

The Cognitive Integration of E-Memory

Robert W. Clowes

Published online: 26 January 2013
© Springer Science+Business Media Dordrecht 2013

Abstract If we are flexible, hybrid and unfinished creatures that tend to incorporate or at least employ technological artefacts in our cognitive lives, then the sort of technological regime we live under should shape the kinds of minds we possess and the sorts of beings we are. E-Memory consists in digital systems and services we use to record, store and access digital memory traces to augment, re-use or replace organismic systems of memory. I consider the various advantages of extended and embedded approaches to cognition in making sense of E-Memory and some of the problems that debate can engender. I also explore how the different approaches imply different answers to questions such as: does our use of internet technology imply the diminishment of ourselves and our cognitive abilities? Whether or not our technologies can become actual parts of our minds, they may still influence our cognitive profile. I suggest E-Memory systems have four factors: totality, practical cognitive incorporability, autonomy and entanglement which conjointly have a novel incorporation profile and hence afford some novel cognitive possibilities. I find that thanks to the properties of totality and incorporability we can expect an increasing reliance on E-Memory. Yet the potentially highly entangled and autonomous nature of these technologies pose questions about whether they should really be counted as proper parts of our minds.

1 E-Memory and Cognitive Hybridity

We appear to be living through an historical moment of rapid development and deployment of new memory technologies, based upon a linked series of innovations in networked digital technologies such as immensely cheap digital memory and massively powerful recording technologies such as mega-pixel cameras and the vast array of sensors and cheap memory included in electronic goods. From iPods to

R. W. Clowes (✉)
IFL, Faculdade de Ciências Sociais e Humanas, Universidade Nova de Lisboa,
1069-061 Lisboa, Portugal
e-mail: robertc@sussex.ac.uk

tablets, from SenseCams to digital voice recorders to the now burgeoning market in “apps” for our mobile gadgetry: New E-Memory technologies are becoming pervasive. Since ever more of this technology is linked via the mobile internet to increasingly massive personal data repositories existing in the cloud it becomes an ever-present background to our lives. As memory technology goes through a revolutionary period, does this have implications for our minds?

According to Andy Clark, we are beings whose nature it is to create and then bond with our technologies; we are, he claims, *Natural Born Cyborgs* (2003). If this is correct, then the potential incorporability of technology is a deep factor in how our brains and minds work of which the new technology is just taking advantage. On this controversial analysis our minds are not merely evolved and composed of our organic brains but are a heterogeneous bunch of sub-systems partly organic, partly composed of technological extensions: Our minds are hybrid.

The notion of hybrid minds has been developed by Menary (2010)—alongside related notions like cognitive integration (Sutton 2010)¹—in an attempt to understand some of the complexes that our minds forge through this reliance upon, and potentially, incorporation of, a heterogeneous array of artefacts. The idea appears to rest on, and was inspired by, the idea of the extended mind, or the Hypothesis of Extended Cognition (hereafter HEC) (Clark and Chalmers 1998; Donald 1993) which holds that environmental resources, although they have a significantly different organic composition and functional profile from our brains, can nevertheless count as proper parts of our minds. But more interestingly, even if we subscribe to the less radical “Brainbound”² Hypothesis of Embedded Cognition (hereafter HEMC (Rupert 2004)), which holds that our technology only situates our minds—rather than being part of it—there may still be important implications for our minds from the use of new technology. Our minds may change their processing profile to accommodate extra-mental cognitive artefacts, even if we do not consider them actually to incorporate technological artefacts as proper parts.

Memory technologies have emerged as an important test case in the discussion of the Extended Mind. First because the idea of extended memory has become something of a test case for the extended mind more generally (e.g., Adams and Aizawa 2001; Rupert 2004).³ Second, because, if there is any realm of the human mind that has been revolutionized by our creation and use of technology, it is memory (e.g., Goody 1977; Ong 1982; Yeo 2008). Historically, technology has played a central role in how human memory has been organised (Donald 1991; Gregory 1981; Luria and Vygotsky 1992) and new technologies may already be having important effects on the way in which our memories are organised. Thirdly, because how we remember and

¹ In fact both Menary’s and Sutton’s texts were in circulation much earlier in the community: a version of Sutton’s text being presented in 2001 and circulated in draft from 2005. For an early discussion of related themes see Sutton (2003).

² To use Andy Clark’s label (from Clark 2008) to refer to outlooks that hold that mind cannot extend into the world beyond the brain.

³ Adam & Aizawa’s argument on memory focuses on organic memory having a particular fine-grained functional profile implementing features such as recency and priming effects. Extended memory systems are very unlikely to have similar functional profiles to organic systems at very fine-grained levels but this seems unlikely to be a decisive point; see arguments in Clark (2008) and discussion of the complementarity principle (see Section 3 below)

how we forget may turn out to be of central importance to the nature of human beings as such (Clowes 2012; Mayer-Schönberger 2011).

Despite much discussion of the extended mind thesis, there is much less apparent interest—at least in the philosophical community—in how and whether new regimes of cognitive technologies might already, now and in the future, be playing important roles in the qualitative nature of human minds. Insofar as there has been serious discussion of these questions, it seems to be relatively insulated from the theoretical debate around HEC. It may even be that the ontological discussion about how we might identify the boundaries of mind is actually blocking serious consideration of what minds might come to be as they incorporate or come to widely rely upon new regimes of cognitive technology.⁴ This article seeks to begin the task of remedying this occlusion.

To this end this paper will, in the next section (Section 2), first discuss some of the technological specificities of what I (and others) have called E-Memory technologies. I examine these with an eye to any novel properties which these technologies, either severally or individually, might have for our minds. I settle on four properties which seem to capture some of the most important dimensions: totality, incorporability, autonomy and entanglement. Section 3, building on what we currently know about the early usage of E-Memory technologies, explores some of the patterns of how E-Memory and O-Memory interpenetrate in real patterns of cognitive usage. I focus on some practical implications of the idea that our minds are hybrid in the context of E-Memory technologies in order to examine ways in which we may incorporate, or at least utilize, the new technologies in ways which imply changes to the qualitative character of human cognition. Section 4 pulls back to reflect on whether the original conditions used to set limits on extended mind from the original paper (Clark and Chalmers 1998) continue to do good service in the light of the technological advances of E-Memory. The analysis reveals that some new work may be needed to prevent cognitive bloat; (the tendency for ever more environmental paraphernalia to appear to be parts of our minds). Section 5 then treats a prominent contemporary argument which holds that we are being diminished by our engagement with internet technology and particularly seeks to explore how the debate between HEC and HEMC can situate this discussion in a more nuanced light. Section 6 then reflects on what this discussion has revealed for debate between theorists who posit our minds can be extended by technology and those who hold it is rather better understood as deeply embedded in and effected by its technological surroundings.

2 Four Factors of E-Memory Technologies and Their Cognitive Implications

Exactly which aspects of our ‘native’ cognitive profiles might depend on cognitive technologies remains deeply controversial (see e.g., the discussion of the role of writing in cognition: Goody 1977; Halverson 1992; Olson 1994). Nevertheless the idea that our understanding of our minds cannot proceed without an account of our

⁴ Although the notion of cognitive technology is in itself contentious I hope not to beg any questions by using it here. Cognitive technology for these purposes might be thought of as any technology which we make extensive use of in tasks considered cognitive. I discuss the matter further here: (Clowes 2012)

use of technology is an idea which has been defended convincingly by a variety of authors (Clark 2003; Gregory 1981; Luria and Vygotsky 1992; Mithen 1996). In order to think about what might be happening to human mind and memory as we make use of new ‘memory’ technology we need to theorize the various components of the relationship. One common and problematic way of conceptualizing the relationship is to ask how it is that Electronic Memory *impacts* upon natural human memory. One problem here is that human memory is in important respects non-natural memory anyway. Our biological memory systems are already embedded in a cultural and technological matrix. Indeed from an historical perspective all human memory systems heavily depend on their cultural and technological embeddings; this makes it difficult to identify properties of human memory which are not culturally or technologically inflected. Equally the ways in which we go about developing and deploying any cognitive technology needs to be understood against a cultural background. It might be better to consider how one culture of technologically based memory activity is being augmented, altered or supplanted by another. Still any given technology can be understood as having relatively determinate properties within a cultural matrix and it is those we are interested in.

Several researchers have hit upon some similar theoretical divisions in making sense of the new territory. Kalnikaite and Whittaker (2008) discuss organic memory by contrast with “prosthetic” memory; (they contrast pen-and-paper and digital note taking with unaided organic memories). Kalnikaite et al. (2010) extend this usage in a paper that investigates how organic memory relates to Lifelogging.⁵ Both papers use the term organic memory as a synonym for “unaided memory” although as already mentioned this is problematic in general⁶ in the sense that organic memory is generally aided, but just by a different technological/cultural matrix. Bell and Gemmill (2009) speak about the effects and interrelations between electronic memory and ‘bio-memory’, while Smart (2012) speaks about bio-memory and its interactions with various digital systems, especially internet knowledge. All of these distinctions are roughly equivalent and workable as long as we remember that we are really attempting to grasp how an existing technologically supported memory culture is shifting through its adoption of new technologies.

In a previous first cut at the distinctions discussed here (Clowes 2012), I defined O-Memory (Organic Memory) as “undoubtedly heterogeneous set of systems and processes which underlie the ways in which human beings and their brains retain and organise knowledge during episodes of experience which they can later bring to mind.” (NB—this definition would in the main exclude an important form of memory: Working memory, as it is conceived by cognitive psychology; we’ll return to this in a moment). E-Memory was then defined as “a heterogeneous bunch of devices and systems which fulfil similar functions either by replacement, extension or augmentation”. I will stick with these broad definitions here although as will become clear it is important not to obscure the various subsystems that compose organic

⁵ Lifelogging is the movement among heavy users of digital media to attempt to produce a total digital capture of one’s life with digital recording devices, more or less as life happens, e.g. (Bell & Gemmill, 2009; Kalnikaite et al. 2010; Sellen and Whittaker 2010). I will return to this in a moment in discussion of the E-Memory property of totality.

⁶ If not in the context of the specific experiments cited here.

memory. The elements stored by such E-Memory systems I will refer to as E-Memory Traces or simply *Memory Traces*.

A final prior problem here is that for cognitive psychology and neuroscience, memory is not typically thought of as a single system at all but rather a series of interlocking systems, themselves composed of processing subsystems. A central division is made between short-term systems (dealing with holding information over a few seconds) and long-term memory systems. Working memory (Baddeley 1992) is really part of the transitory microstructure or workspace of cognition, to be understood somewhat by analogy with the RAM (working memory) of a computer. My main concern here will be to look at longer term systems: These are often divided between explicit and implicit systems (Schacter 1987) although I will primarily be concerned with explicit longer term stores. Of central importance here is the division made in explicit memory between, on the one hand, episodic (and autobiographical⁷) systems which store personal memories, and on the other, semantic systems, which store information about meaning and factual knowledge (Tulving 1972). These distinctions are important because the adoption of E-Memory technologies may affect these different sub-systems in differing ways for instance having implications for semantic memory but not episodic memory. To an extent working memory systems have already been widely treated by those interested in embedded/embodied approaches to mind because the use of the immediate environment can be treated as a sort of extended working memory (see, e.g., Clark 2008, ch. 3.1 & ch. 4; Donald 1991; Rowlands 1999).

An attempt was made (in Clowes 2012) to identify the main factors or dimensions of E-Memory technologies with novel cognitive potential. By novel, what is meant here is that they provide capacities which by support, extension or replacement have distinctive quantitative or qualitative properties with respect to current O-Memory systems and their adjuncts in pre-digital technology. Although there are potentially many ways in which these technologies can be considered novel, four dimensions were identified which conjointly, can be argued to offer novel properties with respect to previous regimes of memory technology:

Totality E-Memory promises to record our everyday activities on a scale and with a fidelity and completeness that would have been practicably unimaginable under previous regimes of memory technologies (mem-tech). This dimension was previously called *Capaciousness and Comprehensiveness*. The use of the term totality better reflects the terminology as used in the contemporary discussion (see, e.g., Bell and Gemmell 2009; Sellen and Whittaker 2010).

Practical Cognitive Incorporability E-Memory technologies—leaving aside for the moment whether they should count as actual parts of our minds—possess a power, portability and an readiness to be rapidly customized that renders them apt for cognitive deployment. They also potentially possess a transparency of use that makes

⁷ For the purposes of this article we will treat episodic and autobiographical systems as doing much the same job, whereas, more precisely episodic memory is construed as the capacity to consciously remember episodes in one's life, whereas autobiographical memory is construed as dealing with knowledge of one's life more generally including certain kinds of semantic knowledge (such as one's date of birth or nationality). In fact the exact definition of autobiographical memory is something of a moving target (Hoerl 2007). I will have a little more to say about this distinction in the discussion of *totality* below.

them competitors (or complements) with certain of our internal resources. E-Memory technologies are thus poised for deep and pervasive integration with O-Memory systems.

Autonomy E-Memory repositories increasingly do not merely store data but actively process it. Thanks to tagging, indexing and artificial intelligence-based retrieval systems we can expect E-Memory systems to restructure, reconstitute and re-present memory traces in ways that may complement our native cognitive profile, but function relatively independently of ourselves as individual human agents.

Entanglement E-Memory traces are often best understood not as the memory traces of individuals but of traces of the interaction of individuals or of groups. The form and content of the data that composes many E-Memory stores is inherently relational, tracking interactions between people, or people and organisations, and is often used by multiple users for different purposes.

I previously suggested (Clowes 2012) that by virtue of these novel properties, which in turn offer novel cognitive affordances, the extended/hybrid cognitive systems that we are likely to form or are already forming could have rather different cognitive profiles from those that human beings had under the previous printed and text-based regime of memory technology. I will now reflect on each of these proposed properties in a little more detail at each point relating them to different types of O-Memory system which might be their counterpart.

The idea of totality is present in claims that current trends in *lifelogging* tend toward or aim at ‘total recall’ (Bell and Gemmell, 2009) or ‘total capture’ (Sellen and Whittaker 2010). Lifelogging is in essence the idea that digital technologies can be used to log the totality of, or at least a very wide sample of, the digital traces of someone’s experiences, or life as lived. Gordon Bell, Jim Gemmell and their colleagues have performed what is probably the most extensive attempt at lifelogging so far, with Bell as the guinea-pig for their work which attempts to build a total digital record of his life. They have built systems that can record all of Bell’s conversations (and not just on the telephone). Thanks to the SenseCam technology (Hodges et al. 2006) which Bell wears around his neck and which snaps and stores an image every couple of seconds or whenever a face comes into view, they produce what approaches a complete visual record of Bell’s everyday life. They use other computer systems which store copies of and cross-reference between every document, webpage and even book which Bell reads or writes. These systems are then (largely automatically) tagged and stored in a series of cross-referenced databases which Bell can and does use in a way similar to which many of us currently use Google, but to search for information about his own life. With these technologies Bell can not only search for and access memory traces but can review episodes of his life in novel ways. He can for instance pick a random day from the past year and watch it as a vastly speeded-up stop-motion movie. Such systems portend ever more and heterogeneous ways that we might access the memory traces of our personal past. Together these systems—which they call MyLifeBits (Gemmell et al. 2006)—add up to an E-Memory database that aims to record, store and allow retrieval of memory traces relating to Bell’s everyday life and personal history with a completeness and flexibility that has scarcely been imagined before.

In general lifelogging—insofar as it seeks to capture episodes or simply information about an individual person's life—can be seen as roughly the E-Memory counterpart of our O-Memory episodic and autobiographical memory systems. The tendency is toward an ever-increasing range of modalities and situations in which we are able to record the correlates of sensory experience more or less as they happen. Whereas once we struggled to choose what to attempt to remember, arguably with E-Memory it now requires more effort to forget (Mayer-Schönberger 2011). In this, lifelogging only brings out what is already implicit in much mass usage of E-Memory: “*Record*” has become the default setting. It is worth pointing out that some of the information that one might seek to capture, for instance heart-rate, would not be part of one's typical episodic memories but could be regarded as part of an extended autobiographical memory system. The scope of both autobiographical and episodic E-Memories can thus go well beyond their organic counterparts at least in certain respects.

Yet despite the rhetoric of “total-recall” there are some problems with the general idea. While it is perhaps imaginable that we could produce and store complete records of aspects of our lives with digital technology, how would we ever be able to retrieve the mass of this data? In fact significant empirical work is now starting to be done on how technologies like the SenseCam might help certain types of retrieval (Berry et al. 2007; Hodges et al. 2006). But perhaps the very notion of *total* recall goes against the nature of recall as it happens in humans. Recall seems by its nature partial and selective if only because we have only a finite amount of time to spend recalling events that are past. In addition E-Memory recollection of the type Bell has in mind is arguably different in type to O-Memory recollection which is generally held to be reconstructive (Loftus and Loftus 1980; Loftus and Palmer 1974), rather than any sort of verbatim retrieval. This has led others to talk of “total-capture” (Sellen and Whittaker 2010). Exactly how ‘the totality’ of E-Memory traces captured by lifelogging might really be brought back and indeed to what extent it is desirable or useful to aspire toward total recollection is still an open question (see: Clowes 2012; Mayer-Schönberger 2011). A steady stream of empirical work that investigates how E-Memory may in various ways augment or aid O-Memory (rather than simply replace it) in any case points toward ways in which it maybe best not to think of these systems in competition but as complementary (see Section 3).

E-Memory technologies capture, store and potentially retrieve ever more and varied representations of our lives—in the form of digital memory traces—with a fidelity and comprehensiveness that has little comparison in any previous technological regime. But these technologies do not only relate to our personal recollections of events (the counterpart of Episodic O-Memory). They are increasingly playing a role better understood as the counterpart of *semantic memory*. While lifelogging per se is still a rather minority pursuit, technologies like Google and Wikipedia have already achieved great penetration into the everyday lives of millions of people. For many people sat at desks, or using their mobile devices, it may be simply easier to use Google Search than attempt to recall much factual information using O-Memory. The reasons for this are in part to do with the sheer weight of information available through search engines but also its easy access, with Google searches taking fractions of a second to return results.

Other important factors concern our increasing facility and familiarity with these systems which bring us to our next factor.

The concept of Practical Cognitive Incorporability as originally developed, (in Clowes 2012)⁸ was supposed to capture how extensively a given cognitive technology is practically infolded into our cognitive life. This is not supposed to beg questions about whether such technologies are really part of our minds or not, but rather serve as an ecumenical notion which might be deployed by either a HEMC or a HEC theorist. It is supposed to reflect the way in which E-Memory technologies (and other cognitive technologies) are increasingly embedded in or accessed through everyday objects, such as mobile phones, that we can learn to interact with in an almost thoughtless way.

The idea was at first developed mainly in terms of transparency-in-use; that notion having its roots in Heidegger's (1927) conception of how a hammer is "ready-to-hand". The idea is that a tool can be used in such an accomplished and practiced way that the it becomes is lost to conscious reflection, conveying the subject's attention to the task in hand. Pen and paper, are especially transparent technologies for many of those schooled in the pre-computer age. They are artefacts we just use with scarcely a thought about the medium itself but that we put to use in whatever task we set ourselves to accomplish.

A very transparent technology may approach invisibility (Norman 1999) and this is often regarded by HEC theorists as being a cardinal sign that a technology can be regarded as a part of the agent. Regardless of these claims, transparency-in-use appears to be a central dimension of how cognitively incorporated a technology is, at least in a weak sense. We will just tend to use technologies we find transparent-in-use whenever they are available and the opportunity presents itself. The notion of transparency-in-use can be generalized to any technology and will depend upon a number of factors including our skills at using that technology and the way the technology has become shaped to our use.⁹ There is perhaps nothing especially novel about the transparency of E-Memory technologies in themselves and it would certainly be putting it too strongly to suggest that these technologies currently approach the transparency-in-use of pen and paper for those practiced in it use. This however, is not the whole story.

In a recent article Heersmink (2012) has suggested that the criteria by which tools and technology might count as extensions of mind depend on a number of relatively independent factors. Heersmink's proposed factors are: Reliable access; durability; trust; transparency; individualization and entrenchment; bandwidth; speed of information flow; distribution of computation; and cognitive and artefactual transformation. It remains something of an open question as to which of these factors, or group of factors, would weigh most in implying we should count a given agent/instrument interaction as intra-cognitive; i.e. as taking place within the agent instead of being part of its set of environmental interactions. Further practical work and conceptual analysis will be required to attempt to sort out which of these dimensions are of the greatest importance. For reasons I have already partly indicated transparency and, as I shall go on to discuss, trust, may turn out to be of central importance where E-Memory is concerned.

⁸ In that paper it was just termed incorporability.

⁹ See the discussion of cognitive dovetailing in Clark (2003)

The nomenclature of incorporability suggests that the concept is supposed to pick out the tendency of technologies to become one with the human mind. But as I have emphasized, it can be understood merely as something ripe for transparent cognitive usage rather than actually committing us to views about the technology being a part of our minds. Moreover, incorporability should not be understood as binary but instead as a graded concept allowing us to imagine patterns of incorporability of greater or lesser density. Heersmink's factors could be considered as different dimensions of incorporability¹⁰; namely the potentiality of any artefact to achieve a deep and pervasive, or indeed a shallow and transitory, bond with the human mind. Any technology could be understood as having an incorporation profile. Incorporability can then be understood as a continuum starting with dense interaction and deep incorporation on the one hand, whilst on the other there is the occasionally utilized cognitive instrument¹¹; (see also the discussion in Wilson and Clark 2009).

The incorporation profile of E-Memory appears novel on several grounds. As much internet access coalesces around mobile devices some of the contours, at least for the near term, become apparent. E-Memory devices will be ever-present in our lives, tethered to huge memory stores to which they are connected by ultra-fast mobile connections. We will carry devices that connect us to these memory stores and interact with in various fluid ways. We will stroke and touch them and they will interpret many of our gestures. We will speak to them and they will understand much of what we say; at least well enough to trigger whatever function or app we wish to access. They will track our movements via GPS throughout our daily lives and this also will be stored in the cloud.

Rather than relying on a special representation code (such as the symbols of writing) these devices will store massive amounts of iconic representations (such as digital pictures or voice recordings) as well as traces of the ways we interact with devices and the informational cloud beyond. Much of what we see and hear, along with records of physiological activities will be stored away automatically as we progress through our everyday lives. The control we have over many of these systems will be an open question but there is a sense in which these technologies rather than being merely transparent to us, will make us transparent to them. This of course raises many questions about how we will choose to use this technology but, to return to our main concerns here, one central one is what will happen with all the masses of data that will be collected and stored in this way. This question nicely brings us to our next factor: *The autonomy of E-Memory*.

The autonomy of E-Memory hinges on the way that the memory traces that we will store in massive online repositories will not simply held inertly but will increasingly be processed and repurposed in multiple ways which will often go well beyond any purpose we originally intended for them. E-Memory repositories do not merely

¹⁰ Indeed Heersmink analysis—in the same 2012 paper—of artefact/organism interactions into different levels of information flows appears to be compatible with the approach to different densities of interaction developed here.

¹¹ To be clear terms like cognitive instrument, cognitive technology or even cognitive interaction are not supposed to imply that said instruments, technologies or interactions are necessarily part of anyone's mind. They can be thought of as merely having important cognitive implications in the way that technologies were analysed for example in Hutchins (1995). None of this is supposed to beg the question against the HEMC theorist.

store memory-traces but index, tag, process, re-present and, in a variety of ways, reorganise them. Oral poetry might partially reorganise itself independently of the teller as other singers of songs re-use material, but it is the active processing nature of many E-Memory systems that is really particular. For example, a library may be re-ordered between visits. But this is nothing compared to the growing tendency of E-Memory systems to re-order themselves in a way which is cognitively opaque or even invisible to the user; (notice the sense of cognitive opacity can go hand-in-hand with the sense of transparency-in-use I have already discussed. Transparent tools—especially when they are also cognitive technologies—will often make their workings opaque.) This active or autonomous nature of E-Memory in some ways mirrors aspects of human memory. We do not simply store away inviolable memories but rather new knowledge continually permeates into what we knew before and what we remember now.

There are, however, two related paradoxical movements here that have to be noted. On the one hand, one of reasons we might use E-Memory—as we have seen in the discussion of totality—is that E-Memory repositories can potentially store and index huge amounts of data in ways that are not subject to degradation and remain potentially accessible forever. The SenseCam can take photo-image every three seconds throughout a day, which can later be indexed and retrieved through interfaces like Gordon Bell's MyLifeBits system (Gemmell et al. 2006). Despite this potentiality to collect and retain really inviolate records, the way that current stores seem to be evolving is toward the re-deployment of collected data in labile ways. Those who warehouse much of the data we are currently storing away are continually involved in the task of making it accessible in novel forms. While in some sense the original records may be retained the interfaces through which we access them continue to go through protean change.

The Entanglement of E-Memory systems concerns the ways in which E-Memory traces, and the systems that store them, are not best understood as fully personal but are inherently relational tracking the interactions between people. One aspect of what is new is seen in systems like facebook—one of the most widely used social media systems—where memory traces of an individual's life are stored, processed and presented in their relation to the memory traces of others. Everything that happens on facebook: status updates, likes, comments and shares, is conceived of as an 'edge'. An edge relates a user to the creator of a piece of data via a weighted sum whose main terms are thought to be an affinity score (relation between users of the system), a time decay factor (more recent news is more salient), and a community saliency factor (essentially how often that item or the creator of that item has been looked at or clicked upon). The salience of any memory trace in such a system is inherently relational based on a history of interactions. Moreover ownership of data in such a system is always contestable: if someone 'likes' my status update, is the new edge that is created mine or hers? Facebook uses the *edgerank* algorithm to determine exactly which updates are shown to an individual user and what is shown is essentially dependent on previous interactions (or lack thereof). Such systems, when viewed from the way they used by individuals can a potentially can be seen as the E-Memory counterparts of organic autobiographical memory or episodic memory, yet from the point of view of how they stored, organise and retrieve memory traces, look more like collective memory systems. This trend appears to be quite general across social media, and algorithms like *edgerank* are used to produce views into a dataset

which while particular to a system user, are in fact composed from data about interactions between users. Arguably, systems like facebook are typically used for tracking the ongoing activity of those we are interested in rather than as any sort of memory. Yet, if we look at the way people tag and revisit photographs—which are effectively then stored in perpetuity in such systems—it is clear that facebook and other social media systems are already serving as the contemporary equivalents of photo-albums and in some cases diaries. As the algorithms which lie behind facebook become ever more dynamic and sophisticated they are also likely to become more autonomous and entangled. Really it is unknown the extent to which the use of semi-public media to store and reflect on our intimate private lives may alter the way we think about ourselves and our past; although some are already worried about the effects (Mayer-Schönberger 2011; Turkle 2011).

Both Wikipedia & Google can also be viewed as forms of hyper entangled semantic memory insofar as they are systems to store knowledge which rely upon the aggregation of widely parallel activities of either making links between webpages which are then aggregated (it is this structure which Google uses to seed its ranking algorithms), or people—in a vastly distributed way—making edits to Wikipedia pages (Leadbeater 2008). In fact such systems are more obviously entangled although arguably in a way which is quite similar to traditional libraries. In some respects entanglement is a very old feature of memory systems and technology: e.g. public libraries store knowledge for all. Entanglement looks more novel in regard to episodic and autobiographical memory systems.

Our discussion of these four factors, when taken together, but especially totality and practical incorporability, suggests that E-Memory may have a novel incorporation profile with implications for how human memory will develop in the future. Still the ways in which we may take up and use these systems is largely open and it is difficult to do more than project trends forward. In the next section I shall explore some of the factors that may influence the ongoing incorporation of E-Memory resources into various cognitive activities. We'll also consider what resources the extended mind debate may give us to help us think about this.

3 Hybrid Minds and Their Principles of Assembly and Refinement

The idea of Cognitive Hybridity suggests that a mind's basic organic composition is extended by the deep incorporation of systems whose properties differ in type from related organismic ones. The so-called *Complementarity Principle*¹² (hereafter CP) suggests that we will cognitively factor in ambient resources insofar as they complement, rather than replace, our native cognitive profile (Sutton 2010). Some of the most compelling examples of the hybridity and complementarity of E-Memory can be found among those using the new resources to compensate for O-Memory deficits.

¹² The complementarity principle is first coined in Sutton (2010) as a sort of antidote to the parity principle's (Clark and Chalmers 1998) tendency to too strictly make novel cognitive technologies need to conform to prior organic cognitive systems. The parity principle is discussed in detail in Section 4.

One of the more interesting uses of the SenseCam has been to study whether those suffering from serious memory deficits might be able to improve their ability to recall salient episodes of their lives with technological aid. Some researchers (Berry et al. 2007; Hodges et al. 2006) have begun to investigate how those with severely compromised O-Memories might use devices like the SenseCam in order to recollect events in their everyday life otherwise lost to them. One study explored how Mrs B and her husband use a SenseCam to capture everyday salient events and then download them to a standard PC to review the pictures together at the end of day. The combination of the images and discussion with her husband significantly improves Mrs B's ability to remember events in her life beyond what a paper diary would achieve (Berry et al. 2007). On the (admittedly provisional) evidence of such studies it seems E-Memory can significantly support O-Memory systems.

Deacon Patrick Jones, who suffers from traumatic brain injury (Marcus 2008), uses the EVERNOTE and CURIO software systems on his iPad and iPhone to do some of the work of his deeply damaged organic episodic and working memory. Jones has deeply compromised anterograde, retrograde and working memory but, with extensive use of these technologies, is able to offset some of his memory disabilities. Rather than supporting O-Memory encoding and retrieval—as in the case of Mrs B—Deacon Jones seems to have extensively incorporated E-Memory into cognitive operations that would largely be carried out internally by O-Memory systems in the rest of us. Jones just thinks of these technologies as part of his mind. The question is whether these are exceptional cases or whether those of us without such compromised O-Memory systems will follow the path that Mrs B and Deacon Jones have trail blazed.

In his book *Supersizing the Mind* (Clark 2008) Clark suggests that minds organise problem-solving by what he calls a *Principle of Ecological Assembly*: “According to the PEA, the canny cognizer tends to recruit, on the spot, whatever mix of problem-solving resources will yield an acceptable result with minimum effort” (Clark 2008, p.13.). If this is right, canny cognizers like us should—when engaged in some cognitive task—factor in whatever the current ambient array of potential cognitive tools makes available. Thinking of the PEA in the context of hybridity suggests that, at least synchronically speaking, cognitive agents like us simply co-ordinate whatever group of external props and internal resources are most convenient. But this is a simplification. Agents will have preferred modes of problem-solving and preferred tools which will count in what particular cognitive assembly is deployed at any given moment and this will change over time, especially as the users of a given technology develop greater facility with its use. It remains an empirical question as to how canny we are in Clark's sense. Under what circumstances do we incorporate technologies like E-Memory into problem-solving or other cognitive activities?

There is a limited amount of empirical evidence that illustrates how some of the predictions of the PEA might be manifest in the ways our minds already factor internet technology into the way they store and retrieve information. Sparrow et al. (2011) set out to examine how this ever present source of information might affect the way we store knowledge and make use of our organismic systems. They first explored how much thoughts of searching for information on the internet might be factored into everyday demands on semantic memory. For instance, when asked questions such as “are there countries with only one colour in the flag?”, would our

first thought be to turn to the internet? By using priming studies they found that subjects do indeed tend to think about the internet when presented with moderately difficult trivia questions. In another experiment they were able to show that participants tended to remember certain information less well if they thought that information was stored in a readily accessible computer file. In other experiments participants who thought they would later be able to retrieve information from a computer system, were shown to be more likely to remember where to access information rather than the information itself; (all of which may well generalize to our use of the internet). These results tend toward the hypothesis that, when we have the opportunity to store information externally we will tend to do so and that our organic systems will adapt accordingly.

In another set of experiments (Kalnikaite and Whittaker 2008) examine how subjects perform in remembering details of stories in conditions where they may either rely on organic memory, use standard pencil and paper note-taking, or a hybrid digital note-taking system (ChittyChatty) which both records the stories and allows hand-written annotations. Not only do the experiments show significant positive effects for cuing recall when using the note-taking systems—the ChittyChatty system performing somewhat better than pencil and paper—but that subjects tended to use the equipment strategically, relying on it more in cases where they are uncertain about organic memory. This is important because it suggests that the distribution of resources between O and E-Memory is indeed influenced by meta-knowledge (the authors refer to this as meta-memory) about which of our cognitive resources (E or O) most ably fits a given task. The PEA therefore might best be understood less as an automatic adaptive process of the brain but at least in part organised dependent upon a subject's self-knowledge

Kalnikaite et al., (2010) also explore how subjects use lifelogging technology—in this case a SenseCam and GPS together with some sophisticated data visualisation systems—to attempt to recall details of everyday life some weeks hence. Participants are asked to recollect a day in their life, recorded previously, with the aid of one of several different visualisation tools which presented back the digital memory traces collected using the various lifelogging technologies. The first visualisation tool *snaps* presented sequential images collected by the sensecam. The second *tracks* used locational information to present their paths across a map. A third visualisation tool allowed the different sets of information to be fused. One of the interesting findings is that users who just have access to the viewer of SenseCam images (*snaps*) tend to say they have rich visual and contextual memories—suggesting episodic memories—while users of the GPS viewer (*tracks*) tend to say they ‘know’ or ‘guess’ (can reconstruct) the events. As the authors discuss, this appears to demonstrate that different styles of lifelogging, i.e. recording large sets of images vs tracking location information imply the triggering or tokening of differing memory types, or at least different types of self-report.

One important point here is that users of the experimental systems tend to think of the usage of E-Memory technology as providing a cue to remembrance rather than actually counting as first-class remembrances in themselves. Such parlance seems to support a general (folk?) commitment to a HEMC rather than a HEC interpretation of E-Memory usage. By contrast Deacon Jones who employs a much more deeply integrated E-Memory system in his daily life (and outside of experimental situation)

speaks about his E-Memory systems as actually being parts of mind and their use as acts of remembrance. This implies that further empirical investigation is needed to get at exactly which aspects of E-Memory technology and usage might imply very deep integration. However we can predict that such factors as constancy and prolongation of usage, user control of apparatus, types of memory trace stored, along with means of accessing those traces all seem to be matters of importance in both a subject's quality of recollection, and for understanding of whether such 'recollection' might count as real remembrance as opposed to some type of inference.

Sparrow, Liu et al., cast their work into the framework of *transactive memory* (Wegner 1987), that is to say, we tend to remember what we need to based on expectations that other trusted sources will remember for us; especially when we expect them to have special knowledge and competence in a domain. The transactive memory framework was developed to describe how people tend to rely on others who they trust to remember certain things on their behalf. The use of term transactive memory in relation to E-Memory does imply, if taken straightforwardly, that we think of E-Memory stores as though they were other (trusted) agents, rather than as parts of ourselves. An important implication here is that as we find E-Memory systems ever more available, trustworthy and easily accessible that they will tend to be factored ever more heavily into ongoing episodes of cognition; at least in the absence of other tendencies. One possible further implication is that with the very incorporable E-Memory resources that mobile internet applications are making available we shall (organically) remember less of the information we think we can readily access through our gadgetry. Instead we will tend to remember how to access it in ways that line up with the predictions of Clark's PEA.

The PEA appears to offer different predictions from the CP and both might have an element of idealization about them. According to the PEA, brains coordinate whichever range of internal resources and ambient equipment happen to be most useful in achieving the task goal. But, the notion of usefulness here might partially obscure the way in which agents will tend to lean upon and use a favoured set of devices and instruments. The CP, on the other hand, implies that equipment is only incorporated if it offers something that complements existing resources, yet this might depend upon rather fine-grained matters such as how fast or convenient it is to access a given online tool (i.e. precise matters of the incorporation profile of a technology). This may undergo rapid change as an agent becomes increasingly skilled in the use of, or more reliant upon, particular external props; or the technology itself is refined; or an agent develops new expectations about the competences of particular technologies with respect to her organic systems. What really gets incorporated (either in a deep or shallow way) is likely to broadly depend on matters of habit and skilled usage as much as the raw properties of a cognitive technologies.¹³

Let us extend this question into more speculative territory. Consider an E-Memory system whose interface works like Google Search but accesses tagged and indexed E-Memory traces that have been continuously recorded by devices like the SenseCam (Berry et al. 2007; Hodges et al. 2006). Such a system could be accessed through a smart phone and its data-stores accessed through the cloud; let us call it E-Recall.¹⁴

¹³ Related inferences are drawn by Smart (2012)

¹⁴ The system is modelled on a next generation version of MyLifeBits as presented in Bell & Gemmill (2009); for further discussion see Clowes (2012).

Users of E-Recall will progressively develop a growing facility with the technology and we can also expect its user interface to evolve over time, with the general tendency that it will be easier and more convenient to use. If extrapolation from the PEA is correct, it is likely that users of E-Recall will progressively factor the technology into cognitive episodes. If the tendency is general we may expect reliance on E-Recall to start to replace some aspects of O-Memory systems that previously would have played similar roles.

E-Memory systems (and related cognitive technologies) will likely continue to become more potentially incorporable as they become more user-friendly and, importantly, as we become more skilled in using them. Over time the balance of our assemblage of tools will change essentially with our culture of usage. If the PEA is correct, there is a strong possibility that we will increasingly favour our extended toolkit over at least some organic resources. We can easily imagine that as we become habituated to the use of E-Memory technologies we will increasingly factor them into cognitive routines and become more likely to deploy them in future. For, as we become more skilled users of search technologies, and they become ever more transparent to us, we will tend to rely on them more. Complementarity on the other hand suggests that we will factor in new technologies insofar as they offer us novel capabilities. These principals although having different points of emphasis are not in deep contradiction. Insofar as the PEA predicts we use whatever are the most useful cognitive resources, whether organic or environmental, available and the CP predicts we will use resources which have complementary properties, we can use them both to predict greater reliance on E-Memory. From what we have already said about the totality of E-Memory it may be controversial as to whether it is really complementary but not that it does not offer new properties. We can thus see how we might start to distribute an ever increasing amount of this “knowledge” in E-Memory resources.

However this may not help us with the ontological argument over whether these systems should ever properly count as part of our minds or merely their environmental embedding. The problem here is that it is difficult to see how mere empirical evidence about deep integration (in the sense we have developed) is going to convince a theorist that chooses to dig in her heels. She may argue that only parts of our organic brain are truly mental because only they have underived intentionality (Adams and Aizawa 2001). If so, she is unlikely to be convinced by ever more detailed demonstrations of how deeply an agent comes to rely on “cognitive technologies”. Similarly theorists who are committed to the idea that it is the fine-grained cognitive profile of cognitive apparatus that matters (Rupert 2004) will be similarly unimpressed by our deep reliance on cognitive technologies which have very different functional profiles. Nevertheless theorists from all sides seem to owe us some account of what the implications for us are of our potential increasing reliance on these technologies. We shall attempt to tease out some of the implications in the next section.

Yet here it becomes clear why the face-off between HEMC and HEC may be in danger of generating more heat than light as we seek to understand the cognitive and ethical implications of E-Memory technologies. The terms of the debate seem to press us toward primarily recognising the uses of E-Memory technologies in ways which are directly comparable to the way that O-Memory functions; and especially its

functions under text based regimes of E-Memory technology. Such comparisons can be useful but they can also obscure the novel properties of hybrid systems that we could be investigating. A better understanding of what is really new here may require us to focus our attention on the use of E-Memory technologies in their own terms and as new patterns of usage and cognition arise.¹⁵ Some of the experimental work we have focused on in this section hopefully gives an indication of how an examination of the use of E-Memory technologies as a practical/experimental inquiry may go forward. Nevertheless, given that E-Memory technologies, and the uses we are putting them to, already seem to have some implications for the more ontological debates about whether such technologies should ever count as proper parts of our mind we shall return to this question more explicitly in the next section.

4 The Limits to Incorporability: Trust and the Autonomy of E-Memory Systems

Clark and Chalmers' original (1998) Extended Mind paper advanced the *parity principle* which held that if a process taking place in the world should count as cognitive were it in the head, then it should count as a cognitive process. Moreover, if that process had the right sort of interaction with the agent, it should also count as part of the agent's mind. It also advanced four criteria (or conditions) on the extensibility of the mind that were designed to see-off the dangers of cognitive bloat—the tendency to see an ever-increasing collection of props as proper parts of the mind—and help identify cases where technologies, tools and other environmental props might usefully be considered mind extenders.¹⁶ The Conditions were:

(Constancy) - A cognitive technology should be considered a constant in an agent's life; when information it contains would be relevant, the agent rarely takes action without consulting it.

(Facility) - Information the technology contains (or makes available) is directly available without difficulty.

(Trust) - Upon retrieving information from the device the agent automatically endorses it.

(Prior Endorsement) - The information a cognitive technology presents has been consciously endorsed at some point in the past and is there as a consequence of this endorsement.

In the same paper Clark and Chalmers also remarked that the Internet (circa 1998) seemed unlikely to meet the conditions and thus count as part of our extended minds: “The Internet is likely to fail on multiple counts, unless I am unusually computer-

¹⁵ This discussion in part re-iterates points made by some of those theorizing a “second wave” approach to the extended mind (Menary 2010; Sutton 2010) where—as we have seen—the emphasis is placed on understanding the dynamics of potential new cognitive systems and what is distinctive about them. As Sutton (2010, p. 41) wrote “in extended cognitive systems, external states and processes need not mimic or replicate the formats, dynamics, or functions of inner states and processes”. Insofar as the ontological discussion tends to obscure this sort of investigation then it will tend to block understanding of the new “kinds of minds” we may be developing.

¹⁶ NB—I have slightly amended these so they can be applied to technologies more generally rather than specifically referring to Otto's notebook.

reliant, facile with the technology, and trusting.” (Clark and Chalmers 1998, p. 17). Since 1998 many of us have however become far more computer reliant, trusting and accomplished users of the technology. Writing much more recently (Smart et al. 2009) have argued that the information contained in the typical web-page is not well poised to form a part of our cognitive economies. This is because web-pages contain irrelevant and difficult to manipulate information which as Smart et al. (2009) point out are often not readily usable to “guide daily sequences of thought and action in a manner that is functionally indistinct from the way in which internally generated information does.” While this is astute observation in regard to how webpages are accessed through standard computer terminals, the variety of ways that mobile devices allow us to access timely information through apps may already be significantly changing the cognitive poise of information accessible through gadgetry. It is unlikely the webpage will remain the privileged unit of information on the internet or the web-browser the main mechanism for access for very long. Mobile apps like Google Maps already seem well poised for relatively effortless cognitive usage.

So while the standard web (accessed through a computer terminal) might not easily meet Clark and Chalmers’ original conditions, today’s Internet, increasingly embodied in a host of mobile devices and their associated applications, appears to meet the criteria of constancy and facility at least, rather more readily. In part this is because mobile apps tend to present pre-packaged chunks of information fitted to certain tasks and these appear more ready to meet the demands of everyday cognitive episodes. Indeed if we are interested in finding devices and systems that are poised to bond with us—as I have begun to argue—we might be better looking at these. David Chalmers has recently (2007) written that his iPhone meets, at least superficially, the criteria set out for mind extenders. He notes he carries his iPhone with him everywhere, relies on it and automatically endorses at least some of its contents and, as a result, his brain may cease to contain some of the information it once might have. Thus he writes: “My iPhone is not my tool, or at least it is not wholly my tool. Parts of it have become parts of me.” His reasons closely track the argument from the parity principle.

But is Chalmers correct in thinking that the properties of devices like his iPhone really press us to consider them as proper parts of our minds, or should we consider them still just smart tools? To attempt to answer this question let us return to the four factors of E-Memory with which I opened this discussion: Totality; incorporability; autonomy and entanglement. How do they weigh in the discussion of whether these technologies should now be considered as deeply integrated parts of our minds?

E-Memory systems which implement the property of totality, tend to offer novel and in certain senses complementary resources to O-Memory as the previous section explored. Similar things can be said about practical incorporability. Because of the heavy emphasis on human computer interaction in the design of artefacts like the iPad, along with the way we quickly develop interaction skills with such artefacts, and the vast range of E-Memory applications that they are helping to proliferate, mobile devices tend to become ever more incorporable. In addition, when we consider the capaciousness of the cloud infrastructure that stands behind them, the possibilities of meeting the Clark and Chalmers’ first two criteria look very high.

Thus, the mobile internet makes condition 1 (constancy) readily beatable by almost any of the current generation of smart phones and many other devices as, battery or network failures aside, smart phones are a constant in our lives. Condition 2 (Facility) closely relates to my own second factor; incorporability. If what has been said so far is correct we can expect that E-Memory and other cognitive devices are undergoing a constant process of becoming more transparent and facile in usage. E-Memory devices will increasingly convincingly meet that condition.¹⁷ Thus the Totality and Incorporability Profiles of E-Memory point toward these technologies rapidly being deeply incorporated in many of our cognitive episodes.

However, the last two conditions of trust and previous endorsement seem actually to be challenged by aspects of E-Memory and the new cognitive technology in general. This is in part because the autonomy of E-Memory, along with its entanglement, can look from a certain viewpoint as being a form of memory tampering. Many recent works of popular fiction have dealt with the theme of memory tampering and indeed it comes up in the original extended mind article. In fiction such as *Before I go to Sleep* (Watson 2011) and, of course, the film *Memento* (Nolan 2000), the recurrent theme is that of an amnesiac (generally with some anterograde and retrograde memory problem) struggling with attempts to use some extended paraphernalia to retain memory and hold together their sense of self and identity. The playing out of the plot in these works always turns on how far the subject can trust that the external trace has not been tampered. This suggests that the possibility of memory tampering may act as a fundamental constraint on memory extension. When one cannot be sure that a purported extended memory system is free from tampering, one should not be ready to count it as part of one's own memory.

The problem with E-Memory technologies is related to their autonomy, and especially the way in which the computational systems which underlie them are subject to change which can be invisible to users. Should this rule out such systems from being considered as deeply incorporated? An autonomous E-Memory technology might subtly change the memory profile of an agent without their even noticing. Similar things can be said about entangled memory systems where changes take place invisibly due to another agent's activity. Even if unintended this looks very like memory tampering when considered from the perspective of an individual agent.

So even if we leave aside deliberate memory tampering, factors of autonomy and entanglement may operate as a fundamental constraint on what *should* really count as part of our minds. Very entangled and autonomous cognitive technology is likely to be technology that we cannot entirely trust to be operating in our interests. Moreover it would not make much sense to say we have endorsed the content of such systems, especially if that content was subject to change in terms of the algorithms that present it. The best we could manage is some sort of endorsement of the *process* or perhaps the company that was offering the technology; like the way some may say today: "I trust Apple, but not Google".

¹⁷ One factor that may count against this is that software companies may sometimes upgrade their software in ways that disturb a user's pattern of smooth use, at least whilst one is adjusting to a new interface (Thanks to Ron Chrisley in personal communication for pointing this out). Personal experience of upgrades in Microsoft Office mean I tend to put off using upgraded software when working on important projects. Nevertheless the general tendency does seem toward more user-friendly gadgetry.

Much might come down to the issue of trust; especially if we can really trust those who manage our purportedly extended memory systems, or those companies or persons with whom our E-Memory systems are entangled. Even if in every other respect an E-Memory system tends toward deep incorporability, we might continue to distrust it by virtue of its tendency to change its underlying processing architecture and hence treat it as a cognitive adjunct rather than part of our minds.

Of course it is always possible that despite the fact that we *should* distrust a particular cognitive system we will nevertheless actually trust it. This brings us to what we might call the paradox of credulity. The more credulous a given subject is, the more likely he is to trust systems he should not, and potentially treat those systems as part of his mind. This seems to imply that the more credulous a subject is about the sources of E-Memory, the more his mind seems to expand, incorporating systems that might well be downright working against his interests. A more suspicious and sceptical agent would tend to treat sources that were potentially open to tampering, or just run according to the interests of others, as unreliable and would take the appropriate stance toward them: They are just potentially unreliable sources of information. The inappropriately credulous agent, however, might nevertheless trust them.

This might be taken as a *reductio ad absurdum* of the trust constraint. It doesn't seem to make sense that a more credulous agent is more cognitively capacious than the cognitively suspicious one, although one has to add that such an agent's set of beliefs often actually compasses a good deal of unreliable information. Perhaps the epistemic manner in which an agent treats his tools should not here count as a decisive issue and instead we take into account not only that an agent acts as though he trusts his extended cognitive equipment, but also that the information contained therein is actually trustworthy. (It is important to remember here that in many attempts to test Wikipedia, for instance, it performs well in matters of reliability (Leadbeater 2008)).

Roughly speaking, agents need to trust that the components of their minds are doing their proper job. Consider, for instance, a neural prosthetic that made timely suggestions about where the cheapest place for lunch was. Although the information from the device was generally reliable, it would also—about one time in 10—include some ad-supported prompts that were not flagged as such. One might rely on such a system generally to find somewhere to have lunch, after all nine times out of ten it would be correct, but one would be unlikely to treat such a system as really a proper part of oneself. Why? We generally assume that parts of our mind are operating in our own interest. If we have reason to believe they are not we are unlikely to treat them as parts of *our* minds. This implies that human minds that are potentially distributed have—if they are to be considered a systematic whole—to develop intra-systemic trust relationships.

A problem here is that we do not always expect all of our own cognitive processes to operate in our own best interests or at least, even if they are self-destructive, self-deceptive or in other ways appear to operate against our own best interests, we do not generally disown them¹⁸; although we may disavow them. Indeed disowning thoughts produced by our minds or ascribing their agency to others can be a cardinal

¹⁸ In the technical sense of not feeling ourselves to be the owner of those thoughts (Campbell 2002).

sign of mental illness such as schizophrenia. However this may be because we operate on a general pre-theoretical assumption that a thought produced by our own mental economy is still ours even if it is dissonant with other thoughts or self-destructive in some respect. It is unlikely we will adopt so charitable an outlook where other agencies would *ex hypothesi* partially determine the content of thoughts; especially if the systems that produce it are very autonomous or entangled. Thus, trust—broadly construed—may be a fundamental factor limiting the extension of mind into gadgetry and cloud based-systems. Even if a system that might potentially be deeply incorporated were to meet all other criteria, the subject's lack of confidence in the reliability of the potentially incorporated system would likely stall its deep incorporation (or should do). Even very high-bandwidth, transparent and otherwise deeply integrated systems may under some circumstances not count as part of an agent's mind if they are untrustworthy or simply too subject to the control of other agencies.

E-Memory systems seem to be tending toward a generally dense incorporation profile affording characteristics that are broadly complementary to our O-Memory systems. For instance those properties of E-Memory systems we discussed under the banner of totality seem precisely to be different in kind to what O-Memory systems afford. Yet there is good reason to think those E-Memory systems will become factored into an ever more expansive range of cognitive operations.

By contrast, when we consider autonomy and entanglement, things start to pull the opposite way. Some aspects at least of very entangled and autonomous E-Memory systems are likely to make the prior endorsement condition difficult to fulfil in a straightforward sense. Likewise the sorts of deep entanglement that many E-Memory systems imply mean that it will be difficult for sensible agents to trust those systems are serving their own best interests. Very autonomous and entangled E-Memory systems are thus unlikely to be treated as part of the agent's mind by the agent (or at least should not) because they are too open to what is effectively memory tampering. This implies important constraints on the concrete incorporability of E-Memory systems: Trust relations appear to operate as a fundamental limit on real (concrete) incorporability. Still, given the very labile nature of these systems and the tendencies that drives E-Memory toward deep incorporation we should keep an open mind on these issues.

5 Extending or Diminishing the Subject?

A prominent and fashionable critique of the cognitive implications of the internet holds that it is a primarily destructive force, disorientating our minds and diminishing us. According to Carr (2008, 2010) it is making us distracted, lazy-brained, satisfied with whatever 'knowledge' is spoon-fed to us and undermining our ability to think critically. The auto-completions of Google Search distract us from our original thoughts; tabbed browsing encourages us to open and read the first few lines of a dozen articles we will never return to; and hypertext encourages us to click around a bunch of links, never finishing a clear path of research and even finding it difficult to maintain a linear train of thought. Moreover sites like Wikipedia, Carr claims, are encouraging us to become intellectually lazy, uncritically accepting the first things we find through a Google Search rather than checking sources and critically questioning

them. Others argue that the Internet limits our ability to think creatively (Lanier 2010) or is undermining our sense of self (Greenfield 2008; Mayer-Schönberger 2011; Turkle 2011).

It is worth arguing with this vision of the Internet, and especially search technology, as *essentially* distracting and disorienting. In part this is because the empirical basis for these claims is quite slight; but it is also dangerous to generalize from the precise shape of the internet today to where we will be tomorrow (see also: Smart 2012). It is possible Carr has identified historically specific and non-necessary aspects of networked media to do with the precise way many or most of the widely used Internet services are embedded in the market system.¹⁹ What I am more concerned with here is whether what I have previously claimed are rather more *essential* factors or tendencies in the way that we use mobile internet technologies are likely to have more general implications for human cognition. One could, for instance, argue that search is already deeply embedded in our mental lives and not obviously to our detriment.

On a traditionalist (vehicle internalist) account of the mental—and assuming Carr is right about the Internet's effect on our brains—things look quite bad for the heavy internet user and potential users of E-Recall. Why? As we interact with and form an ever-deeper reliance on networked digital media, a fundamental redistribution may occur between the knowledge that we expect our minds to hold in our brains, and that which is now dealt with by environmental paraphernalia such as search and other technologies. From the internalist (and embedded) vantage-point it is as if our minds are being steadily off-loaded and dissipated into our tools, potentially with the result that we become less autonomous, less knowledgeable and arguably less interesting sorts of creatures in the process.²⁰ But it is possible to contest all of this.

Consider the case of the satnav (satellite navigation) systems here. As I (2012) have previously noted, it is unlikely that our organic systems will be crowded out in any absolute sense as we accommodate—in shallow or deep fashion—the new environment of cognitive tools. Simply having an ambient environment of a certain class of cognitive tools will not stop our organic resources working. So someone might start to use a satnav device to navigate an unfamiliar city, where the satnav acts as a sort of virtual memory. The satnav user may not have the usual knowledge of the city in other respects which might come if the knowledge were developed through driving and painstakingly learning the way using directions or a traditional map. However, using a satnav system does not imply that one straightforwardly stops encoding any new memories or learning the route one uses. Memory is not directly

¹⁹ For instance, Carr claims that internet reading is necessarily disrupting and diffuse, but it is difficult to see why the use of a specialised reading device like the Kindle—already in mass usage—may not be invented that could switch off hyperlinks when needed (an easy way to facilitate internet reading?). In principle web content need not be much more distracting and dissipating than reading a book. It is possible to imagine the construction of web-reading software or hardware which might minimize the tendency of this technology to distract us.

²⁰ For Carr the mass of internet users (and particularly the so-called digital natives who have never known any other kind of intellectual culture), it is simply as though our minds are bleeding away in the machines leaving us as almost sub-human Eloi. In fact the science fiction writer Dan Simmons (Simmons 2010) develops just such a Wellsian scenario in his book *Ilium*. In the book the semantic memory of human beings has atrophied to the point that they know almost nothing. They rely almost entirely on a future internet to browse whatever facile knowledge they need.

traded against search in this way and we can expect organic memory systems to continue to work as complements to E-Memory systems even if they appear quite redundant. It is likely our internal systems will adapt to what environmental cuing makes available yet they will also continue to do this in a richly-redundant way as we spontaneously learn from our experiential flow.

But let us recast this argument in the context of the face-off between HEC and HEMC. For HEMC (or Brainbound) theorists, it follows that if our minds are being bled out into a series of environmental resources we are diminished in the process; (or so it seems at first pass). For HEC theorists this does not follow, for rather than being dissipated into our tools, we are incorporating them into us.

Consider again David Chalmers and his deeply integrated iPhone. It is far from obvious that even if—as Chalmers contends—part of the function of his brain has been taken over by the iPhone that we need to see this as a necessary diminishment of Chalmers or even of his brain. In fact the total agent Chalmers + (including deeply incorporated equipment) may have enhanced cognitive abilities. Previous functions are now handled by the extended ecology of his hybrid mind and in the process it is at least possible that some of the functions of his brain will have been freed-up to play new roles. If the HEC theorist is right, rather than diminishing him, the total cognitive agent Chalmers + may be enhanced and expanded; even if it can be shown that there is movement of some of his knowledge out into his good devices.

The HEC theorist would thus argue that the diminishment argument rests on a misconception about a proper extent and boundaries of a cognitive agent. When using the internet one has to look at how functions which were previously carried out by the agent's brain are now incorporated by virtue of their incorporation into a range of ambient devices. Brainbound theorists get the limits of the agent wrong and in so doing can miss interesting aspects of the total agent.²¹

If HEC theorists are correct then the diminishment argument appears to lose much of its force, for even if we accept that an agent's brain in some sense holds less knowledge or is even less autonomous as a result of the *incorporation* of internet technology, this cannot be simply judged as good or bad for the agent. The Extended theorist is interested in the wider ensemble of human + that is, the agent plus the motley collection of equipment which has been densely incorporated (and some of that which is only shallowly incorporated too). What is going on in the brain is not the only factor of importance; we should care about the profile of the agent as a whole. The relevant contrast cases for trying to understand what sorts of humans we want to be are not the *brains* of subjects who have interacted deeply with the mobile internet and those who have not, but humans beings using pre-Internet and post-Internet technologies. This does not of course mean that we can ignore the brain, or that it lacks importance, but that in order to understand it we need to look at the whole agent—wherever its boundaries lie—and the dynamics of his/her interaction with the wider array of embeddings. Only in the round can we properly understand and weigh changes that are happening in the brain.

²¹ One argument in favour of the HEC outlook is that it allows us to recognise more interesting systems and regularities in the world than can be recognised by HEMC theories. Theorists who really insist on confining their attentions to the cognitive operations of the brain may find themselves restricted to the analysis of increasingly partial cognitive systems as we practically incorporate E-Memory technologies.

It is also worth drawing out what a sophisticated Embedding theorist (such as Rob Rupert) might say here. Rupert (2009), while rejecting the full-blown Extended framework, still endorses the idea that the *embedding* of the mind could still have profound effects on cognitive economy. So, one could even agree with Carr that certain functions of mind were being taken over by E-Memory systems without necessarily holding that this causes diminishment. It may be that the particular functions which are being taken over are not particularly interesting. For instance pocket calculators take over some of the work needed to do mathematics but this need not be understood as diminishment. (But this introduces a dilemma for the HEMC theorist. He might have to say that cognition is diminishing overall in the system but the environment is doing more work. But more work of what sort, cognitive work?)

HEC theorists offer reasons to think that, as the principal concern is with the cognitive potential and actuality of the agent, there is no obvious need to fear that those who densely incorporate E-Memory systems will be diminished by the process. Even subscribers to HEMC theories need not see our use of these media as diminishing, but if HEMC builds in the presumption that it is the real non-extended mind which is the locus of our concern and attention, then it does seem this viewpoint may have limited scope to capture the nature of the changes that are taking place.

Another of Carr's arguments about the supposedly numbing effect of technology derive from Marshall McLuhan (2001 [1964]). According to McLuhan, tools end up "numbing" whatever part of the body they "amplify". But this discussion of numbing now appears as far too simplistic. Deacon Jones is using E-Memory to attempt to organise otherwise scattered thoughts and integrate himself more coherently over time. Miss B similarly is able to review episodes of her life which would otherwise be lost to her. Both Deacon Jones and Mrs B moreover seem to be able to maintain connections with love-ones, and sustain a fuller life, in virtue of the use of their E-Memory systems. Neither seemed to be numbed but rather have a fuller use the faculties. Who is to say that for those of us without severe memory deficits something similar is possible?

To summarize, there is no real reason to think that our purported extension by technology should lead to a necessary diminishment or numbing, nor that the internet in its various guises need act as an engine of distraction, nor that our agency need necessarily be undermined by its use.

What we can see here is that the problem of cognitive diminishment is not at all straightforward and requires some precise workings out of the practical implications of the theories. Working out some of these implications in the context of technologies like E-Memory may turn out to be a particularly fruitful exercise in isolating the most important aspects of the debate between HEMC and HEC. It starts to become apparent that these positions—perhaps unexpectedly—present ethical implications (in the wide sense) for how we valorise cognitive changes that may take place and have implications for how we consider human cognitive futures. This brings us to one reason to prefer HEC over HEMC. HEMC seems to offer us a rather limited viewpoint on the incorporation of E-Memory which points toward (if not necessitating) cognitive diminishment; HEMC seeming to imply the deep incorporation of at least E-Memory technologies involves knowledge being bled out of agents. HEC theorists on the other hand have a more nuanced set of positions they can use to qualitatively assess, e.g. what boundaries make an agent look most coherent? As the

HEC approach could make available an expanded set of theoretical tools to examine the new unities which might arise around E-Memory it appears to be a more useful lens. What needs to happen though is that we must seriously address ourselves to this more qualitative and ethically inflected discussion and the ontologically lead discussion between HEMC and HEC does not as yet seem to have offered much of a lead on these points.

6 Present and Future Unities of Mind

So where does this leave us? One great advantage of the HEC approach, especially when used in the light of the complementarity principle, is that it allows us to spot potentially new agential systems without overly privileging the sorts of cognitive agents that have gone before. This might turn out to be important as, for as our analysis has shown, there are reasons to think that agents operating in a world with pervasive E-Memory may rapidly start to look different and need to be investigated in ways which are at least open to the possibilities of the establishment of new agential boundaries. Certainly spotting and explaining new and interesting regularities here might be problematic if we are to rule out altogether the extended perspective.

One difficulty for the HEMC theorist is that a prior commitment to what is important about the agent implies that the significant regularities and by extension what should be valued is limited to organismic systems. Everything else is part of the environment and of lesser importance to the cognitive agent as such. There certainly seem to be cases, such as Otto from the original thought experiment, or indeed Deacon Jones where the extended organismic agent + set of cognitive technologies seem to add up to the more coherent agent than when understood without them. Ruling out Deacon Jones' extended systems because they are composed of the wrong material, or have an unusual causal profiles, or do not have original intentionality all seem attempts to pre-judge an important question. Does Jones make more sense as a coherent agent when we include the E-Memory paraphernalia he claims to be part of his mind, or without it?

In similar ways, in attempts to make sense of cognitive diminishment the HEC theorist seems to have more options than HEMC theorists; and this looks important. There are interesting ethical questions over what we really care about with mind and our future use of technology. The HEMC theorist seems committed to something like the Carr picture where the spread of knowledge and resources beyond the organismic boundary—other things being equal—look like diminishment. Rather like the case of Deacon Jones this seems to prejudge too much the interesting regularities and properties we might discover. Forms of unity of extended cognitive systems may just be more interesting and weigh more heavily than organismic integrity.

So what then for the claims that E-Memory technology should count as proper parts of our minds? Of the four factors or tendencies in E-Memory that we have discussed, two: totality and practical cognitive incorporability, seem to pull toward counting for deep incorporation whilst the other two: entanglement and autonomy, for a more shallow incorporation. (These latter two may even suggest—as we have

explored—some new problems for the HEC theorist as such). This leaves the question of whether E-Memory should now and in the near future count as actual proper parts of our minds rather than a significant environmental tool which changes some of the processing profile of our brain, a rather finely balanced question and one I think it would be premature to attempt to rule on here. Rapid technological change will almost certainly mean those committed to drawing the boundaries of mind and organism with the same line may find their job ever more problematic.

One approach is to posit that minds themselves may fall along a scale between the densely incorporated (traditional human minds) and the more sparsely integrated (human minds heavily reliant on E-Memory and other novel cognitive technologies). In these cases it seems that we should generally want to keep an open mind²² on what constitutes the most sensible bounds of an agent to meet a number of explanatory and interpretative tasks. In this light the HEC theorist at least has the advantage of being able to compare the different sorts of unity which might arise and consider the arguments for whether any of them might comprise the sort of unity we should really care about.

Acknowledgments Robert Clowes would like to express his thanks to three anonymous reviewers and especially the editors John Sutton and Kourken Michaelian whose insights over several drafts of this paper much contributed to the paper's scope and hopefully interest; their thoughtful comments have also given its author food for thought for future work. Of course its remaining failings are my responsibility alone. The author would also like to gratefully acknowledge a Portuguese Fundação para a Ciência e Tecnologia grant/BPD/70440/2010 that supported writing this paper.

References

- Adams, F., and K. Aizawa. 2001. The bounds of cognition. *Philosophical Psychology* 14(1): 43–64.
- Baddeley, A. 1992. Working memory: The interface between memory and cognition. *Journal of Cognitive Neuroscience* 4(3): 281–288.
- Bell, C., and J. Gemmell. 2009. *Total recall: How the E-memory revolution will change everything*. Dutton.
- Berry, E., N. Kapur, L. Williams, S. Hodges, P. Watson, G. Smyth, et al. 2007. The use of a wearable camera, SenseCam, as a pictorial diary to improve autobiographical memory in a patient with limbic encephalitis: A preliminary report. *Neuropsychological Rehabilitation* 17(4–5): 582–601.
- Campbell, J. 2002. The ownership of thoughts. *Philosophy, Psychiatry, & Psychology* 9(1): 35–39.
- Carr, N. 2008. Is Google making us stupid? *Yearbook of the National Society for the Study of Education* 107(2): 89–94.
- Carr, N. 2010. *The shallows: How the internet is changing the way we think, read and remember*. London: Atlantic Books.
- Chalmers, D. 2007. Forward to supersizing the mind. In *Supersizing the mind: Embodiment, action and cognitive extension*. Oxford: Oxford University Press.
- 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*. Oxford: Oxford University Press.
- Clark, A., and D. Chalmers. 1998. The extended mind. *Analysis* 58: 10–23.
- Clowes, R. W. 2012. Hybrid memory, cognitive technology and self. In ed. Y. Erdin and M. Bishop. Proceedings of AISB/IACAP World Congress.
- Donald, M. 1991. *Origins of the modern mind*. Cambridge: Harvard University Press.
- Donald, M. 1993. Precipice of the origins of the modern mind: Three stages in the evolution of culture and cognition. *The Behavioral and Brain Sciences* 16: 737–791.

²² No pun intended.

- Gemmill, J., G. Bell, and R. Lueder. 2006. MyLifeBits: A personal database for everything. *Communications of the ACM* 49(1): 88–95.
- Goody, J. 1977. *The domestication of the savage mind*. Cambridge: Cambridge University Press.
- Greenfield, S. 2008. *ID: The quest for identity in the 21st century*. London: Sceptre.
- Gregory, R.L. 1981. *Mind in science: A history of explanations in psychology*. Cambridge: Cambridge University Press.
- Halverson, J. 1992. Goody and the implosion of the literacy thesis. *Man* 27(2): 301–317.
- Heersmink, R. 2012. Mind and artifact: A multidimensional matrix for exploring cognition-artifact relations. In *Proceedings of AISB/IACAP World Congress 2012*.
- Heidegger, M. 1927. *Being and time*. Oxford: Basil, Blackwell.
- Hodges, S., Williams, L., Berry, E., Izadi, S., Srinivasan, J., Butler, A., et al. 2006. SenseCam: A retrospective memory aid. *UbiComp 2006: Ubiquitous computing*, 177–193.
- Hoerl, C. 2007. Episodic memory, autobiographical memory, narrative: On three key notions in current approaches to memory development. *Philosophical Psychology* 20(5): 621–640.
- Hutchins, E. 1995. *Cognition in the wild*. Cambridge: MIT Press.
- Kalnikaite, V., and S. Whittaker 2008. *Cueing digital memory: how and why do digital notes help us remember?* Paper presented at the Proceedings of the 22nd British HCI Group Annual Conference on People and Computers: Culture, Creativity, Interaction-Volume 1.
- Kalnikaite, V., Sellen, A., Whittaker, S., and D. Kirk. 2010. *Now let me see where I was: Understanding how lifelogs mediate memory*. Paper Presented at the Proceedings of the 28th International Conference on Human Factors in Computing Systems.
- Lanier, J. 2010. *You are not a gadget: A manifesto*. London: Allen Lane.
- Leadbeater, C. 2008. *We-think: Mass innovation not mass production*. Profile.
- Loftus, E.F., and G.R. Loftus. 1980. On the permanence of stored information in the human brain. *American Psychologist* 35(5): 409.
- Loftus, E.F., and J.C. Palmer. 1974. Reconstruction of automobile destruction: An example of the interaction between language and memory. *Journal of Verbal Learning and Verbal Behavior* 13(5): 585–589.
- Luria, A.R., and L.S. Vygotsky. 1992. *Ape, primitive man and child: Essays in the history of behaviour*. New York: Simon and Schuster.
- Marcus, G. 2008. What if HM had a Blackberry? Coping with amnesia, using modern technology. *Psychology Today*.
- Mayer-Schönberger, V. 2011. *Delete: The virtue of forgetting in the digital age*. Princeton: Princeton Univ.
- McLuhan, M. 2000 [1964]. *Understanding media: The extensions of man*. London: Routledge.
- Menary, R. 2010. Cognitive integration and the extended mind. In *The extended mind*, ed. R. Menary, 227–244. London: Bradford Book, MIT Press.
- Mithen, S. 1996. *The prehistory of the mind*. London: Thames Hudson.
- Nolan, C. 2000. *Memento*. Los Angeles: Newmarket Films.
- Norman, D.A. 1999. *The invisible computer: Why good products can fail, the personal computer is so complex, and information appliances are the solution*. Cambridge: MIT Press.
- Olson, D. 1994. *The world on paper: The conceptual and cognitive implications of writing and reading*. Cambridge: Cambridge University Press.
- Ong, 1982. *Orality and literacy: The technologizing of the word*. Methuen.
- Rowlands, M. 1999. *The body in mind: Understanding cognitive processes*. Cambridge: CUP.
- Rupert, R.D. 2004. Challenges to the hypothesis of extended cognition. *Journal of Philosophy* 101: 389–428.
- Rupert, R.D. 2009. *Cognitive systems and the extended mind*. USA: Oxford University Press.
- Schacter, D.L. 1987. Implicit memory: History and current status. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 13(3): 501.
- Sellen, A.J., and S. Whittaker. 2010. Beyond total capture: A constructive critique of lifelogging. *Communications of the ACM* 53(5): 70–77.
- Simmons, D. 2010. *Ilium*. Gateway.
- Smart, P.R. 2012. The web-extended mind. *Metaphilosophy* 43(4): 446–463.
- Smart, P. R., Engelbrecht, P., Braines, D., Strub, M., and J. Hendler. 2009. *Cognitive extension and the web*. Paper presented at the WebSci'09: Society, Athens, Greece.
- Sparrow, B., J. Liu, and D.M. Wegner. 2011. Google effects on memory: Cognitive consequences of having information at our fingertips. *Science* 333(6043): 776–778.
- Sutton, J. 2003. Constructive memory and distributed cognition: Towards an interdisciplinary framework. In *Constructive Memory*, ed. B. Kokinov and W. Hirst, 290–303.

- Sutton, J. 2010. Exograms and interdisciplinarity: history, the extended mind, and the civilizing process. In *The extended mind*, ed. R. Menary, 189–225. London: Bradford Book, MIT Press.
- Tulving, E. 1972. Episodic and semantic memory. In *Organisation of memory*, ed. E. Tulving and W. Donaldson. New York: New York: Academic Press.
- Turkle, S. 2011. *Alone together: Why we expect more from technology and less from each other*. New York: Basic Books.
- Watson, S.J. 2011. *Before I go to sleep*. London: Random House.
- Wegner, D.M. 1987. Transactive memory: A contemporary analysis of the group mind. In *Theories of group behavior*, ed. B. Mullen and G. R. Goethals, 185–208. New York: Springer-Verlag.
- Wilson, R.A., and A. Clark. 2009. How to situate cognition: Letting nature take its course. In *The Cambridge handbook of situated cognition*, ed. M. Aydede and P. Robbins, 55–77. Cambridge: Cambridge University Press.
- Yeo, R. 2008. Notebooks as memory aids: Precepts and practices in early modern England. *Memory Studies* 1(1): 115.