motivated by inseparatism, and OPAs are paradigmatic cognitive states, it would seem as if PIT is committed to the notion that all OPAs are phenomenal. Second, the proponent of PIT cannot hold that some OPAs are non-conscious without also claiming that the OPAs in question have either derived intentionality or lack intentionality altogether. But it is wildly implausible that OPAs lack intentionality altogether, and it is even difficult to see how they could have derived intentionality. To say that an OPA has derived intentionality is to say that its intentionality is explainable in terms of some further intentional system. The problem is that one cannot easily understand the intentionality of OPAs in terms of some further intentional system; at least not in the same way that the intentionality of things like highway signs and linguistic objects can be understood in terms of the intentional mental states of the humans responsible for

¹² I am indebted to an anonymous reviewer for bringing my attention to this potential dialectical move by the defender of PIT.

fixing their respective meanings. These reasons suggest that PIT is committed to SP3, which means that SP3 does not beg the question against PIT.¹³

The most controversial sub-premise of the self-stimulating loop argument is undoubtedly SP1. The primary way to reject SP1 is to follow Robert Rupert (2004, 2009a, 2009b, 2011) in distinguishing the *extended* mind thesis from the *embedded* mind thesis and claim that the self-stimulating loop case is best interpreted through the framework of the embedded model. This gives rise to the following objection:

Objection 3 (O3) SP1 is false because self-stimulating loop cases are best construed though the framework of embedded cognition over the framework of extended cognition.

The embedded mind thesis advocates for an internalist vision of cognition while upholding that cognitive systems are nevertheless deeply interactive with and perhaps even casually reliant upon external entities. Rupert defines embedded cognition thusly: “Cognitive processes depend very heavily, in hitherto unexpected ways, on organismically external props and devices and on the structure of the external environment in which cognition takes place” (Rupert, 2004, p. 393). The thesis of embedded cognition is significantly more conservative than the extended view, which presents an externalist vision of cognition according to which cognitive systems can temporarily extend beyond the body into the outer world. Rupert argues that the embedded model is preferable over the extended model because the former possesses more explanatory power and is less metaphysically profligate than the latter. He thus denies that self-stimulating loop cases are genuine cases of mind extension. While cognition may intimately depend upon external artifacts during loop-related activity, the artifacts are not themselves a part of cognition, according to Rupert.¹⁴

¹³ Upon first glance, it might seem as if Angela Mendelovici (2018, 2020) defends a view of OPAs which rejects SP3. Mendelovici is one of the leading proponents of PIT in the field and advocates for a view of OPAs, which she calls *self-ascriptivism*, according to which “propositional attitudes are derived representational states, deriving their contents and their attitude types from our self-ascriptions” (Mendelovici 2020: 1). Since derived representational states lack intentionality altogether on Mendelovici’s view, it might seem as if she not only rejects SP3 or the assertion that *all* OPAs are source intentional, but that she also rejects the more moderate assertion that *some* OPAs are source intentional. However, Mendelovici’s self-ascriptivism only accounts for the *alleged* contents of propositional attitudes and does not account for the *immediate* contents of such attitudes. The immediate contents of propositional attitudes are the contents that are directly accessible via introspection. In contrast, the alleged contents are the contents that folk psychology takes propositional attitudes to have or the contents that we normally attribute to ‘that-clauses’. Mendelovici identifies the immediate contents of propositional attitudes with *phenomenal contents*, meaning that she believes that OPAs contain a phenomenally conscious aspect. She says, “we think in phenomenal tags that we take to stand for more complex and sophisticated contents” (Mendelovici 2018: 154). This filled-in picture of Mendelovici’s view of OPAs allows us to see that she does not actually reject P* but instead holds that OPAs have both a phenomenal intentional aspect (i.e. immediate contents) and a non-phenomenal, non-intentional aspect (i.e. alleged contents).

¹⁴ This criticism is similar in many respects to the famous ‘coupling-constitution’ fallacy promulgated by Adams and Aizawa (2001), which holds that the extended mind thesis is the result of a fallacious inference from the fact that external entities are causally coupled to cognitive processes to the conclusion that such entities are constitutive components of cognitive processes. Casual relevancy, according to Adams and Aizawa, does not entail constitutive role-playing.

Rupert's main contention in favor of the embedded model is that internal cognitive processes and their external functional analogs possess such different causal profiles that one should conceptualize them as distinct natural kinds.¹⁵ Regarding self-stimulating loops in particular, Rupert (2009a, 2009b) avers that outer contributions to the production of intelligent behavior during loop-related activity asymmetrically depend on the contributions of the internal cognitive architecture. For example, the sketchpad in the aforementioned self-stimulating loop case is reliant upon some set of neural activity to help drive cognition, but the relevant neural systems are not in turn reliant upon the sketchpad given that these systems alone, in the absence of any external scaffolding, are sufficient for the realization of cognitive processes. Asymmetric relations like these between internal cognitive mechanisms and external cognitive artifacts lead Rupert to conclude that there is no compelling theoretical reason to regard self-stimulating loops as instances of extended cognition. Worse still, Rupert (2009a) argues that adopting the extended mind framework in self-stimulating loop cases comes at a significant scientific cost. He stresses that widespread acceptance of self-stimulating loops as extended cognitive systems would inhibit our ability to locate persisting biological subjects for psychological and cognitive scientific research.¹⁶ For all of these reasons, Rupert alleges that external entities implicated in loop related activity are properly construed as cognitive scaffolds rather than cognitive extensions. Hence O3.

Clark (2008) presents a multifaceted response to O3 in defense of SP1. First, he repudiates the idea that external entities involved in self-stimulating loops should not be regarded as cognitive simply because they do not share the same fine-grained causal profile as their neural counterparts. Clark says this idea is the product of unjustified anthropocentrism and neurocentrism and results from a misinterpretation of the parity principle. The parity principle, according to Clark, does not mandate a fine-grained identity of causal profile but is rather "a call for sameness of opportunity, such that bio external elements might turn out to be parts of the machinery of cognition even if their contributions are unlike (perhaps deeply complementary to) those of the biological brain" (Clark, 2008, p. 135). Furthermore, Clark denies that relations of asymmetric dependence are relevant to the question of individuating cognitive systems. He supports this denial by noting that the internal neural architecture is composed of various mechanisms, some of which are themselves asymmetrically dependent on others for their role in cognition. Since these internal neural mechanisms clearly qualify as cognitive despite being asymmetrically dependent,

¹⁵ Rupert (2004) draws upon psychological research on working memory to make this contention, showing that internal memory processes possess significantly different fine-grained causal profiles than forms of external memory storage. For example, unlike external memory, internal working memory is integral to the successful execution of verbal exchanges, and is associated with particular kinds of interference patterns in paired-associates experiments.

¹⁶ For example, with respect to developmental psychology, Rupert maintains that the extended mind thesis "offers developmental psychologists no more reason to be interested in, for example, the series of temporal segments we normally associate with Sally from ages two-to-six than it offers to be interested in, say, Sally, aged two, together with a ball she was bouncing on some particular day, Johnny, aged five, together with the book he was reading on some particular afternoon, and Terry, aged seven, plus the stimulus item he has just been shown by an experimenter" (Rupert 2009a: 15).

Clark avers that external entities implicated in loop-related activity which are similarly asymmetrically dependent should also be eligible to be granted the status of ‘cognitive.’

Finally, in response to Rupert’s assertion that self-stimulating loops are not instances of extended cognition because interpreting them as such comes at a high scientific cost, Clark distinguishes between *organism centered* cognition and *organism bound* cognition. He maintains that the extended mind thesis does not diminish the scientific concept of a stable, persisting biological subject given that it preserves the vision of cognition as organism centered even though it denies that cognition is organism bound.¹⁷ Clark thus disaffirms that conceptualizing self-stimulating loops as extended cognitive systems fails to dovetail with the methodology of cognitive science. Quite to the contrary, he insists that the embedded mind thesis is the one in danger of falling victim to a nonscientific conceptual paradigm, avowing that the embedded model “threatens to repeat for outer circuits and elements the mistake that Dennett (1991) warns us against with regard to inner circuits and elements. It [the embedded mind thesis] depicts such outer resources as doing their work only by parading structure and information in front of some thoughtful inner overseer” (Clark, 2008, p. 153). According to Clark, once we dispense with the scientifically flawed notion of a privileged inner homunculus that directs the flow of all bio external information coming into the brain, the extended mind interpretation of self-stimulating loop cases becomes significantly more credible.

Suppose now that Clark’s response to O3 is successful, meaning that the OPAs involved in loop-related activity are cases of extended cognition. Even if this is true, and one acknowledges that all OPAs have source intentionality, the objector might nevertheless insist that self-stimulating loops fail to demonstrate that some source intentional states are extended. In other words, the objector might concede that all of the premises of the self-stimulating loop argument are true but maintain that the argument is invalid because the premises do not jointly imply P1. The idea, in particular, is that P1 does not follow from the conjunction of SC1 and SP3. According to this final possible objection, one cannot use the extended mind thesis to infer extended intentionality because it is a thesis about representational vehicles and not representational contents. Hence,

Objection 4 (O4) The self-stimulating loop argument is invalid because the extended mind thesis concerns representational vehicles, not representational contents.

When discussing the concept of mental representation, philosophers of mind often distinguish between the *vehicle* of representation and the *content* of representation (Hurley, 1998). The vehicle denotes the physical substrate that does the representing, whereas the content picks out that which is represented (i.e. what the

¹⁷ Clark says, “In rejecting the vision of human cognitive processing as organism bound, we should not feel forced to deny that it is (in most, perhaps all, real-world cases) organism centered. It is indeed primarily (though not solely) the biological organism that, courtesy especially of its potent neural apparatus, spins and maintains (or more minimally selects and exploits) the webs of additional structure that then form parts of the machinery that accomplishes its own cognizing” (Clark 2008: 141).

representation is ‘about’). The important point in this context is that intentionality is a feature of representational contents, not vehicles, and the extended mind thesis appears to pertain exclusively to vehicles. Indeed, Clark & Chalmers originally classified the extended mind as a version of *vehicle externalism*. They claim that “under certain circumstances we should see the material *vehicles* that realize the mind as encompassing not just brain activity, but also that of the body and the material environment” (Clark & Chalmers, 1998, p. 1243, emphasis added). If the extended mind thesis does not also encompass representational contents, then it cannot shed light on the prospect of extended intentionality. So, while the relevant OPAs may qualify as instances of extended cognition, they are arguably not examples of extended source intentionality because only the vehicles of these attitudes extend into the environment during loop-related activity. The source intentionality of OPAs is found in their content, but this content does not extend into the environment if O4 is correct.¹⁸

To counter O4 the defender of P1 must reject the notion that the extended mind thesis pertains exclusively to representational vehicles. Fortunately, a strong case can be made that entities in the material environment can partly constitute both vehicles and content. Holger Lyre (2016) contends that the extended mind thesis entails what he calls *active content externalism*, given the plausible assumption that *content supervenes on vehicles*. According to Lyre, the claim that content supervenes on vehicles is a relatively widespread assumption in the philosophy of mind.¹⁹ To say that content supervenes on vehicles is to say that vehicles play a determinative role with respect to content, for determination is the converse concept of supervenience. If vehicles do determine content, and the extended mind thesis is correct in asserting that the vehicles of cognition sometimes include external factors in the environment, then it follows that external factors sometimes partly determine mental content. To illustrate, recall the fictional case of Otto and his notebook. Otto, an Alzheimer’s patient, is taken to possess the extended belief that the MOMA is on 53rd street insofar as his notebook (which contains the directions to the MOMA) is a constitutive component of his cognitive system. Importantly, Otto’s extended belief is true only insofar as his notebook contains the correct address to the museum, meaning that if Otto’s notebook had included an inaccurate address, then his extended belief would be false. The content of Otto’s extended belief state is therefore actively determined by an aspect of the external world; namely, his notebook and the writings therein. Lyre elucidates this point by presenting a twin earth scenario involving Otto: “on twin earth Twotto wants to meet Twinga. But Twotto’s notebook mistakenly displays 51st street as MOMA’s address. So Twotto won’t meet Twinga. This shows [that]....a variation of the external component of the cognitive system, in this case the notebook entry, may change *mental content* in a behaviorally relevant manner” (Lyre, 2016, p. 7, emphasis added).

¹⁸ It is actually unclear whether or not Clark and Chalmers’ original discussion of the extended mind is meant to apply to just vehicles. When invoking the fictional case of Otto to argue for extended beliefs, Clark and Chalmers appear to insinuate that the *content* of Otto’s belief state depends partly on external components outside the brain (i.e. on the notebook itself). I elaborate on this point below.

¹⁹ Moreover, the assumption has recently been explicitly defended by Gottfried Vosgerau (2018).

While Lyre focuses on the case of Otto and extended dispositional states to support active content externalism, the view is also naturally supported by extended OPAs. Recall the fictional case of Eve, whose occurrent act of judging her parents extends into the environment via the self-stimulating feedback loop that she forms with her notebook. The propositional content of Eve's occurrent judgment is also actively shaped by the written outputs that she jots down onto the page, for these outputs "feed back into Eve, influencing her overall evaluative perspective as she continues to be engaged in the activity" (Colombetti & Roberts, 2015, p. 1258). The cases of Eve and Otto examples show that O4 is mistaken in suggesting that the extended mind thesis does not pertain to representational contents. By professing that content supervenes on vehicles, the advocate of the self-stimulating loop argument for P1 can uphold that the pertinent OPAs represent instances of extended source intentionality, and not merely extended cognition. Put differently, P1 does follow from the conjunction of SC1 and SP3 assuming the extended mind thesis implies active content externalism.

The aim of this section has been to motivate the self-stimulating loop argument for P1 and defend the argument against anticipated objections. I have shown that P1 is highly plausible given the following philosophical assumptions: (a) some OPAs extend into the environment via self-stimulating feedback loops, (b) all OPAs possess source intentionality, and (c) the extended mind thesis pertains to both vehicles and contents. I turn now to P2 of the extended mind argument against PIT. If it is assumed for the sake of argument that P1 is true (i.e. some source intentional states are extended), and, as has already been established, source intentional states can be extended on PIT only if consciousness can be extended, then one of two things must be the case: either PIT is false or some conscious states are extended. P2 rejects the latter disjunct by affirming that no conscious mental states are extended.

4 Premise 2: no conscious states are extended

In this section I present and motivate an argument for P2 deriving from Chalmers (2018), which I call *the direct access argument against extended consciousness*. The argument invokes the idea that cognitive extension requires sensorimotor interaction in addition to functional parity and that consciousness requires direct access to information for global control. If the argument is sound and extended consciousness is impossible, then the extended OPAs involved in self-stimulating feedback loops must be non-conscious, meaning that these propositional attitudes represent instances of source intentionality in the absence of phenomenal consciousness (meaning that PIT is false). After unpacking the direct access argument, I briefly consider two possible objections: one based on so-called external circuit cases and the other based on the enactive theory of consciousness. The central problem with these objections, as I explain, is that even if they succeed in undermining the direct access argument (which is highly controversial), they nevertheless fail to undermine P2 under a proper construal of the premise. To effectively counter P2, the objector must show that some currently existing conscious states are extended. In particular, the objector must illustrate that the OPAs featured in self-stimulating loop cases are

instances of extended consciousness. However, external circuit cases at best establish the possibility of extended consciousness in high-tech science fiction scenarios, whereas enactivism at best illustrates that perceptual conscious states are extended, not OPAs.

Chalmers' argument against extended consciousness can be represented as follows:

4.1 The direct access argument against extended consciousness

QP1 A mental state *X* is extended only if *X* involves sensorimotor interaction with external entities.

QP2 A mental state *X* is conscious only if *X* has direct access to information for global control.

QC1 Therefore, a mental state *X* is an extended conscious state only if (i) *X* features sensorimotor interaction with external entities *and* (ii) *X* has direct access to information for global control.

QP3 (i) and (ii) are incompatible.

QC2 Therefore, extended consciousness is impossible.

Consider first QP1. The immediate thing to notice about this premise is that it is at odds with the extended mind thesis as previously formulated. Clark and Chalmers (1998), to recall, conceptualize extended mentality in terms of the parity principle. Call this conception the 'original definition of extension':

Original Definition of Extension (ODE) An external entity *X* counts as an extended mental state of subject *S* just in case *X* plays the right functional role in *S*'s overall cognitive system.

Chalmers (2018) takes issue with ODE on the grounds that it is *objectionably uncontroversial* or 'too weak to be interesting.' In particular, he concurs with Farkas (2012), who contends that ODE is problematic because it includes external circuit cases within its purview. *External circuit cases* are instances where parts of the brain are substituted with functionally isomorphic silicon parts located in the outer environment and connected directly to internal neural mechanisms by wiring or radio transmitters. Clark (2009) offers the following example of an external circuit case:

But now imagine a case in which a person (call her Diva) suffers minor brain damage and loses the ability to perform a simple task of arithmetic division using only her neural resources. An external silicon circuit is added that restores the previous functionality. Diva can now divide just as before, only some small part of the work is distributed across the brain and the silicon cir-

cuit: a genuinely mental process (division) is supported by a hybrid bio-technological system (Clark, 2009).

The case of Diva qualifies as an example of extended cognition according to ODE because the external silicon circuit in question satisfies the parity principle. This implication of ODE leads Farkas to conclude that ODE is objectionably weak and needs to be replaced with a more robust version of the extended mind thesis. The extended mind thesis is supposed to be a compelling, controversial hypothesis, not a statement of unquestionable veracity. The problem is that the thesis is rendered utterly uncontroversial when interpreted through the lens of ODE, for even the most ardent opponents of extended cognition (e.g. Adams & Aizawa, 2008; Rupert, 2004) concede that cognitive processes can become extended in high-tech science fiction scenarios like the case of Diva. Chalmers agrees with Farkas that the bar for extended cognition needs to be set higher. He says, “Andy and I could stand our ground and stick with our stipulated definition of the extended mind thesis, so that Adams, Aizawa, Farkas, and Rupert all count as supporters of the thesis. That would be a little akin to the US declaring victory in Vietnam and going home. I think it makes more sense to find a stronger formulation of the extended mind thesis that captures what is really at issue in the debate” (Chalmers, 2018, p. 4).

Chalmers strengthens the extended mind thesis by adding a ‘sensorimotor requirement’ to the conception of cognitive extension. This requirement affirms that mental states must involve *sensorimotor interaction* with entities in the external environment to count as extended. Chalmers’ revised formulation of the thesis is comprised of both this sensorimotor requirement as well as the original parity principle:

Revised Definition of Extension (RDE) An external entity X counts as an extended mental state of subject S just in case (a) S interacts with X via perception and action, and (b) X plays the right functional role in S’s overall cognitive system.

Chalmers believes that RDE, which is essentially synonymous with QP1, gets the best of both worlds. On the one hand, it accommodates trademark cases of extended mental processes, like the fictional case of Otto (extended dispositional states) and the fictional case of Eve (extended occurrent states). On the other hand, RDE bypasses Farkas’ ‘too weak to be compelling’ criticism of the extended mind. External circuit cases do not involve sensorimotor interaction given that the circuits in question are connected directly to the brain via radio transmitters or wiring. The case of Diva therefore fails to meet the relevant criteria for extended cognition, according to RDE.²⁰

²⁰ To be clear, RDE does not hold that mentality in general requires sensorimotor interaction. The idea is not that sensorimotor interaction is a necessary condition for mindedness at large; it is rather the weaker claim that sensorimotor interaction is a necessary condition for *extended* mindedness. So, extended circuit cases do qualify as mental according to RDE, they just do not qualify as *extended* mental processes on the view.

Crucially, the adoption of RDE over ODE in no way affects the strength of Clark's self-stimulating loop argument for P1. OPAs bound up in self-stimulating loops with external entities satisfy the sensorimotor interaction condition for cognitive extension. For example, consider the previously discussed case of Eve and the self-stimulating loop that she forms with her notepad during the occurrent act of judging her parents. Eve engages with her notepad via perception and action; she looks at the notepad, turns its pages, and writes feverishly inside it. The relationship between Eve and her notepad is markedly different from the relationship between Diva and her external neural prosthetic device. Unlike Diva, Eve qualifies as a cognitive extender from the perspective of RDE because her cognitive artifact necessitates sensorimotor engagement instead of bypassing her perceptual faculties. The upshot of this is that the proponent of the extended mind argument against PIT can invoke RDE to help justify P2 without undercutting the self-stimulating loop argument for P1.

Turn now to QP2, the second key premise of Chalmers' argument against extended consciousness. QP2 states that a mental state is phenomenally conscious only if it has direct access to information for global control. Call this *the direct access condition*.²¹ Chalmers (1996, 1998) originally proposes the direct access condition as a 'pre-experimental bridging principle' that neuroscientists can use to help isolate the neural correlates of consciousness, or the set of mechanisms in the brain jointly sufficient for conscious experience. Discovering the neural correlates of consciousness is a notoriously daunting scientific task since consciousness is not directly observable or measurable. However, Chalmers indicates that consciousness can be indirectly measured by first propounding a functional property associated with conscious experience and then identifying the neural correlations of this functional property. The relevant functional property is regarded as a 'pre-experimental bridging principle' because it must be presupposed prior to any experimental research into the nature of consciousness. Chalmers says that the veracity of this bridging principle is not something that can be determined via scientific experimentation but rather must be established based on phenomenological or conceptual considerations.²² The direct access condition, according to Chalmers, is strongly

²¹ Chalmers (2008) and Clark (2009) consider an alternative version of QP2 which invokes the notion that consciousness requires 'high bandwidth' access to information. Extended consciousness is impossible according to this line of thought because extended processes involving perception and action provide comparatively low-bandwidth access to information. However, as Karina Vold (2015) has contended, perception actually relays external information to the brain at a relatively quick rate and so does not qualify as a low bandwidth process. Vold explains that the bandwidth of perception is roughly similar to the bandwidth of internal neural mechanisms: "non-neural processes must be constantly reporting information back to the brain...at least as quickly as neural processes can operate" (Vold 2015: 21). In light of this point from Vold, Chalmers goes on to endorse the direct access condition over the high-bandwidth access condition.

²² As Chalmers explains, "Experimental research gives us a lot of information about processing; then we bring in the bridging principles to interpret the experimental results, whatever those results may be. They are the principles by which we make inferences from facts about processing to facts about consciousness, and so they are conceptually prior to the experiments themselves. We cannot actually refine them experimentally (except perhaps by first-person experimentation!), because we have no independent access to the independent variable. Instead, these principles will be based on some combination of (1) conceptual judgments about what counts as a conscious process, and (2) information gleaned from our first-person perspective on our own consciousness" (Chalmers 1998: 3).

supported by phenomenological reflection, for it introspectively seems as if information in the brain is conscious if and only if it is directly accessible to guide different behavioral responses like verbal report and motor action. He elaborates:

A correlation between consciousness and global availability (for short) seems to fit the first-person evidence—that gleaned from our own conscious experience—quite well. When information is present in my consciousness, it is generally reportable, and it can generally be brought to bear in controlling behavior in all sorts of ways. I can talk about it, I can point in the general direction of a stimulus, I can press bars, and so on. Conversely, when we find information that is directly available in this way for report and other aspects of control, it is generally conscious information (Chalmers, 1998, p. 5).

Chalmers also emphasizes that the direct access condition is either implicitly or explicitly assumed to be true by many prominent neuroscientists and philosophers of mind working in the field, for the mechanisms that researchers put forward as candidates for the neural correlates of consciousness typically subserve the functional property of global availability. For example, Crick and Koch (1990) famously propose that 40-Hz oscillations in the cerebral cortex constitute a neural correlate of consciousness because these oscillations play a pivotal role in integrating information related to working memory and making it directly accessible for global control. Other notable theories of consciousness which endorse the direct access condition include (but are not limited to) Bernard Baars' (1988) Global workspace theory, Michael Tye's (1995) PANIC theory, and Martha Farah's (1994) High-quality representation theory. The fact that the direct access condition is widely adopted among empirical consciousness researchers and independently substantiated by phenomenological insight lends a high degree of plausibility to the principle, according to Chalmers.

After motivating RDE (i.e. QP1) and advocating for the direct access condition (i.e. QP2), Chalmers proceeds to explain why the conjunction of these two principles entails the impossibility of extended consciousness (i.e. QP3). The fundamental problem is that extended processes cannot supply direct access to information for global control because these processes are mediated by perception and action. Information can only be made directly available for global governance when located in the brain, not in the external environment. This is because information from the external environment must pass through multiple subsystems to reach the internal control system in the head. Chalmers elaborates: “these processes [cognitive processes that are extended via sensorimotor interaction] support information that is only indirectly available for global control: in order to be used in control, it must travel causal pathways from object to eye, from eye to visual cortex, and from visual cortex to the loci of control” (Chalmers, 2018, p. 10). In essence, then, because sensorimotor interaction with the environment is necessary for a conscious mental state to count as extended, but such interaction fails to satisfy the direct access condition, it follows that extended consciousness is impossible.

The direct access argument against extended consciousness is philosophically and empirically formidable, but it is not foolproof (as Chalmers himself concedes). There are three possible ways to object to the argument: (a) reject the direct access condition, (b) reject RDE, or (c) reject the idea that mental processes involving sensorimotor interaction with the environment cannot support information that is directly accessible for global control. The soundness of the direct access argument entails the veracity of P2, which is to say, the impossibility of extended consciousness implies that no actual conscious states are extended. However, it is important to recognize that the inverse entailment relation does not hold: it does not follow from the fact that extended consciousness is possible to the conclusion that some actual conscious states are extended. This is an essential point because it means that an objection to the direct access argument does not necessarily amount to an objection to P2. For example, one might reject RDE on the grounds that external circuit cases illustrate the possibility of extended consciousness. This objection to the direct access argument requires resisting Farkas' 'case against ODE and Chalmers' contention that RDE is preferable to ODE. Suppose one assumes that sensorimotor interaction is not necessary for extended cognition (i.e. RDE is false) and further stipulates that the brain processes implicated in external circuit cases constitute part of the neural correlates of consciousness. It then becomes natural to regard external circuit cases as instances of extended consciousness, for it seems that substituting the neural correlates of consciousness with functionally isomorphic silicon circuits located in the outer environment would extend conscious states.

There are at least two problems with this objection. First, it is unclear whether consciousness is even preserved in external circuit cases since it is currently unknown whether synthetic consciousness is possible. If it turns out that a cognitive system must be composed of carbon substrate in order to instantiate phenomenal consciousness, then consciousness will not be sustained when carbon-based brain bits are substituted with silicon-based functional duplicates. More to the point though, even if it were granted that external circuit cases demonstrate the *possibility* of extended consciousness and thereby undercut the direct access argument, such cases would still not compromise P2 given that this premise concerns the *actuality* of extended consciousness. Thus, even if it is possible for consciousness to extend into the environment in various high-tech science fiction scenarios (e.g. external circuit cases), this provides no reason to believe that any currently existing conscious states are extended.

The most common strategy for defending the actuality of extended consciousness is to endorse *enactivism*, a theory of perceptual consciousness championed by O'Regan and Noe (2001), Noe (2004). The enactivist conceptualizes perception as an active process that necessitates sensorimotor interplay with the world.²³

²³ Sensorimotor enactivism and the Revised Definition of Extension (RDE) are alike in that they both invoke the notion of 'sensorimotor interaction.' However, it is important to recognize that the two views address fundamentally different things: sensorimotor enactivism is a theory of perceptual experience whereas RDE is a theory of extended mentality. Put differently, RDE maintains that sensorimotor interaction is necessary for a mental state *to count as extended*, whereas sensorimotor enactivism claims that such interaction is necessary for a perceptual state *to be phenomenally conscious*.

The view stands in contrast to internalist accounts of perception, which hold that perceptual experience is fundamentally passive and relies solely on the instantiation of internal mental representations. Enactivism sees consciousness as a kind of ‘doing’, and explains the relationship between the perceiving subject and the world in terms of the possession of sensorimotor knowledge.²⁴ Various proponents of the enactivism (e.g. Noe, 2004; Ward, 2012, Pepper, forthcoming) avow that the theory entails extended consciousness. The guiding notion is that it follows from the thesis that perceptual consciousness implicates active engagement with elements in the external environment to the conclusion that these external elements form part of the physical substrate that realizes conscious experience. This leads to the following possible objection to P2:

Objection 5 (O5) P2 is false because the theory of enactivism is true.

The immediate problem with O5 is that enactivism appears to be at odds with the direct access condition for consciousness. As previously discussed, cognitive processes involving sensorimotor interaction cannot support information directly accessible for global control given that information from the environment must pass through two separate processing stages to be used by internal control mechanisms: a stage of processing by the perceptual system, and a second stage of processing to be globally broadcast. Enactivism is thus arguably false because the direct access condition is true. In response to this rebuttal the enactivist must either reject the direct access condition or reject the proposition that enactivism is incompatible with the direct access condition. Kiverstein and Kirchhoff (2019) appear to pursue the latter strategy in their book *Extended Consciousness and Predictive Processing: A Third Wave View*.²⁵

²⁴ As Adrian Downey elaborates, “there is a law-like relation between an organism’s movements and its visual stimulation— when an organism moves closer to an object the object looms in the visual field, when it gets further away the object appears smaller, and so on. On sensorimotor enactivism an organism is thought capable of perceiving only when it *understands* this relation between sensory stimulation and movement” (Downey 2017: 2). The relevant sensorimotor knowledge is varied, as each sensory modality is subject to its own particular law-like relations.

²⁵ Kiverstein and Kirchhoff observe that Chalmers’ direct access argument against extended consciousness derives from the particular conception of ‘sensorimotor interaction’ at play in the Otto notebook case, and specifically, from the role that perception and action play in the retrieval of information from the notebook. They agree that this extended mind-based conception of ‘sensorimotor interaction’ is incompatible with the direct access condition but allege that the enactivist operates with a more sophisticated version of the concept which is mutually consistent with the direct access condition. The kind of sensorimotor interaction pertinent to enactivism, according to Kiverstein and Kirchhoff, involves what they call ‘dynamic entanglement’ and ‘unique temporal signature.’ Dynamic entanglement denotes a reciprocal causal relationship between the agent’s sensory inputs and her motor outputs, while unique temporal signature refers to the idea that brain states must unfold over time to give rise to consciousness, and that this temporal unfolding requires environmental causes. They argue that conceptually engineering these two principles into the enactivist’s understanding of sensorimotor interaction allows the enactivist to avoid the direct access concern and infer the existence of extended consciousness.

The main problem with O5 is that enactivism pertains to perceptual states and not propositional attitudes. Even if enactivism is the correct account of perceptual experience and entails the actuality of extended consciousness (thereby disproving the direct access argument), the view still does not compromise P2 under a proper construal of the premise. The self-stimulating loop argument for P1 concerns extended propositional attitudes; specifically, it holds that OPAs are source intentional states that extend into the environment when coupled with external entities to form a self-stimulating feedback loop. Given this premise, it follows that the kinds of conscious mental states pertinent to P2 are OPAs and that objecting to P2 requires demonstrating that OPAs implicated in loop-related activity are instances of extended consciousness. Establishing that the extended consciousness thesis applies to some existing class of mental states which are not OPAs does not help the objector dismantle the extended mind argument against PIT. As long as P1 is true, and the relevant OPAs are not instances of extended consciousness, the defender of the argument can maintain that PIT leads to the aforementioned contradiction and is therefore false. The scope of P2 should therefore be restricted in the following manner such that it only targets conscious propositional attitudes and not all conscious mental states:

P2* No conscious propositional attitudes are extended.

This modification of P2 clarifies why an enactivist defense of extended consciousness fails to subvert the premise. The fundamental issue is that enactivism is a theory of perceptual consciousness and perceptual states are not types of propositional attitudes. There are, to be sure, certain versions of representationalism which regard perceptual states as propositional attitudes, but enactivism is traditionally put forward as a competitor to these representationalist theories of mind and denies that perceptual states have propositional content.²⁶ Thus, even if enactivism implies that perceptual conscious states are extended, it supplies no reason to think that any conscious OPAs extend into the environment.

5 Conclusion

The extended mind argument against PIT is motivated by what I called the ‘bipartite extension intuition’ (i.e. BEI), the claim that some source intentional states extend into the environment, but no conscious states do. Put another way, BEI holds that the extended mind thesis is true but the extended consciousness thesis is false, where ‘the extended mind thesis’ is understood to encompass some OPAs and ‘the extended consciousness thesis’ is stipulated to pertain exclusively to OPAs. BEI is inconsistent with PIT because the extended mind thesis entails the extended consciousness thesis if PIT is true. This paper has offered

²⁶ See Crane (2009) for an argument as to why perceptual states do not possess propositional content.

a philosophical roadmap for how BEI can be justified in a non-question begging manner. The spirit of my discussion has not been to prove with certainty that the argument against PIT is sound but rather to demonstrate that the argument is creditable given certain philosophical assumptions about the nature of consciousness and extended cognition. There are undoubtedly other ways to validate P1 and P2 (i.e. BEI) that have not been considered here. It is possible that both Clark's self-stimulating loop argument and Chalmers' direct access are flawed but that P1 and P2 are nevertheless true, for the truth of the premises might be grounded in reasons that have nothing to do with the existence of self-stimulating feedback loops or the notion that consciousness requires direct access to information for global control.

More broadly, I think that instead of drawing upon *extended* mental states and the extended mind literature to argue against PIT, one might alternatively seek to undermine PIT by drawing upon *collective* mental states and the collective intentionality literature (e.g. Gilbert, 1989; Pettit, 2003; Tuomela, 1992). Consider how a group of scientists can, throughout time, collectively understand some natural phenomenon that no individual scientist can understand. This case appears to pick out an example of distributed cognition which involves the tokening of a collective mental state, namely, a collective attitude of understanding. Moreover, the collective propositional attitude in question seems to be occurrent rather than dispositional since the process of reaching a shared scientific understanding of any natural phenomenon is one that actively unfolds over time as empirical experiments are conducted and the different stages of the scientific method are completed. If one additionally assumes that there is no phenomenally conscious group agent present in this scenario that is the bearer of the collective scientific understanding, then the scenario arguably represents an instance of a non-conscious OPA (i.e. an instance of source intentionality in the absence of phenomenal intentionality). The argument from collective intentionality that I am envisioning here is structurally similar to the extended mind argument against PIT and might be formally devised as follows:

5.1 The collective intentionality argument against PIT

P1 Collective mental states exist.

P2 Some collective mental states are OPAs.

P3 All OPAs possess source intentionality.

C1 Therefore, some collective mental states possess source intentionality.

P4 No collective mental states are phenomenally conscious.

C2 Therefore, some source intentional states lack phenomenal consciousness (meaning that PIT is false).

Determining whether this argument is tenable would mandate a detailed investigation into the collective intentionality and group agency literature. I raise the argument here only to illustrate that the extended mind argument represents one species of a more general argumentative strategy against PIT. This general argumentative strategy consists in *identifying a counterexample* to PIT by demonstrating that some class of source intentional states is devoid of phenomenal consciousness (and therefore phenomenal intentionality). The extended mind argument avers certain extended states (i.e. extended OPAs) represent counterexamples to PIT, whereas the collective intentionality argument regards certain collective states (i.e. collective OPAs) as counterexamples. The key to making this general argumentative strategy persuasive is to *avoid begging the question against PIT*. The challenge is to show that there is a class of non-conscious, source intentional states without presupposing the falsity of PIT. As discussed in sections II and III, the argument presented here does not presuppose the falsity of PIT and so bypasses any question-begging concern. If the extended mind argument against PIT is sound, then the prevailing sentiment that intentionality is grounded in consciousness is mistaken.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Adams, F., & Aizawa, K. (2001). The bounds of cognition. *Philosophical Psychology*, 14, 43–64.
- Adams, F., & Aizawa, K. (2008). *The bounds of cognition*. Wiley-Blackwell.
- Armstrong, D. M. (1968). *A materialist theory of Mind*. Routledge.
- Baars, B. J. (1988). *A cognitive theory of consciousness*. Cambridge University Press.
- Baars, B. J., Newman, J. B., & Taylor, J. G. (1998). Neuronal mechanisms of consciousness: A relational global workspace approach. In S. R. Hameroff, A. W. Kaszniak, & A. C. Scott (Eds.), *Toward a science of consciousness II* (pp. 269–278). MIT Press.
- Battaly, H. (2018). Extending epistemic virtue: extended cognition meets virtue-responsibilism. In J. A. Carter, A. Clark, J. Kallestrup, S. O. Palermos, & D. Pritchard (Eds.), *Extended epistemology*. Oxford University Press.
- Bobrow, S. A., & Bower, G. H. (1969). Comprehension and recall of sentences. *Journal of Experimental Psychology*, 80(31), 455.
- Bordini, D. (2017). Is there introspective evidence for phenomenal intentionality? *Philosophical Studies*, 174(5), 1105–1126.
- Burge, T. (1979). Individualism and the mental. In P. French, T. Uehling, & H. Wettstein (Eds.), *Midwest studies in philosophy 4* (pp. 73–121). University of Minnesota Press.
- Chalmers, D. J. (1996). *The conscious mind. In search of a fundamental theory*. Oxford: Oxford University Press.

- Chalmers, D. J. (1998). On the search for the neural correlate of consciousness. In Stuart R. Hameroff, Alfred W. Kaszniak & A.C. Scott (eds.), *Toward a Science of Consciousness II*. MIT Press, pp. 2–219.
- Chalmers, D. (2018). Extended cognition and extended consciousness. In M. Colombo, E. Irvine, & M. Stapleton (Eds.), *Andy Clark and his critics*. Wiley.
- Clark, A. (2008). *Supersizing the mind: Embodiment, action, and cognitive extension*. Oxford University Press.
- Clark, A. (2009). Letter to the editor. *London Review of Books*, 31, 6.
- Clark, A. (2010). Coupling, constitution and the cognitive kind: A reply to Adams and Aizawa. In R. Menary (Ed.), *The extended mind*. MIT Press.
- Clark, A., & Chalmers, D. J. (1998). The extended mind. *Analysis*, 58(1), 7–19.
- Cohen, L. J. (1992). *An essay on belief and acceptance*. Clarendon Press.
- Colombetti, G., & Roberts, T. (2015). Extending the extended mind: The case for extended affectivity. *Philosophical Studies*, 172(5), 1243–1263.
- Crane, T. (2009). Is perception a propositional attitude? *Philosophical Quarterly*, 59(236), 452–469.
- Crick, F., & Koch, C. (1990). Towards a neurobiological theory of consciousness. *Seminars in the Neurosciences*, 2, 263–275.
- Dennett, D. C. (1991). *Consciousness explained*. Penguin Books.
- Downey, A. (2017). Radical sensorimotor enactivism & predictive processing—Providing a conceptual framework for the scientific study of conscious perception. In T. Metzinger & W. Wiese (Eds.), *Philosophy and predictive processing*. MIT Press.
- Dretske, F. (1981). *Knowledge and the flow of information*. MIT Press.
- Dretske, F. (1995). *Naturalizing the mind*. MIT Press.
- Farah, M. J. (1994). Visual perception and visual awareness after brain damage: A tutorial overview. In C. Umilta & M. Moscovitch (Eds.), *Consciousness and unconscious information processing: Attention and performance*. MIT Press.
- Farkas, K. (2008). Phenomenal intentionality without compromise. *The Monist*, 91(2), 273–293.
- Farkas, K. (2012). Two versions of the extended mind thesis. *Philosophia*, 40(3), 435–447.
- Fodor, J. A. (1987). *Psychosemantics: The problem of meaning in the philosophy of mind*. MIT Press.
- Fodor, J. A. (1990). *A theory of content and other essays*. MIT Press.
- Forrest, P. (2017). Can phenomenology determine the content of thought? *Philosophical Studies*, 174(2), 403–424.
- Gertler, B. (2001). Introspecting phenomenal states. *Philosophy and Phenomenological Research*, 63(2), 305–328.
- Gilbert, M. (1989). *On social facts*. Routledge.
- Goldin-Meadow, S. (2003). *Hearing gesture: How our hands help us think*. Harvard University Press.
- Graham, G., Horgan, T. E., & Tienson, J. L. (2007). Consciousness and intentionality. In M. Velmans & S. Schneider (Eds.), *The Blackwell companion to consciousness* (pp. 468–484). London: Blackwell.
- Horgan, T., & Graham, G. (2012). Phenomenal intentionality and content determinacy. In R. Schantz (Ed.), *Prospects for meaning*. De Gruyter.
- Horgan, T., & Kriegel, U. (2008). Phenomenal intentionality meets the extended mind. *The Monist*, 91(2), 347–373.
- Horgan, T., & Tienson, J. (2002). The intentionality of phenomenology and the phenomenology of intentionality. In D. J. Chalmers (Ed.), *Philosophy of mind: Classical and contemporary readings* (pp. 520–533). Oup Usa.
- Hurley, S. (1998). Vehicles, contents, conceptual structure, and externalism. *Analysis*, 58(1), 1–6.
- Kirchhoff, M. D., & Kiverstein, J. (2019). *Extended consciousness and predictive processing: A third-wave view*. Routledge.
- Kriegel, U. (2011). *The sources of intentionality*. Oxford University Press.
- Kriegel, U. (2013). The phenomenal intentionality research program. In U. Kriegel (Ed.), *Phenomenal intentionality*. Oxford University Press.
- Lackey, J. (2014). Socially extended knowledge. *Philosophical Issues*, 24(1), 282–298.
- Loar, B. (1987). Subjective intentionality. *Philosophical Topics*, 15(1), 89–124.
- Loar, B. (2003). Phenomenal intentionality as the basis of mental content. In M. Hahn & B. Ramberg (Eds.), *Reflections and replies: Essays on the philosophy of Tyler Burge* (pp. 229–258). MIT Press.
- Ludwig, K., & Jankovic, M. (Eds.). (2018). *The Routledge handbook of collective intentionality*. Routledge.
- Lyre, H. (2016). Active content externalism. *Review of Philosophy and Psychology*, 7(1), 17–33.
- McNeill, D. (2005). *Gesture and thought*. University of Chicago Press.

- Mendelovici, A. (2018). *the phenomenal basis of intentionality*. Oxford University Press.
- Mendelovici, A. (2020). Propositional attitudes as self-ascriptions. In L. R. G. Oliveira & K. Corcoran (Eds.), *Commonsense metaphysics essays in honor of Lynne Rudder Baker*. Oxford: Oxford University Press.
- Mendelovici, A., & Bourget, D. (2014). Naturalizing intentionality: Tracking theories versus phenomenal intentionality theories. *Philosophy Compass*, 9(5), 325–337.
- Millikan, R. G. (1984). *Language*. MIT Press.
- Montague, M. (2016). *The given: Experience and its content*. Oxford: Oxford University Press.
- Müller, V. C. (2012). From embodied and extended mind to no mind. In A. Esposito, A. M. Esposito, R. Hoffmann, V. C. Müller, & A. Viniciarelli (Eds.), *Cognitive behavioural systems* (pp. 299–303). Springer.
- Myin, E., & Loughlin, V. (2018). Sensorimotor and enactive approaches to consciousness. In R. J. Gennaro (Ed.), *Routledge handbook of consciousness*. London: Routledge.
- Neander, K. (1996). Swampman meets swampcow. *Mind and Language*, 11(1), 118–129.
- Noë, A. (2004). *Action in perception*. MIT Press.
- Oregan, J. K. (2009). Sensorimotor approach to (phenomenal) consciousness. In B. Tim, C. Axel, & W. Patrick (Eds.), *The Oxford companion to consciousness* (pp. 588–593). Oxford University Press.
- O'Regan, J. K., & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, 24(5), 883–917.
- Papineau, D. (1984). Representation and explanation. *Philosophy of Science*, 51(December), 550–572.
- Pepper, K. (2021). Do sensorimotor dynamics extend the conscious mind? *Adaptive Behavior*, 22, 99–108.
- Pettit, P. (2003). Groups with minds of their own. In A. I. Goldman & D. Whitcomb (Eds.), *Social epistemology: Essential readings*. Oxford University Press.
- Phillips, I. (2016). Consciousness and criterion: On Block's case for unconscious seeing. *Philosophy and Phenomenological Research*, 92(3), 1.
- Pitt, D. (2004). The phenomenology of cognition, or, what is it like to think that P? *Philosophy and Phenomenological Research*, 1, 1–36.
- Putnam, H. (1975). The meaning of meaning. *Mind, language and reality; Philosophical papers* (Vol. 2, pp. 215–271). Cambridge University Press.
- Quine, W. V., & Orman, . (1960). *Word and Object*. MIT Press.
- Quinton, A. (1975). Social objects. *Proceedings of the Aristotelian Society*, 76, 1–27.
- Rowlands, M. (2009). Extended cognition and the mark of the cognitive. *Philosophical Psychology*, 22(1), 1–19.
- Rupert, R. D. (1999). The best test theory of extension: First principle(s). *Mind and Language*, 14(3), 321–355.
- Rupert, R. D. (2004). Challenges to the hypothesis of extended cognition. *Journal of Philosophy*, 101(8), 389–428.
- Rupert, R. D. (2009a). *Cognitive systems and the extended mind*. Oup Usa.
- Rupert, R. D. (2009b). Innateness and the situated mind. In P. Robbins & M. Aydede (Eds.), *The Cambridge handbook of situated cognition* (pp. 96–116). Cambridge University Press.
- Rupert, R. D. (2011). Cognitive systems and the supersized mind. *Philosophical Studies*, 152(3), 427–436.
- Schneider, S. (2019). *Artificial you: AI and the future of your mind*. Princeton University Press.
- Schweikard, D. P. & Schmid, H. B. (2013). Collective intentionality. In: E. N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/archives/sum2013/entries/collective-intentionality/>. Accessed 6 April 2021.
- Searle, J. R. (1980). Minds, brains and programs. *Behavioral and Brain Sciences*, 3(3), 417–457.
- Searle, J. R. (1990). Consciousness, explanatory inversion and cognitive science. *Behavioral and Brain Sciences*, 13(1), 585–642.
- Searle, J. R. (1992). *The rediscovery of mind*. MIT Press.
- Siewert, C. (1998). *The Significance of Consciousness*. Princeton University Press.
- Smith, D. (2011). Nibbanic (or pure) consciousness and beyond. *Philosophia*, 39(3), 475–491.
- Stampe, D. W. (1977). Towards a causal theory of linguistic representation. *Midwest Studies in Philosophy*, 2(1), 42–63.
- Strawson, G. (1994). *Mental reality*. MIT Press.
- Taraborelli, D., & Mossio, M. (2008). On the relation between the enactive and the sensorimotor approach to perception. *Consciousness and Cognition*, 17(4), 1343–1344.
- Theiner, G. (2018). Groups as distributed cognitive systems. In M. Jankovic & K. Ludwig (Eds.), *Routledge handbook on collective intentionality* (pp. 414–442). New York: Routledge.

- Thompson, E. (2005). Sensorimotor subjectivity and the enactive approach to experience. *Phenomenology and the Cognitive Sciences*, 4(4), 407–427.
- Tuomela, R. (1992). Group beliefs. *Synthese*, 91(3), 285–318.
- Tye, M. (1995). *Ten problems of consciousness: A representational theory of the phenomenal mind*. London: MIT Press.
- Vold, K. (2015). The parity argument for extended consciousness. *Journal of Consciousness Studies*, 22(3–4), 16–33.
- Vosgerau, G. (2018). Vehicles, contents and supervenience. *Filozofija i Društvo*, 29, 473–488.
- Ward, D. (2012). Enjoying the spread: Conscious externalism reconsidered. *Mind*, 121(483), 731–751.
- Weiskrantz, L. (1986). *Blindsight: A case study and implications*. Oxford University Press.
- Weiskrantz, L. (1997). *Consciousness lost and found*. Oxford University Press.
- Wheeler, M. (2005). *Reconstructing the cognitive world: The next step*. London: Bradford.
- Wilson, R. A., & Foglia, L. (2017). Embodied Cognition. In: E. N. Zalta (ed.), *The Stanford encyclopedia of philosophy*. <https://plato.stanford.edu/archives/spr2017/entries/embodied-cognition/>. Accessed 14 December 2020.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.