Extended cognition and the explosion of knowledge
David Ludwig

Abstract: The aim of this article is to show that externalist accounts of cognition such as Clark and Chalmers' (1998) "active externalism" lead to an explosion of knowledge that is caused by online resources such as Wikipedia and Google. I argue that externalist accounts of cognition imply that subjects who integrate mobile Internet access in their cognitive routines have millions of standing beliefs on unexpected issues such as the birth dates of Moroccan politicians or the geographical coordinates of villages in southern Indonesia. Although many externalists propose criteria for the bounds of cognition that are designed to avoid this explosion of knowledge, I argue that these criteria are flawed and that active externalism has to accept that information resources such as Wikipedia and Google constitute extended cognitive processes.

1. Introduction

According to the extended mind thesis, human cognition extends beyond skin and skull and is partly constituted by the environment. The aim of this article is to argue that an active externalism as it has been proposed by Clark and Chalmers (1998) implies an explosion of dispositional beliefs and knowledge that is caused by digital information resources such as Wikipedia or Google. My argument is based on the common observation that any form of externalism faces a demarcation problem: if human cognition extends beyond the boundaries of the organism, we need criteria to distinguish extended cognitive processes from non-cognitive processes in the environment. I will argue that demarcation criteria as they have been proposed by Clark and Chalmers' (1998) or Rowlands (2009) are not suited to exclude digital information resources such as Wikipedia or Google from constituting extended cognitive processes.

The possibility of a "web-extended mind" (Smart, 2012; Smart, Engelbrecht, Braines, Strub, & Hendler, 2009) or "e-memory" (Clowes, 2013) that is constituted by digital and mobile technologies is an important and potentially radical consequence of externalism that has also been acknowledged by Clark: "perhaps external representations on the Web, when integrated appropriately into the processes that govern an agent's behaviour, may count as parts of that agent's cognitive architecture"
(Halpin, Clark, & Wheeler, 2010). Although Clark seems to take the possibility of a "web-extended mind" seriously, he also proposes criteria (e.g., Clark, 2008, p. 80) that are designed to prevent at least contemporary online resources from constituting parts of the cognitive realm. I argue that these criteria are flawed and that externalists have to accept extended cognitive processes that are partly constituted by digital information resources. The consequence is a staggering explosion of dispositional beliefs and knowledge. Internalists may consider the conclusion of this article a reductio of the very idea of an extended mind while externalists will have to develop strategies that incorporate an explosion of knowledge into their theories.

2. The Demarcation Problem of Externalism

In their article "The extended mind," Clark and Chalmers present the now-famous thought experiment of Otto who suffers from a mild case of Alzheimer's and relies on a notebook as a substitute for his biological memory (1998, pp. 12–17). When Otto wants to keep important information, he writes it down in his notebook and is therefore able to access it later. In Clark and Chalmers' example, Otto wants to go to the Museum of Modern Art (MoMA) in New York. He consults his notebook and retrieves the information that the museum is on 53rd street. Clark and Chalmers compare Otto to Inga who also wants to go to the MoMA but does not need a notebook because she has the address stored in her biological memory. Of course, there are many differences between Otto's and Inga's information retrieval but Clark and Chalmers argue that they are analogous

in all functionally important aspects. For example, both Otto and Inga have access to a reliable and portable information resource that allows them to quickly retrieve the information that the MoMA is located on 53rd street. Why are many of us inclined to consider the notebook merely a tool for Otto's cognitive system while accepting Inga's biological memory as a part of her cognitive

system? According to Clark and Chalmers, there is no good reason to exclude external objects such as notebooks as potential parts of cognitive systems. Access to an external information resource can play the same functional role in a cognitive process as access to biological memory. Therefore, cognitive processes and systems are not always solely realized by the brain and the intuition that Otto's notebook is merely a tool for his cognitive system turns out to be a flawed prejudice.

As Clark elaborates in his book Supersizing the mind, Otto's incorporation of an external object in cognitive routines is by no means an isolated case. On the contrary, the example of Otto illustrates a ubiquitous phenomenon that can be described as "cognitive niche construction." Many species alter their environments to their own benefits. Common examples include spiders and beavers who create physical objects that become crucial parts of their environments. Some changes, however, serve distinctively cognitive purposes and can be described as "cognitive niche construction" which Clark defines "as the process by which animals build physical structures that transform problem spaces in ways that aid (or sometimes impede) thinking and reasoning about some target domain or domains" (Clark, 2008, p. 62; compare Stotz, 2010).

Cognitive niche construction can be the work of an individual. For example, I may use a notebook to be able to remember tasks such as grocery shopping. Or I may put an empty pasta box on my kitchen table to remember to buy pasta the next day. Or I organize my desk in order to minimize sources of distraction and to make sure that I have all my cognitive tools (books, laptop, calculator, coffee, etc.) available. Many of the most powerful cognitive niches, however, are social institutions. Consider, for example, libraries or online resources such as Wikipedia and the countless ways in which they transform problem spaces in ways that aid thinking and reasoning about target domains.

Although the ubiquity of cognitive niche construction provides a good reason for cognitive scientists to look beyond the brain, it also offers an important challenge for proponents of extended cognition: when does an external process qualify as part of a cognitive process? Where is the line between genuine extended knowledge and noncognitive background processes that my affect cognitive processing but are not themselves part of the cognitive machinery (Clark, 2008, p. 80; compare Allen- Hermanson, 2012; Rupert, 2004)? According to Clark and Chalmers, the information in Otto's notebook meets the requirements for extended cognition and Otto actually knows the location of the MoMA despite the fact that this information is stored in an external device instead of his brain. But what about information in Otto's favorite library? It would be absurd to claim that every piece of information that is available in the library constitutes a part of Otto's belief system. For example, Otto does not know a lot about rural communities in Papua New Guinea despite the fact that the library recently acquired a book on this topic.

If not all parts of the cognitive niche qualify as parts of cognitive systems or processes, proponents of extended cognition face a demarcation problem (or the problem of "cognitive bloat"; Clark, 2008, p. 80; Varga, 2013): in order to make a convincing case for extended cognition, externalists have to distinguish genuine cases of extended cognition from implausible candidates such as Otto's alleged beliefs about the rural communities in Papua New Guinea. Although demarcation issues are often raised as problems for proponents of extended cognition (e.g., Marsh, 2010, p. 312; Rupert, 2009, pp. 15–18), Clark and Chalmers do not seem to be too worried about the demarcation problem. Instead, they suggest criteria that are supposed to distinguish genuine cases of extended cognition from implausible candidates such as libraries. In the remainder of this article, I will argue that this strategy is not satisfying and that proponents of extended cognition have to get comfortable with a far-reaching "explosion of knowledge" through external media.

3. Clark and Chalmers' Demarcation Criteria

How can we determine whether an external process qualifies as part of a cognitive process? Clark and Chalmers' general answer to this question is their parity principle: 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. (1998, p. 8)

The parity principle, however, is not proposed as a precise criterion but rather as a convenient "informal test" (Clark, 2010, p. 53) while more specific criteria would have to be informed by research in cognitive science. Although Clark and Chalmers avoid commitment to one precise answer to the demarcation problem, they offer the following "rough-and-ready set of additional criteria to be met by non-biological candidates for inclusion into an individual's cognitive system" (Clark, 2008, p. 80):

(1) That the resource be reliably available and typically invoked. (Otto always carries the notebook and won't answer that he "doesn't know" until after he has consulted it.)
(2) That any information thus retrieved be more or less automatically endorsed. It should not usually be subject to critical scrutiny (e.g., unlike the opinions of other people). It should be deemed about as trustworthy as something retrieved clearly from biological memory.
(3) That information contained in the resource should be easily accessible as and when required.
(4) That the information in the notebook has been consciously endorsed at some point in the past and indeed is there as a consequence of this endorsement.

The criteria (1)–(4) provide prima facie attractive tools for distinguishing genuine cases of extended cognition from implausible candidates such as "beliefs" that are realized by libraries. Contrary to Otto's notebook, a library does not satisfy any of these additional criteria: (1) It is only available at a specific place during opening hours. (2) Not all the information in a library will be "automatically endorsed" by Otto. (3) The information is not always easily accessible. (4) Most of the information in the library has not been endorsed by Otto in the past and it

certainly isn't "there as a consequence of this endorsement." It seems, then, that these four criteria do a pretty good job in explaining why Otto has the dispositional belief that the MoMA is on 53rd street and why he does not have beliefs about rural communities in Papua New Guinea despite a book on this topic in his favorite library. Unfortunately, the situation becomes more complicated when we turn to other examples. Imagine Otto buying his first cell phone with Internet access. During the first month, he carries both his notebook and cell phone around but his cell phone successively takes over the functions from his old notebook. For example, Otto figures out that he only has to type the name of an institution such as "MoMA New York" in the search bar of his cell phone and Google will display the address as the first search result. Furthermore, he realizes that this cell-phone-based information access is faster and more reliable than his old notebook method. As a consequence, he gets rid of his old notebook and the next time he wants to go to the MoMA, he simply accesses the information on Google.

Is this still a case of extended cognition and should we consider Google's database an external memory storage in the same way as Otto's notebook? It certainly seems that there are good reasons to consider both cases to be on par. It would be strange to claim that Otto forgot the location of the MoMA and other addresses because he switched to a more efficient method of accessing information. Furthermore, Clark and Chalmers introduce the example of Otto because they want to show that function and not location matter in cognitive processing. However, it seems that Otto's access to the address through Google is functionally equivalent to his access to the address through his notebook.

Of course, there is one obvious difference between Otto's information access from his notebook and from his cell phone. Otto added the information– e.g., the address of the MoMA– to the notebook himself while he did not add the information to Google's database. Given Clark and Chalmers' fourth criterion, one can argue that the information from Google's database cannot constitute extended beliefs in the same way as information from Otto's notebook. In the following section, I will argue that this strategy of excluding digital media will not work as

the fourth criterion actually turns out to be incompatible with Clark and Chalmers' active externalism.

4. Why Active Externalism Implies an Explosion of Knowledge

If we drop the fourth criterion, we seem to open "the floodgates to what many would regard as an unwelcome explosion of potential dispositional beliefs" (Clark, 2008, p. 80). At the same time, it is far from clear why one should accept (4) and treat conscious endorsement and causal responsibility as relevant for extended cognition. Consider my example of Otto starting to retrieve address information from a search engine on his cell phone instead of his notebook. (4) implies that Otto forgot the location of theMoMA and other addresses because he switched from his old notebook method to a new cell phone method of accessing information. Certainly, there appears to be a tension between this implication and any externalism that relies on functional equivalence in its arguments. Furthermore, it seems that the notebook method and the cell phone method can be not only functionally but also phenomenologically equivalent. Otto's use of his cell phone can become part of his second nature in the same way as the notebook used to be part of his second nature. Consider Clark's description of Otto's use of the notebook:

> Otto is so accustomed to using the book that he accesses it automatically when biomemory fails. Calls to the notebook are . . . deeply and subpersonally integrated into his problem-solving routines. The notebook has become transparent equipment for Otto, just as biological memory is for Inga. And in each case, doesn't it add needless and psychologically unreal complexity to introduce additional beliefs about the book or biological memory into the explanatory equations? (2008, p. 80)

Exactly the same could be said about Otto's access to address information through a search engine on his cell phone. How should externalists react to this tension

between the implications of (4) and externalism? One option is to simply bite the bullet of accepting implications such as the claim that Otto forgets the addresses of the MoMA when switching from his old notebook method to his new cell phone method of accessing information. Indeed, this implication may be counterintuitive but externalists never claimed to offer a conceptual analysis of the concept of cognition. A slightly counterintuitive revisionist stance that requires (4) in addition to functional and phenomenal equivalence may therefore be justified by the methodological advantages even if it is prima facie hard to see why Otto should forget the address of the MoMA when switching to an equivalent or even more efficient way of accessing the address information. In this section, I will argue that this justification of (4) won't be successful as (4) turns out to be incompatible with the overall claims of an active externalism.

Consider Laura who constantly uses Wikipedia on her cell phone to check basic facts such as biographical data or geographical coordinates. For Laura, mobile access to Wikipedia has become part of her second nature and she accesses the online encyclopedia automatically when her biological memory fails. Does Wikipedia constitute parts of Laura's belief and knowledge system? If so, we face a literally breathtaking explosion of knowledge (Ludwig, forthcoming). Laura would know the biographical data of hundreds of thousands of people. She would have memorized millions of geographical coordinates. She would know the latest election results of every major country in the world, the population estimates of thousands of cities. And so on.

Even many proponents of extended cognition will be reluctant to accept this consequence and look for ways to exclude Wikipedia from Laura's belief system. However, criteria (1) and (3) are met: information from Wikipedia is reliably available, typically invoked, and easily accessible on Laura's cell phone. Often, (2) is also met: Laura (rightly) believes that Wikipedia is more reliable than her biological memory with regard to basic facts such as the population of a city, the geographical coordinates of an island, or biographical data of a scientist. Therefore, she automatically endorses

this information when she retrieves it from Wikipedia. One may object that at

least contemporary online access is not quick and efficient enough to meet the third condition. As Smart (2012) explains in his discussion of the possibility of a "web-extended mind":

> Think about the problem of accessing factual information from a Web-accessible resource, such as Wikipedia. Even if the delays associated with document retrieval and presentation are resolved, the user is still confronted with the onerous task of surveying the document for relevant information content. In most cases, this requires the user to scroll through the Web page and process large amounts of largely irrelevant content in order to identify the small amount of information that is actually needed. This is a very inefficient means of information access. (Smart, 2012, p. 454)

There is certainly some truth to this objection. Often, information access from an online resource requires too much time and cognitive effort to satisfy Clark and Chalmers' third condition for extended cognition. Think of a rather simple example such as the belief that Charles Darwin became the Secretary of the British Geological Society in 1838. It would be odd to ascribe Laura this belief only because the Wikipedia article on Darwin contains this information and because Laura has reliable access to Wikipedia. Laura would still have to scroll through the article and process large amounts of information in order to find out that Charles Darwin became the Secretary of the British Geological Society in 1838. Therefore, Clark and Chalmers' third condition is not met and this piece of information does not constitute an extended belief.

But even if Smart's objection is successful in this case, it won't be successful with respect to many other examples of online access to information. Consider the belief that Charles Darwin was born in 1809. Laura does not have to scroll through the article and process large amounts of information to access this information. As an experienced Wikipedia user, Laura knows that she'll find Darwin's birth date in the first sentence of the Wikipedia article on Darwin. There is no reason to doubt that

Clark and Chalmers' third condition is met. In the case of Darwin's birth date,

Laura's access to information from Wikipedia will be at least as quick and reliable as Otto's access to information from his notebook. Furthermore, there are countless examples that are analogous to the case of Darwin's birth date. Even if we only consider a very limited set of information such as biographical data and geographical coordinates, access to Wikipedia seems to ensure that Laura has literally millions of unexpected dispositional beliefs about birth dates of scientists and politicians, geographical coordinates of villages, islands, monuments, and so on.

If we only consider Clark and Chalmers' conditions (1)–(3), then we have to accept an explosion of knowledge and the fact that Laura actually has millions of unexpected dispositional beliefs that are constituted by information from Wikipedia and other online resources. If we want to avoid this conclusion, we have to endorse the fourth condition that requires some form of authorship or causal responsibility for information (compare Theiner, 2011, pp. 69–71). Most of us are not authors or causally responsible for any information on Wikipedia. Therefore, (4) seems to offer an attractive way of excluding millions of Wikipedia-based dispositional beliefs about birth dates, geographical coordinates, and so on.

Unfortunately, a slight variation of the thought experiment illustrates that the fourth criterion is incompatible with Clark and Chalmers' active externalism (see also Rupert, 2004 for similar worries about the fourth criterion). Let us assume that Laura is also a passionate editor of Wikipedia and that she specializes on updating large chunks of data. Whenever she finds a new and trustworthy data set, she adds the missing information to Wikipedia. As she often adds large chunks of data that require an update of hundreds of articles, she does not keep track of what articles she has edited. For example, she has updated the population of Tijuana but not the population of Guadalajara. She has also updated the geographical coordinates of the Belizian island Long Cay but not of Hick's Cay. If we accept (4), we have to conclude that Laura knows the population of Tijuana and the geographical coordinates of Long Cay but not the population of Guadalajara or the geographical coordinates of Hick's Cay.

However, this is unacceptable for proponents of an active externalism: according to

our thought experiment, Laura does not even remember what articles she has updated. If she needs the information at some later point, her access to the information about Tijuana or Long Cay is functionally and phenomenologically equivalent her access to the information about Guadalajara or Hick's Cay. If she wants to know the population of Tijuana and Guadalajara or the geographical coordinates of Long Cay and Hick's Cay, she will look up the Wikipedia articles and her information retrieval will be exactly the same no matter whether she added the information or not.

Clark and Chalmers stress the active character of their externalism:

> In the cases we describe . . . , the relevant external features are active, playing a crucial role in the here-and-now. Because they are coupled with the human organism, they have a direct impact on the organism and on its behaviour. In these cases, the relevant parts of the world are in the loop, not dangling at the other end of a long causal chain. Concentrating on this sort of coupling leads us to an active externalism, as opposed to the passive externalism of Putnam and Burge. (Clark & Chalmers, 1998, p. 13)

The problem is that (4) undermines the active character of their externalism as the difference between Laura's access to the information about Tijuana/Long Cay and the information about Guadalajara/Hick's Cay is entirely historical while the relevant roles "in the here-and-now" are exactly the same. Another way of making the same point is to invoke Twin-Earth thought experiments. Imagine Twin-Laura who updates Twin-Wikipedia on Twin-Earth. Twin-Earth is an exact replica of our Earth with only one difference:

- When Twin-Laura edited Twin-Wikipedia, the data of Twin-Guadalajara was outdated, so she updated the article.
                    versus
- When Laura edited Wikipedia, the data of Gudalajara was not outdated so she did not update the article.

If we endorse (4), we have to conclude that Twin-Laura knows the population of

Twin-Guadalajara (she updated the article) while Laura does not know the population of Guadalajara (she did not update the article). However, this is incompatible with an active externalism as Laura and Twin-Laura can be in exactly the same situation when they look up the population of Guadalajara on a later occasion. There is no functional or phenomenological difference between Laura's and Twin-Laura's information access. The only difference is historical which is why the endorsement of (4) would transform an active externalism into a passive Putnam-Burge-style externalism.

5. Responses

The examples of the last section suggest that (4) is not only counterintuitive but also at odds with the very idea of an active externalism. Obviously, one can react to this situation in different ways. Internalists are in the most convenient situation as they can simply argue that the explosion of knowledge through digital media creates a further problem for or maybe even a reductio of externalism. Externalists can react to this situation in three different ways: they can (i) endorse an at least partly passive externalism that sticks with (4); (ii) formulate different criteria that exclude media such as Wikipedia; or (iii) embrace the explosion of knowledge.

5.1 (Partly) Passive Externalism

One option is to endorse an at least partly passive externalism which accepts that historical aspects such as the difference between Laura and Twin-Laura are crucial for distinguishing cognitive and non-cognitive processes. The importance of historical aspects does not contradict the importance non-historical functional aspects– both of them may be necessary but only jointly sufficient. By combining functional and historical criteria, externalists may hope to formulate a position that is "active" enough to meet Clark and Chalmers' intuitions while still accepting (4) and therefore the following claims:

(a) Laura knows the population of Tijuana and the geographical coordinates of

Long Cay but she does not know the population of Guadalajara or the geographical coordinates of Hick's Cay.

(b) Twin-Laura knows the population of Twin-Guadalajara but Laura does not know the population of Guadalajara.

I do not think that this move suits externalists who share Clark and Chalmers' motivation and I assume that any position that relies on historical criteria such as (4) will end up in conflict with the intuitions that fuel active externalism. The very point of a Clark-Chalmers-style externalism is to reject artificial boundaries between the cognitive and non-cognitive that do not reflect important differences. However, the boundaries that are implied by (a) and (b) are at least as artificial and dubious as the internalist boundary between brain-bound and external processes. If externalists reject the latter boundary because it is functionally uninteresting, they should certainly not endorse the former boundaries that are vulnerable to the same kind of objections. Or, to put it differently: of course, we can add a historical criterion such as (4) as a necessary condition for cognitive processes (who could stop us?), but Clark and Chalmers' active externalism is built on the idea that the distinctions we draw between cognitive and non-cognitive processes need justification. We need to be able to show that our distinctions correspond with meaningful differences that are, for example, in line with the parity principle or in some other way interesting for cognitive scientists. However, my thought experiments of Laura suggest that (4) will imply distinctions such as (a) and (b) that are not justifiable by pointing towards the parity principle or any other meaningful differences in Laura's cognitive routines.

The dubious status of the distinctions that are implied by (4) also becomes clear by considering possible scientific applications. One the one hand, we can imagine a meaningful internalist framework in psychology that rejects Laura's knowledge of the population of Tijuana and Guadalajara as she only has access to this information through Wikipedia. On the other hand, we can also imagine a meaningful externalist framework in psychology that accepts Laura's knowledge of the population of Tijuana and Guadalajara as her information access is in many

ways equivalent to her access to brain-bound information. However, it is very hard to imagine any scientific use of a

framework that treats Laura as knowing the population of Tijuana but not of Guadalajara. Such a framework does not seem useful because it implies distinctions that do not correspond with the kind of differences (e.g., behavioral, neural, phenomenal) that cognitive scientists are usually interested in. The very point of an externalism in the spirit Clark and Chalmers is to allow cognitive scientists to draw the line between cognitive and non-cognitive processes along differences that actually matter for their research instead of forcing them to accept criteria such as "inside the head versus outside the head" that seem arbitrary from a research perspective. Unfortunately, the distinctions that are implied by (a) and (b) will appear equally arbitrary as they do not correspond to differences that are of importance for cognitive scientists.

One may object that historical criteria such as (4) can be important in science as a look at other disciplines in the life sciences illustrates. For example, many accounts of species in contemporary biology are at least partly historical as they postulate monophyletic origin as a necessary condition for species membership (compare Ereshefsky, 2000). However, there is an important difference between partly historical accounts of species and partly historical accounts of cognition in the sense of (4). In the case of species, shared ancestry is often suggested as a criterion because it is of utmost importance for the understanding of biological differences between organisms. The situation is different in the case of (4) as it is far from clear what explanatory benefits cognitive scientists would gain from distinctions as they are implied by (a) or (b). Of course, this does not prove that it is empirically or even metaphysically impossible that cognitive scientists could develop interests that turn (4) into an interesting criterion. But such a proof is not necessary and also not available with respect to countless other arbitrary criteria. Externalism in the spirit of Clark and Chalmers avoids the problem of arbitrary criteria by urging us to look at the actual explanatory interests of cognitive scientists. This suggestion brings us back to the observation that distinctions as they are implied by (a) and (b) do not seem to correspond to any differences that are of actual

importance in contemporary cognitive science.

5.2 Alternative Criteria

Externalists do not have to accept Clark and Chalmers' criteria and one may argue that the explosion of knowledge can be avoided by considering alternative externalist responses to the demarcation problem. Rowlands (2009), for example, suggests an alternative set of criteria in his discussion of the problem of cognitive bloat. According to Rowlands (2009, p. 8), a process P is a cognitive process if and only if the following four conditions are met:

> (1) P involves information processing– the manipulation and transformation of information-bearing structures.
> (2) This information processing has the proper function of making available either to the subject or to subsequent processing operations information that was (or would have been) prior to (or without) this processing, unavailable.
> (3) This information is made available by way of the production, in the subject of P, of a representational state.
> (4) P is a process that belongs to the subject of that representational state.

Although Rowlands' criteria clearly exclude many background processes from counting as cognitive processes, it is far from clear that they also exclude digital information resources that are accessed through a mobile device. Let us stick with the example of Laura's use ofWikipedia to access the information that Darwin was born in 1809. I assume that there can be little doubt that this example satisfies Rowlands' criteria (1)–(3). Clearly, Laura's information access "involves information processing" in the sense of (1) as Laura's cell phone use involves the manipulation and transformation of an information-bearing structure. Furthermore, the function of this process is to make information available that was previously unavailable and therefore satisfies Rowlands' condition (2). Finally, it is equally obvious that the condition (3) is satisfied through Laura's representational state that Darwin was born in 1809.

Whether Rowlands' criteria avoid an explosion of knowledge therefore depends on the interpretation of his fourth criterion, i.e., the idea that the process must belong

to the cognitive subject. Although Rowlands admits that the criterion of ownership is rather vague, he attempts to clarify it through an analogy with ownership of noncognitive digestive processes (2009, pp. 16–17). Usually, digestive processes in my body belong to me while digestive processes in someone else's body do not belong to me. However, we can construct science fiction thought experiments in which my digestive processes are located in someone else's body if they are properly integrated with my bodily functions in the sense that they, for example, process my food and release energy into my body. Rowlands suggests that we can understand ownership of cognitive processes in an analogous way. Integration and not location is the crucial feature– in order for a process to be my cognitive process, it must be properly integrated in my cognitive system.

Rowlands' criteria avoid an explosion of knowledge only if one can show that information access from digital resources such asWikipedia is not properly integrated in cognitive systems. Unfortunately, I do not see why Otto's access to Google or Laura's access to Wikipedia should not be properly integrated in their cognitive systems. Again, consider the case of Otto switching from a notebook-based method of information access to a cell-phone-based method of information access. While Otto used to retrieve the information that the MoMA is located of 53rd street from his notebook, he now retrieves it from a search engine on his cell phone. I have argued that both methods can turn out to be functionally and phenomenologically equivalent: in both cases, Otto has quick and reliable access to information and one can imagine that both forms of information access can become part of Otto's "second nature" in the sense that he automatically accesses and endorses the information when his biological memory fails. Both forms of information access therefore show the same degree of integration, and Rowlands' criteria do not lead to an externalism that avoids the explosion of knowledge.

Even if Clark and Chalmers' criteria as well as Rowlands' criteria imply an explosion of knowledge, externalists may still hope to come up with an alternative set of criteria that leaves room for extended cognition while excluding information resources such as Google or Wikipedia. For example, Kaplan has suggested an account "according to which cognitive boundaries are determined by relationships

of mutual manipulability between the properties and activities of putative components and the overall behavior of the cognitive mechanism in which they figure" (2012, p. 545). Weiskopf has proposed a "view of systems demarcation on which cognitive systems are sets of

mechanisms for producing cognitive processes that are bounded by transducers and effectors" (2010, p. 313). Even if both proposals are very interesting, they will clearly not solve the externalists' troubles with Wikipedia. Kaplan's "mutual manipulability" provides functional conditions for cognitive processes that will not distinguish between Laura's use of Wikipedia and typical externalist examples such as Otto's use of a notebook. Weiskopf follows Pylyshyn's (1984) characterization of transducers and effectors which leads to an internalist framework that prevents an explosion of knowledge but also excludes all other common externalist examples.

Another potential source for demarcation criteria are "second wave externalisms" that do not rely on considerations of parity and the similarity of internal and external processes (Menary, 2010; Sutton, 2010). Instead of parity, second wavers stress "complementarity" and the functional entanglement of the brain with external media:

> Brains like ours need media, objects, and other people to function fully as minds. Seeing the brain as a leaky associative engine, its contents flickering and unstable rather than mirroring the world in full, forces attention to our reliance on external representations in the technological and cultural wild. (Sutton, 2010, p. 205)

Although second wave externalism is attractive in many ways, it does not offer a simple solution to the problem of an explosion of knowledge. On the contrary, it seems obvious that a consideration of functional entanglements and complementarity will not lead to a distinction between standard cases of extended cognition and my examples involving Wikipedia or Google because of their functional equivalence.

Of course, there could still exist some set of externalist criteria that successfully prevents an explosion of knowledge, but it is difficult to see what such criteria

could look like. Furthermore, the functional as well as phenomenological equivalence of Otto's methods of information retrieval suggest that it will be very difficult to formulate criteria that avoid an explosion of knowledge while still being consistent with the basic intuitions that motivate externalism.

5.3 Embracing the Explosion of Knowledge

If externalists give up (4) and cannot come up with alternative criteria that exclude digital media, they have to pay the price of an explosion of knowledge. People such as Otto and Laura who have made mobile access to Google or Wikipedia part of their second nature have a truly impressive amount of factual knowledge. Otto does not only know the address of the MoMA but the address of almost every major museum in the world. Laura does not only know the population of Guadalajara but of every major city in the world. And so on. Quick and reliable mobile access to information resources such as Google or Wikipedia literally creates millions of dispositional beliefs.

Many internalists will be inclined to interpret this result as a reductio of the very idea of extended cognition. However, an explosion of knowledge may actually not be as absurd as it appears on first sight. People who incorporate digital and mobile technologies in their everyday life and their cognitive routines alter their cognitive economy profoundly. The thought experiment of Otto's cell phone is just one example of the transformation of analog into digital cognitive environments that is more than common at the beginning of the twenty-first century. Maybe it is not that implausible to claim that this transformation comes with an explosion of dispositional beliefs and knowledge about simple facts such as addresses or city populations.

Note that an explosion of knowledge in this sense does not imply an explosion of understanding of complex situations. Laura, for example, may know a lot of basic facts about Charles Darwin, but that does not mean that she has actually understanding of Darwin's scientific work. She may know many demographic facts about Mexico, but that does not mean that she has a profound understanding of the demographics of Mexico (demographic dynamics, their underlying causes,

their effects on society, etc.). Although this kind of in-depth information may be available online, its retrieval would require too much time and cognitive effort to satisfy Clark and Chalmers' third condition and to constitute extended beliefs.

Given this distinction between knowledge of simple facts such as addresses, biographical data, or population numbers and understanding of complex situations, the explosion of knowledge through digital media may actually sound less absurd and become part of a convincing interpretation of the impact of digital media and mobile technologies on our cognitive economies (see also Ludwig, forthcoming).

References

- Allen-Hermanson, S. (2012). Superdupersizing the mind: Extended cognition and the persistence of cognitive bloat. Philosophical Studies, 158(1), 1–16.
- Clark, A. (2008). Supersizing the mind. Oxford: Oxford University Press.
- Clark, A. (2010). Memento's revenge. In R. Menary (Ed.), The extended mind (pp. 43–66). Cambridge, MA: MIT Press.
- Clark, A., & Chalmers, D. (1998). The extended mind. Analysis, 58, 7–19.
- Clowes, R. (2013). The cognitive integration of e-memory. Review of Philosophy and Psychology, 4, 107–133.
- Ereshefsky, M. (2000). The poverty of the Linnaean hierarchy: A philosophical study of biological taxonomy. Cambridge: Cambridge University Press.
- Halpin, H., Clark, A., & Wheeler, M. (2010, April). Towards a philosophy of the web: Representation,enaction, collective intelligence. Paper presented at the 2010WebSci ACMconference, Raleigh, NC.
- Kaplan, D. M. (2012). How to demarcate the boundaries of cognition. Biology & Philosophy, 27, 545–570.
- Ludwig, D. (forthcoming). Extended cognition in science communication. Public Understanding of Science.
- Marsh, L. (2010). Introduction to the special issue "Extended Mind".

Cognitive Systems Research, 11, 311–312.

- Menary, R. (2010). The extended mind in focus. In R. Menary (Ed.), The extended mind (pp. 1–26). Cambridge, MA: MIT Press.

- Pylyshyn, Z. (1984). Computation and cognition. Cambridge, MA: MIT Press.

- Rowlands, M. (2009). Extended cognition and the mark of the cognitive. Philosophical Psychology, 22 (1), 1–19.

- Rupert, R. (2004). Challenges to the hypothesis of extended cognition. Journal of Philosophy, 101, 389–428.

- Rupert, R. (2009). Cognitive systems and the extended mind. Oxford: Oxford University Press.

- Smart, P. (2012). The web-extended mind. Metaphilosophy, 43, 446–463.

- Smart, P., Engelbrecht, P., Braines, D., Strub, M., & Hendler, J. (2009, March). Cognitive extension and the web. Paper presented at the 2009 WebSci ACM conference, Athens, Greece.

- Stotz, K. (2010). Human nature and cognitive-developmental niche construction. Phenomenology and the Cognitive Sciences, 9, 483–501.

- Sutton, J. (2010). Exograms and interdisciplinarity: History, the extended mind, and the civilizing process. In R. Menary (Ed.), The extended mind (pp. 189–226). Cambridge, MA: MIT Press.

- Theiner, G. (2011). Res cogitans extensa: A philosophical defense of the extended mind thesis. Bern: Peter Lang.

- Varga, S. (2013). Frames of cognition. Cognitive Systems Research, 25-26, 54–60.

- Weiskopf, D. A. (2010). The Goldilocks problem and extended cognition. Cognitive Systems Research, 11, 313–323.