

Dimensions of integration in embedded and extended cognitive systems

Richard Heersmink

Published online: 1 March 2014

© Springer Science+Business Media Dordrecht 2014

Abstract The complementary properties and functions of cognitive artifacts and other external resources are integrated into the human cognitive system to varying degrees. The goal of this paper is to develop some of the tools to conceptualize this complementary integration between agents and artifacts. It does so by proposing a multidimensional framework, including the dimensions of information flow, reliability, durability, trust, procedural transparency, informational transparency, individualization, and transformation. The proposed dimensions are all matters of degree and jointly they constitute a multidimensional space in which situated cognitive systems can be located and have certain dimensional configurations. These dimensions provide a new perspective on the conditions for cognitive extension. They are, however, not meant to provide a set of necessary and sufficient conditions, but to provide a toolbox for investigating the degree and nature of the integration of agent and artifact into “new systemic wholes”. The higher a situated system scores on the proposed dimensions, the more functional integration occurs, and the more tightly coupled the system is.

Keywords Situated cognition · Distributed cognition · Extended mind · Complementarity principle · Second-wave extended mind theory · Cognitive integration · Cognitive artifacts · Information · Function

1 Introduction

Human agents use their social and material environment to help them perform their cognitive tasks, resulting in cognitive systems that are situated in a certain environment. Extended mind theory (EMT) is helpful to better understand the nature of this situatedness. It claims that human thought is, under certain conditions, distributed across an embodied agent and cognitive artifacts or other external resources (Clark 1997, 2008b; Sutton 2006, 2009b, 2010; Menary 2007). Such environmental resources are then not seen as external aids or scaffolds for thinking, but are sometimes proper

R. Heersmink (✉)

ARC Centre of Excellence in Cognition and Its Disorders, Macquarie University, Sydney, Australia
e-mail: richard.heersmink@gmail.com

parts of a distributed or extended cognitive process. Cognitive states and processes are thus conceptualized as hybrids or amalgamations of neurological, bodily, and environmental objects and processes (Wilson 2004; Rowlands 1999, 2010). On this view, some of the vehicles of cognitive processes are sometimes not located in the brain, but in the cultural and artifactual environment. It is thus a metaphysical claim about the location of human cognition, implying rather significant epistemological and methodological consequences for the philosophical and scientific study of human thought.

1.1 Parity and complementarity

John Sutton (2010) has identified two distinct, but historically overlapping, waves in EMT. The first wave is mostly based on the parity principle and is advocated by Andy Clark and David Chalmers (1998), Mike Wheeler (2010, 2011), and others. The parity principle is as follows:

“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).

The parity principle focuses on functional isomorphism between internal states and processes and external states and processes. By stressing functional isomorphism, it downplays differences between internal and external states and processes, in that way implying that the specific properties of cognitive artifacts are not important for better understanding situated cognitive systems. As long as there is a relevant (and rather generic) similarity with internal states or processes, then we do not need to pay attention to the specific properties of the external resource. This is unfortunate as the particular properties of the external resource are often important for better understanding how the cognitive task is performed (Norman 1993; Zhang 1997; Vorms 2012).

Stressing isomorphism is, moreover, not fruitful as there are differences between the internal and external. In case of memory, for example, there are differences in how internal memories and external memories are stored and processed. Internal memories are stored in neural networks that are subject to blending and interfering, which means they are shaped and updated according to other previously stored and new incoming information. Whereas external memories are stored in discrete iconic or symbolic format that are static, have no dynamics, and are not automatically integrated with other information (Sutton 2010). There are thus relevant differences in the informational properties of internal and external memories.

These and other drawbacks of the parity principle have led some theorists to argue for a different route to extended minds. This route is based on the complementarity principle, which downplays functional isomorphism and emphasizes complementarity between the internal and the external. On such a view, cognitive artifacts need not be similar to internal states and processes, but often complement internal states and processes with *different* properties and functions. In fact, complementing brain functions is often the point of deploying cognitive artifacts: so that they can perform functions the brain cannot do or cannot do well. There is often no point in externally replicating what the brain can already do. Jointly, brain-plus-artifact is often a more

versatile and powerful problem-solving system than an embodied brain alone. On a complementarity view, then, we should study the variety of situated cognitive systems by focussing on the particular properties of both embodied agents and cognitive artifacts (Sutton 2010) and by focussing on the functional integration of the two (Menary 2007, 2010a, b; Greenwood 2013a). Wilson and Clark defend a similar view. They write:

“Contrary to any requirement of fine-grained similarity then, what the friends of extended cognition actually expect, and study, are hybrid processes in which the inner and the outer contributions are typically highly distinct in nature, yet deeply integrated and complementary” (Wilson and Clark 2009, p. 72).

Wilson and Clark thus claim that internal and external parts of situated cognitive systems are deeply integrated and exhibit complementary properties. They further point out that “Tracing and understanding such deep complementarity is, we claim, the single most important task confronting the study of situated cognition” (Wilson and Clark 2009, p. 70). However, what they do not point out is how and how deeply the internal and external components are integrated, i.e., they do not look at the processes and degrees of integration. Given that the complementary properties and functions of cognitive artifacts and other external resources are integrated into the onboard cognitive system to varying degrees, it is important to have the tools to conceptualize the degree and varieties of integration. The goal of this paper is therefore to develop some of those tools by proposing and further developing a number of dimensions along which to conceptualize integration.

1.2 Method, motivation and organization

Sutton (2006; see also Sutton et al. 2010), Wilson and Clark (2009), Kim Sterelny (2010), and Richard Menary (2010a) have articulated the idea of a dimensional analysis of the coupling between agents and external resources. By synthesising and building on their work, I develop a multidimensional framework for conceptualizing the different kinds of cognitive interactions and complementary integration between agents and artifacts. This framework consists of the following dimensions: information flow, reliability, durability, trust, procedural transparency, informational transparency, individualization, and transformation. See Heersmink (2012) for a preliminary version of this framework. These dimensions are all matters of degree and jointly they constitute a multidimensional space in which situated cognitive systems can be located and have certain dimensional configurations. The higher a particular system scores on these dimensions, the more deeply the functional integration is.

There are at least two reasons why we need such a multidimensional framework. First, because it provides a toolbox for a detailed study of conceptual and empirical cases of situated cognitive systems, in that way providing us with a richer understanding of the variety of complementary relationships between agents and artifacts. This toolbox is not only helpful for philosophers interested in situated cognition, but also for psychologists, anthropologists, Human-Computer Interaction theorists, design theorists, media theorists, and others interested in cognition-technology relations. It may guide empirical research by providing the dimensions along which to study and observe

agent-artifact interactions. Importantly, although this toolbox is developed primarily to analyse agent-artifact relations, it should be equally useful to analyse agent-agent systems. Conceptual and empirical research on socially embedded or distributed cognitive systems (e.g. Tollefson et al. 2013; Theiner 2013) would thus also benefit from this framework.

Second, conceiving of situated cognitive systems in terms of dimensions that are matters of degree and differ in each case, provides a more realistic view on such systems. Seeing situated systems as either embedded or extended is not a particularly fruitful way of conceptualizing such systems, as some may be more embedded or extended than others (Sutton et al. 2010; Michaelian and Sutton 2013; Arnau et al. 2013). The nature of their embeddedness or extendedness depends on the degree of integration and, consequently, there is a grey area between systems that are embedded or extended. So, rather than providing a set of necessary and sufficient conditions for cognitive extension and distributed cognition, this framework provides a toolbox for investigating the degree and nature of integration of agents and artifacts.

The paper is organised as follows. It starts by concisely outlining complementarity-based EMT (Section 2). It then presents the dimensions for conceptualizing the complementary integration between embodied agents and cognitive artifacts (Section 3). Next, it discusses different perspectives on functions of artifacts and of larger integrated systems (Section 4). It ends by investigating the relations between some of the dimensions and briefly looks at how situated cognitive systems can change their dimensional configurations and shift to different regions in the multidimensional space (Section 5).

2 The complementarity principle

Given the drawbacks of the parity principle mentioned in the introduction (see also Rupert 2004; Sutton 2010; Walter 2010), Sutton has identified and articulated a distinctive route to extended cognition based on what he refers to as the complementarity principle:

“In extended cognitive systems, 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 (enduring or temporary) 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).

This principle downplays functional isomorphism between inner and outer states and processes, and argues for complementary properties and functions of cognitive artifacts or other external resources. Human brains are not good at storing large chunks of information for a long period of time in discrete manner and therefore offload information onto the environment that supplements our internal memory and information-processing capacities. “Biological traces are typically integrative, active, and reconstructive, but in using them we hook up to more enduring and transmissible exograms, mostly of our own making, which supplement and extend our powers”

(Sutton 2010, p. 205). So external information systems exhibit different properties as compared to internal information systems, in that way complimenting biological memory. Cognitive artifacts, however, do not just complement our memory, but a variety of cognitive capacities. Consider two brief examples.

First, mental imagery is rather constrained in that it is relatively fixed, i.e., it is difficult to see new forms and components of a mental image, and it is limited, which means there is only so much we can imagine internally. By sketching, we offload information onto the environment, which may not only prompt and generate new forms and components, but also dramatically increases the amount of information an artist or designer can work with, as it provides much more detail and precision (Clark 2001; Tversky and Suwa 2009). These properties of externalized information rather straightforwardly complement the constraints and limitations of our mental imagery, making us more powerful and creative designers and artists.

Second, internal reasoning processes have limited capacities, in part because we can perform only certain operations on internal information such as e.g. propositions, and in part because that internal information has a limited capacity in terms of the amount of information it can contain at any given moment. Thus when we develop an argument or line of thought, we can do only so much internally. But when using written language, e.g. with the aid of a word-processor, we can formulate much better and elaborate arguments or ideas. This is so because the representational malleability that word-processors afford, complements internal information-processing systems by allowing text to be erased, copied, restructured, reformatted, or otherwise manipulated. It thus allows us to perform operations on the offloaded information that are very hard, if not impossible, to perform in the brain. It also allows us to manipulate and work with significantly larger amounts of information. These functionalities complement the reasoning capacities of embodied agents. So, despite Socrates' well-known worries about the effects of writing on memory, it does allow us to significantly augment our onboard cognitive capacities.

These brief examples show that embodied agents deploy the functional and informational properties of cognitive artifacts to complement their onboard cognitive capacities. The human brain develops in an information-rich environment and has to learn how to complement its shortcomings by using the functionalities of artifacts to perform cognitive tasks in such a way that it does not waste expensive internal resources (see also Kiverstein and Farina 2011). The idea is that why perform a cost-expensive memory or information-processing task, when it can be delegated to an artifact?¹ A complementary integrated agent-artifact system is usually more adaptable and powerful in problem-solving than an embodied agent without such artifacts.² "Brains like ours need media, objects, and other people to function fully as minds" (Sutton 2010, p. 205). On a complementarity view, artifacts or other resources do not just augment existing brain functions by externally replicating them, but add new

¹ See also Clark's (1997) "007 principle" and Mark Rowlands' (1999) "Barking Dog principle".

² A reviewer pointed out that, in some cases, using a cognitive artifact to perform some task does not always result in a more adaptable and powerful problem-solving system. For example, when the artifact provides too much information, making it hard to decide which information is relevant for performing the task, or when the information is not sufficiently transparent (see Section 3.6). So not all agent-artifact systems are necessarily more adaptable.

functionalities to existing ones by integrating them (in varying degrees) into a plastic neural architecture.

The complementarity principle does not have the structure of a conditional statement. Unlike the parity principle, it does not connect two conditional premises with an “if P, then Q” structure. Thus, it does not argue that if a cognitive artifact has a complementary role or function to inner processes or states, then it counts as a constitutive part of an extended cognitive process or system. Many things external to the brain somehow have complementary roles to internal cognitive states and processes. Formulating complementarity in terms of a conditional statement would leave it vulnerable to the cognitive bloat objection (Adams and Aizawa 2008). So rather than trying to stipulate what exactly demarcates an embedded system from an extended or distributed system, Sutton (2006, 2008, 2010) provides a taxonomy of the dimensions that are relevant for better understanding such systems. In the next section, we will look at these and other relevant dimensions.

3 Dimensions for integration

Having briefly outlined complementarity-based EMT, let me now present the dimensions for conceptualizing the integration between embodied agents and artifacts into “new systemic wholes” (see Clark 2007). In their article, *The Extended Mind*, Clark and Chalmers (1998) put forward a number of conditions that describe the coupling between agents and cognitive artifacts. They argued that trust, reliability, accessibility, and past endorsement are necessary conditions for cognitive extension. Clark & Chalmers acknowledged that past endorsement as a condition for cognitive extension may be questionable, but trust, reliability, and accessibility are crucial. These much debated conditions³ are now referred to as “trust and glue”. If one of these three conditions is not sufficiently satisfied, then cognition is not extended. So it provided a fairly clear set of criteria to demarcate between an external resource that is a proper part of an extended cognitive process and one that has a mere causal influence on an internal cognitive process.

Roughly a decade later in the extended mind debate, a number of theorists articulated the idea of a dimensional analysis of the cognitive relations between agents and external resources (Sutton 2006; Sutton et al. 2010; Wilson and Clark 2009; Sterelny 2010; Menary 2010a). Wilson and Clark identify two dimensions. First, the nature of the external resources, which may be natural, technological, or socio-cultural. Second, the durability and reliability of the overall situated cognitive system. Sutton et al. (2010) take the dimensions of durability and reliability as well as the dimensions of “trust and glue” mentioned earlier by Clark and Chalmers (1998), and also briefly develop the dimension of transparency. Sterelny discusses three dimensions: trust, individualization, and individual versus collective use. And, finally, Menary focuses on the dimensions of manipulation and transformation. All these dimensional approaches are perceptive and clearly help us in better understanding the cognitive

³ For critical discussions of these conditions, see Rupert (2004), Rockwell (2010), Roberts (2011), Adams and Maher (2012), Palermos (2014), Smart (2012), Michaelian (2012), Halpin (2013), Arango-Muñoz (2013), and Ludwig (2014).

relationship between agents and external resources, but they focus on certain dimensions while overlooking others. Moreover, it is not at all clear whether these are all dimensions in the same sense of the word. Wilson and Clark, for example, refer to the nature of the external resource (i.e. either natural, technological, or socio-cultural) as a dimension. Sutton et al. and Sterelny refer to trust as a dimension. And Menary refers to manipulation of external resources as a dimension. These theorists thus seem to use the notion of a dimension to indicate quite different aspects of situated cognitive systems.

In this section, therefore, I first briefly clarify the notion of a dimension. I then refine and synthesize some of the above mentioned dimensions into a coherent and systematic multidimensional framework, add a number of dimensions to the framework, and examine where and how some of these dimensions overlap and interact. The framework I am proposing is not meant to be exhaustive, but compared to the above mentioned frameworks, the one developed in this paper is more elaborate in that it contains more dimensions and is more integrative in that it examines how the dimensions relate to each other and interact.

I now clarify and distinguish two elements that are relevant for better understanding the underlying conceptual structure of the dimensions. These elements are: (1) the cognitive profile of the embodied agent and (2) the informational and functional properties of the cognitive artifact. Each dimension emerges out of the interaction between these elements. The dimensions are thus relational in the sense that they never depend on only one of those elements. So, in contrast to Wilson and Clark, who refer to the nature of external resources as a dimension, I use a different notion of dimension which is more in line with Sutton et al. (2010) and Sterelny (2010), i.e., a dimension describes certain properties of the *relation* between an embodied agent and an external resource, not about intrinsic properties of either. Before presenting the dimensions, it is helpful to point out that the eight dimensions are presented in a non-hierarchical order. It may be the case that some dimensions turn out to be more important than others in certain situated systems, which will be investigated in future case studies.

3.1 Information flow

The first dimension concerns the direction and kind of information flow between the components of situated cognitive systems. Before distinguishing different kinds of information flow, it is helpful to first briefly introduce the notion of an epistemic action (Kirsh and Maglio 1994; contrast Loader 2012). An epistemic action is an action whose function is to improve cognition by changing the input to an agent's information-processing system, in that way making certain cognitive tasks easier, faster, more reliable, or possible at all. So when epistemic actions are performed, information is propagated within a situated cognitive system and, therefore, epistemic actions are necessary conditions for information flow. To better understand this information flow, we should conceptualize how and why agents interact with their artifacts. In Michaelian and Sutton's words: "The cognitive scientists' aim then is to examine the microprocesses of interaction across the diverse components of these distributed and heterogeneous cognitive systems, tracing for example the propagation and transformations of particular representational states across distinct (internal and external) media" (Michaelian and Sutton 2013, p. 5). Below I identify four kinds of information flow,

including one-way, two-way, and reciprocal information flow, as well as information flow in larger systems comprising more than one agent and more than one artifact.⁴

The first kind is characterized by one-way information flow from artifact to agent. Examples include clocks, compasses, road signs, maps, dictionaries, newspapers, websites, textbooks, graphs, diagrams, manuals, timetables, etc. Such artifacts are usually open access systems in that they are open to anyone able to interpret the information. Further, and this is essential, the agent typically does not have any influence on the informational content. Such artifacts and the information they provide are designed by other agents, including writers, journalists, designers, engineers, etc. One can thus argue that such artifacts mediate information flow between the creators of the information and its users.

The second kind is characterized by two-way information flow, i.e., from agent to artifact and then from artifact to agent. Humans frequently offload information onto their environment to relieve their memory burdens, thereby creating cognitive artifacts such as post-it notes, diary entries, and shopping-lists.⁵ Artifacts in two-way relations are often tailored for individual use and are frequently not part of publicly available representational systems. They are typically closed systems in that the artifact is meant for an individual agent who has designed the informational content of the artifact for individual use. Once the information is offloaded, it remains fixed and is thus not transformed during a task. In one-way systems, designers outside the system have designed both the physical structure and informational content of the artifact. But in two-way systems, designers have designed the physical structure of the artifact, e.g. the physical structure of a notebook, but the informational content is designed by the user, which is distinctive from one-way systems.

The third kind is based on reciprocal information flow. Cognitive artifacts are sometimes integral parts of ongoing information-processing tasks. Writing an academic paper, solving a difficult calculation with pen, paper, and numerals, or designing an architectural blueprint, often involve small incremental steps. We do not have a finished paper, calculation, or blueprint in our head and then fully offload it onto the artifact. Rather, we offload small bits of information onto the artifact, and the nature and content of the offloaded information contributes to and partly determines the next step in the overall process. For example, when writing an academic paper one often starts with a rough outline, which may prompt ideas about how to fill in the details. Filling in the details may then prompt an adjustment of the outline, which may in turn prompt further details. This process may continue for a number of cycles. Each step in the overall process builds and depends on previous steps. Agent and artifact continuously exchange information and so there is a reciprocal and cumulative information flow that constantly transforms the situated cognitive system. There is “continuous reciprocal causation” between agent and artifact (Clark 1997).

⁴ The dimension of information flow is somewhat related to Sterelny's (2010) dimension of individual versus collective use, but is more fine-grained.

⁵ A reviewer pointed out that, whether a situated system should be classified as two-way or reciprocal, depends on how the artifact in question is used. For example, in case of a shopping list, it can be classified as a two-way system when its user writes down the needed items and then disposes the list when it is no longer needed. However, it can also be classified as a reciprocal system when its user changes the list while doing the shopping and, consequently, turning a shopping list into a more thorough tool for the next use.

Like information flow in two-way systems, reciprocal information flow often takes place in a closed system in that the artifact is meant for a single agent who has designed the informational content of the artifact for individual use. Moreover, each cycle of offloading and intake depends on the outcome of the previous one. The cycles are thus interdependent. For this reason, the functional and informational integration between agent and artifact is significantly denser than in one-way or two-way systems. It is not a mere exchange of information between two entities, as in two-way systems, but what is offloaded onto the artifact in a given cycle depends on what is offloaded in the previous cycle(s) and, therefore, the degree of integration is considerably higher.

The above three kinds of information flow concern situated cognitive systems that are comprised of an embodied agent and a (set of) artifact(s). However, information quite often flows in systems that are comprised of more than one agent and more than one cognitive artifact. These are the cases in which Edwin Hutchins (1995a, b), Nancy Nersessian (2005, 2006), Evelyn Tribble (2011), and others are primarily interested. Examples include a team of navigators trying to navigate a ship through a harbor, pilots in the cockpit of an airplane, researchers in a scientific laboratory, or actors interacting with theatre architecture and other artifacts. These are cases of distributed cognition in which there is a collective of agents that tries to solve a particular problem or perform a certain cognitive task by using a variety of cognitive artifacts. Within those distributed systems, information flows in many directions (sometimes simultaneously) and is often transformed and reformatted.

It could be argued that within those larger distributed cognitive systems, there are often agents interacting with single artifacts, for example a navigator using an alidade to measure the bearings relative to the ship's head, a pilot reading a radar, a scientist using a flow loop to generate representations of blood cells, or an actor interacting with a playbook. Single agent-artifact systems are thus often the building blocks of larger distributed cognitive systems, implying that a better understanding of agent-artifact systems can have a trickledown effect on better understanding the larger distributed system. Furthermore, the way agents interact with those artifacts is influenced by the history and social organization of the larger system, which should be taken into account when analysing single agent-artifact subsystems within larger systems.

Finally, like information flow in computer networks, information flow in situated cognitive systems has a certain bandwidth, which is the amount of information that is interpreted or offloaded per unit of time.⁶ So, the more information is interpreted or offloaded in a given amount of time, the higher the bandwidth, which is often an important aspect for realizing a cognitive purpose. When playing Tetris, for example, a user only has a few seconds to decide in which orientation the zoid should be dropped into a lower template (Kirsh and Maglio 1994). Likewise, pilots landing an airplane need to perform various cognitive tasks within a certain time frame (Hutchins 1995b). In such cases, a high bandwidth is essential for performing the cognitive task.

The bandwidth depends on both the informational properties of the artifact and the cognitive profile of its user. A well-designed map, for example, potentially has a greater bandwidth than a poorly designed map, because it is easier and quicker to interpret. Conversely, the interpretative skills of the agent determine how quick information is

⁶ See also John Haugeland (1998) on bandwidth between two entities, and Clark's (2007) discussion of Haugeland.

interpreted which differs from agent to agent. For example, some people read quickly, while others need more time to read a text. Likewise, some people interpret a map in a glance, whereas others have to study it closely before they know where to navigate. Such interpretative skills often depend on experience, training, and sometimes also on one's age.

3.2 Reliability

Reliable access to external information is essential for how and how often an epistemic action unfolds. There are several features relevant to reliable information access that are often not discussed in the literature. First, the cognitive profile of the human agent partly determines the necessity for information access. For certain cognitive tasks we highly depend on cognitive artifacts. Finding your way in an unfamiliar city without the support of road signs, maps, or navigation systems; multiplying $3,567 \times 5,674$ without the help of pen and paper or a calculator; or remembering all your appointments for the entire month without a diary is for most people impossible. We highly depend on cognitive artifacts to help us perform those cognitive tasks and without such devices our cognitive capacities would just not be the same. The degree of dependency partly depends on one's cognitive profile. Some people, like Otto (Clark and Chalmers 1998), have bad memory capacities and therefore rely and depend on memory aids such as notebooks, post-it notes, diaries, and other reminders. Other people have bad mathematical skills and rely and depend more on calculators or perform calculations with pen and paper. While yet other people have bad navigation skills and rely and depend on navigation aids such as road signs, maps and navigation systems. There are also people who have better memory, mathematical, and navigational skills and do not or rely less on external artifacts. So whether or not one needs to use a cognitive artifact depends on one's cognitive profile.

Second, reliability depends on the kind and properties of the artifact. Due to their physical properties, some artifacts provide better information access than others. Take diaries, for instance. As long as one does not forget to bring one's analogue diary when needed, it provides reliable access to the information in it. In contrast, digital diaries embedded in one's smart phone, tablet, or other electronic device, in one sense, provide less reliable access, because they are inaccessible without electricity. So one not only needs to remember to bring the device when needed, but also to charge it when the battery is empty. Further, digital cognitive artifacts can potentially malfunction in more ways than analogue ones. So next to battery issues, there may be numerous software and hardware issues that prevent one from accessing one's digital diary. Software and hardware issues are irrelevant for analogue diaries. But, conversely, digital diaries such as Google Calendar are online systems that store information in the cloud and are therefore less susceptible for theft, loss, or damage than analogue diaries. Even if one loses one's wearable computing device or if it gets stolen, the information is still available in the cloud. Analogue diaries lack these properties.

The degree of portability is also an important physical property for information access. Cognitive artifacts like smart phones, slide rulers, compasses, or watches are worn or carried on one's body and are thus very portable. As a result, they are (when fully operative) easily, repeatedly, and reliably accessible. The physical design of these artifacts is such that they are small, light, and (ideally) ergonomic, because the

epistemic actions that are performed with these artifacts require them to be small, light, and ergonomic. If not, then they are not suitable for their function. A large and heavy compass, for example, is, in virtue of its physical properties, not particularly portable and thus non-functional for hiking. One could imagine a fully functioning fifty kilogram compass that would be able to fulfil its representational function, but due to its physical properties, unable to be of any help for bushwalkers and so in practice non-functional for the cognitive purpose of measuring one's position in relation to the north pole. Tablets, laptops, analogue diaries, and books are still rather portable but often not (though in some cases they are) worn or carried on one's body and thus a bit less (often) accessible. Desktop computers, DNA-sequencers, and radio-telescopes are non-portable and only accessible at their fixed location. The point of these examples is that partly due to their physical properties (i.e. weight, size, design), cognitive artifacts are able to fulfil their function.

3.3 Durability

There are degrees in the durability of our relation to cognitive artifacts. Wilson and Clark (2009) introduce a trichotomy between one-offs, repeated, and permanent relationships to cognitive artifacts. Shopping-lists are typically one-offs. Abacuses or compasses are frequently re-used because they are devices that are utilized many times for the purpose of calculating or navigating. But some cognitive artifacts enter into permanent and highly durable relationships with their users. Otto and his notebook, a carpenter and his ruler, and an accountant and her calculator enter into long-lasting relationships.

A further distinction can be made between the duration of the temporary coupling of a specific agent-artifact system and the repeatability of the same system. We frequently use calculators to help us solve difficult calculations, which is often a short process. As soon as we have the answer, the system is decoupled. However, throughout our lives we repeatedly and consistently use calculators, so there are many cycles of coupling, decoupling, and recoupling with certain cognitive artifacts distributed over long periods of time. Conversely, there are also cognitive artifacts with which we couple only once or perhaps only a few times. A compass, for instance, is an artifact most people have probably used once or twice. But only bush walkers, scuba divers, and sailors use it quite often. Most people do not need a compass to get around in the world and so there are situated cognitive systems that are not repeated for reasons of necessity.

3.4 Trust

The notion of trust in the literature on EMT concerns our attitude towards the truth-value of information. When we trust information, we typically think it is true. When we distrust information, we either think it is false or we are not sure whether it is true. There are at least two ways in which we can come to (dis)trust information. First, by consciously evaluating it and then come to the conclusion that the information is either trustworthy or not trustworthy. Second, by assuming it is either trustworthy or not trustworthy without consciously evaluating it. The first type of trust may be called explicit (dis)trust and the second may be called implicit (dis)trust.

Some information we trust implicitly because we have endorsed it somewhere in the past and wrote it down because of this endorsement. This is true for Otto's notebook, diary entries, shopping-lists, and the information on post-it notes, as we usually (though not necessarily) do not consciously evaluate information we wrote down ourselves. Other information we trust implicitly because many people rely on it for their actions. This is true for timetables of trains, dictionaries, encyclopaedias, and maps, which are used and shared by many humans. Because these artifacts are shared with many others, and many people rely on them for their actions, there is often no reason to think that the information they contain is false or incorrect (Sterelny 2010, contrast Sterelny 2004). But there are certainly exceptions: Wikipedia, for instance, is used and shared by many people, but given its great variability and constant change, it is in some cases still not a trustworthy source of information and caution with accepting the information as true is often sensible (Magnus 2009; Simon 2013). A further reason to distrust Wikipedia articles is that the reputability of authors is unknown, whereas in case of the before mentioned artifacts like timetables, dictionaries, and so on, the consumer of the information typically knows the information is created by a reputable and reliable source. So, in addition to shared use, the authority and reputation of authors are relevant for establishing implicit trust.

Four other reasons for explicit trust in external information are pointed out by Santiago Arango-Muñoz:

“coherence (the fact that it is in accordance with some of her beliefs), consensus (the fact that most of the people endorse it), intelligibility (the fact that it is easy to understand), and relevance (the fact that it increases the likelihood of attaining her goals)” (Arango-Muñoz 2013, p. 147).

Thus, if external information is coherent, endorsed by most people, intelligible, and relevant for the task at hand, is sufficient reason to trust it.

In two-way and reciprocal situated cognitive systems, we implicitly trust the information because we have endorsed it in the past and because we offloaded it ourselves, but we also implicitly trust it because we believe the information is private and has not been tampered with. Consider a brief example. In Australia, there is a TV commercial for smart phones in which a parent goes shopping with a shopping-list made on a smart phone. The application is connected in real-time to the desktop at home where a child deliberately changes the digital shopping-list as to include items he desires. This example shows that with new digital cognitive artifacts exhibiting networking abilities such as smart phones and tablets, informational privacy and security become increasingly important for trust in information. Privacy and security issues are less likely to emerge when using analogue shopping-lists, which are identifiable by means of one's handwriting (Parsell 2007). So the nature and properties of the artifact partly determine how relevant informational privacy and security are for establishing a trust relation with the artifact and the information it contains.

3.5 Procedural transparency

There are two types of transparency that are relevant in relation to cognitive artifacts: procedural and informational transparency. Procedural transparency (see also Sutton

2009a) concerns the effortlessness and lack of conscious attention with which an agent deploys a cognitive artifact. Otto, for example, is so familiar to using his notebook that he will consult it automatically when he needs to do so. His perceptual-motor processes are proceduralized to such an extent that he does not consciously think about how to retrieve information from his notebook. So the retrieval process is not a two-step process in which Otto first believes that the address of MoMa is in his notebook and then looks up and interprets the information to form his second belief that MoMa is at 53rd street. It is a proceduralized and transparent process. In Clark's words: "the notebook has become transparent equipment for Otto, just as biological memory is for Inga" (Clark 2008b, p. 80).

Likewise, maps, textbooks, compasses, abacuses, computer applications, and countless other cognitive artifacts, typically do not need conscious attention in order to effectively use the artifact. Although interacting with such artifacts clearly require cognitive processing, such processing typically remains below the threshold of consciousness. It must be noted, though, that achieving procedural transparency, in most cases, needs training and may take a considerable amount of time. Novice cognitive artifact-users typically lack procedural transparency. The first time I pick up a compass or abacus, it is difficult for me to effectively use it, because there is no procedural transparency. But as my compass and abacus-using skills develop, I become more fluent in successfully using the artifact.

3.6 Informational transparency

Informational transparency concerns the effortlessness with which an agent can interpret and understand information. Heersmink (2013) developed a taxonomy of cognitive artifacts in which artifacts were categorized based on their informational properties. First a distinction between representational and ecological cognitive artifacts was made and then further distinctions within those two categories were made. I first focus on representational cognitive artifacts. Using the work of Charles Sanders Pierce (1935), a distinction was made between cognitive artifacts containing icons, indices, and symbols.

Icons are defined by exhibiting isomorphism with their target. Generally, but not necessarily, the higher the isomorphism between icon and target, the higher the informational transparency. Photographs, for example, display a high degree of isomorphism with their target and are therefore highly transparent, whereas an fMRI-scan is also isomorphic to its target but needs more interpretative skills and training to meaningfully interpret. For this reason, it is less transparent, at least for people who lack those interpretative skills. Both the skills of the agent and informational properties of the icon determine informational transparency. The easier it is for an agent to understand the salient properties of an icon, the more transparent it is.

Indices have a direct causal connection between index and target. The degree of informational transparency depends on an understanding of this causal connection, i.e., the better the causal connection between target and index is understood by an agent, the more transparent the index is. So, for example, the causal connection between the state of a weathervane and the direction of the wind is for most people quite easy to understand and therefore quite transparent. The causal connection between an electrical current and the reading on a voltmeter is most likely a bit less easy to understand, as

electricity is a more complicated phenomenon than wind direction. And the causal connection between a gel electrophoresis pattern and protein structure and protein size is, unless you are a molecular biologist, quite difficult to understand and so it is close to opaque. Thus, again, one's interpretative skills and informational properties of the artifact determine the degree of transparency.

Symbols such as words in language and numerals in mathematics obtain their meaning through shared use, convention, and logical rules. The degree of informational transparency depends on the degree of understanding of the rules and conventions and the amount of shared use that determine the meaning of symbols. For someone who speaks English, for example, English is transparent because the rules and conventions that determine the meaning of language (i.e., syntax and semantics) are sufficiently understood and, furthermore, language is typically shared with a community of language-users. Conversely, for a non-English speaking person, English is opaque as the rules and conventions are not understood and the language is not shared. The same applies to mathematics, so for someone who has not been exposed to or trained in mathematics, mathematical symbols and numbers are opaque (see also Woelert 2012). Note that the degree of transparency in relation to symbols should be seen as a spectrum from opacity to transparency with varies levels in between.

Icons, indices, and symbols have targets or content, i.e., they *represent* particular objects, structures, or properties of the world. A defining feature of these representational artifacts is that they stand in for something else (Pierce 1935). Ecological cognitive artifacts are different, as they do not stand in for something else. Examples include consistently leaving car keys on a certain spot in your apartment so that you know where they are, leaving a rented DVD on your desk as a prompt to bring it back to the video store, or putting an article you have to read on top of the pile on your desk. By placing artifacts in locations that are either purposely usual or purposely unusual, we intentionally encode information into the artifact and its location (Kirsh 1995; Heersmink 2013), in that way complementing our memory. When we encounter the artifact in that usual or unusual location, it prompts a memory and we know what it means, as we placed the artifact there ourselves. So, because we are the creators of the ecological artifact, it is informationally transparent, unless we have forgotten what it is supposed to mean.

3.7 Individualization

Sterelny (2010; for discussion see Greenwood 2013b) presents a spectrum in which some cognitive artifacts are interchangeable, others are individualized, and yet others are entrenched. For Sterelny, individualization is changing, adjusting, or fine-tuning the informational properties of the artifact such that its use is more effective and efficient for realizing its cognitive function. He argues that most of the books in his professional library are interchangeable, but some of them are individualized with underlining, highlighting, and comments. These adjustments essentially make sense to Sterelny and are less useful and valuable to others. Similarly, Otto's notebook is highly individualized and is useful only for Otto, although others may still be able to read the notebook, only Otto uses it to aid his memory and structure his actions. My tablet computer is fairly individualized: it has applications that I have downloaded to fit my specific needs such as the weather forecast and train timetables for Sydney, and specific

websites, documents, and eBooks. But although it is individualized, most applications are still easily usable by others. In contrast, road signs, speed dials, weathervanes, newspapers, and textbooks are not individualized (and thus interchangeable) and accessible for most people. Like transparency, individualization of cognitive artifacts often takes a certain period of time.

Entrenchment of cognitive artifacts implies a close equilibrium between agent and artifact in which *both* have been transformed in order to ensure the best possible fit between agent and artifact. Sterelny acknowledges that his individualized books are not entrenched in the sense that his professional routines and habits have not been adjusted to those books in the same way as those books have been adjusted to Sterelny. So, he has individualized his books, but his books have not individualized him, or at least not sufficiently. But, according to Sterelny, there may still be cases of entrenchment concerning books. For a Locke scholar, Locke's oeuvre may have transformed the routines of the scholar in the same way as he or she has transformed Locke's oeuvre in the sense of highly individualizing his works by underlining, highlighting, comments, and so on. A more obvious and clear example of an entrenched cognitive artifact is Otto's notebook. The information in his notebook is only meant for Otto himself and is specifically geared to his needs and desires, so it is highly individualized, and his behavioural and cognitive routines are sculpted by his notebook, so it is entrenched as well.

3.8 Transformation

The use of cognitive artifacts transform the representational and cognitive capacities of the human brain. Clark (1997), Stanislaw Dehaene et al. (1999), Helen de Cruz (2008), Menary (2010a), and Michael Kirchhoff (2011; see also Menary and Kirchhoff 2013), amongst others, have argued that external representational systems transform the brain's representational capacities. During ontogenetic development, we interact with public representational systems such as language and mathematics. By so doing, we soak up and learn to think in terms of those systems and the brain takes on the representational properties of those systems. In a sense, this is the reversed version of the parity principle. So external states and processes, rather than being functionally isomorphic to internal ones, are soaked up by the brain which then takes on the properties of those external representational systems. In other words, we internalize external representations which subsequently transform the brain's representational capacities, i.e., we are enculturated.

In case of mathematics, for example, we are born with a capacity to recognise quantity (Dehaene et al. 1999), which, given the right training, can be transformed into a system that is capable to perform mathematical calculations such as multiplication, divisions, and so on. We first learn to do mathematics on the page and typically in an educational setting, before we can do the same operations in our head. Thus, in Menary's words: "The deeply transformative power of our learning histories in the cognitive niche is one that reformats the representational capacities of the brain in terms of public symbol systems" (Menary 2010a, b, p. 576).

Language and mathematics are examples of external representational systems with which we interact substantially for a long period of time, both phylogenetically and ontogenetically. In ontogeny we call this period education. A considerable amount of

research has been done on the transformation effect of those systems on our brain and cognition (e.g., Dehaene et al. 1999; Dehaene 2004). Other cognitive artifacts and representational systems such as road signs, maps, graphs, diagrams, computer applications, and design programs have presumably also a transformation effect on our representational and cognitive capacities. For example, after navigating a city with a map for a certain period of time, the interaction with the map and the city has changed our internal spatial representation, i.e., our cognitive map, of parts of the city. At a certain point, we no longer need the actual map to navigate and we have to a certain degree internalized the information of the map. Likewise, interacting with computers for many hours a day almost certainly transforms our neuronal structures and cognitive capacities. Engineers, for example, spend many hours a day designing objects and structures with design programs. It is not unlikely that after a certain period of training and practice their brains take on the representational properties of the design program, i.e., engineers learn to think in terms of those representational systems. Such transformations seem to be a consequence of long-term interaction with cognitive artifacts over ontogenetic time.

Interacting with cognitive artifacts does not only transform the brain's representational properties. It also transforms our embodied interactive skills. An obvious example is a personal computer. An agent who has never used a personal computer will have trouble interacting with it. The mouse, keyboard, and screen will be difficult to use and it takes a certain amount of time to learn to (fluently) use a computer. Our body schema and motor programs need to adapt and transform in order to be able to successfully interact with computers. Likewise, compasses, rulers, navigation systems, voltmeters, and countless other cognitive artifacts require for their use certain embodied interactive skills that are developed and transformed during the use of such artifacts.

It is, however, not only the human component of the situated system that transforms its representational properties and capacities. The artifactual component transforms its representational properties, too, i.e., cognitive artifacts are often not static and fixed but active and dynamic. The representational properties of post-it notes, slide rulers, and textbooks, for instance, are fairly stable and fixed, but smart phones, tablets, laptops, and other computing devices are very dynamic in their representational properties. We can transform and adjust their representational properties to our own needs and desires, and it is frequently *because* we act on those artifacts and the information they contain that they have dynamic and changing representations. Thus when talking about the dimension of transformation, one can make a distinction between neural, skill, and artifactual transformation.

4 Micro, macro, and system functions

The degree to which each component in a situated cognitive system actively contributes to solving a problem depends on the distribution of computation and on how the overall system is structurally and functionally organized. The structural organisation is how the situated system is physically organised, whereas functional organisation is how its functions are integrated and organised. Consider the following examples from Heersmink ([forthcoming](#)). If I need to perform a difficult calculation, I may use pen, paper, and external numerals, an abacus, a pocket calculator, a spreadsheet program, a

slide ruler, or even Charles Babbage's difference engine. These artifacts have rather different physical and representational structures, but on a course-grained or macro functional level of abstraction, they have the same function, which is to help us perform calculations. So it may be argued that their macro-function is to help us calculate. Therefore, like most artifact functions (see, e.g., Preston 2009, 2013), cognitive functions are, at least on a macro functional level of abstraction, *multiple realizable*, i.e., different physical and representational structures can be used for realizing the same cognitive function.

On a fine-grained or micro functional level of abstraction, however, there are differences in how these devices perform their function. When using pen, paper, and numerals, for example, most of the computation or information-processing is performed by an embodied brain. The external numerals function as to aid or supplement working memory and to structure the task space by decomposing the cognitive task in smaller and easier to perform parts. When using an abacus or slide ruler, the computation is done by an agent-plus-artifact system, i.e., manipulating the beads or slides *is* computation. The artifact's function, then, is to facilitate an analog computation performed by an agent-plus-artifact system. Finally, when using a pocket calculator or spreadsheet program, an agent merely provides the artifact with input and the computation is performed by the artifact. On a micro-level, the artifact's function is not to aid working memory, structure the task space, or to facilitate an analog computation, but to semi-autonomously perform a digital computation.

In these three scenarios, the artifactual elements have different micro-functions and the computation is distributed quite differently in each situated cognitive system. In the first scenario, there is mainly mental computation and the artifactual element is merely a medium for information storage; in the second scenario, there is agent-driven analog computation performed by a joint agent-artifact system; and in the third scenario, there is mainly digital computation performed by a computer. Moreover, using these different artifacts also requires a different set of interactive skills and, consequently, the overall structural and functional organization of the situated cognitive system may be quite different. But, although different artifacts may have different micro-functions, may require different interactive skills, and result in different overall functional architectures of agents and artifacts, ultimately, the situated cognitive system has as its function to perform calculations.⁷ One can therefore conceive of such systems as to varying degrees single integrated functional units. Thus when analysing the integration of cognitive functions into the onboard system, we may take a micro, macro, or systems perspective.

5 Relations between dimensions

In discussing a preliminary version of the framework proposed in this paper (see Heersmink 2012), Robert Clowes (2013; see also 2012) claims that the proposed dimensions are relatively independent and that more practical work and conceptual

⁷ See Clark (2008a) for similar ideas on multiple realizability in situated cognitive systems, and see Gallagher and Crisafi (2009) on socially distributed systems exhibiting different composition but with similar functions regarding legal decision-making.

analysis is needed to clarify which dimensions are more important. Clowes himself suggests that transparency and trust are, at least in cases of external memory, of central importance.⁸ My view is that it is problematic to prioritize certain dimensions because it encourages us to overlook other dimensions. If we would focus only on trust and procedural transparency, we would miss out on how information trajectories are established, how artifacts are individualized, how they transform our onboard cognitive system, how the overall situated system is structurally and functionally organized, and so forth. All these things are essential to better understand integration. Thus, what we need, I think, is a picture of integration not in terms of necessary and sufficient conditions (see also Sutton et al. 2010), but as an inherently multidimensional phenomenon without prioritizing any of these dimensions or group of dimensions.

For analytical purposes, I presented each dimension separately, but some overlap and interact. I now look at some of these interactions and point out how the dimensions are related. The dimension of information flow is concerned with how, how fast, or how much information flows between the components of a situated system. Information trajectories supervene on how the overall system is structurally and functionally organized. In other words, how information trajectories are established depends on the material and functional organisation of the integrated situated system.

Reliability and durability are both concerned with information access. If an artifact is not easily and reliably accessible, then it is hard to establish a durable relation to it. Further, reliability and durability often result in individualization. The more often a certain cognitive artifact is used, the more likely it is that it will be individualized and perhaps in some cases even entrenched. But this need not be the case. There are often-used cognitive artifacts that are neither individualized nor entrenched such as clocks and speed dials. Individualization and entrenchment frequently result in cognitive transformation. Again, the more often we use a certain cognitive artifact, the more likely it is that the human brain soaks-up the representational properties of the artifact. This happens with language and mathematics, but arguably also with maps, abacuses, design programs, and perhaps with graphs, pie charts, diagrams and other illustrations as well. Individualization frequently results in both trust and procedural and informational transparency. Individualized cognitive artifacts, including diaries and notebooks, are designed by the user of the artifact and thus almost automatically trusted. They are also procedurally transparent and informationally transparent. We do not need to think about how to use such artifacts, and the information they carry is trusted implicitly because we wrote it ourselves.

Structural and functional organization are interwoven with transformation. So, when components of situated systems are transformed—either embodied brains, skills, or artifacts—the structural and functional organization are typically transformed as well. Thus, when one's brain structures transform due to artifact-use, when new interactive skills are learned, or when the artifact's informational properties are transformed, then the overall structural and functional organisation is usually transformed, too. For example, when a novice computer-user is over time transformed into an expert computer-user, the novice's brain structures and interactive skills are changed such that

⁸ Although Clowes does not distinguish between procedural and representational transparency, I take it he is referring to procedural transparency, as he illustrates his point by making an analogy to Heidegger's notion of a hammer being ready-to-hand.

the computer almost becomes a different cognitive artifact. The interactive skills one has partly determine the affordances of a cognitive artifact. A novice cannot perform the same cognitive tasks with a computer as compared to an expert, due to a lack of skill and thus also a lack of affordances. How the overall system is organised and transformed over time also relates to individualization. Typically, when agents individualize an artifact, they change its functional and informational properties so that it is better equipped to complement an agent's own properties, i.e., to ensure a better fit between the cognitive profile of the agent and the informational and functional properties of the artifact. And, lastly, informational transparency often results in a greater bandwidth. The idea being that the easier information is to interpret and understand, the faster we can take it onboard and process it.⁹ Although there are more interactions between the dimensions, these are the most obvious ones.

Situated cognitive systems can shift from one region of this multidimensional space to another, i.e., the degree and nature of integration may change over time. When an artifact is used for a longer period of time and it becomes gradually more efficient, durable, individualized, transparent, and trustworthy, the overall system becomes increasingly more integrated. As a result, the situated system will shift to a higher region in the multidimensional space and will get a different dimensional configuration. Highly individualized cognitive artifacts are likely to maintain a stable relation with its user and, consequently, populate a given region in the space for a longer period of time, but most situated systems are frequently shifting from one region to another. This is so because many of these systems are quite dynamic, constantly changing their functional and informational properties, and renegotiating existing functional and informational equilibriums. The framework in this paper is able to deal with this dynamic nature of situated cognitive systems.

In one corner of this multidimensional space we find systems with a high-bandwidth reciprocal information flow, where the artifact is reliably available, the relation between agent and artifact is durable, the information it provides is trustworthy, the artifact is procedurally and informationally transparent, it is individualized or possibly even entrenched, and its representational systems has transformed the brain of its user.¹⁰ These are obvious cases for extended and distributed cognitive systems. In another corner of this space we find systems that score low on all these dimensions. These are clear cases for embedded systems. In between those two corners, there is a large space in which it may not always be clear whether a system is embedded or extended. In future work I will try to map various situated systems in this multidimensional space.

6 Conclusion

This paper has further developed complementarity-based EMT by proposing a multi-dimensional framework for conceptualizing the complementary integration of embodied agents and cognitive artifacts or other external resources. It includes the dimensions

⁹ A reviewer pointed out that, in some cases, fast information-processing may not be desirable, as it runs the risk of superficiality and lack of understanding. In an educational context, for example, slow information-processing may be advantageous, as it increases understanding of a problem or concept.

¹⁰ Sterelny (2010) presents a similar approach on a multidimensional space, but with only three dimensions.

of information flow, reliability, durability, trust, procedural transparency, informational transparency, individualization, and transformation. All these dimensions are matters of degree and jointly constitute a multidimensional space in which situated cognitive systems can be located. These dimensions provide a new perspective on the conditions of cognitive extension and cognitive integration. Importantly, the proposed dimensions are not meant to provide a set of necessary and sufficient conditions, but to provide a toolbox for investigating the degree and nature of the integration of agent and artifact into “new systemic wholes”.

Acknowledgments I would like to express my gratitude to John Sutton and Richard Menary for helpful comments on earlier versions of this paper. I would also like to thank the audiences at the Distributed Cognition and Distributed Agency Workshop (Sydney, 2012) and the AISB/IACAP World Congress (Birmingham, 2012) for useful feedback.

References

- Adams, F., & Aizawa, K. (2008). *The bounds of cognition*. Oxford: Wiley-Blackwell.
- Adams, Z., & Maher, C. (2012). Cognitive spread: under what conditions does the mind extend beyond the body. *European Journal of Philosophy*. doi:10.1111/j.1468-0378.2012.00543.x.
- Arango-Muñoz, S. (2013). Scaffolded memory and metacognitive feelings. *Review of Philosophy and Psychology*, 4, 135–152.
- Arnaud, E., Ayala, S., & Sturm, T. (2013). Cognitive externalism meets bounded rationality. In M. Knauff, M. Pauen, N. Sebanz, & I. Wachsmuth (Eds.), *Proceedings of the 35th Annual Conference of the Cognitive Science Society* (pp. 1768–1773).
- Clark, A. (1997). *Being there: Putting brain, body and world back together again*. Cambridge: MIT Press.
- Clark, A. (2001). *Mindware: An introduction to the philosophy of cognitive science*. Oxford: Oxford University Press.
- Clark, A. (2007). Reinventing ourselves: the plasticity of embodiment, sensing, and mind. *The Journal of Medicine and Philosophy*, 32(3), 263–282.
- Clark, A. (2008a). Pressing the flesh: a tension in the study of the embodied, embedded mind? *Philosophy and Phenomenological Research*, 76(1), 37–59.
- Clark, A. (2008b). *Supersizing the mind: Embodiment, action, and cognitive extension*. Oxford: Oxford University Press.
- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 58, 10–23.
- Clowes, R. W. (2012). Hybrid memory, cognitive technology, and self. In J. M. Bishop & Y. J. Erden (Eds.), *Proceedings of the 5th AISB Symposium on Computing and Philosophy* (pp. 4–13).
- Clowes, R. W. (2013). The cognitive integration of e-memory. *Review of Philosophy and Psychology*, 4(1), 107–133.
- De Cruz, H. (2008). An extended mind perspective on natural number representation. *Philosophical Psychology*, 21(4), 475–490.
- Dehaene, S. (2004). Evolution of human cortical circuits for reading and arithmetic: The neuronal recycling hypothesis. In S. Dehaene, M. D. Hauser, & G. Rizzolatti (Eds.), *From monkey brain to human brain* (pp. 133–157). Cambridge: MIT Press.
- Dehaene, S., Spelke, E., Pinel, P., Stanescu, R., & Tsivkin, S. (1999). Sources of mathematical thinking: behavioral and brain imaging evidence. *Science*, 284, 970–974.
- Gallagher, S., & Crisafi, A. (2009). Mental institutions. *Topoi*, 28, 45–51.
- Greenwood, J. (2013a). Contingent transcranialism and deep functional cognitive integration: the case of human emotional ontogenesis. *Philosophical Psychology*, 26(3), 420–436.
- Greenwood, J. (2013b). Is mind extended or scaffolded? Ruminations on Sterelney’s (2010) extended stomach. *Phenomenology and the Cognitive Sciences*. doi:10.1007/s11097-013-9337-8.
- Halpin, H. (2013). Does the web extend the mind? In *Proceedings of the 5th annual ACM web science conference* (pp. 139–147).

- Haugeland, J. (1998). Mind embodied and embedded. In *Having thought: Essays in the metaphysics of mind* (pp. 207–237). Harvard University Press.
- Heersmink, R. (2012). Mind and artifact: A multidimensional matrix for exploring cognition-artifact relations. In J. M. Bishop & Y. J. Erden (Eds.), *Proceedings of the 5th AISB Symposium on Computing and Philosophy* (pp. 54–61).
- Heersmink, R. (2013). A taxonomy of cognitive artifacts: function, information, and categories. *Review of Philosophy and Psychology*, 4(3), 465–481.
- Heersmink, R. (forthcoming). The metaphysics of cognitive artifacts. *Philosophical Explorations*.
- Hutchins, E. (1995a). *Cognition in the wild*. Cambridge: MIT Press.
- Hutchins, E. (1995b). How a cockpit remembers its speed. *Cognitive Science*, 19(3), 265–288.
- Kirchhoff, M. D. (2011). Extended cognition and fixed properties: steps to a third-wave version of extended cognition. *Phenomenology and the Cognitive Sciences*, 11(2), 287–308.
- Kirsh, D. (1995). The intelligent use of space. *Artificial Intelligence*, 72, 1–52.
- Kirsh, D., & Maglio, P. (1994). On distinguishing epistemic from pragmatic action. *Cognitive Science*, 18, 513–549.
- Kiverstein, J., & Farina, M. (2011). Embraining culture: leaky minds and spongy brains. *Teorema*, 32(2), 35–53.
- Loader, P. (2012). The epistemic/pragmatic distinction. *Philosophical Explorations*, 15(2), 219–232.
- Ludwig, D. (2014). Extended cognition and the explosion of knowledge. *Philosophical Psychology*. doi:10.1080/09515089.2013.867319.
- Magnus, P. D. (2009). On trusting wikipedia. *Episteme*, 6(1), 74–90.
- Menary, R. (2007). *Cognitive integration: Mind and cognition unbounded*. Basingstoke: Palgrave MacMillan.
- Menary, R. (2010a). Dimensions of mind. *Phenomenology and the Cognitive Sciences*, 9(4), 561–578.
- Menary, R. (2010b). Cognitive integration and the extended mind. In R. Menary (Ed.), *The extended mind* (pp. 227–243). Cambridge: MIT Press.
- Menary, R., & Kirchhoff, M. (2013). Cognitive transformations and extended expertise. *Educational Philosophy and Theory*. doi:10.1080/00131857.2013.779209.
- Michaelian, K. (2012). Is external memory, memory? Biological memory and extended mind. *Consciousness and Cognition*, 21(3), 1154–1165.
- Michaelian, K., & Sutton, J. (2013). Distributed cognition and memory research: history and current directions. *Review of Philosophy and Psychology*, 4(1), 1–24.
- Nersessian, N. J. (2005). Interpreting scientific and engineering practices: Integrating the cognitive, social, and cultural dimensions. In M. Gorman, R. Tweney, D. Gooding, & A. Kincannon (Eds.), *New directions in scientific and technical thinking* (pp. 17–56). Mahwah: Erlbaum.
- Nersessian, N. J. (2006). The cognitive-cultural systems of the research laboratory. *Organization Studies*, 27(1), 125–145.
- Norman, D. (1993). *Things that make us smart: Defending human attributes in the age of the machine*. New York: Basic Books.
- Palermos, S. O. (2014). Loops, constitution, and cognitive extension. *Cognitive Systems Research*, 27, 25–41.
- Parsell, M. (2007). The cognitive cost of extending an evolutionary mind into the environment. *Cognitive Processing*, 7, 3–10.
- Pierce, C. S. (1935). *The collected papers of Charles S. Pierce vol 2*. Cambridge: Harvard University Press.
- Preston, B. (2009). Philosophical theories of artifact function. In A. Meijers (Ed.), *Philosophy of technology and engineering sciences* (pp. 213–234). The Netherlands: Elsevier.
- Preston, B. (2013). *A philosophy of material culture: Action, function, and mind*. New York: Routledge.
- Roberts, T. (2011). Taking responsibility for cognitive extension. *Philosophical Psychology*, 25(4), 491–501.
- Rockwell, T. (2010). Extended cognition and intrinsic properties. *Philosophical Psychology*, 23(6), 741–757.
- Rowlands, M. (1999). *The body in mind: Understanding cognitive processes*. Cambridge: Cambridge University Press.
- Rowlands, M. (2010). *The new science of the mind: From extended mind to embodied phenomenology*. Cambridge: MIT Press.
- Rupert, R. (2004). Challenges to the hypothesis of extended cognition. *Journal of Philosophy*, 101(8), 389–428.
- Simon, J. (2013). Trust. In D. Pritchard (Ed.), *Oxford bibliographies in philosophy*. New York: Oxford University Press.
- Smart, P. (2012). The web-extended mind. *Metaphilosophy*, 43(4), 446–463.
- Sterelny, K. (2004). Externalism, epistemic artefacts and the extended mind. In R. Schantz (Ed.), *The externalist challenge. New studies on cognition and intentionality* (pp. 239–254). Berlin: De Gruyter.

- Sterelny, K. (2010). Minds: extended or scaffolded? *Phenomenology and the Cognitive Sciences*, 9(4), 465–481.
- Sutton, J. (2006). Distributed cognition: domains and dimensions. *Pragmatics and Cognition*, 14(2), 235–247.
- Sutton, J. (2008). Material agency, skills and history: Distributed cognition and the archaeology of memory. In C. Knappet & L. Malafouris (Eds.), *Material agency* (pp. 37–55). Berlin: Springer.
- Sutton, J. (2009a). The feel of the world: Exograms, habits, and the confusion of types of memory. In A. Kania (Ed.), *Philosophers on memento* (pp. 65–86). London: Routledge.
- Sutton, J. (2009b). Remembering. In P. Robbins & M. Aydede (Eds.), *The Cambridge handbook of situated cognition* (pp. 217–235). Cambridge: Cambridge University Press.
- Sutton, J. (2010). Exograms and interdisciplinarity: History, the extended mind and the civilizing process. In R. Menary (Ed.), *The extended mind* (pp. 189–225). Cambridge: MIT Press.
- Sutton, J., Harris, C. B., Keil, P., & Barnier, A. J. (2010). The psychology of memory, extended cognition, and socially distributed remembering. *Phenomenology and the Cognitive Sciences*, 9, 521–560.
- Theiner, G. (2013). Transactive memory systems: a mechanistic analysis of emergent group memory. *Review of Philosophy and Psychology*, 4(1), 65–89.
- Tollefson, D., Dale, R., & Paxton, A. (2013). Alignment, transactive memory, and collective cognitive systems. *Review of Philosophy and Psychology*, 4(1), 49–64.
- Tribble, E. (2011). *Cognition in the globe: Attention and memory in Shakespeare's theatre*. Houndmills: Palgrave MacMillan.
- Tversky, B., & Suwa, M. (2009). Thinking with sketches. In A. Markman (Ed.), *Tools for innovation* (pp. 75–84). Oxford: Oxford University Press.
- Vorms, M. (2012). Formats of representation in scientific theorising. In P. Humphreys & C. Imbert (Eds.), *Models, simulations, and representations* (pp. 250–273). New York: Routledge.
- Walter, S. (2010). Cognitive extension: the parity argument, functionalism, and the mark of the cognitive. *Synthese*, 177, 285–300.
- Wheeler, M. (2010). In defence of extended functionalism. In R. Menary (Ed.), *The extended mind* (pp. 245–270). Cambridge: MIT Press.
- Wheeler, M. (2011). In search of clarity about parity. *Philosophical Studies*, 152, 417–425.
- Wilson, R. A. (2004). *Boundaries of mind: The individual in the fragile sciences*. New York: Cambridge University Press.
- Wilson, R., & Clark, A. (2009). How to situate cognition: Letting nature take its course. In P. Robbins & M. Aydede (Eds.), *The Cambridge handbook of situated cognition* (pp. 55–77). New York: Cambridge University Press.
- Woelert, P. (2012). Idealization and external symbolic storage: the epistemic and technical dimensions of theoretic cognition. *Phenomenology and the Cognitive Sciences*, 11, 335–366.
- Zhang, J. (1997). The nature of external representations in problem solving. *Cognitive Science*, 21(2), 179–217.