



Why the extended mind is nothing special but is central

Giulio Ongaro¹ · Doug Hardman² · Ivan Deschenaux³

Accepted: 3 June 2022 / Published online: 15 July 2022
© The Author(s) 2022

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. Mind extension to other humans corresponds essentially to the ‘we-mode’ of cognition, the unique power of human minds to be jointly directed at goals, intentions, states of affairs, or values (which, importantly, differs from having a ‘group mind’). Because the capacity for collective intentionality holds evolutionary and developmental 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

✉ Giulio Ongaro
g.ongaro@lse.ac.uk

¹ Department of Anthropology, London School of Economics, London, UK

² Department of Psychology, Bournemouth University, Poole, UK

³ Department of Methodology, London School of Economics, London, UK

1 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 approach and subject matter because the 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 engaging in this debate have glossed over the most important way in which the mind might extend into the world: to other humans in acts of collective intentionality. We argue that mental extension to other humans corresponds essentially to the 'we-mode' of cognition (Tuomela, 2013a), the power of minds to be jointly directed at objects, matters of fact, states of affairs, goals, or values. Extension to epistemic artefacts is simply a derivative of the more fundamental capacity for collective intentionality. The exclusion of this (uniquely human) capacity from the debate is arbitrary and limits the implications of the extended mind thesis for the cognitive sciences.

The idea that the mind primarily engages the world socially, as part of a group – which, importantly, differs from having a 'group mind' – has been advanced in many guises before, without, however, having much impact on the cognitive sciences. The aim of this paper is to revive these views by charting a connection between the 'collective intentionality' and 'extended mind' debates. 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. Having established this connection, we survey how recent approaches to the socially extended mind could be corrected through a focus on collective intentionality. 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'.

2 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 needs 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 & 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 meta-physical 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 & 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, simultaneously opening new horizons 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 implication 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 primarily 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.

This has changed recently, with some philosophers invoking the idea of the ‘socially extended mind’ (Gallagher, 2011, 2013; Slaby & 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 a problem dictated by an individual private intention, without references to a group’s collective intention. As we show later in this essay, recent waves of approaches to the extended mind have significantly enriched the original formulation, and they have done so precisely by paying attention to the person’s sociocultural environment. We note, at the same time, that their import is limited because they remain implicitly wedded to methodological individualism. Indeed, as the Routledge Encyclopaedia of Philosophy states, 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). In what follows, we point out that philosophers debating this issue have overlooked an intimate link between extended mind and collective intentionality, and that acknowledging this link could have a deep impact on the cognitive sciences. To make our case, we will highlight some striking parallels between human-tool coupling in individual tasks and human-human coupling in the context of collective tasks. First, though, an overview of collective intentionality.

3 Collective intentionality

Imagine a group of people unrelated to each other sitting on various places in a park. Suddenly, it starts to rain so they all get up and run towards a centrally located shelter. Now think of a group of people in the same park that gets up and starts running towards the common point, but as part of a choreographed outdoor ballet. The outward behaviour of the two groups might be exactly the same, yet the intentional and emotional state of the people in the two scenarios obviously differs. Searle (1990), who came up with this example, pointed out that the latter case is distinct from the former because people think and act *as part of a group*, where the collective act cannot be explained by a summation of I-referential intentions. People here intend and act with reference to a collective intention, or ‘we-intention’ (Tuomela and Miller, 1988). Although characterisations of slightly different analytical flavours have been proposed (Searle, 1996; Bratman, 1999, 2013; Tuomela, 2007; Gilbert, 2015), philosophers agree that for something like joint action to take place there needs to be a ‘shared point of view’ – the ability to conceive a sense of ‘we’ in relation to which members regulate their own thoughts and actions. Rather than a private intention, it is the collective commitment to a group’s ethos that gives the group’s members reasons to think, act and share emotions in a certain way. The enabling condition of collective intentionality is taken to be ‘mutual trust’ (Schmid, 2013; Tomasello, 2019), i.e. trusting that other members of the group share the group’s collectively accepted goals and ethos.

It is important to stress that the notion of collective intentionality is not tied to any commitment to supra-individual entities like a ‘group mind’. Whether groups can be considered to have a mind of their own, in a literal sense, is a question of a different order than the question concerning the character of ‘we-intentions’, which are held by single persons.¹ Collective intentionality is also not defined by its objects of intention. Examples of collective acts in the literature range from the simple dyadic action of lifting a table together to fighting as part of a nation at war; thus, we-intentions can be held towards both social and non-social objects, and across a range of spatio-temporal scales. What defines collective intentionality, rather, is its specific ‘mode’ of cognition. Tuomela (2007) dubbed it the ‘we-mode’, to contrast it from the ‘I-mode’ of I-referential intentions. He noted that the ‘we-mode’ holds across phenomena as

¹ We think that the ‘group mind’ debate bears similarities to the debate on the ‘mark of the mental’ insofar as they both deal with metaphysical questions that no kind of empirical evidence is likely to settle. By contrast, whether, how and in which contexts individuals engage in collective intentionality *is* an empirical question, as we will see.

different as goals, states of affairs, values, or emotions. A goal, for instance, is collective if and only if, when satisfied for one member, it is satisfied for all. Emotions are collectively experienced (i.e. in the we-mode) if and only if they are induced by the ethos of a collectively recognised group, whereas they are individually experienced if the reasons that induce them are independent of group membership. The ballet dancers in the example above might have a collective experience of ‘flow’, in the we-mode, but if a lightning storm approaches, they might be individually afraid, in the I-mode.

There are in-between cases. Imagine a group of people who jointly agree to take turns to grow fresh flowers in the village commons to keep it beautiful. The intentions that go into this collective action certainly tick all the boxes for the ‘we-mode’ of cognition. But what if an individual decides to perform the same action privately, without being part of a group, while being aware that their action satisfies collectively held ideas of beauty in the community? We might refer, in this case, to a weaker form of collective intentionality, which lacks collective commitment but still retains the adherence to a group ethos or ideal. Tuomela (2013a) termed it ‘pro-group I-mode’ intentionality. Now imagine someone simply driving a car following road conventions. Any sense of jointness seems to have vanished in this case (we are in the pure I-mode), yet the person still acts on the basis of collectively accepted norms and conventions. Philosophers here refer to the ‘background’ (Searle, 1996; Schmitz, 2013), a set of learnt shared dispositions people act on pre-reflectively as they engage with the social world. In short, philosophers agree that there are different layers and scales of collective intentionality (e.g. Carassa & Colombetti 2013). As we move from acting in the we-mode to acting in the I-mode, the conscious sense of group membership recedes into the background, but, even so, certain conditions of collective intentionality (i.e. trusting, implicitly, that other people share the same ethos in given domains of action) are retained. ‘We-mode’ and ‘I-mode’ are thus better conceived as two opposite intentional poles embedded in a background – i.e. a more diluted version of the ‘we-mode’ – where the sense of being part of a group retreats from consciousness. Importantly, though, the ‘background’ still rests on acts of collective intentionality. Money and the norms that surrounds it, for example, ultimately stems from a social joint commitment about the way it is used. There is a link running between joint action and social convention, and it is not coincidental that the same philosophers who write about collective intentionality feel like that the same discussions apply to broader issues of social ontology. At present, philosophers are actively debating how to distinguish ‘we-mode’ from ‘I-mode’ and about how these relate to the ‘background’, and generally about how exactly to define all these terms.

Notably, this debate has been largely unheeded by the mind sciences. The contemporary fields of ‘social psychology’ or ‘social cognition’ are not concerned with the study of what philosophers define as ‘collective intentionality’. This is clear in the very definitions of these disciplines: “What makes social psychology social,” write Hogg and Vaughan, “is that it deals with how people are affected by other people who are physically present or who are imagined to be present” (2018:4). What defines them, in short, is not the mode of psychological states but their object, namely other people or groups. They study how individuals think and act *towards* other individuals, rather than studying how individuals think and act *with* other individuals towards the world.

By default, they study social cognition in the I-mode. When they do consider phenomena that take place in the we-mode, they do not focus on the we-mode per se or on the differences with the I-mode (which is the central question of the philosophical debate). A classic example is the debate on ‘theory of mind’ in the context of face-to-face interactions. Cognitive scientists differ on how to account for intersubjective understanding – 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 I-referential framing is the same: two or more people trying to individually represent and predict the other’s mental states in the I-mode.² An approach that is conscious of the significance of collective intentionality would see things differently. It would not deny that some aspects of face-to-face encounters can take this form, but it would point out that mindreading is embedded the we-intention of having the encounter in the first place. It would consider face-to-face interactions primarily as joint actions. To underscore the central difference once again: whereas contemporary social psychology and social cognition look at how individuals engage the social world, collective intentionality is concerned with how individuals socially engage the world (regardless of whether the worldly object of intention is social or not). The two fields, which, intuitively, should be about the same phenomena, have moved along different tracks.

This wasn’t always the case. Greenwood (2003) chronicled a shift occurring around the 1930s in America where social psychology abandoned a focus on the distinctive social dimension that it previously maintained. For many American social psychologists writing in the 1920s, the difference between individual and socially embedded psychological states was evident, as was evident that ‘social psychology’ was supposed to research the latter. These psychologists’ conceptions of ‘socially embedded psychological states’ bear remarkable similarity (if less sophistication) to the conceptions of collective intentionality of today’s philosophers. Dunlap (1925), for instance, defined “social consciousness” as “consciousness (in the individual, of course) of others in the group, and consciousness of them, as related, in the group, to oneself; in other words, consciousness of being a member of the group” (1925:19), which amounts to an alternative characterisation of the ‘we-mode’ that appeared decades later.³ Empirical research, accordingly, looked specifically at whether and how certain attitudes and behaviour were held socially or individually.⁴

Before long, interest in this question waned. A few reasons for this shift were cultural, due to the rampant growth of moral and political individualism in mid-century America; others had to do with shifts within the discipline of psychology itself., 200 Particularly, a few psychologists who championed the idea of socially embed-

² There is an alternative enactivist position (Fuchs & De Jaegher, 2009) that challenges these views, but it is unclear whether and how they align with the collective intentionality approach. We hope to explore these differences in a future publication.

³ The difference between we-mode and I-mode wasn’t lost to early 20th century phenomenologists either. Heidegger, for instance, made a distinction between the conditions of ‘being-towards’ and ‘being-with’ other people that closely resembles the I-mode/we-mode distinction. Philosophers like Stein and Merleau-Ponty also held sophisticated views of what we now consider ‘collective intentionality’.

⁴ An example among many is Schanck’s (1932) investigation of whether Methodists hold preferences towards various forms of baptism ‘as church members’ or simply out of their ‘private feelings’.

ded psychological states drew gradually closer to the more metaphysically contested idea of ‘group mind’. This ambiguity led to an effective (and not wholly unjustified) backlash against the concept of collective cognition as a whole in the years that followed. What put the final nail in the coffin, however, was the rapid spread of a narrow form of experimentation that has dominated the mind sciences since, namely, the systematic setting up of experiments aimed at demonstrating functional relationships between specific stimulus elements – i.e. independent variables – and specific response elements, the dependent variables. The increased technical focus on the manipulation of independent variables acted as a constrain on the acknowledgement of the social dimension of cognition. This is so because unless the goal of the experimenter is to study the we-mode specifically – i.e. unless the we-mode is taken as the independent variable itself – any kind of orientation of the subjects towards social groups and their ethos becomes a source of confounding that must be eliminated at the outset. Thus, as experiments became more and more technically sophisticated and statistically rigorous, so the social dimension of cognition became increasingly neglected, marking a win of methodology over ontology.

The neglect continues in today’s cognitive sciences. Despite sporadic calls for recognising the we-mode scientifically (e.g. Gallotti & Frith 2013), and recent studies of joint action that offer budding hopes for the future (which are important and which we will discuss later), a fully-fledged science of collective intentionality does not yet exist. There is little doubt that the experimental conditions required to study the we-mode are more challenging than those aimed at studying cognition in the I-mode. Early 20th century social psychologists knew this too (Greenwood, 2003). But such scientific study is by no means impossible. Our argument in this essay is that it is essential, lest we leave the most distinctive dimension of human mental life out of the purview of the cognitive sciences.

To build our case, we exploit the intuitive move implicit in the extended mind thesis. In a stepwise manner, we argue that if one accepts that the mind extends to epistemic artefacts in the pursuit of individual tasks, one must also accept that it extends to other humans in acts of collective intentionality. Our argument, in essence, is that the ‘extended mind’ applied to other humans is tantamount to cognition in the we-mode. Crucially, since the debate on the extended mind is (also) a debate over the unit of analysis of the cognitive sciences, it follows from our argument that cognitive science will provide an accurate picture of mental life only by paying attention to the we-mode. One of its central questions, in any given context, should be whether and how individuals think, act and feel with reference to collective intentions.

4 Charting the link: 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 merely derivative of the capacity for collective intentionality, as illustrated in the former.

4.1 Parallels

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.⁵ 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.⁶ As such, we suggest a homology between the Otto and Elmo cases, with the very important difference that in the Elmo case the accomplished task is not individual but the product of a collective intention. 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 & 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 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

⁵ See Szanto (2013) for a similar application of the parity principle to collectives.

⁶ 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.

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 we-intentions, 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 'reliability' and 'accessibility', we notice that these too can be applied to the Elmo case, even though it quickly becomes apparent that the condition of 'trust' subsumes them both, to the point that 'reliability' and 'accessibility' emerge as unnecessary. Elmo and his fellow hunters are engaging in a different task than Otto: their goal is to acquire new information on the location of a deer, which does not involve information retrieval, as in the Otto case, but rather information acquisition and processing. Consequently, 'reliability' and 'accessibility' pertain in this case to the acquisition and processing of information. All that is required for these conditions to be met and is that Elmo's companions, upon seeing 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. For that matter, Elmo could be equipped with a remote-controlled drone where video recorded from the drone's camera is broadcast in real time to a headset that he is wearing, such that he effectively sees the world through the drone's camera. In addition, a computer-vision algorithm highlights any deer in the video feed. Following the parity principle, extended mind philosophers would have no hesitation in saying that mind extends to the drone. We suggest that in the low-tech example, other hunters are functionally equivalent to the drone. But - and this is important - we note that the conditions of being 'reliably available' and 'easily accessible' are entirely superfluous vis-a-vis Elmo's goal of killing the deer: trusting the other companions to have the same intention of hunting the deer as part of a group (i.e. trusting that they will process information to that end) lead to the same result. 'Reliable availability' and 'accessibility' are conditions that should only apply to tools; 'mutual trust' makes them irrelevant in the context of collective action. Finally, there is a need for the collective intention to be "consciously endorsed at some point in the past", which is Clark's fourth tightening condition. Elmo accomplishes this when he agrees to go hunting with the other people in the band.⁷

⁷ Other parallels include the idioms used to define the two types of coupling. The idioms used to describe the coupling between a subject and epistemic artefacts – "glue and trust" (Clark, 2010:83) – are notably similar to the idioms used to define human-human relations in the collective intentionality literature, strengthening the case for the homology. Tuomela, for instance, employs the same adhesive metaphor when describing participants in joint acts as "glued together" (2013b:22). Also, as extended minds

The homology runs deeper when we consider it from a phenomenological point of view. Central to this is the observation 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 as part of a group, we do not perceive the other member of a group 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 & De Jaegher, 2009). In the condition of ‘being-with’, rather than ‘being-towards’ (Heidegger, 1962), other companions are not experienced as opaque entities whose intentions of movements need to be figured out, but as transparent extension of one’s own activity (Seemann, 2009). This is ultimately an implication of trusting others’ intentions. 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.

4.2 Beyond problem solving

Both the Otto and the Elmo cases feature problem-solving tasks. Problem-solving has dominated the extended mind debate because ‘extension’ has typically been framed as the employment of external epistemic resources or tools to solve practical individual tasks. We drew the above homology based on problem-solving cases in order to establish a point of contact between extended mind and collective intentionality. If one accepts extension in the former case, we argued, one must also accept extension in the latter. But the scope of collective intentionality extends far beyond the domain of problem-solving. As we illustrated earlier, rather than being defined by an object or spatio-temporal scale, collective intentionality is defined by a cognitive ‘mode’ – the we-mode – which can hold across a constellation of phenomena which might have little to do with the type of problem solving exemplified by Otto and Elmo. Think of rituals, play, convivial conversations, political projects, or the pursuing of collective interests. These cannot be framed as problems to be solved. In a context such as ‘dancing together’, other people are typically thought as the very reason why we do it, not tools that allows us to do it. It is hard to draw a comparison here because other people cannot replace any internal function and the parity principle simply does not apply. Yet, the kind of coupling involving interpersonal trust that typifies all the cases explored previously is exactly the same. It involves intending a joint activity with shared goals and norms under the assumption that everyone in the group does, as part of a group (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. 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

employ epistemic tools “on the fly” (Clark, 2008:42, 2017:19), so people are said to move from individual to collective action “on the fly” (Tollefsen, 2014:15).

of mind extension can be at play in non-problem-solving actions with other people. 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.

5 The primacy of collective intentionality

A key reason 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 - constitutes the real 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 deer, they both have to individually have the goal of capturing the deer 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'.

Ontogenetically speaking, evidence shows that the acquisition of collective intentionality is a two-step process. It first starts as joint attention, when babies point out objects to other persons with no motive other than the recipient share in the baby's focus, and develops later into learning by instruction, language and fully-fledged understanding of the social world (Trevvarthen, 1979; Tomasello, 2019). Crucially, the same evidence indicates that collective intentionality is not added on top of other individual cognitive skills – it “shapes cognition all the way down” (Kern & Moll, 2017:327). It permeates the human individual's reasoning and engagement with the world as a whole and not just how they socially interact with others. Children's 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. Mental life is social through and through.

A great many philosophers over time (and across cultures) have made arguments for the primacy of collective intentionality that chime with this line of evidence. Phenomenologists have argued it is mistaken 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: 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 primordially 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). Pragmatists such as Dewey, Mead and James made similar points (see Crippen & Schulkin 2020). They stressed that thoughts and actions are entwined with social life all along and; that the latter is not built from the former but shapes individual cognition in turn. The argument in this paper fully supports these views and seeks to make them relevant to the cognitive sciences.

There is little doubt that someone like Husserl or Dewey would be unimpressed with Clark’s narrow conception of the ‘extended mind’, as they would recoil at the way the cognitive sciences operate today. They would see the cognitive sciences as grounded on an epistemology of isolated individualism that elides sociality from view. A slight problem with these positions, however, is that in declaring the primacy of social cognition they do not discriminate between a person’s modes of social engagement with the world. In other words, they do not dwell on the distinction between ‘we-mode’ and ‘I-mode’ that is central to the collective intentionality debate – a distinction that, we soon suggest, holds the potential of providing the cognitive science with a whole new set of empirical questions. Before getting to this point, though, we briefly review some recent attempts at extending the ‘extended mind’ to the social environment (attempts that fall within the so-called ‘third-wave’ of extended mind theorising), which, while significantly advancing Clark’s account, have the same problem and could be much more incisive were they to reach out to the collective intentionality literature (readers uninterested in this reassessment can jump to section 7 on the ‘scientific implications’).

6 Revisiting ‘third-wave’ views

6.1 Gallagher’s ‘socially extended mind’

The idea of the ‘socially extended mind’ was first introduced by Gallagher (2013). As we noted earlier, we think that Gallagher’s original formulation remains wedded to methodological individualism insofar as other people figure as tools to achieve an individual task rather than a collective task. In the same vein of contemporary social psychology, he approaches the issue from an ‘I-mode’ perspective. Thus, as he argues through one example, the legal system emerges as a cognitive institution that facilitates an individual’s actions in the world and is for this reason constitutive of one’s mind. But to say this, 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. Gallagher (2013:2–3) states that ‘critical reflection’ is simply more cognition that is added to the overall extended cognitive process; hence it should be counted as extended. The problem,

however, is that critical reflection hampers ‘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. We hope to have shown that 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.

Gallagher’s more recent co-authored publications on the ‘socially extended mind’ have taken this idea into novel directions by developing the notion of ‘cognitive institutions.’ According to Petracca and Gallagher, these are defined as institutions that “not just allows agents to perform certain cognitive processes in the social domain but, more importantly, without which some of the agents’ cognitive processes would not exist or even be possible” (2020:1). Their primary example is the market. Markets, they argue, do not simply constrain or enable individual economic reasoning but are also constituted by the actions and interactions of individual agents. Agents and social institutions stand in a mutually constitutive relation. In making this point, Gallagher et al. engage with the debate on collective intentionality, which already produced a large literature on institutions, but take some distance from it due to the internalism of some of its proponents, primarily Searle’s.

Searle is surely known to be an internalist. But we suggest that his avowed internalism has little to do with his characterisation of collective intentionality (indeed, he has never delved into the extended mind debate). Pace Searle’s views on internalism, we argued in this paper that there is direct link between the extended mind thesis and the capacity for collective intentionality that form the basis of social phenomena such as institutions. The social world emerges through ‘joint commitment’ (Gilbert, 2015), which, as we argued at length earlier, entails mental extension to other people. Given that institutions already build on the extended mind, we don’t see the need for the additional term ‘cognitive institution’. ‘Institution’ suffices. In our view, Gallagher’s most insightful contributions on the matter – e.g. on the role of social narratives in establishing norms and cementing collective ethos (Gallagher & Tollefsen, 2019), or the idea that it is processes and interactions rather than models that make up the social world (Gallagher et al., 2019) – should fit squarely within collective intentionality discussions rather than in the extended mind debate.⁸

6.2 Kirchoff and Kiverstein’s ‘extended consciousness’

The idea that the mind extends socially has also been advanced, in a different guise, in Kirchoff and Kiverstein’s recent *Extended Consciousness and Predictive Processing* (Kirchoff & Kiverstein, 2019). In this compelling book and other publications (2019, 2020, 2021), Kirchoff and Kiverstein have put forward

⁸ As a tentative aside, we think that the ‘market’ might be an awkward place to look for collective intentionality since, as Gallagher et al., (2019:9) also indirectly note, it is specially designed to promote I-mode attitudes (at least in its neoliberal form). The main aspect of collective intentionality lies in the collective agreement to act in the I-mode.

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, we focus our discussion on a particular interpretation of cognition Kirchoff and Kiverstein adopt in making their arguments.

Kirchoff and Kiverstein note that the extended mind debate has focused on individual problem-solving. As a result, commentators have approached cognition in a strictly synchronic way; that is, 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 & 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. 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, Kirchoff and Kiverstein suggest that 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," Kirchoff 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 would like to elaborate on Kirchoff and Kiverstein's point by noting that, in fact, 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 & 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, premised as it is on a synchronic reading of cognition, is 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. Particularly, we think that Kirchhoff and Kiverstein’s notion of ‘cultural practice’ can be easily subsumed in Searle’s notion of the ‘background’, namely, a temporally diluted version of the ‘we-mode’ (which, we argued, entails mental extension to other people). Bringing in collective intentionality throws the point Kirchhoff 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. In short, we do accomplish an action like writing a philosophy paper with other people 'in mind'.⁹ 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 from the overall wider process of cognition and would yield a selective view of mental life.

7 Scientific 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. What, then, would a cognitive science that takes the concept of collective intentionality to heart look like?

We wager that it would look a lot different than it does now. As we mentioned earlier, current cognitive science is grounded in methodological individualism: it approaches cognition in the I-mode by default, 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 & 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 problems under investigation. These experiments tend

⁹ 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.

to represent uncommon, socially-isolated and experimenter-defined tasks that do not acknowledge that cognition can take place in the we-mode.

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.

As an illustrative case, consider Loehr et al. (2013) musical ensemble study, which investigates the cognition of a pianist who produces tones in the course of playing a duet with another pianist. In this 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 (i.e. occurring in the I-mode), while a finding that aligns with the second hypothesis would essentially find a cognitive signature of collective intentionality (i.e. occurring in the we-mode). 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).

Studies of joint action such as this one are the very few experiments that find the cognitive signature of the we-mode: they show that it would be inaccurate to study the cognition of the duetting pianist without factoring in collective intentions. But as we mentioned earlier, there are different scales of 'we-mode'. What philosophers call the 'background' can be considered as temporally diluted form of we-mode cognition, where the sense of being in a group recedes from consciousness. If this is so, there is 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 collective 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 (Smith, 2013; Adams et al., 2019).

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 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 '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 example of 'cultural shock' – was originally used to describe the first (way 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 & 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.

8 Conclusions

This paper has circumvented the dispute between internalism versus externalism that fuelled the extended mind debate. It has argued, though, that if one accepts that the

mind extends to epistemic artefacts in the pursuit of individual tasks, one must also accept that it extends to other humans in acts of collective intentionality. To put it succinctly: the socially extended mind equals cognition in the ‘we-mode’. Because collective intentionality holds primacy over epistemic tool use, we call for a reversal of the 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’, as part of a social group, and this is why the extended mind is nothing special but is central.¹⁰ Drawing a link between the extended mind and collective intentionality also implies expanding the scope of the sciences of the mind. A cognitive science that takes collective intentionality to heart will take the ‘we-mode’ as one of its central research objects, and how people act in it one of its central research questions.

Authors’ contributions Giulio Ongaro: Conceptualization; formal analysis; writing – original draft. Doug Hardman: Conceptualization; formal analysis; writing – review & editing. Ivan Deschenaux: Conceptualization; formal analysis; writing – review & editing.

Funding Research related to this work was supported by ESRC Grant number ES/V005731/1 and a Wenner-Gren Hunt Postdoctoral Fellowship (GO) and ESRC Grant number ES/V009664/1 (ID).

Data Availability N/A.

Code Availability N/A.

Declaration

Conflicts of interest/Competing interests The authors have no conflicts of interest to declare.

Ethics approval N/A.

Consent to participate N/A.

Consent for publication N/A.

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/>.

¹⁰ The title of this paper and overall style of argumentation is inspired by Bloch’s ‘Why religion is nothing special but is central’ (2008).

References

- Adams, F. (2019). The Elusive Extended Mind: Extended Information Processing Doesn't Equal Extended Mind. *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>
- Adams, G., Estrada-Villalta, S., Sullivan, D., & Markus, H. R. (2019). The Psychology of Neoliberalism and the Neoliberalism of Psychology: Neoliberalism of Psychology. *Journal of Social Issues*, 75(1), 189–216. <https://doi.org/10.1111/josi.12305>
- 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>
- Carassa, A., & Colombetti, M. (2013). Creating interpersonal reality through conversational interactions. In M. Schmitz, B. Kobow, & H. B. Schmid (Eds.), *The Background of Social Reality* (pp. 91–104). Springer
- 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
- Di Paolo, E. (2009). Extended Life. *Topoi*, 28(1), 9–21. <https://doi.org/10.1007/s11245-008-9042-3>
- Dunlap, K. (1925). *Social psychology*. Williams & Willkins Company
- 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., Mastrogiorgio, A., & Petracca, E. (2019). Economic Reasoning and Interaction in Socially Extended Market Institutions. *Frontiers in Psychology*, 10, 1856. <https://doi.org/10.3389/fpsyg.2019.01856>
- 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; A Journal Devoted To The Study Of The Nervous System And Behavior*, 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>
- 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
- Heidegger, M. (1962). *Being and Time*. New York and San Francisco: Harper Collins
- Hogg, M. A., & Vaughan, G. M. (2018). *Social psychology* (Eighth Edition.). New York: Pearson
- 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
- Husserl, E. (1970). *The Crisis of European Sciences and Transcendental Phenomenology: An Introduction to Phenomenological Philosophy*. (D. Carr, Trans.)
- 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: An Illuminated Factor in Evolution*. PM Press/Kairos
- 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* (333 vol.). 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>
- Petracca, E., & Gallagher, S. (2020). Economic cognitive institutions. *Journal of Institutional Economics*, 16(6), 747–765. <https://doi.org/10.1017/S1744137420000144>
- 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>
- Schanck, R. L. (1932). A study of a community and its groups and institutions conceived of as behaviors of individuals. *Psychological Monographs*, 43(2)

- 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
- Schmitz, M. (2013). *The background of social reality: selected contributions from the inaugural meeting of ENSO*. New York: Springer
- Schütz, A. (1967). *Collected papers I. The problem of social reality*. The Hage: Martinus Nijhoff
- Searle, J. (1996). *The Construction of Social Reality* (Reprint edition.). London: Penguin
- 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 : A Journal Of Medical Care Organization, Provision And Financing*, 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>
- 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>
- 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
- 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
- Tollefsen, D. (2014). A Dynamic Theory of Shared Intention and the Phenomenology of Joint Action. In S. R. Chant, F. Hindriks, & G. Preyer (Eds.), *From Individual to Collective Intentionality* (pp. 12–33). Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199936502.003.0002>
- Tomasello, M. (2019). *Becoming human: a theory of ontogeny*. Cambridge, Massachusetts: The Belknap Press of Harvard University Press
- Trevarthen, C. (1979). Communication and Cooperation in Early Infancy: A Description of primary Intersubjectivity. In M. Bullowa (Ed.), *Before Speech: The beginning of Human Communication*. Cambridge 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). *Social ontology: Collective intentionality and group agents*. Oxford University Press
- 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>

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