CrossMark

RESEARCH REPORT

Identifying and individuating cognitive systems: a task-based distributed cognition alternative to agent-based extended cognition

Jim Davies[1] · Kourken Michaelian[2]

Abstract This article argues for a task-based approach to identifying and individuating cognitive systems. The agent-based extended cognition approach faces a problem of cognitive bloat and has difficulty accommodating both sub-individual cognitive systems ("scaling down") and some supra-individual cognitive systems ("scaling up"). The standard distributed cognition approach can accommodate a wider variety of supra-individual systems but likewise has difficulties with sub-individual systems and faces the problem of cognitive bloat. We develop a task-based variant of distributed cognition designed to scale up and down smoothly while providing a principled means of avoiding cognitive bloat. The advantages of the task-based approach are illustrated by means of two parallel case studies: re-representation in the human visual system and in a biomedical engineering laboratory.

✉ Jim Davies
jim@jimdavies.org

Kourken Michaelian
kourken.michaelian@otago.ac.nz

[1] Institute of Cognitive Science, Carleton University, Ottawa, Canada

[2] Philosophy Department, University of Otago, Dunedin, New Zealand

Introduction

This article argues for a task-based approach to identifying (i.e., singling out) and individuating (tracing the boundaries of) cognitive systems. The "The ontology of cognitive systems" section reviews difficulties faced by traditional intracranialist or brainbound views as well as parity-based and complementarity-based variants of extended cognition with respect to the challenges of "scaling up" (identifying and individuating supra-individual cognitive systems) and "scaling down" (identifying and individuating sub-individual cognitive systems). A distributed approach can in principle successfully meet the challenges of scaling up and scaling down, but distributed cognition, like extended cognition, also faces the challenge of avoiding cognitive bloat, i.e., providing a criterion for the identification and individuation of cognitive systems that does not imply an unacceptable proliferation of ephemeral and explanatorily-inert systems. In order to meet this challenge, "The task-based approach" section develops a task-based (as opposed to agent-based) variant of distributed cognition. According to the task-based approach, cognitive *systems* are defined relative to cognitive *processes*, which themselves defined relative to cognitive *tasks*; the approach makes no essential reference to agents. In order to show that the task-based approach meets the challenges of scaling up and scaling down, "Applying the task-based approach" section applies it to a pair of parallel case studies: re-representation in the human visual system (scaling down) and re-representation in a biomedical engineering laboratory (scaling up). The

🖄 Springer

section provides a brief summary of our argument.

The ontology of cognitive systems

Participants in the debate over extended and distributed cognition are by now no doubt intimately familiar with the details of Clark and Chalmer's thought experiment about Otto, an Alzheimer's patient who relies on a notebook in which he records relevant information in order to compensate for his failing memory. Given its centrality to the debate, it is nevertheless worth quoting the story once more:

> Otto suffers from Alzheimer's disease, and like many Alzheimer's patients, he relies on information in the environment to help structure his life. Otto carries a notebook around with him everywhere he goes. When he learns new information, he writes it down. When he needs some old information, he looks it up. For Otto, his notebook plays the role usually played by a biological memory. Today, Otto hears about the exhibition at the Museum of Modern Art, and decides to go see it. He consults the notebook, which says that the museum is on 53rd Street, so he walks to 53rd Street and goes into the museum.
>
> Clearly, Otto walked to 53rd Street because he wanted to go to the museum and he believed the museum was on 53rd Street. And just as Inga [an agent with a normal memory] had her belief even before she consulted her memory, it seems reasonable to say that Otto believed the museum was on 53rd Street even before consulting his notebook. For in relevant respects the cases are entirely analogous: the notebook plays for Otto the same role that memory plays for Inga. The information in the notebook functions just like the information constituting an ordinary non-occurrent belief; it just happens that this information lies beyond the skin. (Clark and Chalmers 1998, pp. 12–13)

The moral of the story, if Clark and Chalmers are right, is that, contra the traditional intracranialist view, cognition is not "brainbound" [as Clark has subsequently put it (Clark 2008)]: when Otto looks up the museum's address in the notebook, the situation is best described not as an individual cognitive agent retrieving information from an external store, but rather as a single, extended cognitive system retrieving information from an internal store, in a manner that is functionally indistinguishable from that in which a human agent with a normal memory, such as Inga, retrieves information from her internal (biological) memory (Clowes 2013). In other words: cognition "ain't (all) in

the head" (Clark and Chalmers 1998, p. 8): the process of remembering loops out into the world, running from Otto's brain through the notebook and back again.[1]

The fact that the memory process is viewed as starting with the agent and "extending" to include the contribution of the external resource is not incidental to the extended cognition approach. Consider a series of three cases, this time involving mental rotation rather than memory retrieval, that Clark and Chalmers (1998) use to make the same point.

- In case 1, a person mentally rotates a two-dimensional figure displayed on a computer monitor in the standard (onboard) way.
- In case 2, the person has the computer rotate the figure on the monitor itself.
- In case 3, he makes use of a neural implant to (mentally?) rotate the figure.

The thought we are invited to have is that the process unfolding in case 1 is uncontroversially cognitive; most of us would be willing to count the process in case 3 as cognitive; but if we are willing to count the latter process as cognitive, then we should also be willing to count the process in case 2 as cognitive, given that the only relevant difference between cases 2 and 3 is the location of the non-biological resource. The point to note here is that, just as in the Otto example, in all three cases cognition is viewed as *starting* from the human agent. While this approach to identifying and individuating extended cognitive systems—starting with an *agent* who is himself uncontroversially cognitive and then asking whether an external resource counts as engaged in cognitive processing in virtue of its relation to the agent in question—is natural enough, we will argue that it is ultimately unsatisfactory.

Agent-based approaches: intracranialism and (first- and second-wave) extended cognition

The agent-based approach fits comfortably with Clark and Chalmers' explicitly functionalist reasoning: it is in virtue of the fact that an external resource (such as a notebook)

---

[1] Clark and Chalmers argue that both *cognition* and *mind* are extended; the Otto case was originally used to argue that mind, in particular, is extended, but it can be used for either purpose. Our focus here is on the hypothesis of extended *cognition*; hence we will not be concerned with the claim that Otto's *beliefs* are located in his notebook but rather with the claim that Otto and his notebook constitute a single cognitive system. See Huebner (2014) for a defense of distributed mentality. Our focus is also distinct from that of theorists who have argued for extended consciousness (Manzotti 2011; Honderich 2014), and we take no stand on the question of extended or distributed consciousness. A fuller treatment would, however, eventually have to deal with this question, especially in view of recent debates on cognitive phenomenology, many participants in which have argued that cognition has a distinctive conscious character (see, e.g., Strawson 2011; Smithies 2013.)

plays the same functional role as the relevant internal resource that it is legitimate to view the relevant cognitive process as being realized in part by the external resource. This functionalist reasoning is encapsulated in the parity principle (as it has come to be known):

> If, as we confront some task, a part of the world functions as a process which, *were it done in the head*, we would have no hesitation in recognizing as part of the cognitive process, then that part of the world *is* (so we claim) part of the cognitive process. (Clark and Chalmers 1998, p. 8).

Given functionalism, and given that internal and external resources are in fact sometimes functionally isomorphic, we have a powerful reason to accept the extended cognition hypothesis.[2]

It should be emphasized, however, that, in order to derive extended cognition from functionalism, we do need to assume that external resources do in fact sometimes play functional roles that are typically carried out by internal resources, and this assumption can be challenged. Among the most influential arguments against extended cognition and in favor of intracranialism are those provided by Adams and Aizawa (2001, 2008). One of their key arguments (we will not consider their discussion of the "mark of the cognitive") points to apparent functional differences between internal and external resources (cf. Rupert 2004, 2013). Consider, for example, Donald's (1993) discussion of exograms or external memory traces (so-called by analogy with engrams, i.e., internal memory traces). While Donald explicitly refers to exograms as "exact functional analogues" of engrams, he himself points to many functionally relevant differences between exograms and engrams; e.g., in contrast to engrams, which have a labile and distributed character, exograms are designed for stable storage of discrete items of information (Sutton 2010; Michaelian 2012, 2014). The specific disanalogies to which Adams and Aizawa point do not show that internal and external resources are *never* functionally isomorphic, but they do suggest that the parity principle may be unable to support a form of extended cognition which sees the existence of extended cognitive systems as anything more than an exceptional occurrence.

Intracranialists take functional disanalogies between internal and external resources to constitute a reason for rejecting extended cognition, but it is possible for an extended approach to take such disanalogies on board. Indeed, it is (in part) in recognition of the importance of

such disanalogies that extended cognition theorists have tended to shift away from *parity*-based to *complementarity*-based arguments for the extended cognition hypothesis, in what Sutton has referred to as a move from first-wave to second-wave extended cognition (Sutton 2010) [though he argues that the germ of the complementarity-based approach can be distinguished as early as Clarke (1998b)].[3] On the second-wave approach,

> external states and processes need not mimic or replicate the formats, dynamics, or functions of inner states and processes. Rather, different components of the overall...system can play quite different roles and have different properties while coupling in collective and complementary contributions to flexible thinking and acting. (Sutton 2010, p. 194)

The second-wave approach brings extended cognition much closer to distributed cognition (described below), but extended and distributed cognition continue to differ in that, while the latter tends to focus on distributed socio-technical systems which may involve multiple agents and artifacts and which do not necessarily have a clear center, the former continues to focus on systems which are centered on a single agent (see Hutchins 2011). As Clark recently put it, "[i]ndividual cognizing...is organism-centered even if it is not organism-bound" (Clark 2007, p. 176) (cf. Giere 2011, 2012). Thus second-wave extended cognition is continuous with first-wave extended cognition in the sense that it advocates an agent-based approach to the ontology of cognitive systems: we *identify* a cognitive system by focusing on an uncontroversially cognitive agent and *individuate* the system by asking which external resources interact with the agent in the right way.

The challenge of cognitive bloat

A distinct line of objection to extended cognition concerns "cognitive bloat" (Clark 2001; Rupert 2004; Rowlands 2009; Palermos 2014). The worry here is that Clark and Chalmers' original criteria for cognitive extension, the so-called trust-and-glue conditions—according to which the agent (1) must have reliable access to the information contained in the resource, (2) must have reliable access to the resource itself, (3) must tend to endorse information contained in the resource upon retrieval, and (possibly) (4) must have previously endorsed the information (Clark and Chalmers 1998)—are too easily satisfied, with the consequence that extended cognition entails an "unacceptable proliferation" of

---

[2] But see Sprevak (2009) for an attempt to use the relationship between functionalism and the extended cognition hypothesis to ground a reductio of the latter.

[3] For an overview of the transition from first-wave to second-wave extended cognition theorizing, see Kirchhoff (2012).

"extremely short-lived" cognitive systems (Rupert 2004, p. 396).[4] The problem of cognitive bloat, however, may not be as severe as many participants in the debate (on both sides) have assumed. Second-wave extended cognition theorists, in fact, have tended to emphasize the fact that extended systems may in fact often be extremely short-lived—in other words, that we may just need to learn to live with cognitive bloat.

But learning to live with cognitive bloat will not make the problem go away entirely. It may not be a problem to posit extremely *short-lived* cognitive systems, but it is a problem to posit *too many* systems. In order for an approach to the ontology of cognitive systems to have any explanatory traction, it cannot imply that an extended cognitive system comes into existence any time the agent makes use of an external resource in order to accomplish a given cognitive task; we require a way of distinguishing between cases in which an agent merely makes use of an external resource and cases in which the agent and the resource constitute a joint cognitive system [cf. Adams and Aizawa (2008) on the coupling-constitution fallacy]. Second-wave extended cognition theorists thus still owe us a means of identifying and individuating cognitive systems. It is not obvious how the second-wave approach might provide such a means, however, especially as it tends to emphasize that coupling between agents and external resources is a complex, multidimensional affair (Heersmink 2012, 2015), with the result that it becomes difficult to see how the extended cognition theorist might "draw the line" in such a way as to prevent the theory from implying that an extended cognitive system pops into existence every time an agent relies to some extent on an external resource (and pops out of existence again as soon as he ceases to rely on it). Distributed cognition, which is likewise prepared to acknowledge short-lived distributed systems composed of temporarily interacting elements, is, as we will see, in the same boat when it comes to cognitive bloat.

Two additional challenges

In addition to the problem of cognitive bloat, approaches to the identification and individuation of cognitive systems face the interrelated challenges of scaling up and scaling down.

---

[4] In fact, it is unclear to what extent it is compatible with the spirit of first-wave extended cognition to impose criteria in addition to functional isomorphism; in retrospect, the "trust-and-glue" conditions to some extent look like an ad hoc attempt to prevent cognitive bloat (Palermos 2011).

Scaling up

"Scaling up" refers to ability of an approach to acknowledge that an assemblage of distinct entities, including individual agents, might, under the right conditions, itself count as a cognitive system. For straightforward reasons, intracranialism does not acknowledge that such a set of entities might qualify as a cognitive system: if cognition is brainbound by definition, it crosses neither the boundary between agent and artifact nor the boundary between agent and agent. Extended cognition was designed in part to overcome this limitation of intracranialism, but standard forms of extended cognition have difficulty scaling up in certain cases.

Extended cognition can of course acknowledge both agent-artifact and agent-agent systems. But due to its agent-based character, it flounders in cases where the system is not centered on a single human (or other biological) agent and in cases where no human is involved. While the possibility of purely artefactual cognitive systems may be consistent with the letter of extended cognition, extended cognition theorists have in practice focused largely on systems with humans in the loop. And while extended cognition may have no difficulty *identifying* a cognitive system in cases where the system is not centered on a single human agent, it will have difficulty *individuating* the relevant system. For the extended cognition theorist, cognitive systems are identified by starting from a human agent; they are individuated by determining which external resources (possibly including other human agents) are hooked up to the relevant agent in the right way: functional isomorphism (plus satisfaction of the trust-and-glue criteria), in the case of first-wave extended cognition; sufficiently tight coupling, in the case of second-wave extended cognition (see "The task-based approach" section below).

In the discussion to follow, we cannot hope to cover all forms of scaling up. Starting with Hutchins (1995a), distributed cognition researchers have provided detailed investigations of cases involving complex networks of agents and artifacts (see Sutton 2006; Dror and Harnad 2008; Michaelian and Sutton 2013). Here, we will focus on a relatively simple case involving a single agent and multiple artifacts but in which the relevant artifacts appear to be doing the lion's share of the work involved in performing the cognitive task (i.e., a case in which the putative system is not centered on the human agent), as such a case provides an appropriate illustration of the advantages of the task-based approach.

Scaling down

Intracranialism, extended cognition, and many standard versions of distributed cognition fail to meet the challenge

that we refer to as "scaling down": moving below the level of the agent to acknowledge that, in a given case, it may not be the agent as a whole that constitutes the relevant cognitive system but rather some subsystem within the agent. Cognitive neuroscientists and cognitive scientists more generally are often interested in distinguishing which parts of an individual's brain are implicated in a given cognitive process. Many tasks are associated with specific brain regions. These associated regions are inferred through a variety of methods, including lesion studies and neuroimaging. In this context, it often makes sense to view a certain brain area or structure as itself performing a cognitive task, with respect to which other brain areas count as its environment. The limbic system, for example, appears to be the seat of emotion, the frontal lobes the seat of planning, and the hippocampus is centrally involved in the consolidation process responsible for turning labile short-term memories into stable long-term memories.[5] These generalizations about brain areas are, of course, drastic oversimplifications, but it does appear that certain brain areas are much more implicated in particular cognitive tasks than others.

Intracranialism and extended cognition, again, are agent-based and thus lack a principled means of singling out sub-individual cognitive systems. Both tend to view the individual cognizer as a unified whole, a sort of black box, whose cognition may or may not extend to objects normally considered to be part of his environment. But the goal of cognitive neuroscience is precisely to open up the black box and figure out its workings—what its components are and how they interact in cognitive processing. The intracranialist view can do no better than to say that planning, for example, is something done by a cognitive agent, operating on his own. The extended cognition view can do no better than to say that planning is something done by a cognitive agent and which sometimes does and sometimes does not extend into external resources. Hence both approaches fail to scale down. Finally, distributed cognition, with its non-agent-based character, can in principle acknowledge sub-individual systems. As Hutchins recently put it, "the interesting question…is not 'is cognition distributed or is it not?' or even 'is cognition sometimes distributed and sometimes not distributed?' Rather, the interesting questions concern the elements of the cognitive system, the relations among the elements, and how cognitive processes arise from interactions among those elements" (Hutchins 2014). As Theiner (forthcoming)

points out, this means that the perspective of distributed cognition can be adopted with respect to systems at multiple scales; at a microscale, the perspective of distributed cognition might shade into that of embodied cognition, which views cognitive processes as distributed across not only the brain but also the body (Shapiro 2011). In practice, however, the focus of distributed cognition research tends to be firmly on supra-individual cognitive systems. Existing distributed approaches, moreover, fail to provide a principled means of defining sub-individual systems. The task-based approach developed here is distinctive in focusing equally on sub-individual and supra-individual systems and in attempting to provide a principled means of defining both kinds of system.

Distributed cognition as a non-agent-based approach

As noted above, moving to second-wave extended cognition already brings us closer to distributed cognition, and it is ultimately a form of distributed cognition that we want to defend here.

Distributed cognition is *like* extended cognition in that it argues that certain cognitive tasks [involved, e.g., in mathematics (Lave 1988) or reading (Donald 1993)] are accomplished in part by structures external to the agent. Cognition, viewed as the set of processes that accomplish cognitive tasks, is often distributed across richly interacting systems of minds and non-biological resources. This much, extended cognition can agree with. Distributed cognition is *unlike* extended cognition in that it goes further in the same direction, rejecting the agent-based approach to identifying and individuating cognitive systems. Hutchins, for example, famously analyzes ship navigation as being accomplished by a distributed system involving multiple agents and instruments, no one of which is responsible for the overall direction of the system (Hutchins 1995a). Similarly, he argues that the distributed system constituted by an aircraft's cockpit—including human agents, computers, display screens, and sensors—itself is responsible for piloting a plane (Hutchins 1995b).

In virtue of its non-agent-based approach, distributed cognition is better positioned than extended cognition to meet the challenge of scaling up. As noted above, however, standard versions of distributed cognition do have difficulty scaling down. Moreover, the threat of cognitive bloat affects distributed cognition just as much as extended cognition: the distributed cognition theorist must provide a principled approach to the bounds of cognition, a means of identifying and individuating (possibly short-lived) cognitive systems which does not imply an unacceptable proliferation of such systems. The following section of the paper is devoted to articulating a version of distributed cognition

---

[5] Similarly, artificial intelligence researchers will sometimes perform "ablation experiments" on their complicated programs. These are analogous to lesioning experiments in animals in that they remove a part of the program and observe the resulting behavior, allowing the researcher to draw conclusions regarding the functions of specific parts of the program.

intended to provide a means of identifying and individuating cognitive systems that successfully avoids this threat.

The task-based approach

An initial sketch

One point on which (virtually) all parties to the debate over the bounds of cognition can agree is that cognition is a matter of information processing.[6] This implies that cognition is something that happens, a *process* that unfolds. Though it is useful shorthand to describe certain entities as "cognitive" in virtue of their ability to play a role in cognitive processes, strictly speaking it is inappropriate to ask whether a given entity or system is, in itself, cognitive. The right question to ask is whether a given process is cognitive or not; systems themselves are cognitive only in a derivative sense.

What, then, is it for a process to be cognitive? We assume that a given process is cognitive if it is devoted to accomplishing some cognitive (i.e., information processing) *task*. Though there are bound to be gray areas, we can often agree without difficulty that certain tasks (e.g., planning) are cognitive and certain others (e.g., digestion) are not. Identifying cognitive tasks provides a means of identifying cognitive processes. Cognitive *systems* are then individuated relative to cognitive processes: the system responsible for carrying out a given process consists of all and only those entities that contribute (in the right way—see below) to the carrying out of the process. In some cases, the system may be equivalent to an individual agent. In other cases, the system will consist of one or more subsystems of an agent. In yet other cases, the system will consist of an agent together with one or more external resources, multiple agents, or multiple agents interacting with external resources.

To illustrate, let us return to the case of Otto. Do Otto and his notebook constitute a single cognitive system? On the task-based approach, the question has no fixed answer. If Otto happens to be holding his notebook while he watches a movie, the answer is "no": the notebook does not contribute to carrying out the cognitive processes responsible for accomplishing the task of understanding the movie, and there is thus no reason to count it as part of the cognitive system that is responsible for peforming the movie-watching task. If, on the other hand, Otto is

attempting to determine the location of an art exhibit and refers to the notebook for the address of the relevant museum, then the notebook can indeed reasonably be counted as part of the (distributed) cognitive system responsible for carrying out process that accomplishes the task of remembering the location of the exhibit.

Summing up, the core claim of the task-based approach is that cognitive systems are identified and individuated relative to cognitive tasks: once a given task is fixed, this allows us to identify the relevant cognitive process; a cognitive system is then individuated relative to the cognitive process.[7] In short, the claim is that all and only those entities that generate, store, manipulate, or transmit information or representations relevant to the given task are included in the cognitive system picked out by the task. We consider the ability of the task-based approach to meet the challenges of scaling up and scaling down in "Applying the task-based approach" section. The remainder of this section first looks at its ability to avoid unacceptable cognitive bloat.

Avoiding cognitive bloat

Discussing use of the World Wide Web to retrieve information, Clark and Chalmers ask: "Is my cognitive state somehow spread across the Internet?" (Clark and Chalmers 1998, p. 17), suggesting that the answer may be indeterminate. The task-based approach suggests that the answer will vary depending on the task being performed but that it is determinate relative to any given task: if the task at hand involves representations generated, stored, manipulated, or transmitted by components of the internet, then the relevant components should indeed be considered constituents of a distributed cognitive system (DCS) including the agent; those parts of the internet that are not involved in information processing relevant to the task at hand should not.

While its ability to provide a determinate answer in such cases is a virtue of the task-based approach, it might be objected that the approach faces a version of the problem of cognitive bloat yet more severe than other distributed/extended approaches. The thought is the following. As we have seen, agent-based approaches start from an uncontroversially cognitive agent and then ask whether that agent's cognitive processes loop out into (what we would ordinarily take to be) his environment. On the task-based approach, however, we lack an uncontroversially cognitive entity to start from; thus the view, as formulated risks counting as parts of DCSs entities that clearly should not be so counted. Consider a simple example. A person is solving a long division problem with a pencil and paper in a room.

---

[6] We say "virtually" because there are exceptions, e.g., Tallis (2004). While most theorists agree that cognition is a matter of information processing, there is controversy over the nature of information itself. In a recent review, for example, Crnkovic and Hofkirchner (2011) list several senses of information endorsed by scholars. Our conception of information is meant to be very inclusive, neutral among the kinds of information listed.

[7] This raises the question of how cognitive tasks are to be individuated; see "Identifying cognitive systems" section below.

In the next room, behind a locked door, is a whiteboard on which the answer to the problem is written. The worry is that, because the whiteboard encodes a representation relevant to the task performed by the person, the task-based view, as it stands, implies that the whiteboard is part of a DCS including the person (as well as the pencil and paper). This implication is highly implausible, as the person has no access to the answer—there is no *information flow* between the white board and the other components of the putative distributed system, so it should not be counted as a component of the system. Similarly, consider a set of people, in different far-flung locations, all working on the same sudoku puzzle on a Sunday morning. The worry is that the task-based approach will mistakenly classify them all as belonging to the same DCS.

In order to avoid these implications, two additions to the task-based view are required: first, we need a systematic way of identifying an initial cognitive entity (not necessarily an agent, but an entity involved in carrying out the given cognitive task); second, we need a more explicit notion of information flow.

Identifying cognitive systems

In general, we are able to recognize that a cognitive process is unfolding because we recognize a relationship between an informational input and an informational output. An initial cognitive entity can be identified by identifying an entity involved in securing this relationship. To take an easy example, if you ask someone his name (the input) and he responds by uttering a name (the output), you can assume that the production of the name involved a cognitive process. In this case, the only entity that might plausibly have secured the relationship between input and output is the person himself. To take a slightly more difficult—but still simple—case, you might observe a robot maneuvering around an obstacle. Assume that we know that its sensors are taking in data (the input) and that its actuators are responsible for guiding its movements (the output). The robot thus plays the role of the initial cognitive entity.

Of course, in the cases of most interest here, the initial cognitive entity may not be a (human or robotic) agent; we discuss such cases below. The key point, for now, is that identifying a cognitive system does not presuppose first identifying an agent, since we appeal directly to input–output relations. We identify a cognitive task by identifying an informational input–output relation. An initial cognitive entity can be identified by identifying the relevant sensors and actuators. The rest of the system is filled in according to the relevant representations and information flow, to which we turn in the next section. Note that, while our goal here is to provide a means not only of identifying (zeroing in on) but also of individuating (describing the

contours of) cognitive systems, and while we identify and individuate cognitive systems relative to cognitive tasks, it is sufficient for our purposes that we have a means of identifying cognitive tasks; i.e., providing an explicit means of individuating cognitive tasks is unnecessary. The role of tasks, in our framework, is to enable us to initially zero in on a cognitive system. The system is then individuated in terms of relations of information flow, and this does not require us to describe the contours of the initial cognitive task in precise terms, since the latter is simply our starting point.

Individuating cognitive systems

Identifying an initial cognitive entity enables us to identify a cognitive system. Once we have identified an initial entity, we are in a position to ask which other entities, if any, contribute to accomplishing the cognitive task—i.e., to individuate the cognitive system. The key notion here is *information flow*. Two agents should not be considered to be components of a single cognitive system if no information-bearing messages are passed between them. Conversely, if such messages are passed, we may take them to be components of a single system. Applying this approach to the problem case discussed above, the whiteboard containing a representation of the answer to the long division problem, located in the next room, is not included in the cognitive system including the person doing the problem along with the pencil and paper simply because no information is exchanged between the whiteboard and any of the components of the system. The same thing goes for the sudoku puzzle case.

The relevant notion of information is that used in information theory: information must be about something and must be interpretable by some entity. That is, information must *represent* something to some entity. Cognition itself is difficult to define in part because it spans many processes, such as reasoning, perceiving, judging, imagining, and remembering, and in part because we do not have a grip on what it might mean defined without reference to specifically *human* cognition (Poirier and Chicoisne 2006). One common characteristic of all forms of cognition, however, is the generation, storage, manipulation, and transmission of representations, where representation is loosely defined in terms of structures standing for their referents. For a structure to stand for, or represent, a referent, in turn, only makes sense with respect to an entity using the structure in place of the referent in a given task. A Chinese character, for example, may represent something for someone who can read Chinese, but may not represent anything for someone who cannot.

This point can help us to determine which entities count as parts of a cognitive system and which do not. In order to

be considered part of the system performing the task, an entity must be involved in generating, storing, manipulating, or transmitting representations relevant to the given task.[8] If all such entities are located within the brain of an individual, the system can be described in traditional intracranialist terms. If not, however, the relevant system is best viewed as a DCS. The task of addition provides an illustration. For most adults, adding 2 and 2 requires nothing outside the head. That is, the answer is retrieved from memory, and the performance of the task may be described in intracranialist terms. Adding 4123 and 8326, however, is more challenging. If a person uses a pencil and paper in order to calculate the answer, then the situation is optimally described as one in which the task is performed by a cognitive system distributed over the person's brain and the pencil and paper (as well, perhaps, as the person's arm and hand). The intracranialist, of course, will say that the pencil and paper are part of the environment with which the person is interacting (it is less clear, from an intracranialist perspective, what to say about the arm and hand). But the marks on the paper embody representations which are then manipulated by the person, making a DCS description apt. If the person uses a calculator to add 4 123 and 8 326, a DCS description is even more clearly appropriate, as the calculator itself generates, stores, and manipulates representations.

Summing up, the core of the view is that two entities are components of a single (distributed) cognitive system if (1) they are connected by a relation of information flow and (2) they generate, store, manipulate, or transmit representations relevant to a common cognitive task. Note that the task-based approach differs from approaches appealing to the notion of "coupling" invoked in extended cognition (e.g., Palermos 2014), especially its second-wave variant. Coupling refers to information flow, but specifically to bidirectional information flow, whereas on the task-based approach one-way information flow is sufficient to give rise to a DCS. As the discussion of re-representation in "Applying the task-based approach" section below illustrates, this is an important advantage of the task-based approach,

as requiring two-way information flow would rule out the existence of a DCS in many cases in which doing so would prevent us from providing a full explanation of the target phenomenon. At the same time, the task-based approach does provide a means of responding to the challenge of cognitive bloat. The approach does acknowledge the existence of many cognitive systems where we would not intuitively acknowledge their existence—that is, of course, the point of distributed and extended approaches in general. And, in common with second-wave extended cognition, it does acknowledge the existence of many short-lived systems. But it nevertheless provides us with a principled means of distinguishing between cases in which a set of entities constitutes a cognitive system (however long- or short-lived) and cases in which they do not.[9]

Applying the task-based approach

How do we know whether this particular means of avoiding cognitive bloat is preferable to alternatives (such as those based on bidirectional information flow)? We can offer no decisive argument for the task-based approach, but its ability to meet the challenges of scaling up and scaling down provides significant support. In this section, we consider two parallel case studies to show how the approach meets these challenges: re-representation in the human visual system (scaling down) and in a biomedical engineering laboratory (scaling up).[10]

Re-representation in the human visual system

In this section we will describe a theory of mental imagery developed by Kosslyn (1996), Kosslyn and Thompson (2003). While this theory is controversial, and the very existence of mental imagery is debated (e.g., Pylyshyn 2002), we will for present purposes take Kosslyn's influential account of visual mental imagery for granted. On that account, mental imagery in general is a matter of re-experiencing a remembered perceptual experience or some combination of remembered perceptual experiences. *Visual* mental imagery (to which we restrict our attention here) is, more specifically, a matter of the generation of "a pattern of activation in the visual buffer that is not caused by immediate sensory input" (Kosslyn 1996, p. 74). The visual buffer is a set of brain areas grouped as a single

---

[8] This implies that cognition necessarily involves representations. While this view is compatible with most approaches to cognition, it is rejected by radical brands of antirepresentationalism (Beer 1990; Brooks 1999; Gelder 1995; Chemero 2009). Many dynamical systems theorists, for example, reject talk of representations. While this approach has found applications in certain domains in fields such as robotics, dynamical systems theorists have difficulty providing a general account of cognition that does not appeal to representations, since cognition includes, in addition to processes that might be explained in terms of interaction between the cognizer and its environment, processes that refer to entities beyond the cognizer's environment (Clark 1998a). Consider, e.g., the difficulty of providing an account of mental time travel [i.e., remembering the past and imagining the future (Michaelian 2016; Michaelian et al. 2016)] without appealing to representations of past and future events.

[9] One might ask what additional explanatory power is gained by positing the existence of sub- and supra-individual cognitive systems; we return to this question in "Summing up" section.

[10] While we have chosen these cases because they involve analogous cognitive tasks performed at different scales, this should not be taken to suggest that we endorse the functionalist reasoning behind first-wave extended cognition.

functional structure; these areas are topographically organized visual areas in the occipital lobe. During perception, input from the eyes produces a pattern of neuron activation in these areas that spatially matches retinal input (Daniel and Whitteridge 1961). That is, the pattern of activation in the visual buffer spatially matches the pattern of light on the retina. The "attention window" specifies a contiguous set of points in this buffer for deep perceptual processing. The attention-shifting module moves the attention window so the reasoner can focus on different items in the buffer.

Kosslyn's model involves two perceptual modules, associated with specific brain areas: the ventral system, for identifying shape, color, and texture, and the dorsal system, which identifies location and size. Output from these modules is used by the associative memory system to match to stored information (for object recognition, categories, parts of speech, and so on). If the associative memory fails to identify what is needed, the information lookup module collects more information. This process is top-down, based on the already-activated parts of associative memory. These seven subsystems constitute the relevant players in Kosslyn's theory of mental imagery. When a visual memory is sufficiently primed, it activates the visual buffer, forming a mental image. These mental images are perceived with the same areas (the ventral and dorsal subsystems) that are used in normal perception, as brain imaging studies support (Farah et al. 1988). The same visual buffer used in visual perception is used for mental imagery. This re-representation occurs because the visual memories are not stored as patterns of spatially organized dots, but as interpreted structures that can function in many instances as propositions (both sides of the imagery debate agree on this issue; what they disagree on are whether these structured representations are ever turned into mental images).[11] Mental imagery takes this interpreted, or "descriptive," representation and turns it into a "depictive" representation: the mental image.

Of what use is re-representation to the cognizer? Different representational formats affect ease of retrieval of and inference from information (Marr 1982). So, for example, descriptively stored memories are relatively easy to retrieve because a symbol can be used as a cue (e.g., the letter D). A depictive representation, on the other hand, does not represent things like the letter D explicitly, but represents the points that make up the letter visually. If a reasoner has a need to describe the shape of the white space inside a capital letter D, it might prove necessary to generate a mental image of the letter and use the ventral

system to identify the shape of the whitespace: a half-circle.[12] Mental imagery is an instance of representational change: a propositional descriptive representation is transformed into a depictive mental image. The representational change enables the subject to infer information that is difficult to infer from the original representation: the mental image is perceived by perceptual processes in order to generate new beliefs.

Re-representation in a biomedical engineering laboratory

Before showing how the task-based approach applies to re-representation in the visual system, we describe a case of re-representation in a DCS. For 2 years, the first author (JD) was part of a research group conducting ethnographic studies of biomedical engineering laboratories, where distributed systems seem to be responsible for accomplishing many tasks (Nersessian et al. 2003, 2005).[13] Our discussion focuses on ethnographic data gathered during a study of "Lab A".[14] This laboratory has among its goals understanding natural blood vessel and cardiac tissues and engineering artificial blood vessels and cardiac tissue. We will look at a particular task that occurs in the laboratory and examine it with respect to re-representation.

Graduate students in Lab A create "constructs" of endothelial cells. In nature, endothelial cells line blood vessel interiors. In Lab A, these cells are "grown" onto glass rods, resulting in tubes about an inch long and a few millimeters thick. These tubes are cut, resulting in rectangular shapes. They are then put in a device called a "flow loop," which continuously passes media (a red liquid food for cells) over the construct, simulating blood flow in a vessel.

One of the laboratory members, A10, wants to see the result of this "shear stress" on the cells. One of the research themes of the laboratory is that mechanical stimulation can change cell structure. Since cells are microscopic, A10 needs to use a confocal microscope to see the effect of the shear stress on the cells. He remarks:

> "...So then I can tag that antibody, and then look under the confocal microscope and see the pattern of this protein expression on the cells."

---

[11] We use the term "proposition" here in the sense that it has in the knowledge representation literature: propositions are sentence-like representations, as opposed to, for example, images. We do not use it in the logical sense, in that, for example, propositions do not necessarily obey the law of excluded middle.

[12] Though re-interpretation is the major functional reason offered for mental imagery, empirical support for peoples' ability to do it at all has been spotty. See Chambers and Reisberg (1985) and Slezak (1992) for some failed attempts, and Finks et al. (1989) for a response.

[13] Nersessian was the first to suggest that scientific laboratories were distributed cognitive systems. See Giere (2002) for a useful description of this approach.

[14] Aliases are used to protect the anonymity of respondents and their laboratories.

The confocal microscope captures this florescence and outputs a graphic image on a computer screen that A10 looks at. From the image A10 gets a qualitative idea of what has happened to the cells. However, quantitative measures are taken from this image using the confocal microscopes software:

> "...I can trace around the cell and get geometric information from it. And I can use a line and follow the angle the majority of the filaments are in."

This results in numerical measures of cell angle (the shear stress tends to elongate the cells), which is used to test hypotheses regarding the relationship between shear stress and the resulting angle of the cells.

Scaling up and scaling down

An intracranialist, of course, might argue that there is no need, in order to provide an adequate description of the way this task is performed, to invoke a distributed cognitive system: Why can we not simply view A10 as the cognizer and the relevant artifacts (the flow loop, the confocal microscope, etc.) as the environment with which he interacts in performing the flow loop experiment? The performance of certain tasks may be best described in terms of an individual human cognizer interacting with his environment, but other tasks are better described in terms of the operation of a DCS (interacting with *its* environment). If we start with the task being performed and consider the process responsible for performing it, it is clear that the process is carried out by a system including not only the human agent but also the artifacts with which he interacts. Indeed, a task-based approach makes clear that there is no compelling reason to assign the human agent the central role in carrying out the process—the task-based approach, unlike the agent-based approach characteristic of extended cognition, might or might not single A10 out as the initial cognitive entity. A10 remarks:

> "So we use the flow loop as um, a first order approximation of a blood vessel environment, is like, in that, um, as the blood flows over the lumen, the endothelial cells experience a shear stress."

Note how the flow loop represents the natural blood vessel environment for A10. We can infer from this that the construct itself represents a blood vessel. Much like a human mind might manipulate an in-the-head model that represents some physical system, A10, by putting the construct in a flow loop, is manipulating an in-the-world model that represents some other physical system. The construct serves as a memory storage—the intracranialist would view it as an external store, but we treat it as internal to the DCS.

Once the flow manipulation is complete, the DCS must determine the shape and angle of the cells that make it up. The construct is treated with florescent tags and put under the confocal microscope. The microscope takes as input the construct (which, with reference to the DCS, is a representation) and outputs a different representation: a graphic image of the cells in the construct, magnified so that other parts of the DCS perceive the needed information. In this case, the perceptual modules used are the human visual perception system and the software built into the confocal microscope for the determination of cell shape and angle. The confocal and its software may also be classified as components of the DCS because they generate, store, and manipulate relevant representations.

We have described the flow loop experiment from the start of the experiment to the collection of results.[15] The right-hand column in Fig. 1 depicts a part of the process analogous to individual mental imagery, depicted in the left-hand column. Processing modules are contained in boxes; representations are contained in ovals. The labels on the left indicate the categories into which the modules and the representations fit. In the case of mental imagery, the initial representation is the internal descriptive representation, which is analogous to the construct after it emerges from the flow loop. These initial representations are taken as input by visual representation creators—the visualization module, in the case of the visual system, and the confocal microscope, in the case of the lab—which in turn output visual representations. These outputs—the depictive mental image and the graphic image—are then processed by visual inference modules, resulting in the final outputs at the bottom of the figure. Not only do the modules involved in each case behave similarly, but they have similar functions. The purpose of both visual representation creators is to transform representations such that certain kinds of information are easier to infer. Both visual inference modules have the function of extracting facts from data.

In contrast to intracranialism and extended cognition, task-based distributed cognition scales up smoothly. Intracranialism, of course, simply denies the possibility of DCSs. Extended cognition has difficulties with cases in which there is no (human) brain in the loop, but it also provides an implausible description of cases—such as that reviewed here—in which there is a brain in the loop but in which it plays a subsidiary role. In the extended cognition framework, the brain in effect takes credit for the contributions of the artifactual components of the DCS. In cases where artifactual components are doing the lion's share of the work, this leads to a distorted picture of the workings of
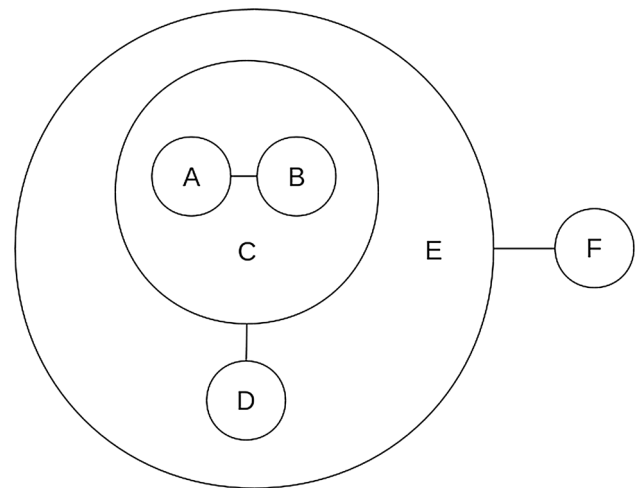
---

[15] After the results are collected, statistics are run on the data. Since software is used to run these analyses, it too is done by a DCS consisting of, at the very least, a person and a computer program.
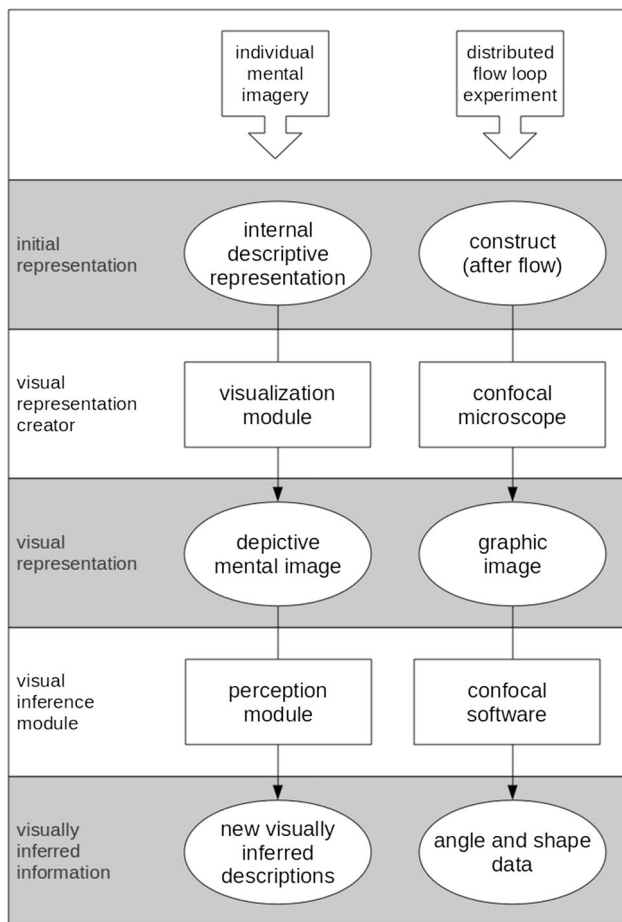
Fig. 1 Comparison of individual and distributed re-representation

Fig. 2 Levels of abstraction in the task-based approach

the cognitive system. Consider another case, running a factor analysis with a statistics software package: making the human the center of the system in this case seems grossly unfair to the software—in practical terms, such analyses are essentially impossible to conduct without the aid of a computer. In the case of the flow loop experiment, the human agent is unaware of, let alone a participant in, the complex representational changes conducted by the confocal microscope. The task-based approach, in contrast, since it regards all entities which contribute to the performance of the relevant task by generating, storing, manipulating, or transmitting representations as equal components of the DCS singled out by the task, is able to be consistent in its treatment of brains, microscopes, computer software, and so on.

The task-based approach likewise scales down more successfully. If the starting point for identifying a cognitive system is, as in intracranialism and extended cognition, the agent or his brain, taken as a unit, it makes little sense to speak of certain brain areas carrying out cognitive processes on their own, as in the case of visual re-representation. The task-based approach, in contrast, handles such

cases comfortably. Because it holds that inclusion in a cognitive system depends simply on the dealing with task-relevant representations, there is no need, when identifying a cognitive system, to start from a whole agent. For example, the frontal lobes (perhaps) carry out the processing involved in planning, making use of long-term memory and whatever other brain areas deal with representations relevant to the task of planning. The brain areas that do not deal with relevant representations are not, on this view, part of the cognitive system defined by the task. In other words, the agent is sometimes too coarse a level of analysis relative to a given task.

In short, the task-based approach is able to accommodate multiple levels of analysis, from brain areas to large-scale distributed systems. Figure 2 illustrates the utility of such a flexible approach. Information processing might be carried out by A, interacting with B as its environment. Relative to another task, it may be appropriate to view A and B together as constituting a cognitive system C, which in turn interacts with D as its environment. Relative to yet another task, it may be preferable to view C and D as constituting E, a cognitive system which interacts with its environment F. And so on. Where, in Fig. 2, is the agent? It might be that A and B are different parts of the human brain, such as the visual system and the hippocampus, which operate together—as C—to encode memories from D, a computer screen. F may be another agent requesting the information that the system E can provide on the basis of its interaction with the screen. We could go further in the same direction, viewing E and F themselves as constituting a DCS, perhaps trying to complete a joint project. In another scenario, A might be an agent and B might be his smartphone. In yet another scenario, A and B might be two agents cooperating. The point is that the elements that we should consider to be part of the cognitive system depend on the task the performance of which we want

to understand. An individual agent may be viewed as anything from black box to a complex system consisting of many interacting components, depending on the task in question. Its ontological flexibility—its ability to accommodate this point by moving flexibly among different levels of analysis—is the key advantage of the task-based approach, relative to extended cognition and other agent-based approaches.

Summing up

One might ask why all this matters: who cares whether we make the unit of analysis the individual plus his environment or a DCS? For one thing, while there are of course certain tasks, such as dreaming, which are entirely brainbound, a great deal of the tasks we perform are performed in a distributed manner; as Clark and Chalmers (1998) point out, limiting ourselves to agent–environment interactions gives rise to overly complicated descriptions. Moreover, the decision about where to draw the line between the agent and his environment has implications for areas such as cognitive modeling. Most cognitive modeling architectures have only very limited support for representation and processing outside the head. To the extent that we accept something like the notion of a task-based DCS, creators of cognitive modeling architectures will have reason to provide additional support for cognitive artifacts in models.[16]

The two key contributions of this paper are (1) the proposal to identify cognitive systems in terms of the cognitive tasks they perform and (2) the proposal to individuate cognitive systems in terms of the entities that deal with task-relevant representations and that are involved in a flow of information with the initial cognitive object. As an illustration of this task-based approach, we have presented a case study of distributed cognition in a biomedical engineering lab, providing a detailed description of the similarities between the roles played by the (human and artefactual) components of the laboratory and those played by the brain structures involved in individual mental imagery.

The extension of the concept of cognition that we are advocating here is in line with certain other views in the literature. In their study of ants, Höldobler and Wilson, for example, claimed that "[t]he colony is a superorganism. It can be analyzed as a coherent unit and compared with the organism in the design of experiments, with individuals treated as the rough analogues of cells"; the hope is that, by regarding the ant colony as an organism, "more general and exact principles of biological organization will be revealed by the meshing of information from insect sociobiology with equivalent information from developmental biology" (Höldobler and Wilson 1990, pp. 2–3). The view that we advocate is that a similar hope is reasonable in the case of cognition. While a task-based view of distributed cognitive systems may be surprising when spelled out explicitly, it is consistent with the practice of a good deal of cognitive science; computational cognitive scientists, for example, regularly provide analyses of cognitive capacities which abstract away from the details of the way in which the capacities in question are implemented in human individuals [see, e.g., Van Rooij (2008) on the tractable cognition thesis], and such analyses can in principle be applied to sub- and supra-individual cognitive systems. Cognitive tasks are performed at multiple scales, ranging from clusters of neurons, to brain areas, to organisms, to suitably-integrated systems of organisms and artifacts, to intelligent (software or robotic) artifacts or systems of artifacts operating independently.[17] Restricting cognition to a single scale threatens to inhibit the transmission of insights among fields ranging from neuroscience to cognitive anthropology. A field as challenging as cognitive science can ill afford to deprive itself of such insights.

References

Adams F, Aizawa K (2001) The bounds of cognition. Philos Psychol 14(1):43–64

Adams F, Aizawa K (2008) The bounds of cognition. Wiley-Blackwell, Hoboken

Beer RD (1990) Intelligence as adaptive behavior. Academic Press, Cambridge

Brooks RA (1999) Cambrian intelligence: the early history of the new AI. MIT Press, Cambridge

Chambers D, Reisberg D (1985) Can mental images be ambiguous? J Exp Psychol Hum Percept Perform 11(3):317–328

Chemero A (2009) Radical embodied cognitive science. MIT Press, Cambridge

Clark A (1998a) Being there: putting brain, body, and world together again. MIT Press, Cambridge

Clarke A (1998b) Author's response. Metascience 7(1):95–104

Clark A (2001) Reasons, robots and the extended mind. Mind Lang 16(2):121–145

---

[16] The architecture CLARION (Sun 2005) supports the modeling of multi-person cognition, though not (or at least not explicitly) of cognitive artifacts.

---

[17] Swarm intelligence (Panigrahi et al. 2011) provides one illustration of the insights that can be obtained when we are prepared to view processes unfolding at different scales—including the scale of multi-organism systems—as engaged in cognition.

Clark A (2007) Curing cognitive hiccups: a defense of the extended mind. J Philos 104(4):163–192

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

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

Clowes RW (2013) The cognitive integration of e-memory. Rev Philos Psychol 4(1):107–133

Daniel PM, Whitteridge D (1961) The representation of the visual field on the cerebral cortex in monkeys. J Physiol 159(2):203–221

Dodig-Crnkovic G, Hofkirchner W (2011) Floridi's "open problems in philosophy of information", ten years later. Information 2(2):327–359

Donald M (1993) Origins of the modern mind: three stages in the evolution of culture and cognition. Harvard University Press, Cambridge

Dror IE, Harnad S (eds) (2008) Cognition distributed: how cognitive technology extends our minds. John Benjamins Publishing, Amsterdam

Farah MJ, Peronnet F, Gonon MA, Giard MH (1988) Electrophysiological evidence for a shared representational medium for visual images and visual percepts. J Exp Psychol Gen 117(3):248

Finks RA, Pinker S, Farah MJ (1989) Reinterpreting visual patterns in mental imagery. Cogn Sci 13(1):51–78

Giere RN (2002) Scientific cognition as distributed cognition. In: Carruthers P, Stich S, Siegal M (eds) Cognitive bases of science. Cambridge University Press, Cambridge

Giere RN (2011) Distributed cognition as human centered although not human bound: reply to Vaesen. Soc Epistemol 25(4):393–399

Giere RN (2012) Scientific cognition: human centered but not human bound. Philos Explor 15(2):199–206

Heersmink R (2012) Mind and artifact: a multidimensional matrix for exploring cognition-artifact relations. In: Proceedings of AISB/IACAP world congress 2012

Heersmink R (2015) Dimensions of integration in embedded and extended cognitive systems. Phenomenol Cogn Sci 14(3):577–598

Hölldobler B, Wilson EO (1990) The Ants. Harvard University Press, Cambridge

Honderich T (2014) Actual consciousness. Oxford University Press, Oxford

Huebner B (2014) Macrocognition: a theory of distributed minds and collective intentionality. Oxford University Press, Oxford

Hutchins E (1995a) Cognition in the wild. MIT Press, Cambridge

Hutchins E (1995b) How a cockpit remembers its speeds. Cogn Sci 19(3):265–288

Hutchins E (2011) Enculturating the supersized mind. Philos Stud 152(3):437–446

Hutchins E (2014) The cultural ecosystem of human cognition. Philos Psychol 27(1):34–49

Kirchhoff MD (2012) Extended cognition and fixed properties: steps to a third-wave version of extended cognition. Phenomenol Cogn Sci 11(2):287–308

Kosslyn SM (1996) Image and brain: the resolution of the imagery debate. MIT Press, Cambridge

Kosslyn SM, Thompson WL (2003) When is early visual cortex activated during visual mental imagery? Psychol Bull 129(5):723

Lave J (1988) Cognition in practice: mind, mathematics and culture in everyday life. Cambridge University Press, Cambridge

Manzotti R (2011) The spread mind: is consciousness situated? Teorema 30(2):55–78

Marr D (1982) Vision: a computational investigation into the human representation and processing of visual information. MIT Press, Cambridge

Michaelian K (2012) Is external memory memory? Biological memory and extended mind. Conscious Cogn 21(3):1154–1165

Michaelian K (2014) JFGI: from distributed cognition to distributed reliabilism. Philos Issues 24(1):314–346

Michaelian K (2016) Mental time travel: episodic memory and our knowledge of the personal past. MIT Press, Cambridge

Michaelian K, Sutton J (2013) Distributed cognition and memory research: history and current directions. Rev Philos Psychol 4(1):1–24

Michaelian K, Klein SB, Szpunar KK (eds) (2016) Seeing the future: theoretical perspectives on future-oriented mental time travel. Oxford University Press

Nersessian NJ, Kurz-Milcke E, Newstetter WC, Davies J (2003) Research laboratories as evolving distributed cognitive systems. In: Markman A, Barsalou L (eds) Proceedings of the 25th annual conference of the cognitive science society. Erlbaum, pp 857–862

Nersessian NJ, Kurz-Milcke E, Davies J (2005) Ubiquitous computing in science and engineering labs: a case study of a biomedical engineering lab. In: Kouzelis G, Pournari M, Stoeppler M, Tselfes V (eds) Knowledge in the new technologies. Peter Lang, Bern, pp 167–195

Palermos SO (2011) Belief-forming processes, extended. Rev Philos Psychol 2(4):741–765

Palermos SO (2014) Loops, constitution, and cognitive extension. Cogn Syst Res 27:25–41

Panigrahi BK, Shi Y, Lim M-H (2011) Handbook of swarm intelligence. Springer, Berlin

Poirier P, Chicoisne G (2006) A framework for thinking about distributed cognition. Pragmat Cogn 14(2):215–234

Pylyshyn ZW (2002) Mental imagery: in search of a theory. Behav Brain Sci 25:157–182 4

Rowlands M (2009) Extended cognition and the mark of the cognitive. Philos Psychol 22(1):1–19

Rupert RD (2004) Challenges to the hypothesis of extended cognition. J Philos 101(8):389–428

Rupert RD (2013) Memory, natural kinds, and cognitive extension; or, Martians dont remember, and cognitive science is not about cognition. Rev Philos Psychol 4(1):25–47

Shapiro L (2011) Embodied cognition. Routledge, London

Slezak P (1992) When can visual images be re-interpreted. In: Proceedings of 14th annual conference of the cognitive science society, pp 124–129

Smithies D (2013) The nature of cognitive phenomenology. Philos Compass 8(8):744–754

Sprevak M (2009) Extended cognition and functionalism. J Philos 106(9):503–527

Strawson G (2011) Cognitive phenomenology: real life. In: Bayne T, Montague M (eds) Cognitive phenomenology. Oxford University Press, Oxford, pp 285–325

Sun R (2005) The CLARION cognitive architecture: extending cognitive modeling to social simulation. In: Sun R (ed) Cognition and multi-agent interaction. Cambridge University Press, Cambridge

Sutton J (2006) Distributed cognition: domains and dimensions. Pragmat Cogn 14(2):235–247

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

Tallis R (2004) Why the mind is not a computer: a pocket lexicon of neuromythology. Imprint Academic, Exeter

Theiner G (forthcoming) Group-sized distributed cognitive systems. In: Routledge handbook of collective intentionality. Routledge

Van Gelder T (1995) What might cognition be, if not computation? J Philos 92(7):345–381

Van Rooij I (2008) The tractable cognition thesis. Cogn Sci 32(6):939–984