Why the extended mind is nothing special but is central

Abstract

The extended mind thesis states that the mind is not brain-bound but extends into the physical world. The philosophical debate around the thesis has mostly focused on extension towards epistemic artefacts, treating the phenomenon as a special capacity of the human organism to recruit external physical resources to solve individual tasks. This paper argues that if the mind extends to artefacts in the pursuit of individual tasks, it extends to other humans in the pursuit of collective tasks. Paradigmatic cases of extended mind in the original literature are only particular manifestations of the more general capacity for collective intentionality, the unique power of human minds to be jointly directed at goals, intentions, or values. Because this capacity holds developmental and diachronic primacy over human-epistemic artefacts relations, the extended mind should not be seen as a special phenomenon, but as a central aspect of the human condition. The original extended mind thesis carried important implications for how the cognitive sciences should proceed. In a version of the thesis that accommodates collective intentionality, these implications would go far deeper than originally assumed.

Keywords

Extended mind; collective intentionality; shared cognition; extended emotions; social ontology

Introduction

One of the strengths of philosophy is its ability to subvert traditional wisdom by unsettling intuitions, making one view the world in a fundamentally new light. Clark and Chalmers' (1998) 'extended mind thesis' is one such example. This suggests that the mind, hitherto conceived as skull-bound, spills over into body and world. Our tools and artefacts are not just connected to us but, like neural mechanisms, can be a constitutive part of our cognitive system. The thesis has stirred important debate on the proper unit of analysis of the cognitive sciences: if it is right, then the cognitive sciences have mistaken their subject matter and approach because the material realisers of the mind extend into the environment.

In this paper, we argue that the thesis has yet to see its full implications. In debating whether the mind extends or not, how, and where to, philosophers have glossed over the most important way in which the mind extends into the world: to other humans in acts of collective intentionality. We argue that paradigmatic examples of the extended mind are only a particular manifestation of the more fundamental cognitive capacity for collective intentionality – the power of minds to be jointly directed at objects, matters of fact, states of affairs, goals, or values. The exclusion of this (uniquely human) phenomenon from the debate is arbitrary and limits the reach of the extended mind thesis. We begin with a standard account of the original extended mind debate. We then draw a parallel between extension to tools in individual tasks and extension to humans in collective tasks, arguing that the capacity for the latter holds primacy over the former. We end by reflecting on what the cognitive sciences would look like if they took to heart the idea that we constantly think and act with other people 'in mind'.

1. The extended mind

In The Extended Mind, Clark and Chalmers appealed to functionalism to argue that:

> 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*… part of the cognitive process. (1998:8)

This has come to be known as the 'parity principle'. In their flagship example, two individuals, Otto and Inga, wish to attend an exhibition at the Museum of Modern Art (MoMA) in New York. Inga retrieves the location of the museum, 53rd Street, from biological memory. Otto, who has onset Alzheimer's, retrieves it from a notebook in which he writes information he may need at a later time. Clark and Chalmers suggest that because the notebook plays the role of the cognitive vehicle for memory it is legitimate to treat it as a constitutive part of Otto's cognitive system. The extended mind thesis claims that some cognitive processes extend beyond the brain to include artefacts in the environment in a way that helps us navigate it more effectively. In response to critique that the thesis takes an excessively permissive position on cognition, leading to 'cognitive bloat' (Rowlands 2010; Rupert 2004), some tightening conditions have been proposed. Thus, the epistemic artefact in question must:

1 Be reliably available and typically invoked.
2 Contain information that is automatically endorsed. It should not usually be subject to critical scrutiny (e.g., unlike the opinions of other people). It should be deemed about as trustworthy as something retrieved clearly from biological memory.
3 Contain information that should be easily accessible as and when required.
4 Contain information that has been consciously endorsed at some point in the past and indeed there is as a consequence of this endorsement.

The extended mind thesis challenges internalism, which posits that psychological processes are confined to the brain. Internalists argue that the input and output of a psychological process may be located in the environment, but the process itself happens inside the head. Modern cognitive science has been particularly internalist in orientation, conducting research primarily in laboratory settings, where cognitive mechanisms are abstracted from their environment. If Clark and Chalmers are right, however, the cognitive sciences have mistaken their subject matter. According to what Clark labels the 'principle of cognitive impartiality' (2007:174), brains accomplish their tasks following cost-functions that are impartial to the nature (motoric, perceptual) or location (in-the-head, in-the-world) of the processing. For cost-benefit reasons, brains happen in many cases to recruit external epistemic resources and it is thus legitimate to say that cognition spans brain and environment. Yet, Clark argues explicitly that cognitive extension only occurs ephemerally and in special conditions, mostly in situations of fast and frugal human-artefacts interaction (Clark, 2007:176).

The extended mind thesis has been much debated. One of the main criticisms is that in order to say that the mind extends, one must first define what mind is; in other words, there needs to be a 'mark of the mental' (Adams 2019; Adams and Aizawa 2001, 2009, 2010). The debate on this question has now reached a deadlock because it has become clear that no empirical finding can adjudicate such a definition. To the extent in which it has been about defining the 'mind', the debate has taken a metaphysical turn (Kiverstein 2018; Sprevak 2010). To some, including us, this debate is unnecessary because the definitions of 'mind' that have been proposed (e.g. Adams and Aizawa, 2001) turn out to be remote from any common usage of the term as well as of little use to cognitive science. They are so because cognitive scientists already work with a rough and ready notion of mind that broadly refers to processes that support intelligent behaviour, and that's enough for cognitive science to proceed. Biologists, for instance, do not need an exact definition of 'life' to study living beings. The value of biological science is not compromised by the fact that biologists also study viruses, which might not be considered 'alive'; at the same time, there is no current risk of seeing biologists studying rocks (Allen 2017). Clark's way out of this metaphysical muddle is to state that "cognition is as cognition does" (Clark 2010:93). Ultimately, the move behind the original extended mind thesis consists in engaging with a common, intuitive understanding of 'mind' and 'cognition' and using philosophical argument to change – in this case, expand – these understandings, while simultaneously opening new horizons and research questions for scientific practice. In other words, cognitive science cannot decide whether the extended mind thesis is true, but the intuition induced by the thesis has the power of changing the approach of the cognitive sciences. The implications of Clark's argument is that a cognitive science that confines itself to internal cognitive activity would be as impoverished as a cognitive science that confines itself to one part of the brain.

In this article, we follow Clark's lead, but reach an entirely different conclusion. We are not interested in delving into the extended mind debate, except to say that it would look entirely different had its discussants considered the following possibility: that the mind primarily extends to other people in acts of collective intentionality, and that it is only because of this that it can extend to epistemic artefacts. In the current debate, the almost exclusive focus on epistemic artefacts has accompanied an individualistic position that precludes the extension to other humans in collective action. In the profuse literature that developed from the original article, the 'part of the world' that mind extends to has mostly been conceived as a scaffold of material objects that aids solitary epistemic action: a notebook helps to overcome *one*'s memory issues; reliable access to a calculator aids *one's* calculations. Clark and Chalmers gesture at the possibility of the mind extending to other human beings – e.g. a waiter at Otto's favourite restaurant might act as a

repository of his beliefs about his favourite meals – but the point has not been elaborated further by the authors. In Clark's writings we find a fascination with gadgets, props, tools, prostheses, implants, etc. When other people are mentioned, it is done so with thinly veiled scepticism. This has changed recently, with some authors invoking the idea of the 'socially extended mind' (Gallagher 2013; Slaby and Gallagher 2015). In certain contexts, these authors argue, other people and even institutions afford extension and can legitimately become part of one's cognitive system, providing that we adopt a more liberal version of the extended mind where the requirements of accessibility, trustworthiness and reliability are relaxed. But even when this possibility is contemplated – the legal system, for instance, could be considered a cognitive institution that facilitates one's actions in the world and is therefore constitutive of one's mind – the epistemic actions in question are undertaken to overcome an individual problem. Gallagher (2011) explicitly likens cognitive institutions to 'tools' which afford individual action and accomplish individual aims.

At root, the extended mind literature has been underpinned by methodological individualism. It has been treated separately from the literature on collective intentionality. For example, The Routledge Encyclopaedia of Philosophy states that the extended mind thesis "is a claim about individual human cognition extending, not about shared or group cognition. [It] may be combined with a variety of other externalist claims about the mind, but it makes a separate, unique claim about how the human mind spreads out into world." (Sprevak 2019). To be sure, philosophers who write about collective intentionality have sometimes suggested that the latter entails mental extension to other human beings. This idea is found in the discussions of distributed cognition (e.g. Theiner et al. 2010), enactivist approaches to 'participatory sense-making' (Thompson and Stapleton 2009; Garavito 2019), and, in similar guises, in past thinkers such as Dewey or Merleau-Ponty. For example, modern enactivists have proposed that a Deweyan concept of cognition is largely convergent with a range of externalist approaches (Dreon 2021); in particular, the importance Dewey gave to the 'situation' – defined as an environing experienced world – in understanding cognition (Dewey 1931, 1982, 1995). On these terms, human cognition is understood through the coupled organism-environment (Gallagher 2017; Crippen and Schulkin 2020).

If collective intentionality researchers have discussed and incorporated ideas about mental extension, however, extended mind researchers have not returned the favour, and, with them, most cognitive scientists. Arguably, this has to do with the fact that the language and conceptual framework of these approaches are so removed from those of the extended mind that philosophers like Clark feel they do not speak much to the debate. Hence there remains a rift. This

paper is an attempt at charting a connection, and is mostly addressed at extended mind philosophers – aware that the individualistic premises that underpin the thesis also hold sway across the cognitive sciences (including social psychology (Greenwood 2003)). Starting from the assumptions of the extended mind thesis, we gradually show that the latter entails extension to other human beings. We argue that a fully-fledged version of the thesis, in which the links with collective intentionality are acknowledged and their wider implications explicated, holds the potential to effect a far more radical shift in our approach to human cognition. To make our case, we will highlight some striking parallels between individual-tool coupling in individual tasks and individual-other human coupling in the context of collective tasks. First, however, a brief overview of collective intentionality.

2. Collective intentionality

Collective intentionality is the capacity of human minds to be jointly directed at objects, goals, matters of fact, states of affairs or values (Schweikard and Schmid 2013), and it is often proposed as an important, perhaps even defining, human cognitive trait (Tomasello 2019). In the philosophical literature, collective intentionality broadly takes the form "I intend that we J" (Bratman 1999). Although there has been substantial debate on how best to define it beyond such formula (Tuomela 2007; Bratman 2013; Gilbert 2015), there is now increasingly widespread consensus that acts of collective intentionality are not reducible to a collection of individual I-referential intentions. Think, for instance, of a group of random people sitting in a park who suddenly get up and run to a central shelter as it starts to rain. Think, now, of a group of people in a park who accomplish the same action as part of a theatrical performance. The outward behaviour of these people might be the same, but their intentional states differ starkly. The latter case evinces a sense of doing something *together*, where the individual intentionality of each person is part of the collective intentionality that they share (this example comes from Searle 1990). Philosophers have termed this state 'we-intention'. Modern debate on we-intention has focussed on Tuomela and Miller (1988) version, in which we-intention consists of an individual having the intention to do their part, a belief others will do their parts, and a belief that this belief is mutual among the relevant others. Tuomela (2013) sees this 'we-mode' as a common-sense alternative to the extremes of individualist-content and collectivist-subject (as well as 'group mind') accounts of collective intentionality. Philosophical dispute notwithstanding, central to collective intentionality is the idea that for something like joint action to take place there needs to be a 'shared point of view' – a sense of 'we'. With this key point in mind, we return to the parallels between the extended mind and collective intentionality.

3. The extended mind and collective intentionality

Consider the following scenario. Coming back to his village after an unsuccessful solitary hunt, Elmo gathers friends and relatives to organise a collective expedition the following day. They set off in the morning, looking for deer. They agree to disband to cover more of the forest, keeping within auditory distance of each other. At one point, someone in the group spots a footprint in the mud and calls Elmo and the other companions to cover that patch of the forest. Before long, another member of the band sees a deer at a distance. He whistles to the others, who quietly come closer. The group encircles the deer and shoots it. Everyone then walks back to the village before sharing the meat equally among each other.

In what follows, we shall describe some striking functionalist and phenomenological parallels between the relationship between Elmo and the rest of the group and that between Otto and his notebook. Our argument is that these parallels are not coincidental. Rather, we hope to show that the 'mental extension' illustrated in the latter case is a token of the capacity for collective intentionality, as illustrated in the former.

3.1. The functionalist parallel

From the viewpoint of the individual in question, both cases present a task-dependent action (reaching MoMA; hunting a deer) which an external part of the world helps achieve (notebook; other hunters). If for Otto the notebook contains information needed to reach MoMA, so for Elmo the band of hunters furnishes the necessary information needed to find the deer. Both cases can be taken as illustrations of the parity principle central to the extended mind thesis. Standing by the original functionalist premise of that principle, both the group of hunters and the notebook replace the individual's cognitive function and could thus be considered part of the individual's mind.[1] The specific cognitive function being replaced is different in the two examples, but this is immaterial. Otto's notebook replaces biological memory whereas the other hunters replace Elmo's perceptual system - the parity principle holds for both.[2] In sum, we suggest a homology between the Otto and Elmo cases, with the very important difference that in the Elmo case the

[1] See Szanto (2013) for a similar application of the parity principle to collectives.

[2] The hunting case also illustrates the 'principle of cognitive impartiality' (Clark, 2007:174), according to which brains typically adopt the most cost-effective task performance strategy regardless of whether, in doing so, they rely on internal or external resources. This is a principle that accords with the grounding premise of evolutionary biology: as a number of evolutionary theorists have argued (e.g. Hrdy 2011; Kropotkin 2021[1902]; Tomasello 2019), to effectively pursue certain tasks, off-loading cognitive action to other humans manages cognitive demand with adaptive efficacy.

accomplished task is not individual but the product of collective intentionality. This is an action that entails simultaneous mind extension from every other hunter: from the point of view of each hunter, the rest of the band functions as the 'part of the world' that helps achieve the goal of hunting the deer. Take your task to be collective rather than individual, and it is apparent that mind extension to other human beings, in the very way laid out by Clark and Chalmers, is the necessary process for achieving that task. In sum, the capacity for collective intentionality is enabled by the same process that underpins classic examples of the extended mind.

At this point, a skeptic will raise the tightening conditions required for an external resource to be considered part of an individual mind: "a high degree of trust, reliance, and accessibility" (Clark and Chalmers, 1998:17). We are drawn to this homology by the recognition that it is exactly these tightening conditions that are, with large measures of congruity, the same necessary conditions for collective intentionality in the example above. 'Trust' is arguably the most fundamental. According to Clark, in order for an artefact to be part of one's cognitive system, it needs to contain information that is "more or less automatically endorsed. It should not usually be subject to critical scrutiny (unlike the opinions of other people, for example)" (Clark 2010:46). Trust is clearly not at play in the act of probing someone's opinions – quite the opposite. But in acts of collective intentionality, trust is indeed the attitude we adopt towards other agents. It is the *sine qua non* condition of collective intentionality, since, for group action to take place, each member must automatically assume that the other members have the same intention and same goal (Searle 1990; Tuomela 2007; Bratman 2013). Philosophers have variably used the terms 'mutual awareness' or 'acceptance' to the same effect. As Schmid (2013) points out, this is the essence of 'mutual trust'. In venturing into the forest, Elmo automatically endorses the collective intention of hunting the deer that, he trusts, is also held by the rest of the band. The difference between the Otto and Elmo cases lies in the fact that while in the former trust applies to the information contained in the external resource that is necessary to achieve a task, in the latter trust applies to the intention of other people of processing information that is necessary to achieve a task.

As we move on to consider the conditions of 'reliance' and 'accessibility', we notice that these too are key conditions for collective intentionality. Otto's task is to retrieve information which he has been presented with in the past, and for whom the notebook therefore replaces biological memory. The reliance and accessibility conditions are ultimately only further specifications of the parity principle. They state that Otto should have *access* to the notebook and *endorse* the information it contains similarly to how another agent would have access to and endorse information retrieved from well-functioning biological memory. Elmo and his fellow hunters are engaging in a different task: their goal is to acquire new information on the location of a deer.

Their task does not involve information retrieval but rather information acquisition and processing. Consequently, 'reliability' and 'accessibility' pertain in this case to the acquisition and processing of perceptual information. Were Elmo hunting alone, he would rely on his own eyes to detect footprints, and upon perceiving one, would infer that a deer is nearby and move in the adequate direction. When Elmo is hunting with his companions, all that is required for the reliability and accessibility conditions to be met and is that his companions, upon perceiving a footprint, process and communicate this information to Elmo so that he can behave in the same way as he would had he perceived it himself. Finally, there is a need for the collective intention to be "consciously endorsed at some point in the past", which is the fourth tightening condition. Elmo accomplishes this when he agrees to go hunting with the other people in the band.

If all this still fails to convince, consider briefly a high-tech version of Elmo, equipped with a remote-controlled drone. Video recorded from the drone's camera is broadcast in real time to a headset that Elmo is wearing, such that Elmo effectively sees the world through the drone's camera. In addition, a computer-vision algorithm highlights any deer in the video feed. Proponents of the extended mind thesis should readily agree that the parity principle and tightening conditions are met in this example, and thus, that Elmo's mind extends to the drone. We propose that, in the low-tech example, other hunters are functionally equivalent to the drone and that Elmo's mind must therefore extend to them.

Having set out our argument for a functionalist parallel between the extended mind and collective intentionality, we now turn to our second, and more important, phenomenological parallel. Indeed, we only used the functionalist parallel to make the case for a homology between classic extended mind examples and one particular case of collective intentionality. But as we show in the following section, it makes more sense to see classic cases of extended mind as particular cases for the capacity of collective intentionality. This is because the same fundamental type of coupling persists even if we drop the conditions of reliability, accessibility, and even the parity principle – conditions that were originally conceived only to fit examples of humans-artefacts extension specifically.

3.2. The phenomenological parallel

Central to our phenomenological argument is the notion that, much like the notebook is to Otto, other people during acts of collective intentionality are, in some important respects, 'transparent' to the subject. In trusting that other people hold the same intention as we do, we do not perceive the other as an object of inquiry, much like Otto doesn't need to critically reflect about the

notebook as he uses it. As phenomenologists put it, interpersonal trust in acts of collective intentionality affords a degree of 'mutual incorporation' (Fuchs and De Jaegher 2009). This idea is grounded on a recent turn to phenomenological approaches to intersubjective understanding that challenge the dominant way in which cognitive science accounts for how we relate to one another. The latter assumes an epistemic gulf between intentional subjects, which can only be bridged by producing a 'theory of mind' in order to ultimately predict the other's behaviour. Views differ on how we do so – do we 'mindread' by inferential interpretation (Leslie et al. 2004), or by subpersonally mirroring (Gallese 2013) or simulating (Goldman 2006) the other's mental states in ourselves? – but the starting individualistic premise is the same in all these cases. The framing is of two or more individuals trying to internally represent and predict the other's mental states. The opaqueness of the other is taken for granted.[3]

The problem faced by 'theory of mind' or 'simulation theory' approaches lies in their non-representativeness. There is no question that some face-to-face encounters, or some aspects of face-to-face encounters, can take this form; especially so when people vie for power, or where there is a power asymmetry between two subjects (a case where much of the interpretative work in figuring out the other's intention falls on the subordinate) (see Dennett 1989:49). But this Machiavellian picture is hardly reflective of the experience of engaging with other people ordinarily and on an equal basis, especially when we have interest in doing so. Defenders of 'theory of mind' would respond by saying that this proposed mechanism operates at the subpersonal level, rather than at the level of experience, but there is something about the experience of intersubjectivity that demands more recognition than what any of these accounts can offer. Crucially, like a collective hunt or a theatrical performance, a face-to-face encounter cannot arise through a summation of individuals independently predicting each others' mental states. In order for a face-to-face encounter to arise at all there needs to be a we-intention of having and pursuing that encounter in the first place. It is the collective intention of having the encounter that produces the kind of alignment, coordination and synchronization of verbal expressions between the participants that are constitutive of mutual understanding. Phenomenologists speak here of 'empathy' (see Stein 1989), or, indeed, of 'mutual incorporation', a situation where the other appears as a transparent extension of one's own intentions (Fuchs and De Jaegher 2009). Understood this way, we access the minds of others directly, through their expressive behaviour and meaningful actions (Zahavi 2014); 'mind-reading' might play a minor role but is itself dependent on the collective intention of having the encounter. Human social cognition is

[3] Not so different in this respect is the approach that Dennett (1989) dubbed as the 'intentional stance'

embedded in shared cognition. If this is a salient in dyadic interactions, i.e. in the experience of *being-towards-others*, it is even more so in the experience of *being-with-others* – in acts such as play, rituals or any kind of engagement in collective action or projects. Here, other companions are not experienced as opaque entities whose intentions or movements need to be figured out, but as a transparent extension of one's own activity (Seemann 2009).

Importantly, in suggesting that the mind extends to other people in acts of collective intentionality we argue something different than what Gallagher (2013) does for his "socially extended mind". As we noted above, Gallagher's framework is individualistic: other people figure as tools to achieve an individual rather than collective task. In his example, the legal system is a cognitive institution that facilitates one's actions in the world and is, for this reason, according to Gallagher, constitutive of one's mind. To make his case, Gallagher is forced to considerably relax Clark and Chalmers' tightening conditions. It would be hard to make the case, for instance, that a person "automatically endorses" the information given by the legal system: usually the engagement with the legal system entails critical reflection. For Gallagher (2013:2-3), 'critical reflection' is simply more cognition that is added to the overall extended cognitive process. The problem, here, is that critical reflection gets in the way of 'trust', a fundamental condition for coupling in the original extended mind thesis. It also precludes drawing the homology at the phenomenological level because reflection is antithetical to transparency. By contrast, it is much easier to draw the homology in the context of collective tasks, where the "glue and trust" (Clark, 2010:83) conditions vis-a-vis other people are persuasively met, and where the experience of the transparency of others is central.[4] Of course, this transparent relationship of trust can break when individuals manifestly deviate from the collective intention. But the same break of reliable coupling can happen to Otto's notebook, if this is lost or ruined, or, for that matter, as extended mind theorists like to point out, to internal cognitive functions.

3.3. Parallels beyond problem solving

In making the case for the functionalist and phenomenological parallels between the extended mind and collective intentionality, we focused on problem-solving tasks that dominate the discussion in the original paradigm. However, the same conditions of coupling, and therefore mental extension, apply to phenomena that go beyond the kind of problem-solving tasks

[4] It is interesting to note that the idioms used to define the coupling between a subject and epistemic artefacts – "glue and trust" (Clark, 2010:83) – are very similar to the idioms used to define human-human coupling in joint action in the collective intentionality literature, strengthening the case for the homology (Tuomela even employs the same adhesive metaphor when describing participants in joint acts as "glued together" (2013b:22)).

exemplified by the Otto and Elmo cases. Collective intentionality, after all, also encompasses a wide range of phenomena in which other people are not typically conceived as 'tools' to solve a problem (Gilbert 2015; Huebner and Hedahl 2018). Think of rituals, play, convivial conversations, political projects, or the pursuing of collective interests. These cannot be framed as problems to be solved. Consider the 'task' of dancing together. Other people are typically thought as the very reason why we do it, not tools that allows us to do it. The parity principle simply does not apply here because other people do not replace any internal function. Yet, the kind of coupling involving interpersonal trust that typifies all the cases explored previously is exactly the same. Because the extended mind literature has focussed on artefacts, which are usually employed to solve practical individual problems, it has glossed over how the very same process of mind extension can be at play in non-problem-solving actions with other people.

Non-problem-solving actions come in a high variety, ranging from joint action with concrete partners one knows (e.g. dancing) to participation in temporally extended collective projects where the mind extends to a whole community. There is certainly room for analysing the differences among all these cases of collective intention in fine-grained phenomenological detail. But what's striking about the collective intentionality debate is that it regards this constellation of phenomena as being regimented by the one same principle: intending a joint activity with shared goals and norms under the assumption that everyone in the group or community does (Gilbert 2015; Butterfill 2017; Tomasello 2019). What we have argued so far is that this assumptive attitude is homologous to the attitude that allows mental extension to epistemic artefacts in classic examples of the extended mind. It is therefore misleading to think of mental extension to epistemic artefacts as a special phenomenon, and it is arbitrary to separate the debate around it from the collective intentionality literature. This is all the more so because, as we argue in the next section, the capacity for the former depends, in some important respects, on the capacity for the latter.

4. The primacy of collective intentionality

4.1. Cognitive development

One of the main reasons why cognitive extension to other people has primacy over individual-artefact extension is evolutionary and developmental. As noted earlier, evolutionary accounts of human cooperation illustrate that the ability to perform important tasks relies on cognitive outsourcing of epistemic action to others under conditions of trust (Tomasello, 2019). For a number of evolutionary theorists, the coupling involving interpersonal trust that allow these

actions – i.e. collective intentionality - is the true mark of humankind (Hrdy, 2009; Tomasello, 2019). Tomasello suggests that this had an evolutionary origin in contexts of food scarcity that required active coordination among members of a group to hunt. For two hunters to capture a stag, they both have to individually have the goal of capturing the stag with the other, and, crucially, they have to have mutual knowledge of the other's goal and awareness that the collective goal can be attained through different individual roles. Agents thus began to relate to each other not only as independent agents, but also as an 'I' to a 'you' in the context of our 'we'. Tomasello argues that distinctively human capacities such as language and morality develop out of this shared point of view. It is only on the basis of joint commitment, which presupposes a shared and equal status between subjects as intentional beings, that the normative idea of fairness can arise. Thus, it feels right for a hunter to share the meat with others, even if he could stealthily keep it to himself. As Engelmann and Tomasello (2018) explain through a range of empirical evidence, all of this is absent among other primates, which are capable of forms of social cognition based on sympathy and strategic cooperation, but not on joint intentionality and commitment.

Ontogenetically speaking, the rudiments of this uniquely human capacity emerge around 9 months of age, when babies point out objects to other persons with no motive other than the recipient share in the baby's attention (something that chimpanzees don't do), and develops later into the capacities of recursive mind-reading, learning by instruction, language and fully-fledged understanding of the social world. Crucially, collective intentionality is not added on top of other individual cognitive skills – it "shapes cognition all the way down" (Kern and Moll 2017:327; see also Tomasello, 2019). It permeates the human individual's reasoning and engagement with the world as a whole and not just how they socially interact with others. Evidence for this claim comes from a suite of empirical studies demonstrating that a person's interaction with the environment is transformed by collective intentionality even outside joint activities with others (Tomasello, 2019). Children do not develop their technical know-how by trial and error or solipsistic hypothesis testing. Their instrumental rationality is shaped by acts of shared agency with adults who show them how to use and craft tools and address instrumental problems. As Moll et al. (2020:172) put it, "human sociality is irredeemably written into humans' technical capacity". The capacity for extending the mind to tools around us, let alone to epistemic artefacts like a notebook (which requires language), is premised on the primal capacity for extending the mind to other humans in acts of shared intention.

A second developmental sense in which collective intentionality has primacy is phenomenological, though here the primacy is not over individual-tool use per se but over the overall individualism that underpins it. As Welsh (2013) suggested, early infancy is not defined by

a well-demarcated individuality but by an 'anonymous collectivity' with no distinction between you and I, and it is thanks to the capacity of sharing an experience, later reinforced through language, that children learn to differentiate themselves from others in the same sphere of experience. In short, the 'we' precedes the 'I'. At a more general level, phenomenologists argue that it is wrong to consider human beings in isolation from the web of social relations they are immersed in because sociality already permeates the world in which we live and act in: the world is intelligible to us the way it is because of our tacit conformity to public norms, and it is this reliance on our shared social background that allows us to be human at all. Essentially, the social world is not something that we make out with our pre-existing cognitive tool-kit, but something that encompasses and shapes this tool-kit. This is what early phenomenologists meant by expressions such as the "ontological primordiality of intersubjectivity" (Husserl 1970), the "*being-with-others* as the existential human condition" (Heidegger 1962) or the "vivid simultaneous presence" of the 'we' in our conscious stream (Schütz 1967). (Incidentally, early phenomenologists also came up with analyses of collective intentionality that foreshadowed the more recent analytic philosophy debate on the subject, and did so, arguably, in more sophisticated ways (Szanto and Moran 2015)).

4.2. Diachronic cognition

The primacy of collective intentionality is also supported by a diachronic conception of cognition. This is fundamental, as it links back directly to the extended mind debate. We are going to engage, here, with what's perhaps the most interesting development in that debate in recent years: Kirchhoff and Kiverstein's *Extended Consciousness and Predictive Processing* (Kirchhoff and Kiverstein 2019). In this book and other publications (2019, 2020, 2021), Kirchhoff and Kiverstein have put forward two important claims about the extended mind: that consciousness as well as cognition extends beyond the brain; and that the increasingly powerful framework of predictive processing is compatible with, and actually mandates, the extended mind (*pace* views to the contrary (Hohwy 2013)). For our purposes, what is important to discuss is an interpretation of cognition Kirchhoff and Kiverstein adopt in making their arguments.

Kirchhoff and Kiverstein note that, as a consequence of being focused on individual problem-solving, commentators on the extended mind debate have approached cognition in a strictly synchronic way. They have treated cognition as if made up of processes that unfold in a linear and stepwise manner over short timescales. Maintaining focus on the here and now, they have considered external elements as constitutive of the cognitive system only if these are wholly present at each instant that the system exists. The implication of this framing is that a person's

history of engagements with cultural practices gets screened off from the analysed process. On a synchronic reading of cognition, cultural practices appear to merely 'set the scene' (Clark 2011:459) for cognitive processes to take place; they are internalised over development and so play an important causal role, but they are not constitutive of the cognitive process in question.

Kirchoff and Kiverstein disagree, but do so by introducing an element that overhauls the terms of the debate. Drawing from an important essay by Van Gelder and Port (1995), they bring forth the challenge that the metaphysics of cognition is intrinsically temporal. The interaction of the organism with the environment in cycles of action and perception is a dynamic process that unfolds over multiple interacting timescales in a way that does not warrant the privilege that the cognitive sciences place on the synchronic here and now. Of course, problem solving can take a synchronic form and can be analysed this way, but this should not rule out cultural practices unfolding over longer timescales from playing a role in the material constitution of people's process of thinking. Based on a *diachronic* understanding of cognition that Kirchhoff and Kiverstein suggest, history and culture are always carried along in the practices and artefacts we engaged in, and entrain what individuals do in faster timescales. Extend the temporal scope of cognition, and the mind comes to encompass wider aspects of one's cultural environment. Consider, for instance, a child that uses pen and paper to do multiplication. Clark would concede that the child makes use of the external scaffolding of pen and paper but that this extension is only temporary and limited to material tools. The dispositions that enable the child to do multiplications are fully internalised. "But to say that a disposition is internalized," Kirchhoff and Kiverstein contend, "is not at all the same as saying that what people know when they take part in cultural practices is fully internalized" (2020:6). The actions she performs are embedded in and organized by the practice of which they are part.

We wish to elucidate what Kirchhoff and Kiverstein mean here by noting that, in fact, the synchronic/diachronic distinction does not tease apart two genuinely different types of cognition. This is because even the problem-solving tasks most amenable to a synchronic reading fundamentally depend on a temporal dimension for their achievement. Consider Otto. However advanced his Alzheimer's, in order for the notebook to play a functional role in his cognitive system, Otto must retain some memory of reading up that MoMA is on 53rd street, otherwise he would never reach his destination. He would constantly look up the notebook and, forgetting what he read straight away, wouldn't go anywhere. This means that one of two propositions must be valid. Either there is an intrinsically diachronic aspect to any act of cognition, including that performed by Otto with his notebook, or the notebook merely 'sets the scene' for the actual (synchronic) cognitive moment. The second proposition would invalidate the central claim made

by proponents of the extended mind. But if the first proposition is correct, as we take it to be, it is arbitrary to include in cognition the notebook but not wider, longer-term cultural practices.

Of course, philosophers who consider 'memory' to be a system of stored representations about the world would take this as a blow to the whole idea of the extended mind because it suggests internalism. But Kirchhoff and Kiverstein's (2019; 2020) position – in line with sensorimotor enactivism – is that the extended mind thesis only holds on non-representationalist grounds, and, more to the point, that it applies to conscious experience (Di Paolo 2009; Silberstein and Chemero 2012; Ward 2012). From an enactive perspective, conscious experience is a process that emerges in interaction with the environment to which the brain is coupled through cycles of action and perception. Kirchhoff and Kiverstein add that this relation of coupling with the cultural environment is one of ongoing 'phenomenal attunement' (2020:2). This is an experience that cannot always be generated solely out of processes unfolding inside a person's brain. To make their point, they ask us to consider the negative corollary of 'phenomenal attunement', which they identify as the experience of 'cultural shock', a situation in which someone is suddenly moved into an unfamiliar cultural environment and experiences alienation as a result. It is impossible, Kirchhoff and Kiverstein argue, to explain such experience only by looking at the person's neural states; it can only be explained by considering the familiar cultural environment as constitutive of the person's conscious experience. Would a person's neural duplicate in a different environment feel the same experience? The answer is negative because this thought experiment is premised on a synchronic reading of cognition and is thus an impossibility. For two people to be neural duplicates, they must also be environmental duplicates. The mind, in short, is partially constituted by cultural practices.

Our comment on Kirchhoff and Kiverstein is that their main argument would gain more traction if placed within the framework of collective intentionality. Bringing in collective intentionality throws the point Kirchoff and Kiverstein are making into sharp relief, with no shortage of more intuitive examples. Consider, for instance, the action of writing a philosophy paper. There is no question that a set of dispositions necessary for doing so (the basic ability to write, argumentative skills, knowledge of various debates) have been internalised over time. Aside from loops of action and perception between brain, upper limbs and laptop, the activity of writing the paper is at first glance all internal. It isn't so, however, once we consider the activity as derivative of the longer-term engagement with the philosophy community (or at least a subset thereof) of which one is part and considered to be part by other members - people who may read the paper and who shares the group ethos. The fact that we are part of this community is constitutive of the act of writing: if we knew that someone else would hijack the authorship of the

paper, thereby violating the premise of collective intentionality, we probably wouldn't write it; if forced to do it under this condition, the experience of writing it would be entirely different. So, we do write with other people 'in mind'; the process of writing the paper hinges on *trusting* the concrete external presence of a philosophical community that entertains shared values and norms. Following the line of reasoning we have undertaken it is legitimate to say that our mind extends to the philosophy community. Take this community away, and the practice vanishes (like Otto's notebook and Otto's capacity to reach MoMA). There is thus a difference between internalised dispositions, which are internal, causal, and cannot be cleared away, and people out there in the world, whose constitutive presence in collective intentions can, hypothetically, be removed asudden. We think that this is what Kirchhoff and Kiverstein mean when they say that dispositions, while internalized, are constrained by norms, rules and principles that operate at the scale of cultural practice which the person must be attuned to.[5]

What Kirschhoff and Kiverstein call 'phenomenal attunement' essentially equals trusting that we are part of a community that share a same set of intentions, whether in the forms of goals, values, norms or emotions. We add that like epistemic objects which must be accessed 'on the fly', so the presence of other people in acts of collective intentionality is automatically assumed. Also, like epistemic objects (Clark 2017), other people holding shared intentions become part of our cognitive system in a flexible and context-dependent manner. My philosophy interlocutors slide away from my mind as I move from writing the paper to playing computer games. The one main difference between epistemic objects in individual tasks and other humans in acts of collective intentionality is that the latter coupling entails a bi-directional impact because a community is not

[5] As an aside, we think that the terms 'culture' and 'cultural practice' get in the way of Kirchhoff and Kiverstein's message because they lack concreteness (see also Facchin and Negro (2020)). They are also not very well defined. While the notion of 'culture' has been endlessly debated in the field of anthropology, it is often employed somewhat uncritically in the cognitive sciences. What the concept usually evokes is a set of meanings, symbols or ideas that is held by a people. 'Cultural practice', in turn, means a practice that employs this set of shared meanings. This conception is certainly close to, and likely derivative of, the definition of 'culture' developed within American anthropology. Geertz defined the latter as "a system of inherited conceptions expressed in symbolic forms by means of which men communicate, perpetuate, and develop their knowledge about and attitudes toward life" (1973). But because 'culture' is about "shared symbols and meanings", anthropologists have often had good reasons to treat it separately from norms or values (concepts that, we note, are central to discussions of collective intentionality). For Schneider, culture "contrasts with norms in that norms are oriented to patterns *for action*, whereas culture constitutes a body of definitions, premises, postulates, presumptions, propositions and perceptions about the nature of the universe and man's place in it" (1976:202; original emphasis). Because of its static overtones, a number of anthropologists (particularly in the British tradition) have rejected the usefulness of the concept in capturing social reality. By contrast, the original notion of *kultur* in the German anthropological tradition meant something altogether different than in the American tradition. For Herder (2002[1772]; 2004[1774]) 'culture' was essentially a field for the pursuit of certain forms of value, values that shaped humans into organisms whose very perceptions were attuned largely to that pursuit (in this sense, we note, the idea of *kultur* is much closer to the concepts of social action and values central to the collective intentionality literature). This brief excursus shows that a) the concept has been polyvalent throughout history and b) has sometimes been rejected because of its failure to capture collective action. As such, it might be inadequate for effectively conveying Kirchhoff and Kiverstein's ultimate claim.

an inert object; it embeds its participants in social practices of coordinated behaviour (Slors 2020). This is what Kirchhoff and Kiverstein mean when speaking of "the co-constitution of agent and environment" (2020:2). Narratives, for instance, play a fundamental role in establishing, not always explicitly, shared norms and values and in cementing a collective ethos (Gallagher and Tollefsen 2019). In short, we do accomplish an action like writing a philosophy paper with other people 'in mind'.[6] The upshot of all this, to anticipate our final section, is that any potential scientific study of my act of writing the paper cannot solely focus on the interaction between brain and computer as if unplugged from the structure of collective intentionality it is placed in. Doing so would abstract away from the overall wider process of cognition, yielding a selective view of mental life.

5. Implications

If what we have suggested is right – if the mind does extend to other people in acts of collective intentionality – then there are important implications for how cognitive science is to proceed. Challenging and expanding the unit of analysis of the cognitive science has been a central consequence and driving thrust of the extended mind thesis. It has also been an oft-stated intimation of phenomenology. According to contemporary proponents, the very same concepts and constructs that are used to describe experience should be central to a proper scientific understanding of the mind (Thompson 2007). If collective intentionality – the condition of *being-with-others* – is the typifying feature of human experience, then it should figure prominently as a concept in how the cognitive sciences study the human mind. What, then, would a cognitive science that takes this concept to heart look like?

It would look a lot different than it does now. Current cognitive science is grounded in methodological individualism and has been almost entirely divorced from the philosophical discussion on collective intentionality outlined above (attempts at drawing links have been few, e.g. Gallotti and Frith 2013). Moreover, the absence of the concept is most obvious in the context where it should play the most central role, namely in social cognition (e.g. in 'theory of mind' debates, as we have shown above). This individualism has informed a kind of methodologism (Teo 2009) whereby cognitive scientists – especially cognitive psychologists – have focussed on the experimental methods of inquiry, rather than selecting suitable methods for the topics and research

[6] In challenging the idea that the 'extended mind' runs counter to common sense understandings of the mind, Clark points out that "the folk grip on mind and mental states […] is surprisingly liberal when it comes to just about everything concerning machinery, location, and architecture" (Clark 2008:105-6). It's worth noticing that the common sense understanding of the mind, even in Western individualistic societies, allows, and, if anything, reinforces, the idea that other people can be part of one's mind.

problems under investigation. These experiments tend to represent uncommon, socially-isolated and experimenter-defined tasks that do not acknowledge structures of collective intentionality within it (let alone the structure of collective intentionality implicated in the experiment itself: as Vygotsky among others pointed out, psychology experiments describe a particular social situation of their own, where subjects share the intention of performing the experiment with the scientist, in line with the norms that regulate the experiment.)

There is one important domain of cognitive science that has incorporated the idea of collective intentionality. This is the domain of joint action. Research in the area is relatively scanty because it does not easily lend itself to laboratory experimentation, but it is significant, if anything because it has offered a proof of the neurocognitive signature of collective intentionality.

Consider Loehr et al. (2013)study of musical ensemble studies, which investigates the cognition of a pianist who produces tones in the course of playing a duet with another pianist. In this typical experiment, there is an outcome to which the pianist's action is directed, the production of a tone or melody; and there is an outcome to which her and her partner's actions are collectively directed, the production of a combination of pitches or harmony. Loehr et al (2013) asked the following question: do pianists monitor their own or their partner's actions with respect to individual action goals (those necessary to achieve each individual's part of the joint action) or with respect to shared action goals (the combined outcome of their coordinated actions)? A result that points towards the first hypothesis would support the premises of individualistic social cognition (e.g. Friston and Frith 2015), while a finding that aligns with the second hypothesis would essentially find a cognitive signature of collective intentionality. One way to investigate this question involves covertly introducing errors. Loehr et al (2013) contrasted two kinds of error: those which were errors relative to the goal of an individual pianist's actions (the pitch) but not relative to the collective goal of the two pianists' actions (the harmony); and those which were errors relative to both. They found neural signatures for both kinds of errors in expert pianists (i.e. pianists were sensitive not only to deviations of the self and other from the desired sound, as would be expected from an individualistic social cognition perspective, but that they also were sensitive to deviations from the joint product in the desired sound). This is evidence that duetting pianists do indeed maintain collective goals. As the authors conclude: "[the] findings indicate that people monitor not only their individual contributions to a joint action, but also their partner's actions and the combined outcome of their coordinated action. […] Successful joint action relies not only on monitoring one's own actions but also the shared goal of coordinated actions" (Loehr et al 2013).

Maintaining a shared goal – rather than individualistically predicting and responding to one's action – has been found to be cognitively parsimonious: it reduces the space of possible predictions about a partner's action (Pesquita et al. 2018). As Butterfill puts it:

> "There are not two processes but one. Anticipation of another's actions and preparation for your own are not two separate things. They are parts of a single process in the same sense that, in preparing to perform a bimanual action, preparation for the actions to be performed by the left hand and anticipation of the movements of the right hand are parts of a single process" (Butterfill 2018:75).

Phenomenologists once again speak aptly here of 'mutual incorporation', or of 'transparency' of other subjects, one that can admit a variety of degrees.

Studies of joint action such as that of the musical ensemble are the very few experiments that find the cognitive signature of collective intentionality: they show that it would be inaccurate to study the cognition of the duetting pianist without factoring in shared intentions. But a key claim of our paper is that, in an important sense, joint action is human cognition writ large – that there should be something equivalent to a shared harmony in most human thought and action in the form of shared goals, norms, values, etc. *And just as it would be inaccurate to consider the duetting pianist's cognitive process without taking into account their shared intention – the harmony – so it is inaccurate to study most cognitive process without factoring in the goals, values, norms, etc that are shared with relevant persons.* A cognitive science that studies cognition as if unplugged from the shared intentionality it is situated in can only offer an impoverished, if not mystified, view of the human organism, one that is especially problematic when it spills over into popular discourse, thereby feeding into individualistic ideologies (Hacking 1995; Smith 2013; De Vos 2013).

The methodological consequences of taking collective intentionality seriously would likely result in a rebalancing of methodological approaches in cognitive science, that arguably reflect its purported interdisciplinary identity. If the mind does extend to others in acts of collective intentionality, and this process informs even individual cognition, then it is likely that experimental approaches would be restricted to answering specific questions that arise from other, more ecologically valid, approaches. One ecologically valid approach to studying cognition that accords with taking collective intentionality seriously is Hutchins' (1994) cognitive ethnography. Here, one makes accurate records of the cognitive aspects of specific instances of human behaviour, using

wider ethnographic observation to inform such data collection. Hutchins has successfully used this approach to develop a rich analysis of cognition that is distributed across place, people, and time – i.e. diachronic – in various environments such as naval ships (Hutchins, 1994) and airplane cockpits (Hutchins 1995). Such an approach has the potential to provide a more accurate functional specification of human cognition, and inform both experimental studies and human systems design.

The incorporation of the idea of collective intentionality into cognitive sciences should also speak directly to researchers of affect or emotion, who have always maintained that there is no such thing as affectless cognition (Colombetti 2017). A significant variety of human experiences – from abject misery to bliss – depends in some fundamental way on the structure of collective intentionality the person is immersed in. In Searle's (1990) classic example, running towards the shelter as part of a theatrical performance is affectively different than running towards the shelter on one's own, even if the outward behaviour is the same. In the first case there is a felt 'sense of 'we", a sense of 'being together' (Searle 1983). The same point can be made by using a negative example. For instance, the work produced by a slave, or by most waged workers in capitalist systems, is only in a reduced sense perceived as a contribution to a collective goal, intention, or value because it is expropriated by the owner. It is precisely for this reason that the experience of performing this work is affectively different than the experience of performing the very same kind of work as part of a community in which that work is perceived as a contribution to a collective goal, intention or value. The concept of 'alienation' – employed by Kirchhoff and Kiverstein in their (somewhat unwieldy) example of 'cultural shock' – was originally used to describe the first (much more common) kind of experience (Marx 1964; see also Graeber 2001). Our argument here accords with recent work within the field of 'situated affectivity' and with recent calls for 'extended emotions' (Slaby 2014; Krueger and Szanto 2016; León et al. 2019). As we did with Kirchhoff and Kiverstein, we suggest that arguments for extension gain more traction when placed within the framework of collective intentionality.

Finally, there are important methodological implications for neuroscience. If the mind extends beyond the skull, then looking at the brain will only tell us a partial story about the mind. The upshot is that the brain should be conceived as a 'mediating organ' (Fuchs 2011) that enables our interaction and coupling with the rest of the (social) world, but cannot, on its own, tells us about consciousness and intentionality, as they are merely processes arising out of this interaction. It certainly remains one of the most fundamental scientific quests to understand how the brain enables this interaction.

Conclusion

The extended mind thesis suggests that the mind spills over into body and world. This has stirred important debate on the proper unit of analysis in the cognitive sciences. In this paper, we argue that the philosophical debate on the extended mind has glossed over the most important way the mind extends: i.e., to other humans in acts of collective intentionality. In arguing that collective intentionality holds primacy over epistemic tool use, we call for a reversal of the individualistic view that sees cognitive extension as something that takes place only in special circumstances. It is not that parts of the world are occasionally and ephemerally involved in one's cognitive system, but the other way around: it is the brain-bound severance from the social world that turns out to be occasional and ephemeral (and something that perhaps only specific cultural practices try to achieve). Typically, we think and act with other people 'in mind', and that's why the extended mind is nothing special but is central.[7] A cognitive science that takes this principle to heart would transform into an altogether different kind of cognitive science than the one we currently have.

[7] The title of this paper and overall style of argumentation is inspired by Bloch's 'Why religion is nothing special but is central' (2008)

Bibliography

Adams, F. (2019). The Elusive Extended Mind: Extended Information Processing Doesn't Equal Extended Mind. In *Andy Clark and His Critics* (pp. 21–31). Oxford University Press.

Adams, F., & Aizawa, K. (2001). The bounds of cognition. *Philosophical psychology*, *14*(1), 43–64. https://doi.org/10.1080/09515080120033571

Adams, F., & Aizawa, K. (2009). Why the mind is still in the head. In M. Aydede & P. Robbins (Eds.), *The Cambridge Handbook of Situated Cognition* (pp. 78–95). Cambridge: Cambridge University Press.

Adams, F., & Aizawa, K. (2010). Defending the Bounds of Cognition. In R. Menary (Ed.), *The Extended Mind* (pp. 67–80). The MIT Press. https://doi.org/10.7551/mitpress/9780262014038.003.0004

Allen, C. (2017). On (not) defining cognition. *Synthese*, *194*(11), 4233–4249. https://doi.org/10.1007/s11229-017-1454-4

Bloch, M. (2008). Why religion is nothing special but is central. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *363*(1499), 2055–2061. https://doi.org/10.1098/rstb.2008.0007

Bratman, M. E. (1999). *Faces of intention: Selected essays on intention and agency*. Cambridge University Press.

Bratman, M. E. (2013). *Shared agency: A planning theory of acting together*. Oxford University Press.

Butterfill, S. (2017). Joint Commitment: How We Make the Social World, written by M. Gilbert. *Journal of moral philosophy.*, *14*(4), 475–478. https://doi.org/10.1163/17455243-01404004

Butterfill, S. (2018). Coordinating joint action. *The Routledge handbook of collective intentionality*, 68–82.

Clark, A., & Chalmers, D. (1998). The Extended Mind. *Analysis*, *58*(1), 7–19. https://doi.org/10.1093/analys/58.1.7

Clark, A. (2007). Curing Cognitive Hiccups: A Defense Of The Extended Mind. *Journal of Philosophy*, *104*(4), 163–192. https://doi.org/10.5840/jphil2007104426

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

Clark, A. (2010). Memento's Revenge: The Extended Mind, Extended. In R. Menary (Ed.), *The Extended Mind* (pp. 43–66). The MIT Press. https://doi.org/10.7551/mitpress/9780262014038.003.0003

Clark, A. (2011). Finding the Mind: Book Symposium on Supersizing the Mind: Embodiment, Action, and Cognitive Extension (Oxford University Press, NY, 2008). *Philosophical Studies*, *152*(3), 447–461. https://doi.org/10.1007/s11098-010-9598-9

Clark, A. (2017). How to Knit Your Own Markov Blanket:How to Knit Your Own Markov Blanket:: Resisting the Second Law with Metamorphic Minds: Resisting the Second Law with Metamorphic Minds. *Philosophy and Predictive Processing*. https://doi.org/10.15502/9783958573031

Colombetti, G. (2017). Enactive Affectivity, Extended. *Topoi*, *36*(3), 445–455. https://doi.org/10.1007/s11245-015-9335-2

Crippen, M., & Schulkin, J. (2020). *Mind Ecologies: Body, Brain, and World*. New York: Columbia University Press.

De Vos, J. (2013). *Psychologization and the subject of late modernity*. Basingstoke: Palgrave Macmillan.

Dennett, D. C. (1989). *The intentional stance*. MIT press.

Dewey, J. (1931). Qualitative thought. In J. Dewey (Ed.), *Philosophy and civilisation* (pp. 93–116). USA: G. P. Putnam's Sons.

Dewey, J. (1982). The pattern of inquiry. In H. S. Thayer (Ed.), *Pragmatism: the classic writings* (pp. 316–334). USA: Hackett Publishing Co, Inc.

Dewey, J. (1995). *Experience and nature*. New York: Dover Publ.

Di Paolo, E. (2009). Extended Life. *Topoi*, *28*(1), 9–21. https://doi.org/10.1007/s11245-008-9042-3

Dreon, R. (2021). Framing cognition: Dewey's potential contributions to some enactivist issues. *Synthese*, *198*(1), 485–506. https://doi.org/10.1007/s11229-019-02212-x

Engelmann, J. M., & Tomasello, M. (2018). The middle step: Joint intentionality as a human-unique form of second-personal engagement. In *The Routledge handbook of collective intentionality* (pp. 433–446). New York, NY, US: Routledge/Taylor & Francis Group.

Facchin, M., & Negro, N. (2020). Predictive processing and extended consciousness: why the machinery of consciousness is (probably) still in the head and the DEUTS argument won't let it leak outside. Preprint. http://philsci-archive.pitt.edu/18547/. Accessed 24 June 2021

Friston, K. J., & Frith, C. D. (2015). Active inference, communication and hermeneutics. *Cortex*, *68*, 129–143. https://doi.org/10.1016/j.cortex.2015.03.025

Fuchs, T. (2011). The Brain -- A Mediating Organ. *Journal of Consciousness Studies*, *18*, 196–221.

Fuchs, T., & De Jaegher, H. (2009). Enactive intersubjectivity: Participatory sense-making and mutual incorporation. *Phenomenology and the Cognitive Sciences*, *8*(4), 465–486. https://doi.org/10.1007/s11097-009-9136-4

Gallagher, S. (2011). The overextended mind. *Versus: Quaderni di studi semiotici*, *113*, 57–68.

Gallagher, S. (2013). The socially extended mind. *Cognitive Systems Research*, *25–26*, 4–12. https://doi.org/10.1016/j.cogsys.2013.03.008

Gallagher, S. (2017). *Enactivist Interventions: Rethinking the Mind* (Illustrated edition.). Oxford, United Kingdom: OUP Oxford.

Gallagher, S., & Tollefsen, D. (2019). Advancing the 'We' Through Narrative. *Topoi*, *38*(1), 211–219. https://doi.org/10.1007/s11245-017-9452-1

Gallese, V. (2013). Mirror neurons, embodied simulation and a second-person approach to mindreading. *Cortex*, *49*(10), 2954–2956. https://doi.org/10.1016/j.cortex.2013.09.008

Gallotti, M., & Frith, C. D. (2013). Social cognition in the we-mode. *Trends in Cognitive Sciences*, *17*(4), 160–165. https://doi.org/10.1016/j.tics.2013.02.002

Garavito, M. C. (2019). Incorporating others: what an extended self tells us about intersubjectivity. *Adaptive Behavior*, *27*(1), 47–59. https://doi.org/10.1177/1059712318789921

Geertz, C. (1973). *The Interpretation of Cultures*. New York: Basic Books.

Gilbert, M. (2015). *Joint commitment: how we make the social world* (Paperback ed.). New York, NY: Oxford University Press.

Goldman, A. I. (2006). *Simulating Minds: The Philosophy, Psychology, and Neuroscience of Mindreading*. Oxford: Oxford University Press.

Graeber, D. (2001). *Toward an anthropological theory of value: the false coin of our own dreams*. New York: Palgrave.

Greenwood, J. D. (2003). *The Disappearance of the Social in American Social Psychology*. Cambridge: Cambridge University Press.

Hacking, I. (1995). The looping effects of human kinds. In *Causal cognition: A multidisciplinary debate*. (pp. 351–394). New York, NY, US: Clarendon Press/Oxford University Press.

Heidegger, M. (1962). *Being and Time*. New York and San Francisco: Harper Collins.

Herder, Johan Gottfried von. (2004). Another philosophy of history for the education of mankind. *ID Evrigenis & D. Pellerin (Eds. & Trans.), Another philosophy of history and selected political writings*, 1À97.

Herder, Johann Gottfried von. (2002). Treatise on the Origin of Language. *Philosophical writings*, 65–164.

Hohwy, J. (2013). *The predictive mind*. Oxford; New York: Oxford University Press.

Hrdy, S. B. (2011). *Mothers and others: the evolutionary origins of mutual understanding*. Cambridge (Mass.): Belknap Press of Harvard University Press.

Huebner, B., & Hedahl, M. (2018). Shared values, interests, and desires. In *Routledge handbook of collective intentionality*. London New York: Routledge, Taylor & Francis Group.

Husserl, E. (1970). *The Crisis of European Sciences and Transcendental Phenomenology: An Introduction to Phenomenological Philosophy*. Northwestern University Press.

Hutchins, E. (1994). *Cognition in the wild*. Cambridge, Mass: MIT Press.

Hutchins, E. (1995). How a Cockpit Remembers Its Speeds. *Cognitive Science*, *19*(3), 265–288. https://doi.org/10.1207/s15516709cog1903_1

Kern, A., & Moll, H. (2017). On the transformative character of collective intentionality and the uniqueness of the human. *Philosophical Psychology*, *30*(3), 319–337. https://doi.org/10.1080/09515089.2017.1295648

Kirchhoff, M. D., & Kiverstein, J. (2019). *Extended Consciousness and Predictive Processing: A Third Wave View* (1st edition.). New York: Routledge.

Kirchhoff, M. D., & Kiverstein, J. (2020). Attuning to the World: The Diachronic Constitution of the Extended Conscious Mind. *Frontiers in Psychology*, *11*, 1966. https://doi.org/10.3389/fpsyg.2020.01966

Kirchhoff, M. D., & Kiverstein, J. (2021). How to determine the boundaries of the mind: a Markov blanket proposal. *Synthese*, *198*(5), 4791–4810. https://doi.org/10.1007/s11229-019-02370-y

Kiverstein, J. (2018). Extended cognition. *The Oxford Handbook of 4E Cognition, Oxford University Press, New York*, 19–40.

Kropotkin, P. (2021). *Mutual Aid: A Factor in Evolution*. Black Rose Books

Krueger, J., & Szanto, T. (2016). Extended emotions. *Philosophy Compass*, *11*(12), 863–878. https://doi.org/10.1111/phc3.12390

León, F., Szanto, T., & Zahavi, D. (2019). Emotional sharing and the extended mind. *Synthese*, *196*(12), 4847–4867. https://doi.org/10.1007/s11229-017-1351-x

Leslie, A. M., Friedman, O., & German, T. P. (2004). Core mechanisms in 'theory of mind.' *Trends in Cognitive Sciences*, *8*(12), 528–533. https://doi.org/10.1016/j.tics.2004.10.001

Loehr, J. D., Kourtis, D., Vesper, C., Sebanz, N., & Knoblich, G. (2013). Monitoring Individual and Joint Action Outcomes in Duet Music Performance. *Journal of Cognitive Neuroscience*, *25*(7), 1049–1061. https://doi.org/10.1162/jocn_a_00388

Marx, K. (1964). *Economic and Philosophic Manuscripts of 1844* (Vol. 333). New York City: International Publishers.

Moll, H., Nichols, R., & Pueschel, E. (2020). Shared intentionality shapes humans' technical know-how. *Behavioral and Brain Sciences*, *43*, e172. https://doi.org/10.1017/S0140525X20000163

Pesquita, A., Whitwell, R. L., & Enns, J. T. (2018). Predictive joint-action model: A hierarchical predictive approach to human cooperation. *Psychonomic Bulletin & Review*, *25*(5), 1751–1769. https://doi.org/10.3758/s13423-017-1393-6

Rowlands, M. J. (2010). *The new science of the mind: From extended mind to embodied phenomenology*. Mit Press.

Rupert, R. D. (2004). Challenges to the Hypothesis of Extended Cognition: *Journal of Philosophy*, *101*(8), 389–428. https://doi.org/10.5840/jphil2004101826

Schmid, H. B. (2013). Trying to Act Together. In M. Schmitz, B. Kobow, & H. B. Schmid (Eds.), *The background of social reality: selected contributions from the inaugural meeting of ENSO* (pp. 37–55). Dordrecht Heidelberg New York London: Springer.

Schneider, D. (1976). Notes toward a theory of culture. In K. Basso & H. Selby (Eds.), *Meaning in Anthropology, University of New Mexico Press: Albuquerque* (Vol. 198). University of New Mexico Press.

Schütz, A. (1967). *Collected papers I. The problem of social reality*. The Hage: Martinus Nijhoff.

Schweikard, D. P., & Schmid, H. B. (2013). Collective intentionality. *Stanford Encyclopedia of Philosophy*.

Searle, J. R. (1983). *Intentionality: An essay in the philosophy of mind*. Cambridge university press.

Searle, J. R. (1990). Collective intentions and actions. *Intentions in communication*, *401*, 401.

Seemann, A. (2009). Joint Agency: Intersubjectivity, Sense of Control, and the Feeling of Trust. *Inquiry*, *52*(5), 500–515. https://doi.org/10.1080/00201740903302634

Silberstein, M., & Chemero, A. (2012). Complexity and extended phenomenological-cognitive systems. *Topics in Cognitive Science*, *4*(1), 35–50. https://doi.org/10.1111/j.1756-8765.2011.01168.x

Slaby, J. (2014). Emotions and the extended mind. In C. von Scheve & M. Salmela (Eds.), *Collective Emotions* (pp. 32–46). Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199659180.003.0003

Slaby, J., & Gallagher, S. (2015). Critical Neuroscience and Socially Extended Minds. *Theory, Culture & Society*, *32*(1), 33–59. https://doi.org/10.1177/0263276414551996

Slors, M. (2020). From Notebooks to Institutions: The Case for Symbiotic Cognition. *Frontiers in Psychology*, *11*, 674. https://doi.org/10.3389/fpsyg.2020.00674

Smith, R. (2013). *Between Mind and Nature: A History of Psychology*. London: Reaktion Books.

Sprevak, M. (2010). Inference to the hypothesis of extended cognition. *Studies in History and Philosophy of Science Part A*, *41*(4), 353–362. https://doi.org/10.1016/j.shpsa.2010.10.010

Sprevak, M. (2019). Extended Cognition. https://doi.org/10.4324/9780415249126-V049-1

Stein, E. (1989). *On the Problem of Empathy*. Washington, D.C: ICS Publications 1917/1989.

Szanto, T. (2013). Shared Extended Minds: Towards a Socio-Integrationist Account. In D. Moyall-Sharrock, A. Coliva, & V. A. Munz (Eds.), *Mind, Language and Action. Proceedings of the 36th International Wittgenstein Symposium*. Kirchberg am Wechsel.

Szanto, T., & Moran, D. (2015). Phenomenological Discoveries Concerning the 'We.' In *Phenomenology of Sociality*. pp. 1-26. Routledge

Teo, T. (2009). Philosophical concerns in critical psychology. In D. Fox, I. Prilleltensky, & S. Austin (Eds.), *Critical psychology: An introduction (2nd ed.)* (pp. 36–53). London: Sage.

Theiner, G., Allen, C., & Goldstone, R. L. (2010). Recognizing group cognition. *Cognitive Systems Research*, *11*(4), 378–395. https://doi.org/10.1016/j.cogsys.2010.07.002

Thompson, E. (2007). *Mind in life: biology, phenomenology, and the sciences of mind*. Cambridge, Mass.: Belknap Press of Harvard Univ. Press.

Thompson, E., & Stapleton, M. (2009). Making Sense of Sense-Making: Reflections on Enactive and Extended Mind Theories. *Topoi*, *28*(1), 23–30. https://doi.org/10.1007/s11245-008-9043-2

Tomasello, M. (2019). *Becoming human: a theory of ontogeny*. Cambridge, Massachusetts: The Belknap Press of Harvard University Press.

Tuomela, R. (2007). *The Philosophy of Sociality: The Shared Point of View*. Oxford England ; New York: Oxford University Press, USA.

Tuomela, R. (2013). Who is afraid of group agents and group minds? In M. Schmitz, B. Kobow, & H. B. Schmid (Eds.), *The background of social reality: selected contributions from the inaugural meeting of ENSO* (pp. 13–35). Springer.

Tuomela, R., & Miller, K. (1988). We-intentions. *Philosophical Studies*, *53*(3), 367–389. https://doi.org/10.1007/BF00353512

Van Gelder, T., & Port, R. F. (1995). It's about time: An overview of the dynamical approach to cognition. *Mind as motion: Explorations in the dynamics of cognition*, *1*, 43.

Ward, D. (2012). Enjoying the Spread: Conscious Externalism Reconsidered. *Mind*, *121*(483), 731–751. https://doi.org/10.1093/mind/fzs095

Welsh, T. (2013). *The child as natural phenomenologist: primal and primary experience in Merleau-Ponty's psychology*. Evanston, Ill: Northwestern University Press.

Zahavi, D. (2014). Empathy and Other-Directed Intentionality. *Topoi*, *33*(1), 129–142. https://doi.org/10.1007/s11245-013-9197-4