

Distributed learning:

Educating and assessing extended cognitive systems

Abstract

Extended and distributed cognition theories argue that human cognitive systems sometimes include non-biological objects. On these views, the physical supervenience base of cognitive systems is thus not the biological brain or even the embodied organism, but an organism-plus-artifacts. In this paper, we provide a novel account of the implications of these views for learning, education, and assessment. We start by conceptualising how we learn to assemble extended cognitive systems by internalising cultural norms and practices. Having a better grip on how extended cognitive systems are assembled, we focus on the question: If our cognition extends, how should we educate and assess such extended cognitive systems? We suggest various ways to minimize possible negative effects of extending one's cognition and to efficiently find and organise (online) information by adopting a virtue epistemology approach. Educational and assessment implications are foregrounded, particularly in the case of Danish students' use of the Internet during exams.

Keywords: Extended mind; Distributed cognition; Learning; Technology; Education; Assessment; Extended knowledge; Virtue epistemology; Intellectual virtues

1 Introduction

In a thoughtful reflection on the role of technology in education from an extended mind perspective, Mike Wheeler (2011) writes:

“Perhaps what we ought to focus on, then, is the education of those hybrid assemblages, a focus which is entirely consistent with the goal of endowing the brain with the skills it needs to be an effective contributor to such assemblages. From this perspective, of course, there are extremely good reasons to support the increased presence of technology in the examination hall.”

Wheeler thus argues that we should focus on educating extended cognitive systems and allow students to use technology when they are completing exams. In this paper, we pick up on Wheeler's suggestion and propose several ways to educate and assess extended cognitive systems. We first outline two versions of extended cognition theory, one based on parity considerations (Clark & Chalmers, 1998) and the other on complementarity considerations (Sutton, 2010), emphasising the priority of the latter. On this view, artifacts have complementary properties that are integrated into the human cognitive system to varying degrees. We then briefly outline the implications of the extended mind thesis for epistemology, arguing that knowledge can under certain conditions be extended when the artifact is appropriately integrated into the agent's cognitive system. Thereafter, we conceptualise how humans learn to assemble extended cognitive systems by internalising cultural norms and practices, along the dimensions of accessibility, trust, procedural transparency and informational transparency. Building on Edwin Hutchins (2011), we argue that humans are enculturated to assemble extended cognitive systems, which happens by being immersed in an (educational) culture. Having a better understanding of how extended cognitive systems are assembled, we suggest several ways to minimize possible negative effects of extending one's cognition and to efficiently find and organise information by adopting a virtue epistemology approach. Finally, we argue that taking the extended cognition perspective seriously requires consideration of assessment regimes, as well as the extent to which

they focus on assessing unaided brains versus sufficiently reflect assessing extended cognitive systems. We specifically look at Danish students' use of the Internet during exams.

2 Extended minds

Human beings have evolved to incorporate informational objects and structures into their cognitive systems, in that way overcoming limitations in the brain's information-storage and processing capabilities (Donald, 1991). We are, in Andy Clark's (2003) words, natural born cyborgs, "joint products of our biological nature and multilayered linguistic, cultural, and technological webs" (p. 195). This means that technologies can be incorporated into our perceptual and cognitive systems as a result of our brain's plasticity. Essentially, our embodied brains are unfinished systems and need cultural artifacts to fully function. Humans have therefore developed a wide variety of informational objects and systems which we use to perform many of our cognitive tasks. We use diaries to help us remember our appointments, pen and paper to perform calculations, maps to navigate, diagrams to make inferences, timetables to plan a trip, word-processors to write documents, interactive whiteboards to learn, and so forth (Heersmink, 2013). Such artifacts have informational properties that complement and (often) enhance the cognitive capacities of unaided brains (Sutton, 2010; Heersmink 2015). Some philosophers have argued that such artifacts are not just aids or scaffolds for cognition, but are literally part of cognition. Thus, when we use such artifacts to complete our cognitive tasks, "the human organism is linked to an external entity in a two-way interaction, creating a coupled system that can be seen as a cognitive system in its own right" (Clark & Chalmers, 1998, p. 8). In the literature, there are two versions or waves of extended mind theory, which we outline below, but we will focus on and emphasize the priority of the second wave.

2.1 Parity-based extended mind theory

Clark and Chalmers (1998) argue that artifacts can become part of the system that performs cognitive tasks. One of their examples is Otto, a man with Alzheimer's disease who uses a notebook to assist his deteriorating biological memory. Information in the notebook (e.g., an address) plays relevantly similar roles as the same information stored in biological memory. It does not matter where the information is stored, what matters is the functional role it plays in thought and action. The notebook is furthermore easily accessible and reliably there when Otto needs it, and the information in the notebook is trustworthy, has been endorsed in the past, and indeed is there because of this past endorsement. Clark and Chalmers emphasize functional parity between inner (or biological) and external (or technological) states and processes. They propose the parity principle as a way of thinking about demarcating extended cognitive systems from non-extended or embedded cognitive systems. According to this principle, if an artifact has similar functional and informational properties as an internal state or process that is clearly recognised as being cognitive, then it is part of an extended cognitive system. This principle is straightforwardly functionalist in nature. The functional profile of information (either internal or external) determines its cognitive status. The parity principle extends the notion of multiple realizability in traditional functionalism as to include external states and processes. Artifacts and external information can thus be part of the physical system that realises cognitive states and processes. This view has been referred to as "extended functionalism" (Clark, 2008, Wheeler, 2010). However, both critics (Adams and Aizawa, 2001; Rupert, 2004) and proponents (Menary, 2010; Sutton, 2010; Heersmink, 2015) of extended mind theory have argued that functional parity does not exist.

2.2 Complementarity-based extended mind theory

So, rather than thinking about the relation between embodied brains and artifacts in terms of functional parity, John Sutton (2010) argues that we should think about this relation in terms of complementarity. Artifacts have informational and functional properties that are quite different from those of the brain. For example, when remembering a list of items, we remember the first

and last item better than those in the middle, which is called the serial position effect. When using a written list of items as external memory, these effects do not occur, as all the items on the list are equally easily readable. Moreover, information stored in biological memory is integrative, dynamic, and not stored in discrete format, whereas information in external memory is – at least in most cases (e.g., notebooks, diaries, etc.) – static and discrete. The functional profiles of internal and external information are thus rather different. An emphasis on functional parity as a condition for extended mind thus fails. Instead, as Sutton points out, we should think about artifacts not as replicating what the brain already can do, but as providing complementary properties to existing brain functions.

Recognizing that artifacts complement brain-based cognitive states and processes is a useful starting point for further conceptualizing the relation between agents and artifacts. One way to think about this relation is in terms of cognitive integration (Menary, 2007, 2010). On this view, human cognitive systems and artifacts are integrated into wider systems that perform cognitive tasks. By synthesising work of Clark and Chalmers (1998), Sutton (2006; Sutton et al, 2010), Wilson and Clark (2009) and Sterelny (2010), we suggest this is a matter of degree and is best seen as a multidimensional phenomenon in which integration varies along several dimensions (Heersmink 2015). These dimensions include the following¹:

- *Information flow* describes the information trajectories between the embodied brain and external artifacts. This may be one-way, where information flows from an artifact to an agent (e.g., when navigating with a map). This may be two-way, where information is first offloaded onto an artifact and then later used to perform some cognitive task (e.g., when writing an appointment in your diary and looking it up at some later point). Information flow can also be reciprocal, where there are many interdependent cycles of offloading and intake (e.g., when making a calculation with pen and paper or writing an article).
- *Accessibility* describes the level of availability of the artifact. Some artifacts are easily available (e.g., one's smartphone), whereas others are not (e.g., a library book in Alaska). Reliable access to external information is essential for how and how often a cognitive task unfolds (Clark & Chalmers, 1998).
- *Durability* describes how often we interact and couple with an artifact. Wilson and Clark (2009) propose a trichotomy between one-offs (e.g., using a shopping list), repeated (e.g., using a map), and permanent (e.g., using a smartphone) relationships to cognitive artifacts.
- *Trust* describes how much trust an agent puts into the information the artifact provides. When we trust information, we think it is true. When we do not trust information, we think it is false or we are not sure whether it is true. Trust is important, as information we do not trust, we typically do not use (Clark & Chalmers, 1998).
- *Procedural transparency* describes the degree of transparency-in-use. The easier it is to use and interact with an artifact, the more procedurally transparent it is. For example, to be able to use a computer, a user must learn how to use a mouse, keyboard, and (touch)screen. This is difficult at first, but becomes easier and more fluent after frequent use, as one's perceptual-motor processes become proceduralised such that one does not have to think about how to use the artifacts.
- *Informational transparency* describes the ease with which information can be interpreted. Some information is opaque, which means we cannot interpret it. Certain scientific symbols or formulas, for example, are for most people opaque as they do not know their meaning. Other information is transparent. One's native language is often fully transparent.

¹ In the original framework (Heersmink 2015), cognitive transformation was also included as one of the dimensions. We decided not to include it here. Heersmink now sees cognitive transformation not so much as a dimension that describes the degree of integration between agent and objects, but as a dimension that describes the change that occurs when we learn to use an object.

- *Personalisation* describes how much the artifact is personalised (Sterelny, 2010). Some artifacts are not personalised and thus interchangeable (e.g., a tourist map of Sydney), whereas other artifacts are highly personalised (e.g., a notebook). Personalisation often streamlines a cognitive task and thus makes performing the task easier.

These dimensions are all matters of degree and jointly constitute a multidimensional space in which both embedded and extended cognitive systems can be located and have certain dimensional configurations (see, e.g., Smart, Clowes & Heersmink, 2017). If the integration is dense, then the artifact is part of an extended cognitive system, whereas if it is shallow, the artifact merely scaffolds an internal cognitive system. Focussing on complementarity and the degree of cognitive integration allows one to conceptualise the rich variety of agent-artifact interaction. On this view, we should focus on the cultural practices of artifact-use and try to better “understand the nature of the integration between these elements of a hybrid process” (Menary, 2010, p. 229).

3 What do we know when we extend our minds?

Recent attention has turned to the epistemological implications of the extended mind thesis, with discussion focusing on the conditions under which one would be said to ‘know’ one’s cognitive extensions (Pritchard, 2010; Palermos, 2014). Such discussion has particularly drawn on virtue epistemology in developing accounts of the role of extensions in cognitive processes. Such views indicate that my knowledge might extend beyond the bounds of the brain as a result of some characteristic of myself or my methods that permits of a reliance on the external resource. For example, knowing that my watch reliably informs me of the time, and having the capacity to identify when the watch is malfunctioning. Duncan Pritchard (2014) adopts such an approach in arguing against a strict epistemic individualism that claims that knowledge cannot extend beyond the individual capacity of the mind qua brain. Such epistemic individualism taken to its extreme precludes use of external artifacts in understanding an agent’s epistemic state. As Pritchard notes, the focus of virtue epistemology is on the development of epistemically virtuous cognitive character, rather than (for example) the accrual of facts. That is not to say that facts don’t matter, but rather that they play an instrumental role in doing and in developing a cognitive integration across external artifacts; clearly, one must know that there is a museum and a notebook and understand a notation system to make use of the external artifact to navigate to that museum.

Pritchard’s discussion describes the common analogy of scaffolding in educational theory. Under one model of scaffolding, the analogy runs that once the scaffold has served its purpose (i.e., when the student is able to complete the task without the structured support of the scaffold), it may be removed. The scaffold thus acts as a crutch for individuals and once it is removed, the individual qua unextended agent can be said to have the capacity that the scaffold supported. Under an alternative conception, the scaffold should be considered much as a support structure in the construction of a bridge over a gorge: the scaffold is in place to support the creation of a permanent new structure that could not otherwise be constructed. This latter conception of scaffolding forges a new capacity for “cognitive processes which extend beyond the skin of the subject and involve ‘external’ technology” (Pritchard, 2014, p. 7). Clark and Wilson provide a useful example of the latter cognitive extension:

“Mathematical notation does not simply feed existing mathematical abilities (though it does that, to be sure), but builds on those abilities to produce an agent with significantly greater mathematical capacities. The difference between the ability to multiply using Arabic numerals, versus that using Roman numerals, serves as a reminder of how much specific forms of writing can contribute to particular abilities here” (2009, p70).

In both cases, technology might play a scaffolding role, but in the former case the technology is temporary and incidental, while in the latter it is central to the developed capacity. On Pritchard’s account, then, educational debates regarding the use of technology ought not to assume *prima*

facie that technologies result in a lessening of the cognitive capacities of the student. Rather, Pritchard argues that under some circumstances, technologies can be said to extend the cognitive capacity of the agent, and that where this is the learning goal, assessments that involve use of technologies are entirely appropriate (see also Wheeler, 2011). Elaborating this perspective, Ben Kotzee (forthcoming) argues that we should educate children to engage in responsible practices of technology-use, inculcating them into virtuous technology-use. To do this, Kotzee suggests, educators must reverse engineer the cognitive integration, making visible the steps that integration has emerged over, and perhaps raising the educator's critical stance towards those technologies. We agree with Pritchard and Kotzee's arguments, and in the following sections briefly outline the kinds of knowledge students might be said to gain. A key distinction in this paper is that while both Kotzee and Pritchard discuss important issues regarding the creditworthiness of a student's extended knowledge, they do not discuss the means through which – educationally – this cognitive integration might occur, or be assessed, which are central concerns in this paper.

It is useful here to introduce a distinction between knowing-how and knowing-that. The latter consists of declarative propositions, while the former consists of procedural and declarative knowledge alongside an ability to perform some task. As Marlene Scardamalia and Carl Bereiter (2003, 2006) note, textbooks, schemes of work, subject tests, etc., focus on what could be characterised as 'knowing-that'. To these terms we might add metacognition, classically thought of as "knowing about knowing" (Metcalfe & Shimamura, 1996) or one's capacity to think about one's own cognition. In this context, metacognition should include all knowledge regarding an agent's awareness of the availability of tools, and the reliability (and reasons for the reliability) of those tools. Developing metacognitive awareness is an important learning aim (Kuhn, 2000); thus, developing awareness of the strategies one might take to develop cognitive integration might be considered an important learning aim. Understanding the particular kinds of knowledge implicated in extended mind and associated learning processes is important for developing understanding of the implications of extended mind for education.

4 Learning - Beyond the brain

4.1 Introduction

In an early exploration of distributed cognition theory for learning, David Perkins (1993) argues that it is not just the unaided brain that learns but brain-plus-artifact systems. He focusses on a student using a notebook as an external medium for thinking and remembering. "We could say that this person-plus system has learned something, and part of what the system has learned resides in the notebook rather than in the mind of the student" (1993, p. 89). Perkins puts forward the "equivalent access hypothesis", which asserts that thinking and learning for the person-plus system depend on the access characteristics of the notebook. This, in turn, depends on the kind of information in the notebook, how it is represented and organised, and how easy it is to retrieve. Very much in the spirit of contemporary extended mind thinking, he argues that the location of the information is not important. If the information is easily available when needed and created by the student herself, "what does it matter whether the ideas lie inside or outside the student's cranium?" (1993, p. 90).

On this view, we should enlarge the unit of analysis in the learning sciences, not just to single persons interacting with artifacts, but also to larger systems. Briefly consider an example of such a larger distributed system. Nancy Nersessian (2006; Nersessian, Kurz-Milcke, Newstetter & Davies, 2003) performed an empirical case study in which she and her colleagues looked at how a biology laboratory, including multiple researchers and various instruments and artifacts, learns to do experiments and generate scientific knowledge. Nersessian writes: "The researchers are PhD and MS candidates, undergraduates, and post-doctoral trainees, all of whom have learning trajectories. These trajectories, in turn, intersect with the developmental trajectories of the diverse technological artifacts and of the various social systems within the laboratory" (2006, p. 129).

Learning trajectories of people and developmental trajectories of artifacts are thus interwoven. Nersessian conceives of the entire laboratory as an evolving distributed cognitive system that continuously learns to perform new procedures and tasks. An important learning trajectory within the overall system is for individual researchers and research groups to learn to use the “flow loop”, which is a device that models the shear stress experienced by cells within blood vessels. Typically, one learns to use the device in an apprenticeship structure. The first flow loop in the laboratory dates back some generations and was an ineffective, clunky device. But each new generation of researchers improved its properties and passed on their knowledge of the device to the next generation. Note that this concerns both knowledge-how and knowledge-that. Learning here is a process transforming the overall distributed cognitive system, including humans and an artifact, and takes place over several generations. These two brief examples show that extended and distributed cognitive systems can learn new knowledge-that and knowledge-how, implying that we should try to better understand the learning processes involved (Sawyer & Greeno, 2009).

4.2 Cultural practices and learning

How does learning relate to the process of cognitive integration? In other words, how do we learn to assemble extended cognitive systems? Clark posits the “principle of ecological assembly”, according to which, “the canny cognizer tends to recruit, on the spot, whatever mix of problem-solving resources will yield an acceptable result with a minimum of effort” (Clark, 2008, p. 13). On this principle, we are thus rather pragmatic, perhaps even opportunistic, in choosing the ecological resources that we use to perform our cognitive tasks. Hutchins (2011) points out that Clark’s use of the word “assembly” is ambiguous, as it can both mean the assembly process or the product, that is, the agent-artifact assembly. Here we will focus on the assembly *process* and argue that it involves learning along various dimensions outlined above. Hutchins draws attention to the importance of cultural practices for better understanding the process of recruiting ecological objects. Hutchins has quite a nuanced and detailed account of cultural practices:

“Cultural practices are the things people do and their ways of being in the world. A practice is cultural if it exists in a cognitive ecology such that it is constrained by or coordinated with the practices of other persons. Above all else, *cultural practices are the things people do in interaction with one another*. Virtually all external representations are produced by cultural practices. All forms of language are produced by and in cultural practices. Speaking is accomplished via discursive cultural practices. Reading and writing are cultural practices *par excellence*. They are fully embodied skills” (2011, p. 440-441, original italics).

Hutchins’ main point, we take it, is that extended cognitive systems are deeply and inherently cultural systems which are always part of a larger cognitive ecosystem, consisting of other people, artifacts, cultural norms, practices, and institutions (Hutchins, 2014). The cultural norms and practices underlying the assembly of integrated and extended cognitive systems are learned. Humans are enculturated to assemble integrated and extended cognitive systems, which happens by being immersed in an (educational) culture. Cultural practices, knowledge, and skills are transmitted from one generation to the next by means of learning, which, in turn, happens through our embodied interactions with the social and material environment. Cultural norms and practices are learned in formal settings such as classrooms and apprenticeships, but also in informal settings such as by interacting with parents, caregivers, friends, and by reading newspapers, watching TV, or surfing the Web.

4.3 Learning and cognitive integration

Above we outlined the dimensions that are relevant for conceptualising the degree of cognitive integration between human cognitive systems and artifacts. Some of these dimensions are directly related to cultural learning. Accessibility, trust, procedural transparency, and informational

transparency are particularly relevant in relation to learning². We learn, in various ways, how accessible an artifact is. For example, through experience and cultural practices, we learn that smartphones and the functionalities they afford such as Internet search engines are highly accessible, whereas other artifacts such as library books are much less accessible.

We also learn how to evaluate information and to decide whether it is trustworthy. How do we do this? In terms of authorship, there are at least two kinds of external information: information that we made ourselves (e.g., notes in a notebook) and information that others have made (e.g., an entry in an encyclopaedia). Usually, we do not evaluate information we created ourselves (unless we know it is outdated). Otto, for example, won't consciously evaluate the information in his notebook. Rather, he automatically endorses it, because the information is in the notebook due to past endorsement. Information made by others is sometimes evaluated based on the reputation of the source. This reputation is learned through experience and enculturation. A reputable source (e.g., a recent textbook on genetic diseases) is more appropriately trusted than a less reputable source (e.g., an online health forum).

Santiago Arango-Muñoz points out four other reasons for trust in external information: "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)" (2013, p. 147). So, on his view, if external information is coherent, endorsed by most people, intelligible, and relevant for the task at hand, then (in most cases) that is sufficient reason to trust it. However, note that not all these features need to be satisfied. I may, for example, read on Wikipedia that Sydney has 4.92 million inhabitants while casually surfing the Web. Say, that I already knew that Sydney is roughly the same size as Melbourne, which I know has roughly 5 million inhabitants. So it is coherent with at least one of my pre-existing beliefs. It also seems intelligible, but I do not know whether it is endorsed by most other people, and because I am just casually surfing the Web, I am not consuming the information to do any particular cognitive task. So, it only satisfies two conditions, but that still seems sufficient reason to trust it. Learning to evaluate information is done both formally (e.g., in school) and informally (e.g., from friends or by trial and error).

Using a cognitive artifact such that it becomes procedurally transparent is essentially a learning process involving knowledge-how. For example, using a pen to write is a difficult embodied skill which may take several years to learn, as it requires highly sophisticated hand-eye coordination skills. Making a simple tool such as a pen procedurally transparent thus needs a lot of training, typically done in a classroom setting. When the pen becomes transparent, it is absorbed in the body schema, which is a subpersonal (or subconscious) representation of the body's size, position, and location in space (Gallagher, 2005). Body schemas are the basis for our motor-programs, allowing us to interact with objects. Because our bodies change over time, our body schemas are updated accordingly. This flexibility allows tools to be incorporated into the body schema. When that happens, the tools are experienced as transparent extensions of the body, fully under control of our agency. When writing with a pen, the perceptual focus is on the pen-paper interface, rather than on the hand-pen interface. The focus is thus on the task, rather than the artifact. Likewise, using a mouse, keyboard, and (touch)screen as to be able to fluently interact with computer systems must be learned. Our body schemas must adjust as to be able to interact with these objects. This also applies to other cognitive artifacts such as compasses, navigation systems, calculators, models, diagrams, and so on. Ideally, these tools become as transparent as possible such that they do not interfere with the task at hand (Norman, 1998). In general, the less we have

² This is not to say that the dimensions of information flow, durability and personalisation are irrelevant for learning processes. Rather, once we have learned how accessible an artifact is, how to evaluate its information, and how to use it and interpret it; information flow, durability, and personalisation are automatically established.

to consciously think about interacting with the tool, the better the task is performed. When we have made the artifact transparent-in-use, we have enlarged our knowledge-how.

In a similar way as one learns how to use an artifact, one also learns how to use an informational system. Being able to interpret external information is essential to extend one's cognition. One can only extend one's cognitive systems when one sufficiently understands the syntax and semantics of the informational system one is using. The paradigm examples in extended and distributed cognition theory are language, number systems, scientific formulas, models, and diagrams. Before we can offload information, we must learn and internalise these systems, resulting in a transformation of our onboard cognitive systems. Building on the work of neuroscientist Stanislaw Dehaene, Richard Menary and Michael Kirchhoff (2014) present an interesting case of learning mathematics. The *unenculturated* brain can discriminate between small sets of items, but is not able to do long multiplication and other mathematical calculations; those must be learned through cultural practices. "The ability to perform exact calculations of mathematics depends upon the public system of representation and its governing norms. We learn the interpretative norms and manipulative norms as a part of a pattern of practices within a mathematics community and these practices transform what we can do" (2014, p. 620). Our learning histories thus reformat existing capacities of the embodied brain in terms of informational systems that we absorb and internalise.

Dwight Atkinson (2010) applies an embodied and extended cognition view to second language acquisition. His view is that: "If cognition is the site of learning, it is extended, embodied cognition that makes learning possible, at least in part" (2010, p. 612). He analyses an empirical case study where a Japanese high-school student (Ako) completes English grammar exercises by interacting with a tutor (Tomo) and an assignment-sheet. The student must convert statements containing time adverbials into "how" questions cued by the adverbial. Thus, when reading the statement "I have been busy for two weeks", the student's task is to formulate the question "How long have you been busy?" The student starts by reading the question out loud. Atkinson argues this actively externalizes her language processing into the sociocognitive problem space, creating shared attention between student and tutor. The tutor then repeats the underlined adverbial, thereby generating alignment between tutor and student and focusing their shared attention on this key part of the prompt. The student writes down the answer, which the tutor then approves. Atkinson concludes that cognition is extended, as it is "circulating across and through Ako, Tomo, and the grammar worksheet-as-cognitive technology" (2010, p. 613). Many learning processes are like this, where students interact with teachers, other students, and artifacts such as textbooks, interactive screens, assignment-sheets, abacuses, multiplication tables, pen and paper, books, (PowerPoint) presentations, diagrams, pictures, ball-and-stick models of molecules, and so on. Our embodied interactions with artifacts and other people are essential to learning and cognitive development.

5 Educating and assessing extended cognitive systems

5.1 Preventing negative consequences

If we couple with artifacts to form integrated and extended cognitive systems, how should such systems be educated? First, we should be aware that not all extended cognitive systems are cognitively beneficial. In some cases, extending one's cognitive system may have detrimental effects. In a careful reflection on the properties of extended cognitive systems, Rob Wilson writes:

"One way for cognitive extension to lead to a reduction in functional capacity is through cognitive clutter: by adding more bells and whistles to an existing cognitive system, we might well cause it to operate less effectively, or even lose certain kinds of functionality; so-called "smart technologies," for all the benefits they bring, are often used in ways that have this effect for particular tasks (e.g., in driving)" (2014, p. 23).

We should educate extended minds such that possible negative effects are prevented or reduced. There are at least two kinds of undesirable consequences of extending one's mind: (1) those that influence the wider cognitive system from effectively performing its cognitive tasks, and (2) those that are moral and social in nature. Examples of the first kind are clutter and information overload, including a messy desk with too many books, articles, and post-it notes or a search engine results page with more than a million search results. In such cases, information is not presented logically, effectively or there is just too much of it, preventing an agent from effectively performing a cognitive task.

An example of the second kind is concerned with privacy issues (see also Carter, Clark & Palermos, forthcoming; Carter & Palermos, 2016). One's extended memory like a notebook or Internet-based application such as Google Calendar can be accessed by others, potentially resulting in an infringement of one's privacy. It is important to educate people to use technologies in the best possible way and to make people aware that their extended minds might be accessed by others (Reiner & Nagel, 2017). We think this should be part of formal school and university curricula, for example as part of courses in digital literacy skills on which we elaborate below.

On an extended cognition view, the kind of tools one uses to perform a task partly determine how the task is done and have different effects on their users. An example is taking notes on a computer versus taking notes with pen and paper. Empirical research shows that "laptop note takers' tendency to transcribe lectures verbatim rather than processing information and reframing it in their own words is detrimental to learning" (Mueller & Oppenheimer, 2014, p. 1159). Making notes on a laptop generates more information because typing is faster than writing, which may be beneficial in some situations. However, because writing is slower, one has to process and think about what one writes down, resulting in a deeper understanding of the topic at hand. Pen and paper are simple, reliable, and effective tools for extending one's cognition; being able to write notes remains cognitively helpful, particularly in an educational context. As Neil Levy (2003) notes, calculators and spellcheckers have for some time raised the concern that the use of these tools will result in poorly developed (or atrophied) capacities to calculate and spell. The argument made here is that this is of concern because the tools are fallible, and the capacity to work without them is a part of intelligence. We are, of course, very much in favor of teaching pupils and students digital literacy skills, but that does not mean we should abandon pen and paper. A healthy mixture of learning to use analogue and digital tools seems preferable.

However, it is important to note that these are not particularly epistemological questions; that is, they are not about quantification of knowing – someone knowing *more* in one context than another – nor, even, questions of whether someone 'knows' in one context but not the other. Rather, they are issues regarding the normative judgements that are made in educational systems: we desire students to know particular things for various normative reasons (for employment, for cultural integration, for moral development, etc.). As such, decisions regarding the ways in which cognitive extensions are taught are in part normative in nature. It may be entirely feasible to outsource our arithmetic knowledge to a device, but we may place some premium on the nature of the knowledge entailed in mental arithmetic.

5.2 Organizing and finding information

"Focussing on the real-world cognitive situations that citizens encounter – situations which are these days laden with technology – is entirely the right approach for our educational policies to take" (Pritchard 2014, p. 50). We agree with Pritchard's observation and suggestion (see also Pritchard, 2013). We should therefore learn how to efficiently find and organise information. This applies to our own writing and notes, but also to information made by others, for example textbooks and online resources. Given the prevalence of the Internet in our information-seeking behaviors, having the skills to efficiently navigate, evaluate, compare, and synthesize online information are very valuable (Heersmink, 2016). Howard Rheingold's (2012) book *Net smart*:

How to thrive online is a helpful resource in this regard, giving suggestions for training your attention, evaluating information, participating in online communities such as Wikipedia, and drawing on online collective intelligence such as Amazon's book recommendation system.

Particularly important is learning how to use search engines, as they are the main portal to online information. Search engines are enormously helpful, but users should be aware that they have epistemically undesirable aspects. Personalised page ranking, for example, may result in confirmation bias and undermine objectivity (Simpson 2012; but compare Smart & Shadbolt 2018), and autocomplete of search terms may suggest search terms that are misleading or false, potentially nudging one towards an epistemically undesirable path of enquiry (Miller & Record 2017). In his recent book, Michael Lynch (2016, p. 66) gives an example of how search engines can suggest and prioritise false information. Google Search provides a "featured snippet" at the top of the search results page, when you ask it a question. When Lynch searched for "what happened to the dinosaurs?", this is what Google provided:

"The Bible gives us a framework for explaining dinosaurs in terms of thousands of years of history, including the mystery of when they lived and what happened to them. Dinosaurs are used more than almost anything else to indoctrinate children and adults in the idea of millions of years of earth history."

If one is unaware of evolutionary history, then this may seem like the truth. However, it is false and manipulative, as search engine optimisation strategies by creationists have led to this snippet being top-ranked for this question. This example suggests that it is tremendously important to teach pupils and students how to use search engines in an epistemically responsible way by being able to define search queries, choose the best search results, and evaluate sources for reliability and validity (see van Dijk & van Deursen 2014).

On Clark's principle of ecological assembly, we are opportunistic assemblers of extended cognitive systems and will use anything that does the job. This makes sense from an evolutionary perspective as one does not always have the time or resources available to select the best information. It certainly seems to apply to our use of online resources, but as van Dijk and van Deursen (2014) argue, it is much more effective to train pupils and students in their digital literacy skills from the outset, rather than to learn less effective ways by trial-and-error, as most of us do. Trial and error methods can be laborious, frustrating, inefficient, and ineffective. Once we have learned an ineffective way of doing an online task, we run the risk of continuing to perform the task in the same ineffective way. We think that digital literacy skills should be taught as part of formal school and university curricula.

A promising approach for teaching digital literacy skills is virtue epistemology, which is a set of approaches in contemporary epistemology, giving epistemic or intellectual virtues a key role (Battaly, 2008). Virtue responsibilism, one of the two main camps in virtue epistemology, emphasises the role of learned cognitive character traits such as open-mindedness, attentiveness, and intellectual autonomy in obtaining knowledge. The underlying idea is that agents who are intellectually virtuous are more likely to obtain true beliefs, knowledge and understanding than agents who are less intellectually virtuous (Zagzebski, 1996). Analogous to moral virtues, intellectual virtues require the right motivation, action, and affective response. So, an intellectually virtuous agent is intrinsically motivated to seek knowledge for its own sake, will actively seek knowledge and use appropriate strategies to do so, and feels rewarded when knowledge is obtained and disappointed when it is not obtained. Further, like moral virtues, intellectual virtues are a mean between two vices. For example, the virtue of open-mindedness is a mean between the vices of naivety and dogmatism. A naïve agent will consider too many options, whereas a dogmatic agent will consider too few.

The goal of virtue epistemology is to provide a framework for living an epistemically or intellectually virtuous life. It is thus a very promising approach to optimize our epistemic interactions with the Internet and other cognitive technologies (Heersmink, 2018; Michaelian & Arango-Muñoz, 2018). It can therefore provide the normative guidance that traditional extended mind theory lacks. Extended mind theory and the more naturalistically inclined distributed cognition theory typically focus on conceptualizing *how* agents interact with artifacts, not how they *should* interact with artifacts. However, as shown above, there are good and bad ways of using cognitive technology.

Heersmink (2018) adopted a virtue-epistemic approach to analyse how we might improve our interactions with Internet search engines. Building on the work of Jason Baehr (2011), he outlines nine intellectual virtues (i.e., curiosity, intellectual autonomy, intellectual humility, attentiveness, intellectual carefulness, intellectual thoroughness, open-mindedness, intellectual courage, and intellectual tenacity) and then explores how these virtues should be deployed when using Google Search. The analysis results in various suggestions for an epistemically virtuous use of search engines. Let us briefly highlight three of them. An intellectually open-minded agent will consider several alternative views, and if these views are more accurate, then the agent is willing to change one's mind. Personalised page ranking may result in confirmation bias and therefore undermine objectivity and open-mindedness, because it results in presenting an agent with information that is consistent with one's existing beliefs. To overcome this issue, an open-minded agent could turn off personalisation or use a search engine that doesn't personalise (Simpson 2012). An attentive agent pays close attention and has a sustained focus on the cognitive task at hand. The Internet is an informational environment that promotes cursory reading, skimming of information, and distracted thinking. Rheingold (2012) suggests mindfulness techniques to train oneself to become aware of distractedness and to force oneself to stay focused. One can also use software programs that block certain webpages (such as social media) at certain times of the day, in that way delegating attentiveness to a software system. Lastly, an intellectually autonomous agent can think for oneself, is cognitively capable, and has a certain degree of scepticism. When using search engines, such an agent will interpret Google's featured snippets, page ranking, and autocompleted search terms with a healthy dose of scepticism³.

A reviewer suggested to clarify whether an intellectual virtue is a trait of an individual agent or a trait of an extended cognitive system. On a responsibilist view, an intellectual virtue is a cognitive character trait that is truth-conducive and minimizes error, including open-mindedness, attentiveness, and intellectual autonomy. Such character traits are dispositions to think and act in certain ways. On our view, dispositions and intellectual character traits can be influenced and improved by environmental structures such as artifacts and other people (see also Alfano & Skorburg, forthcoming), but are not necessarily extended or distributed in the same way as cognitive states and processes are. On an extended and distributed cognition view, cognitive states (e.g., memory states) and processes (e.g., calculating, reasoning, and navigating) often include information stored in artifacts and other people. When the relation between the embodied agent and the external information ranks high on the dimensions outlined in section 2 and is thus densely integrated, we should think of the embodied agent and external resource as one cognitive system. Whilst technology may nudge one to be more attentive, intellectually careful, intellectually thorough; it seems difficult (but perhaps not impossible) to think of cases where technology becomes reciprocally integrated into cognitive character traits.

³ Note that in each of these cases, in Pritchard's terms as outlined in section 3, the virtuous integration of the technology goes beyond the mere crutch metaphor. So, the search engine is not simply providing support towards a capacity to be attained unsupported. Instead, the search engine becomes a part of the bridge into new cognitive domains, in which the more or less virtuous integration plays a core feature in how the cognitive capacity is developed, for example with or without diverse sources of high quality information, and towards the accessing of more advanced forms of information (see also Heersmink & Sutton, accepted).

Various philosophers have proposed teaching students to be intellectually virtuous (Baehr, 2013; Pritchard, 2013; Battaly, 2016). On this view, the goal of education is not (only) to provide students with information, but (also) to teach them to think for themselves. Heather Battaly (2016) suggests teaching intellectual virtues in three steps: first, explaining what intellectual virtues are by means of formal instruction; second, further clarifying individual responsibilist virtues by using several exemplars, i.e., giving examples of agents who are intellectually virtuous; third, providing opportunities to practice identifying intellectually virtuous actions and applying them in classroom settings. Battaly suggests including this teaching strategy in lower-division courses in logic or critical thinking. We think this is a good step towards educating students to become more intellectually virtuous. But given our focus in this paper and the prevalence of the Internet in our epistemic and cognitive practices, we think it is desirable to design and implement university courses in digital information literacy that include teaching intellectual virtues as they pertain to the Internet and other cognitive technologies. Ideally, other parts of the curriculum should also include teaching intellectual virtues. Moreover, because becoming intellectually virtuous is partly a matter of habit, it therefore seems preferable to start teaching these skills as soon as possible, ideally starting at primary school and continuing into secondary and tertiary education. It takes time to develop good epistemic habits, so the sooner one begins learning these skills, the better.

5.3 Assessing extended minds – the Danish example

We take it that assessments in formal educational contexts are conducted for a variety of purposes, each of which would be characterised as assessing the knowledge state of a learner; that is, assessment has an epistemological flavour (Davis, 1998). Distinctive aims under this general purpose include a desire to assess whether a learner has met some pre-specified criterion (i.e., criterion-based assessment), to rank the knowledge states of learners (i.e., norm-referenced assessment), and perhaps even to assess the quality of educational provision to the learner via the proxy measure of their learning. Across these aims the particular object of assessment might range from the skills or capacities of the learner to their accrual of particular propositional facts and their application via, for example, written essays. Given that, as David Boud (2000) flags, education should aim at alignment between educator and student understanding of assessment, it is important that contemporary educational debates are informed by a well-grounded understanding of the nature of knowledge and mind.

What we assess may seem remarkably intuitive: we assess the knowledge we have taught, which fits into the various subject domains. That knowledge is a mixture of knowledge-that and knowledge-how required in the specific discipline, a mixture of ‘facts’ and means to manipulate those facts in meaningful disciplinary ways. If we take the extended mind thesis seriously, which we take as a given in this paper, then we should take seriously the possibility that given the right kind of learning or enculturation, knowledge may extend beyond the bounds of the brain. As such, we should consider the implications of the extended mind thesis for how we understand education and the assessment of what is learned through that education. As outlined above, in considering the extended mind and learning, consideration must be paid to how students learn to integrate technology into their cognitive capacities. We thus distinguish three contexts:

1. The assessment of situated or distributed learning, in which external resources act as a crutch to support knowledge that could not otherwise be displayed, but perhaps where the use of these crutches is acceptable (or desirable) as an intermediary step in learning.
2. The assessment of knowledge-as-extended, in which external resources are considered part and parcel of the knowledge state of an agent.
3. The assessment of knowledge-de-extended, in which access to technology is deliberately inhibited in order to artificially assess an agent (perhaps to understand their extension-making ability). Note that this is not the same as removing a crutch; it is closer to knocking

the bridge down and seeing how people attempt to cross it (presumably using whatever crutches are available).

Clearly, within existing assessment regimes, judgements are made about what is, and is not, an acceptable use of external resources. We permit open book exams and use of calculators in some exams (and of mathematical notation in all exams), but not others. Indeed, in a recent paper, Helenrose Fives et al (2017) outline some ways in which an educator's implicit or explicit beliefs regarding the nature of knowledge (i.e., their epistemic cognitions) might play out in their treatment of assessment. However, as Simon Knight and Karen Littleton (2017) note, discussion of teacher's perspective on knowledge – their epistemic cognitions – has not drawn on the kinds of contemporary epistemological discussions flagged in this paper. Yet, it is apparent that these issues are implicated in many current debates around educational assessment.

What is distinctive about our account of assessment of extended minds is, firstly, that our primary focus is on assessing knowledge as encompassing of cognitive extensions, rather than of knowledge being simply an internal capacity (situated or otherwise). That is, we see the inclusion of external resources as fundamental to the knowledge state of the agent, not as an optional extra. Secondly, in the virtue epistemological account that we have outlined (see also Heersmink, 2018), clear constraints are introduced that provide limits to reliable and responsible cognitive extension.

In fact, we need not turn to science fiction for an example of an assessment regime in which tools, beyond the calculator or open book examination, are made available. In Denmark, a pilot study (which was subsequently rolled out) was conducted in which students were given access to the Internet during their high school examinations (Cunnane, 2011). Students in a variety of subjects, including Danish language and mathematics, are thus given access to most Internet websites (although not those which could be used to communicate with other students) in order to support their assessment. This is a natural extension of earlier Danish examinations which had allowed access to a variety of media sources via CD-ROM. For example, in Danish language exams, students might be given access to video, audio, webpages, etc., in order to push them to work across and integrate materials from political speeches, documentaries, movie clips, and so on. These exams, then, are designed to evaluate the student's ability to interpret, analyse, and evaluate media sources within the context of their wider knowledge regarding the subject content (and use of media environments). Similarly, in mathematics the students can use the Internet both to find formulae and other resources that one might typically find in a textbook, although of course they must know what resources to look for, and to check the results they obtain via calculation websites.

Thus, in Denmark, students are expected to have a certain degree of knowledge-that, a level which is suitable to assist their tasks. They also must use a rather high level of knowledge-how to manipulate information and process it effectively, as well as metacognitive strategies to make effective use of the tools to deploy their knowledge. The Internet provides potential for students to 'check' their answers, demonstrating a particular metacognitive knowledge. Across these cases, students must know how to conduct calculations or interpret historical texts. They must also know what to look for, who the key actors might be, which formulae might be relevant, thus requiring a certain degree of intellectual virtue. A component of these assessments is still the recall of factual knowledge-that; students can use the tools to augment this, for example, where they have failed to learn some key historical context or mathematical procedure, but to do so falls outside the description of cognitive extensions we have provided. In addition to this, the assessments also focus on the manipulation of novel information; in this case, the tool is not acting as crutch; it is a key component of the process through which the students address the problem. The tools can provide facts that they could retain internally (closer to being a crutch), but the key concern is how they integrate these with existing knowledge, how they seek and evaluate the

information in discipline-based normatively grounded ways, and so on. As such, a normative judgement has been made that the students must have knowledge-that regarding particular background knowledge and knowledge-how with regards to deploying that knowledge. To succeed in the assessments, the students must also demonstrate a degree of cognitive integration with the tools available.

6 Conclusion

Human beings have evolved to incorporate external objects and information into their cognitive systems. Such objects are then constitutive parts of the physical supervenience base of cognitive systems. In this paper, we have explored the implications of this view for learning, educating, and assessing such extended and distributed cognitive systems. We conceptualised how we learn to assemble extended cognitive systems through internalising cultural norms and practices. We have argued for the need to educate extended cognitive systems, including ways to minimize possible negative effects of extending one's cognition, and to efficiently find and organise information by adopting a virtue epistemology approach. We also argued that current assessment regimes should be adjusted, as they focus too much on assessing unaided brains and do not sufficiently reflect assessing extended cognitive systems. The Danish case of using the Internet during certain exams is a promising start to assess our extended cognitive systems.

Acknowledgments

Heersmink presented this paper at an invited seminar organised by the Eidyn research centre, University of Edinburgh and at the Epistemic Angst and Extended Knowledge workshop at Macquarie University. We would like to thank the organisers and audience for helpful feedback. Knights' contribution draws on work begun as a Masters student at the University College London Institute of Education, and developed in a talk and chapter for the Institute of Network Cultures Society of the Query conference. Knight extends his thanks to the organisers and participants at that event, and particularly to his supervisor, Dr. Jan Derry, at the Institute of Education. Lastly, we like to thank the two anonymous reviewers for very helpful comments.

References

- Adams, F., & Aizawa, K. (2001). The bounds of cognition. *Philosophical Psychology*, 14(1), 43-64.
- Alfano, M., & Skorburg, J.A. (forthcoming). The embedded and extended character hypotheses. In J. Kiverstein (Ed.), *Philosophy of the social mind*. London: Routledge.
- Arango-Muñoz, S. (2013). Scaffolded memory and metacognitive feelings. *Review of Philosophy and Psychology*, 4(1), 135-152.
- Atkinson, D. (2010). Extended, embodied cognition and second language acquisition. *Applied Linguistics*, 31(5), 599-622.
- Baehr, J. (2011). *The inquiring mind: On intellectual virtues and virtue epistemology*. Oxford: Oxford University Press.
- Baehr, J. (2013). Educating for intellectual virtues: From theory to practice. *Journal of Philosophy of Education*, 47(2), 248-262.
- Battaly, H. (2008). Virtue epistemology. *Philosophy Compass*, 3(4), 639-663.
- Battaly, H. (2016). Responsibilist virtues in reliabilist classrooms. In J. Baehr (Ed.), *Intellectual virtues and education: Essays in applied virtue epistemology* (pp. 163-183). London: Routledge.
- Boud, D. (2000). Sustainable assessment: Rethinking assessment for the learning society. *Studies in Continuing Education*, 22(2), 151-167.
- Carter, J.A., Clark, A., & Palermos, S.O. (forthcoming). New humans? Ethics, trust and the extended mind. In J.A. Carter, A. Clark, J. Kallestrup, S.O. Palermos & D. Pritchard (Eds.), *Extended epistemology*. Oxford: Oxford University Press.
- Carter, J.A., & Palermos, O. (2016). Is having your computer compromised a personal assault? The ethics of extended cognition. *Journal of the American Philosophical Association*, 2(4), 542-560.

- Clark, A. (2003). *Natural-born cyborgs: Minds, technologies, and the future of human intelligence*. New York: Oxford University Press.
- Clark, A. (2008). *Supersizing the mind: Embodiment, action, and cognitive extension*. Oxford: Oxford University Press.
- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 58(1), 10-23.
- Cunnane, S. (2011). The Danish gambit: online access, even during exams. Retrieved from www.timeshighereducation.com/news/the-danish-gambit-online-access-even-during-exams/416090.article.
- Davis, A. (1998). Understanding and holism. *Journal of Philosophy of Education*, 32(1), 41-55.
- Donald, M. (1991). *Origins of the modern mind: Three stages in the evolution of cognition and culture*. Cambridge: Harvard University Press.
- Fives, H., et al. (2017). Teachers' epistemic cognition in classroom assessment. *Educational Psychologist*, 52(4), 270-283.
- Gallagher, S. (2005). *How the body shapes the mind*. Oxford: Oxford University Press.
- Heersmink, R. (2013). A taxonomy of cognitive artifacts: Function, information, and categories. *Review of Philosophy and Psychology*, 4(3), 465-481.
- Heersmink, R. (2015). Dimensions of integration in embedded and extended cognitive systems. *Phenomenology and the Cognitive Sciences*, 14(3), 577-598.
- Heersmink, R. (2016). The Internet, cognitive enhancement, and the values of cognition. *Minds and Machines*, 26(4), 389-407.
- Heersmink, R. (2018). A virtue epistemology of the Internet: Search engines, intellectual virtues and education. *Social Epistemology*, 32(1), 1-12.
- Heersmink, R. & Sutton, J. (accepted). Cognition and the Web: Extended, transactive or scaffolded? *Erkenntnis*.
- Hutchins, E. (2011). Enculturating the supersized mind. *Philosophical Studies*, 152(3), 437-446.
- Hutchins, E. (2014). The cultural ecosystem of human cognition. *Philosophical Psychology*, 27(1), 34-49.
- Knight, S. (2014). Finding knowledge: What is it to 'know' when we search? In R. König & M. Rasch (Eds.), *Society of the query reader: Reflections on web search* (pp. 227-238). Amsterdam: Institute of Network Cultures.
- Knight, S., & Littleton, K. (2017). Socialising epistemic cognition. *Educational Research Review*, 21, 17-32.
- Kotzee, B. (forthcoming). Cyborgs, knowledge and credit for learning. In J.A. Carter, A. Clark, J. Kallestrup, S.O. Palermos & D. Pritchard (Eds.), *Extended epistemology*. Oxford: Oxford University Press.
- Kuhn, D. (2000). Metacognitive development. *Current Directions in Psychological Science*, 9(5), 178-181.
- Levy, N. (2003). Cyborgs-R-Us. In J. Weckert & Y. Al-Saggaf (Eds.), *Selected papers from the ninth international conference on computers and philosophy* (pp. 13-17). Canberra: Australia.
- Lynch, M. (2016). *The Internet of us: Knowing more and understanding less in the age of big data*. New York: W.W. Norton and Company.
- Menary, R. (2007). *Cognitive integration: Mind and cognition unbound*. London: Palgrave MacMillan.
- Menary, R. (2010). *Cognitive integration and the extended mind*. In R. Menary (Ed.), *The extended mind* (pp. 227-243). Cambridge: MIT Press.
- Menary, R., & Kirchhoff, M. (2014). Cognitive transformations and extended expertise. *Educational Philosophy and Theory*, 46(6), 610-623.
- Metcalfe, J., & Shimamura, A.P. (Eds.). (1996). *Metacognition: Knowing about knowing*. Cambridge: MIT Press.
- Michaelian, K., & Arango-Muñoz S. (2018). Collaborative memory knowledge: A distributed reliabilist perspective. In M. Meade, C. B. Harris, P. van Bergen, J. Sutton & A. J. Barnier (Eds.),

- Collaborative remembering: Theories, research, applications* (pp. 231-247). Oxford: Oxford University Press.
- Miller, B., & Record, I. (2017). Responsible epistemic technologies: A social-epistemological analysis of autocompleted web search. *New Media and Society*, 19(12), 1945-1963.
- Mueller, P.A., & Oppenheimer, D.M. (2014). The pen is mightier than the keyboard: Advantages of longhand over laptop note taking. *Psychological Science*, 25(6), 1159-1168.
- Nersessian, N. J. (2006). The cognitive-cultural systems of the research laboratory. *Organization Studies*, 27(1), 125-145.
- Nersessian, N.J., Kurz-Milcke, E., Newstetter, W.C., & Davies, J. (2003). Research laboratories as evolving distributed cognitive systems. In R. Alterman & D. Kirsh (Eds.), *Proceedings of the 25th Annual Conference of the Cognitive Science Society* (pp. 857-862). Cambridge, MA: USA.
- Norman, D. (1998). *The invisible computer: Why good products can fail, the personal computer is so complex, and information appliances are the solution*. Cambridge: MIT Press.
- Palermos, S.O. (2014). Knowledge and cognitive integration. *Synthese*, 191, 1931-1951.
- Perkins, D.N. (1993). Person-plus: A distributed view of thinking and learning. In G. Salomon (Ed.), *Distributed cognitions: Psychological and educational considerations* (pp. 88-110). Cambridge: Cambridge University Press.
- Pritchard, D. (2010). Cognitive ability and the extended cognition thesis. *Synthese*, 175, 133-151.
- Pritchard, D. (2013). Epistemic virtue and the epistemology of education. *Journal of Philosophy of Education*, 47(2), 236-247.
- Pritchard, D. (2014). Virtue epistemology, extended cognition, and the epistemology of education. *Universitas: Monthly Review of Philosophy and Culture*, 478, 47-66.
- Reiner, P.B., & Nagel, S.K. (2017). Technologies of the extended mind: Defining the issues. In J. Illes & S. Hossain (Eds.) *Neuroethics: Anticipating the future*. Oxford: Oxford University Press.
- Rheingold, H. (2012). *Net smart: How to thrive online*. Cambridge, MA: MIT Press.
- Rupert, R. (2004). Challenges to the hypothesis of extended cognition. *The Journal of Philosophy*, 101(8), 389-428.
- Sawyer, R.K., & Greeno, J.G. (2009). Situativity and learning. In M. Aydede & P. Robbins (Eds.), *The Cambridge handbook of situated cognition* (pp. 347-367). Cambridge: Cambridge University Press.
- Scardamalia, M., & Bereiter, C. (2003). Knowledge building. In *Encyclopedia of Education*, 2nd ed. (pp. 1370-1373). New York: Macmillan Reference.
- Scardamalia, M., & Bereiter, C. (2006). Knowledge building: Theory, pedagogy, and technology. In K. Sawyer (Ed.), *Cambridge handbook of the learning sciences* (pp. 97-118). New York: Cambridge University Press.
- Simpson, D. (2012). Evaluating Google as an epistemic tool. *Metaphilosophy*, 43(4), 426-445.
- Smart, P., Clowes, R., & Heersmink, R. (2017). Minds online: The interface between Web science, cognitive science, and the philosophy of mind. *Foundations and Trends in Web Science*, 6(1-2), 1-234.
- Smart, P., & Shadbolt, N. (2018). The world wide web. In J. Chase & D. Coady (Eds.), *Routledge handbook of applied epistemology*. New York: Routledge.
- 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. (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(4), 521-560.

- van Dijk, J., & van Deursen, A. (2014). *Digital skills: Unlocking the information society*. New York: Palgrave MacMillan.
- Wheeler, M. (2010). In defense of extended functionalism. In R. Menary (Ed.), *The extended mind* (pp. 245-270). Cambridge: MIT Press.
- Wheeler, M. (2011). Thinking beyond the brain: Educating and building, from the standpoint of extended cognition. *Computational Cultures*, 1.
- Wilson, R. (2014). Ten questions concerning extended cognition. *Philosophical Psychology*, 27(1), 19-34.
- 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). Cambridge: Cambridge University Press.
- Zagzebski, L. (1996). *Virtues of the mind: An inquiry into the nature of virtue and the ethical foundations of knowledge*. Cambridge: Cambridge University Press.