

# Cognition and the Web: Extended, transactive, or scaffolded?

Accepted for publication in *Erkenntnis*

Richard Heersmink  
Department of Politics and Philosophy  
La Trobe University  
Melbourne, Australia

John Sutton  
Department of Cognitive Science  
Macquarie University  
Sydney, Australia

## Abstract

In the history of external information systems, the World Wide Web presents a significant change in terms of the accessibility and amount of available information. Constant access to various kinds of online information has consequences for the way we think, act and remember. Philosophers and cognitive scientists have recently started to examine the interactions between the human mind and the Web, mainly focussing on the way online information influences our biological memory systems. In this article, we use concepts from the extended cognition and distributed cognition frameworks and from transactive memory theory to analyse the cognitive relations between humans and the Web. We first argue that while neither of these approaches neatly capture the nature of human-Web interactions, both offer useful concepts to describe aspects of such interactions. We then conceptualize relations between the Web and its users in terms of cognitive integration, arguing that most current Web applications are not deeply integrated and are better seen as a scaffold for memory and cognition. Some highly personalised applications accessed on wearable computing devices, however, may already have the capacity for deep integration. Finally, we draw out some of the epistemic implications of our cognitive analysis.

Keywords: Extended cognition, Transactive memory, Cognitive integration, Knowledge, Virtue epistemology, External memory; Cognitive scaffolding

## 1 Introduction

In his book *Origins of the Modern Mind*, Merlin Donald writes: “The globalization of electronic media provides cognitive scientists with a great future challenge: to track and describe, in useful ways, what is happening to the individual human mind” (1991, p. 359). When Donald’s book was published, the Web had only just emerged, and its transformative effects on human cognition seemed to many observers likely to be marginal. The Web, however, has expanded substantially, providing many people with more or less constant access to a vast amount of digital information. It has been claimed that this constant access to external information is transforming the way we think and remember (Sparrow, Liu & Wegner 2011). Whilst we have always offloaded information-storage functions to external media (Donald 1991; Clark 2003), the Web differs from previous external information systems in its scope, constant availability, tendency to absorb and transform other media, and its multifunctionality. This suggests a need to analyse any specific and distinctive cognitive capacities and practices arising from our ongoing and increasingly pervasive interactions with the Web.

We have previously argued that metaphysical debates about extended or distributed cognition, which seek necessary and sufficient conditions by which to assess whether some external resource is part of an agent's cognitive system, are less productive than a focus on method (Sutton et al 2010). We defended a view of distributed or scaffolded cognition intended to be neutral on the metaphysics of extended mind, yet far from trivial in its implications for cognitive scientific practice. We have continued to develop this approach, both by identifying a wider range of dimensions on which we can empirically study variation across cases (Heersmink 2015), and by arguing that the notion of 'scaffolded cognition' in no way privileges the naked brain or the unscaffolded mind as the basic explanandum for cognitive science (Sutton 2015a). Here we apply this framework to one specific and significant case study, the case of the Web. Arguing that some existing philosophical work on the topic has got stuck on the metaphysical issue of whether or not the Web extends cognition, we offer an empirically productive alternative direction. By seeing the Web as a scaffold for our biological memory, we can ask a range of theoretically significant questions which link up to empirically tractable actual or possible research programs.

The analysis unfolds as follows. First, we describe the Web's status and nature as a cognitive technology. Second, we review some recent literature on extended cognition and the Web, suggesting that rather than focusing on metaphysical issues about whether the Web extends our cognitive system, we should focus more on psychological, epistemological and methodological issues. Third, we review recent claims that the Web is a transactive memory system, arguing that there are some differences between richly interactive social transactive memory systems, which involve ongoing directory updating, information allocation, and information coordination across distinct members of a small group, and most current human-Web interactions. Neither of these theoretical frameworks quite captures the particular features of such interactions. But fourth, building on concepts of the extended cognition and distributed cognition frameworks and transactive memory theory, we then offer an alternative way of looking at the interactions between embodied agents and online information. This is in terms of the cognitive integration of embodied agent and online resource, along a number of dimensions. This approach also makes firmer contact with empirical research, as we suggest in sketching its application to three specific and distinct applications of the Web, and suggests new directions for such research. We end by conceptualizing some of the epistemic implications of our cognitive analysis.

## **2 The Web as cognitive technology**

The Internet is a global system of numerous connected computer networks, storing, processing, and providing vast amounts of information. A distinction is made between the Internet and the World Wide Web (Berners-Lee et al 2006; Smart, Heersmink & Clowes 2017). The Web facilitates a way of accessing information over the medium of the Internet by using the HTTP protocol, which is one of the artificial languages used on the Internet. Web-browsers (such as Google Chrome or Safari) are used to access Webpages that are linked to each other via hyperlinks. The focus in this paper is on the cognitive relation between embodied agents and online content. Given this focus, it is helpful to briefly look at how this content is generated, navigated, and accessed.

Online content is generated in a variety of ways. A characterizing property of the Web is that it absorbs pre-existing representational media. Newspapers, magazines, books, TV programs, movies, photographs, maps, dictionaries, encyclopaedia, databases, and scientific journals used to be distinct media channels, but are now also available online. The Web has absorbed and transformed these media channels, thereby making information accessible in a way not seen before, making it a very powerful cognitive technology. David Bolter and Richard Grusin (1999) refer to this phenomenon as "remediation", the expression of one informational medium in another (see also Erll & Rigney 2009; Hoskins 2009). When media systems are remediated, their cognitive affordances change. Because online information is digital, it affords different search functions, as compared to information in an analogue medium. For example, when looking for a keyword or phrase in a digital text, we use the search function of our Web-browser. Consequently, we do not have to scan the entire text for what we are looking for, in that way changing our reading practices (Hayles, 2012). Some paper-based texts also have search functions in the form

of an index, but these are more time-consuming to use and contain only a limited number of keywords. For example, one might be looking for articles on “distributed cognition” in a journal. The journal’s search bar allows one to find all articles with that phrase within seconds, also those that do not have it listed as a keyword. Depending on the journal, it would have taken many hours, perhaps even days, to scan all the physical issues. So, the way we interact with information represented digitally differs from the way we interact with analog information.

The two most prominent ways to navigate online information are by search engines and hyperlinks. Hyperlinks allow a user to go to another (part of a) Webpage by clicking on it. Hyperlinks can be made between any Webpages with a URL, and can thus direct a user to any part of the Web, in that way allowing non-linear browsing strategies. There are various search engines such as Bing, Wolfram Alpha, and Yahoo, but Google is the most popular one. Google Search uses various algorithms to rank search results. These algorithms take into account the keyword usage on Webpages and in the URL, the structure of Webpages, the amount of time users have spent on Webpages, page loading speed in HTML, the number of inbound hyperlinks, the quality of inbound hyperlinks, and other factors. Because the Web is constantly changing, a Webpage that is ranked first today might not even be on the first page next week. Page ranking is cognitively and epistemically relevant as empirical research using eye-tracking technology has shown that Web-users implicitly trust Google’s ranking in that they prefer to click on links in higher positions even when the abstracts are less relevant to the task they are doing (Pan et al 2007). Another epistemic aspect of Google Search is its autocomplete system (Miller & Record 2017). While typing a query in the search bar, the system suggests two of the user’s previous search queries and two popular search queries based on a statistical analysis of other people’s search queries in one’s geographical location. The suggested queries are updated in real-time, so when a user is typing, new words and phrases are continuously suggested.

Online content is accessed with several devices, currently including desktop computers, laptops, tablets, smartphones, and smartwatches. These devices use Web-browsers to access Webpages such as, for example, Wikipedia or the Internet Movie Database (IMDB). Smartphones are now the most popular way to access the Web (Pew Research Center 2015), giving users reliable access to online the online world. The Web thus obtains its informational power and cognitive efficacy by absorbing other media, by introducing new ways to navigate information such as with search engines and hyperlinks, and allowing easy and reliable access to online information. Below in section 5, we analyse how Wikipedia, IMDB, and Google Search are integrated into their users’ cognitive systems. Wikipedia and IMDB are among the most used online encyclopaedic sources. Google Search is the most used search engine and thus the main portal to the online world for most Web-users<sup>1</sup>. These applications are thus exemplary for our cognitive relation to the Internet. Moreover, these Web applications have received attention from philosophers (Smart 2012; Ludwig 2015; Miller & Record 2017) and so this paper further builds on and is continuous with previous philosophical analyses.

### **3 Does the Web extend our cognition?**

The theory of extended cognition is a view about the location of the physical substrates of cognitive states and processes. The claim is that these are, under certain circumstances, not located exclusively in the brain, but distributed across brain, body, and environment (Clark, 1997, 2008; Menary, 2007; Rowlands, 2009; Sutton, 2010a). This is a metaphysical approach to cognition in that it moves beyond an individualist form of cognitivism and towards a picture that involves brain, body, and environment. But it also has methodological consequences, which have perhaps more often been highlighted in the closely related ‘distributed cognition’ framework (Hutchins 1995; Michaelian & Sutton 2013; Sutton 2015b): rather than merely focussing on cognitive processes realised in the brain, it advocates a focus on the relations between brain, body, and environment, both conceptually and empirically. Various authors have recently used an extended or distributed cognition framework in asking whether the Web can ever be part of the

---

<sup>1</sup> <https://www.alexa.com/topsites>.

physical substrate that realises certain cognitive states and processes. Below we outline and evaluate their claims.

### 3.1 Andy Clark

In their article *The Extended Mind*, Andy Clark and David Chalmers ask the following question: “Is my cognitive state somehow spread across the Internet?” (1998, p. 17). They then introduce a number of conditions by which to examine whether specific external resources count as proper parts of an extended cognitive system. These conditions are reliability, trust, accessibility, and past endorsement. They speculate that past endorsement as a criterion for belief is questionable: but later Clark (2008, p. 80) pointed out that not including it would lead to an undesirable explosion of potential dispositional beliefs. If one of these four conditions is not sufficiently satisfied, then cognition is not extended. So, they provide a fairly clear set of criteria to distinguish between an external resource that is a constitutive part of an extended cognitive system and one that is scaffolding cognition. Applying these conditions to the Internet, Clark and Chalmers point out that “The Internet is likely to fail on multiple counts, unless I am unusually computer-reliant, facile with the technology, and trusting, but information in certain files on my computer may qualify” (1998, p. 18). Clark later (2008) returned to this point, arguing that mobile access to Google fails the trust and past endorsement criteria. In a later article, however, Clark does suggest that online information can be part of an extended cognitive system. In a co-authored paper with Harry Halpin and Mike Wheeler, he says: “Perhaps external representations on the Web, when integrated appropriately into the processes that govern an agent’s behaviour, may count as part of that agent’s cognitive architecture” (Halpin, Clark & Wheeler, 2014, p. 24).

### 3.2 Paul Smart

Smart (2012, 2017) says that current HTML-based Webpages make it hard for users to quickly obtain the information they are looking for, in that way making it difficult to fulfil Clark and Chalmers’ (1998) availability criteria. For example, when looking for a specific item of information on a Wikipedia page, it

“requires the user to scroll through the Webpage and process large amounts of largely irrelevant content in order to identify the small amount of information that is actually needed. This makes it difficult to see how current forms of Web-based content could have the kind of functional poise sufficient to count as part of our personal body of knowledge and beliefs about the world” (2012, p. 452-453).

So, the way most information is organised on the Web, constrains efficient user-Web interaction. Smart then goes on to suggest several useful improvements that need to be made before the Web can be seen as a part of an extended cognitive system. These suggestions focus on making online information more accessible and more suitably poised to influence and scaffold our everyday thoughts and actions.

### 3.3 David Ludwig

Ludwig (2015) argues that Clark and Chalmers’ (1998) active externalism “implies an explosion of dispositional beliefs and knowledge that is caused by digital information sources such as Wikipedia or Google” (2015, p. 355). He argues that the past endorsement condition is incompatible with active externalism and including it as a condition for cognitive extension would “transform an active externalism into a passive Putnam-Burge-style externalism” (2015, p. 362). To make his point, he gives an example of someone updating Wikipedia entries. Laura updates fact A, but not fact B on a Wikipedia entry. After a while, Laura forgot which facts she has updated and when retrieving either fact A or fact B “her information retrieval will be *exactly the same* no matter whether she added the information or not” (2015, p. 362, original italics). In both cases, her information access is functionally and phenomenologically equivalent. The difference between fact A and fact B is entirely historical while the relevant roles in the here-and-now are identical, undermining the active character of active externalism. We should therefore, Ludwig argues, remove the conscious endorsement condition. But doing so then leads us to an explosion of

knowledge, as Wikipedia and other online sources fulfil the other three conditions (reliability, trust, and accessibility) of cognitive extension.

Based on his analysis, Ludwig claims that “Internalists may consider the conclusion of this article a reduction of the very idea of an extended mind while externalists will have to develop strategies that incorporate an explosion of knowledge into their theories” (2015, p. 356). In section 5, we respond to this claim not by incorporating an explosion of knowledge into extended cognition theory, but by proposing other dimensions for thinking about when an external resource extends our cognition that exclude Wikipedia and most (but not all) other online resources as part of extended cognitive systems.

### **3.4 Discussion**

What Clark, Smart, and Ludwig have in common is that they work with Clark and Chalmers’ (1998) original criteria of reliability, trust, accessibility, and past endorsement, and treat these as something like necessary conditions for cognitive extension. If one of these conditions is not sufficiently satisfied, then the Web cannot extend the mind of its user. Clark (2008) argued that information on the Web is not automatically endorsed, is not there because of prior endorsement, and is often not as trustworthy as information retrieved from biological memory. For these reasons, it is not part of an extended cognitive system. In a co-authored paper (Halpin, Clark & Wheeler, 2014), Clark later adjusted his view and said that when integrated appropriately into the processes that govern an agent’s behaviour, online information may be part of an extended cognitive architecture. Smart claims that current Web technology does not satisfy these conditions, because most information is not easily accessible, whereas Ludwig claims that some online information does satisfy all the relevant conditions, except past endorsement, which he sees as not relevant anyway.

This philosophical debate on human-Web relations, as outlined above, focusses on metaphysical issues concerning the constitution of extended cognitive systems and whether the Web satisfies Clark & Chalmers’ conditions of cognitive extension. In doing so, the debate overlooks psychological, epistemological and methodological aspects of human-Web interactions. It does so, arguably, because of the primary concern of these authors with identifying sharp *criteria* by which to identify whether a particular external resource is part of an extended cognitive system or not. But the specific criteria under discussion do not in fact lend themselves naturally to this kind of analytical project. They are matters of degree, and are likely realised to different degrees across a continuous range of cases. They are not so much sharp criteria as continuous dimensions of variation (Sutton et al 2010; Heersmink 2015). We will argue below that a fuller multidimensional approach, embracing such continuous variation and seeking to identify a wider range of relevant ways that distinct cases can vary, can be productively applied to the case of cognition and the Web. But first, we look at a field in which the cognitive consequences of using the Web for our goals have been studied, in the work of cognitive psychologists who focus on the Web as a transactive memory system.

## **4 The Web as a transactive memory system?**

### **4.1 Social transactive memory systems**

Transactive memory theory, as developed by Daniel Wegner and colleagues (Wegner, Giuliano, & Hertel, 1985; Wegner, 1986, 1995; Wegner, Erber & Raymond, 1991; see also Sutton, Harris, Keil & Barnier, 2010; Tollefsen, Dale & Paxton, 2013; Theiner, 2013), describes how social groups process and structure information, in that way developing a transactive memory system (TMS). A TMS is a cognitive system comprising people in close relationships in dyads or larger groups who engage collaboratively in encoding, storing, and retrieving information. Transactive processes most obviously occur at the encoding and retrieving stage. In transactive encoding, group members may discuss incoming information, negotiating who has to store it and in which form. In transactive retrieval, group members need to know where relevant information is stored. Both encoding and retrieval may be the topic of explicit negotiation and strategies, or alternatively can

over time become more or less delegated to structural features of the system, and to implicit interactive processes. In either case, sometimes this involves the combination or interplay of different units of information retained in or across different brains.

Wegner (1986) emphasizes that a TMS cannot be reduced to individual memory. It is a group property that exists only when people interact and communicate in a particular way, i.e., it is more than the sum of its parts. An effective TMS typically takes a fair bit of time to develop. Over time people learn what others know, what others do not know, and take responsibility for storing certain kinds of information. Wegner writes: “transactive memory can be built because individuals in a group accept responsibility for knowledge” (1986, p. 194). If people do not take this informational responsibility, the TMS does not work properly.

Deploying an analogy between memory in social groups and computer networks, Wegner (1995, p. 336) points out that for cooperative memory capacities to function, we need to develop an organisation, channel information to the right places, and use effective strategies for accessing that information:

“First, the formation of human transactive memory systems depends on processes whereby people learn what others are likely to know about (directory updating). Second, the formation of these systems also requires processes whereby information coming into the group is communicated to individuals whose expertise is likely to facilitate its storage (information allocation). Third, the operation of human transactive memory depends on each person in the group having a retrieval plan for any topic based on the relevant expertise of self and the others in the group such that the sequence of locations for any information search can be determined (retrieval coordination)” (Wegner, 1995, p. 326).

Directory updating, information allocation, and retrieval coordination are ongoing processes and are “necessary for the formation and successful operation of transactive memory” (1995, p. 320). In an earlier article, Wegner, Erber & Raymond (1991) discuss three kinds of directory updating. First, through negotiated responsibility, where group members explicitly negotiate who has to remember what. Second, through perceptions of relative expertise of oneself and the others in the group, which often happens by having conversations about certain topics, during which one learns what the other knows. Third, through knowledge of the other members’ past access to information. A group member might, for example, know that another group member did see that documentary of David Attenborough on birds of paradise and thus knows more about its content.

#### **4.2 Online transactive memory systems**

Transactive memory theory has recently been extended from social groups to human-Web systems. In a series of four experiments, Sparrow, Liu & Wegner (2011) examined the effects of external information-storage on remembering the content of various trivia statements and the locations where they were stored. The authors show four things:

- (1) When we do not know something, we are primed to use the computer to look it up;
- (2) When we know information is accessible externally, we put less effort into encoding it internally;
- (3) When we know information will not be accessible externally, we put more effort into encoding it internally and have better recall of that information;
- (4) When we know information is stored externally, we often know where it is stored, but not the exact content of that information.

Extrapolating from these results, the authors suggest that human biological memory is adapting to Web technology, incorporating it into our practices of remembering. They further argue that relying on online information for memory depends on the same transactive processes that occur in social TMSs. Just as people learn about what other people know in a group, we learn what the

computer “knows” and where to find it. The authors point out that we are developing a symbiotic relationship with computer systems, storing less knowledge internally and developing ways to navigate external information stored on the Web. Whilst this may be true, it is important to realize that the experiments done by Sparrow et al (2011) do not involve using the Web. The trivia statements were stored in folders on a desktop computer in a psychology laboratory. Whether that is relevantly similar to using the Web in real-world situations as to justify the claim that the Web is a transactive memory system needs more justification (Heersmink 2016).

Building on the above research, Adrian Ward (2013) further theorizes about the Web as a transactive memory partner. He argues that the Web satisfies the three criteria of a TMS, which he defines as availability of information, relative expertise, and accessibility of information. We take it that availability and accessibility are basically identical, reducing the criteria to access and expertise. Ward argues that the Web is virtually always available and that the information on the Web is often more accurate and elaborate than information in people’s individual memory, such that the Web should be seen as a transactive memory partner. Moreover, because the Web is readily accessible and holds an enormous amount of factual knowledge on the Web, Ward sees it as a superior transactive memory partner.

### **4.3 Discussion**

There is no doubt the Web plays a significant and transformative role in our practices of remembering, but it is not obvious whether the concepts of TMS theory naturally extend to human-Web interactions. Wegner (1995) argued that directory updating, information allocation, and information coordination are the necessary building blocks of transactive memory systems. How do these communicative processes apply when interacting with the Web? Sparrow et al claim that: “Relying on our computers and the information stored on the Internet for memory depends on the same transactive memory processes that underlie social information-sharing in general” (2011, p. 778). Below we evaluate this claim.

#### **4.3.1 Directory updating**

Directory updating can in this context be characterized as learning what the Web “knows”, that is, learning which information is available on the Web. In social transactive memory systems, there are three kinds of directory updating: (a) negotiated responsibility; (b) perceptions of relative expertise of self and others; and (c) knowledge of other members’ past access to information. How do these apply to the Web? Option A does not seem to apply. We do not intentionally negotiate with the Web regarding responsibility for storing information. In dyads and groups, responsibility for storing certain kinds of information is often delegated to experts on those topics. In case of human-Web interactions, there is only one agent involved and thus the Web cannot take responsibility for storing information in the same way as a person can. Option C might in some cases apply. A Web-user may, for example know, that Google Books has access to most books but not to ancient Sumerian manuscripts that are only available in a library. Option B seems most relevant here. We learn about the Web’s informational content and relative expertise from teachers, friends, and colleagues, but most obviously from our history of browsing. Most people have spent thousands of hours on the Web, looking for information on news sites, Wikipedia, social media, blogs, and many other websites. Our history of browsing results in developing numerous metamemories regarding online content, i.e., memories not of informational content itself, but of where to find information. It also results in beliefs about what information is likely to be online. For instance, I may have never seen the Wikipedia entry on Barack Obama but know it is very likely it exists. Because most people have spent so many hours online they have an intuitive sense of which information is online.

#### **4.3.2 Information allocation**

In social transactive memory systems, incoming information is delegated to relevant experts in the group. In a dyad, for example, one person may be responsible for the couple’s joint appointments. So, when a friend makes an appointment to have dinner on Friday, the information is delegated to the individual responsible for appointments. This can operate more or less

implicitly, or can be an intentional communicative process whereby one individual tells or asks another to store certain information and to retrieve it when required. Sometimes we also delegate information-storage functions to the Web, for example when we use Google Calendar as our diary or IMDB to list the movies we have seen or want to see. Information is then offloaded to the Web in a personalised application with the intention to retrieve it at some later point to guide action or decision-making. Such applications can be seen as the relevant experts on certain specialised topics. It is, however, important to note that only a small part of our embodied interactions with the Web concern the offloading and subsequent intake of information. The vast majority of our interactions with the Web involve one-way informational exchanges: from the Web to a user. The point here is that most incoming information is not delegated in any kind of interactive fashion to the Web, either intentionally or non-intentionally, as it is already there. The Web is not like a notebook in which we intentionally write information. There are, of course, personal information management systems such as Slack, Trello, and others, where one uses a cloud-based application to manage one's appointments and tasks. For those who use them, such applications can potentially be seen as external memory systems when used in the right sort of way.

### 4.3.3 Information coordination

Information coordination is the retrieval of information based on one's own knowledge and on that of the Web. Given that there is considerably more information available online, as compared to biological memory, we quite often turn to the Web for our informational needs. It is helpful to distinguish between two kinds of information retrieval: (1) the retrieval of information others have put online (e.g., Wikipedia) and (2) that of information we have put online ourselves (e.g., Google Calendar). How do we coordinate the retrieval of this information? Generally, information we have offloaded ourselves is retrieved via bookmarked websites<sup>2</sup> and information others have offloaded is typically retrieved with search engines, either high-level ones such as Google, Bing, or Wolfram Alpha's search engine or low-level ones such as the search bar in Wikipedia. There are also cases where we do not use search engines but go directly to the website of interest. Most Web-browsers store previously visited URLs and automatically finish the URL when typing the first letters into URL bar.

Wegner, Giuliano & Hertel (1985) repeatedly emphasize the cognitive *interdependence* in dyads, which means that both agents rely on each other for their cognitive performance. In case of human-Web systems, only the agent depends on the online information, not vice versa. In this sense, there is no interdependency, but only dependency. A key example of interdependency is when dyads perform interactive cueing, in which case none of the agents knows the answer, but by giving each other cues they come to the right answer. Consider an example of interactive cueing from Harris et al (2010). In this example, a long-married couple recalls the name of the show they saw on their honeymoon more than forty years ago. Neither of them seems to know the answer, but by giving each other cues, they jointly construct the answer.

Wife: And we went to two shows, can you remember what they were called?

Husband: We did. One was a musical, or were they both? I don't ... no ... one ...

Wife: John Hanson was in it.

Husband: *Desert Song*.

Wife: *Desert Song*, that's it, I couldn't remember what it was called, but yes, I knew John Hanson was in it.

Husband: Yes

Note that this conversation has a reciprocal information flow structure driving the required cross-cueing. The wife asks something, the husband is then triggered by what she says and responds to it, the wife is in turn triggered by what the husband says, and so on. This process has various iterative cycles of informational exchanges in which each cycle depends on what is said before. This kind of interactive cueing does not occur when using the Web. Whilst some Websites give

---

<sup>2</sup> An interesting phenomenon is social bookmarking where different users can add, annotate and share bookmarks with other Web-users.



suggestions or cues to facilitate information retrieval, for example when Google's autocomplete system suggests search queries or when Amazon suggests books related to your current search term or search history, it is somewhat different, in interesting ways, from human interactive cueing. In case of Google, a search engine responds to the informational input of a user by giving suggestions that are partly based on a statistical analysis of other people's search queries relative to your geographical location.

For instance, when typing "cog" into Google's search bar, it suggests the following terms: "cognitive virtues", "cognition and material culture", "cognitive", and "cognitive dissonance". The first two are based on one of the authors' (RH) search history, and the latter two on a statistical analysis of other people's search queries in relation to RH's geographical location. The cues Google gives and the cues other people give differ in several ways. The Desert Song example is based on a shared experience and has a strong autobiographical component, whereas Google's cues are based on one's search history, statistics, and location. Moreover, in case of Google's cues: we start typing, Google then suggests a number of alternative phrases which may or may not be relevant, and we may then decide whether or not to use the alternative search query. This is not a reciprocal but a two-way process. So, it seems that current agent-Web interactions do not exhibit the fuller reciprocity that can exist between members of social transactive memory systems.

#### **4.4 Some concluding remarks**

Each of the above transactive processes occurs when interacting with the Web. However, they never jointly occur when interacting with one specific Webpage. We learn what the Web "knows" from our peers and from our history of browsing, in that way updating our directories. Information is allocated to the Web only in case of personalised applications such as Google Calendar or IMDB. Information retrieval from the Web mainly concerns one-way information flow structures: from the Web to an agent. This happens fairly often with the aid of cues in which an application responds to the informational input of a user by giving suggestions that are based on one's own search history, a statistical analysis of other people's search queries, and one's location.

Given the differences between social and online memory, can we still see the Web as a transactive memory partner? We think current Web applications do not sufficiently exhibit the three necessary transactive processes to count as a genuine transactive memory system. A somewhat similar conclusion is made by Kyle Lewis & Benjamin Herndon regarding information systems in organizations. They write: "these systems do not effectively emulate or facilitate transactive processes" (2011, p. 1262). Whilst these systems differ from the Web, they also have many similarities. "Incorporating the functions of transactive processes in information technology involves modelling the transactive aspects of learning, storage and retrieval (i.e., cognitive abilities that result from interactions between people or between a person and a system)" (2011, p. 1262). Current information technology, including the Web, is unable to do that (see also Huebner 2016).

Overall, then, even though some of the concepts of TMS theory are helpful in identifying big-picture parallels between cases of socially distributed cognition and cases in which agents are interacting with the Web, it is not entirely natural to push the analogy too far. It seems awkward to treat the (micro)processes of interaction and communication between humans and online systems as occurring within fully integrated transactive systems. The main reason for this, we suggest, is that in general these interactions are unidirectional, lacking the more deeply integrated two-way or reciprocal information flow which characterises some socially distributed cognitive systems. This point takes us on to a broader way of trying to capture the variety of relations in play, which we can then apply to distinctive cases of agent-artifact relations.

### **5 A multidimensional approach to situated cognitive systems**

Current theoretical frameworks describe aspects of, but do not fully capture, human-Web interactions. Halpin, Clark and Wheeler (2014), Smart (2012), and Ludwig (2015) provide valuable claims about the Web and extended cognition. But instead of focussing on the

metaphysical issue whether or not the Web extends our minds, it is perhaps more fruitful to focus on what the Web is actually doing for and to our cognitive system. Sparrow et al (2011) and Ward (2013) do focus on what the Web is doing for and to our cognitive system, but their theoretical apparatus does not seem to generalize neatly or fully to the Web. That is, it is difficult to see the Web as a transactive memory system in the sense defined by Wegner and colleagues (Wegner, Giuliano & Hertel, 1985; Wegner, 1986, 1995), because there is no cognitive interdependency: there is only information allocation in some cases as most information is already online, and informational exchanges are mainly one-way, from the Web to an agent. So, we suggested above, the interactive dynamics of interactions with the Web differ from that of social TMSs.

## **5.1 Introducing dimensions of integration**

In this section, we further conceptualize human-Web interactions by investigating how embodied agents and online information are integrated into larger cognitive systems along various dimensions. Our approach builds on Sutton (2006; Sutton et al 2010), Wilson and Clark (2009), Sterelny (2010), and Menary (2010). These authors have suggested analysing the interactions between agents and external resources in terms of dimensions, which are not seen as necessary conditions for cognitive extension, but are matters of degree (Heersmink 2015). Below we refine and synthesize some of the proposed dimensions into a systematic framework, add several dimensions to the framework, and then apply them to conceptualize the degree of integration in three case studies: Wikipedia, IMDB, and Google Search.

### **5.1.1 Information flow**

In situated cognitive systems consisting of an embodied agent and an external resource, the propagation of information between agent and resource may be one-way, two-way, or reciprocal. One-way information flow goes from resource to agent, such as using traffic signs to navigate in an unfamiliar city. In case of two-way information flow, an embodied agent first offloads information onto the environment and then takes it onboard again at some later point. There is then one cycle of informational offloading and intake, which happens, for instance, when you write an appointment in your diary and look it up at some later point. Finally, when information flow is reciprocal, it concerns more complicated and elaborate informational exchanges between agent and resource. When that happens, the external information is part of an ongoing cognitive task and often changes its content during the task. The line between two-way and reciprocal information flow may be fuzzy, especially when considering exchanges which have more complicated temporal structures: but the general distinction is clear enough to be useful. Consider, for example, making a difficult calculation with pen and paper, reorganising Scrabble tiles, doing a preliminary sketch of complex structures, or writing an academic paper. In these cases, there is not one cycle of offloading and intake of information, as is the case in two-way interactions, but a number of incremental and interdependent cycles of offloading and intake. To use Clark's (1997) phrase, there is "continuous reciprocal causation" between agent and resource, which results in an "incremental cognitive self-stimulation" (Clark, 2008). Situated cognitive systems exhibiting reciprocal information flow are often seen as the paradigm cases of extended cognitive systems. In the transactive memory literature a clear example of reciprocal information flow is interactive cueing (e.g. Harris et al, 2010).

### **5.1.2 Accessibility**

Informational accessibility is central to the degree of integration: information that is not reliably available when needed cannot be used for performing some cognitive task. So, we suggest, the more reliably task-relevant information is available to its user, the easier it can be integrated into the user's cognitive system. Informational accessibility has strong effects on the way we think and remember. Sparrow et al (2011) demonstrated that when agents know information is reliably available in some external artifact or media, they tend to put less effort into encoding it internally. So, knowledge of accessibility transforms our cognitive strategies.

### **5.1.3 Durability**

There is a great variety in the durability of the relation between agent and artifact. Wilson and Clark (2009) make a tripartite distinction between one-offs, repeated, and permanent relationships to external informational resources. An online recipe, for example, is often used only once, an online TV guide might be used several times, and an online calendar such as Google Calendar is used on an ongoing basis. In general, the more often we use information, the more deeply it is integrated into our cognitive practices.

### **5.1.4 Trust**

Whether we trust external information to be true is important, because trustworthy information we may use for some cognitive task, whereas information that we think is untrustworthy we typically do not use. As Halpin, Clark and Wheeler argue “A central challenge will be to analyse the conditions under which users trust, by responding unreflectively and uncritically to, the collectively retrieved information from the Web” (2014, p. 23). There are at least two ways we can come to trust external information: either after consciously evaluating it (i.e., explicit trust) or without consciously evaluating it (i.e., implicit trust). This is particularly important in relation to the Web, as some online information is of questionable quality. For this reason, it is important to develop strategies to be able to consciously evaluate online information (see also Smart 2017).

### **5.1.5 Procedural transparency**

The easier it is to use and interact with an artifact, the more procedurally transparent it is. Heidegger (1962) conceptualized the relation between a hammer and a skilful carpenter as potentially transparent-in-use, when the carpenter does not focus on the hammer but on the task at hand. Heidegger calls this “ready-to-hand”. This phenomenon also occurs with cognitive artifacts. For example, skilful computer-users do not consciously think about how to use the keyboard and mouse, but operate automatically, as the artifacts withdraw from attention and become invisible. Cognitive artifacts, such as computers, should be as invisible as possible to ensure a high use-efficiency (Norman, 1998). When cognitive artifacts become procedurally transparent, our intentional stance is changed, because such artifacts enable us to perform actions that we would otherwise not be able to do, i.e., our action repertoires are enhanced (Kirsh, 2013).

### **5.1.6 Informational transparency**

Information can be transparent or opaque. When it is transparent we understand what it means, when it is opaque we do not, with various grades in between. So, the easier it is to interpret and understand external information, the more informationally transparent it is. Some informational media, like language and mathematics, are semantically and syntactically complex and take a long time to learn and make transparent. Other media, such as maps, are in some cases and in some respects less complex and easier to learn and to render transparent in use. So, the degree of transparency often depends on the complexity of the information as well as the learning history of the user. There is a relation between the degree of informational transparency and ease of use. If information is opaque (e.g., a manual or Webpage in a language one does not understand), it cannot afford its intended informational use. Note that informational transparency is not an intrinsic property of cognitive artifacts, but partly depends on the cognitive profile and capacities of the interpreting agent.

### **5.1.7 Individualization**

Some cognitive artifacts are interchangeable and used by many people, whereas others are individualized and used by only one person (Sterelny, 2010). Interchangeable artifacts include timetables, maps, clocks, and textbooks. Individualized artifacts include diaries, notebooks, to-do lists, and lists of bookmarked websites<sup>3</sup>. Individualisation is important for the degree of integration because it often streamlines a cognitive task, making it easier and faster to perform. So, in a sense, individualisation is not important in itself, but only as it contributes to forms of

---

<sup>3</sup> A reviewer pointed out that contemporary Web-browsers (for example, Chrome) store bookmarked Webpages on the cloud and can therefore be accessed on any device, making access to bookmarked Webpages easier.

transparency. A key example of a highly individualized cognitive artifact in the extended mind literature is Otto's notebook (Clark & Chalmers, 1998). The information in the notebook is created by Otto and for Otto to help him get around in the world and guide his actions. Whilst the information in the notebook is interpretable and transparent to others, it is only helpful to Otto.

### 5.1.8 Transformation

Interacting with external information transforms our onboard cognitive system and cognitive practices in various ways. During our lifetime, we learn to internalize publicly available informational systems such as language and number systems. The brain absorbs these informational systems which thereby transform the brain's cognitive capacities (Menary, 2010). Another way external information transforms our cognitive abilities is by changing our memory strategies. Knowing that information is accessible in some external artifact transforms the way we think, in that it changes our memory strategies. The informational properties of cognitive technologies also transform over time. This is particularly true for the Web. Webpages are constantly added, deleted, updated, and because search results are personalised and depend on many factors, the page ranking is also constantly changing.

### 5.2 Case studies: Wikipedia, IMDB and Google Search

Jointly, the above dimensions form a multidimensional space in which situated cognitive systems can be positioned. Systems ranking high on most dimensions are deeply integrated and therefore clear candidates for extended or distributed cognitive systems (Heersmink 2015; compare Clowes, 2015). Those ranking low on most dimensions are shallowly integrated and are cases of embedded or scaffolded cognitive systems. To briefly demonstrate the utility of the proposed framework, we now examine the uses of Wikipedia, IMDB, and Google Search as case studies. There is, of course, a lot of variety in how different people use these Web-applications. Gender, age, education, and cultural background all play a role in the way people interact with the Web (van Deursen & van Dijk 2009). In the case studies below, we take a college student as typical Web-user, partly because most empirical research in psychology and cognitive science is done on college students and partly because students often have relatively good Internet literacy skills and so would show the framework in a charitable way. College students are, of course, not a homogenous group and the level of Web-skills varies across this group (Hargittai 2010). The case studies should, therefore, be seen as idealisations. It is important to point out that the goal is not to provide a precise analytical analysis. Rather, the goal here is to apply the dimensional framework as a heuristic tool to explore and conceptualize the degree of integration in an idealised group of Web-users. We are aware that college students are not representative for the population at large. They are what Joseph Henrich, Steven Heine & Ara Norenzayan (2010) refer to as Western, Educated, Industrialized, Rich, and Democratic (WEIRD) people. These constitute only a subset of the world's population and therefore the outcome of the analysis cannot be generalised.

When using Wikipedia, the flow of information is one-way: from a Website to a user. For a typical student, the information flow is one-way, but as a reviewer helpfully pointed out, a small group of users actively discuss and edit the content of entries, in which case there is reciprocal information flow between various agents mediated by an online resource. The informational accessibility is high as it can be viewed on smart phones, tablets, laptops, and desktops. As many people, particularly students, have at least one of these devices - the Pew Research Center (2015) shows that 68% of Americans have a smartphone, 45% have a tablet, and 73% have a desktop or laptop computer - they have a high degree of access to Wikipedia<sup>4</sup>. How durable the relation to Wikipedia is, depends on the student. Some use it often, others rarely. Alison Head and Michael Eisenberg (2010) show that 52% of students use Wikipedia frequently, while 22% use it rarely, if ever<sup>5</sup>. So,

---

<sup>4</sup> The Pew Research Center (2011) shows that 53% of Americans use Wikipedia to look for information. According to Alexa (see: <https://www.alexa.com/topsites>), Wikipedia is the 5<sup>th</sup> most visited Webpage in the world and has a daily pageview per visitor of 3.3.

<sup>5</sup> Furthermore, students majoring in architecture, engineering, or the sciences use Wikipedia more often than students in the humanities (Head & Eisenberg 2010).

the durability varies from one-off to permanent. The trust students put in the truth-value is probably in general medium to high. Most students are aware that information on Wikipedia might be incomplete or outdated, but research has shown that overall it is quite reliable (e.g., Fallis 2008). They may or may not explicitly evaluate the claims, so trust may be implicit or explicit, but most information is most likely not consciously evaluated.

The retrieval of information from Wikipedia is a fully procedurally transparent process: many students do it so regularly that they do not have to think (a lot) about doing it. Search bars, hyperlinks, sitemaps, portals, and scroll bars are all transparent in their function to most regular Web-users and utilizing these tools to find information on Wikipedia is quite easy. Empirical research by Head and Eisenberg (2010) shows that 64% of US college students find it easy to use Wikipedia and to obtain the information they are looking for. The informational transparency is typically also quite high, since the entries are often written such that they are understandable for lay people, although there are of course (parts of) entries that are too difficult for all readers to fully understand. Wikipedia is not individualised: students do not gear it for personal use, although they may bookmark certain entries to read later. Having constant access to it is transforming memory practices in that Web-users including students store less information internally because we know it is available online. This is a consequence of the Web in general, but given that it ranks 5<sup>th</sup> in terms of the most visited Webpages, Wikipedia probably plays an important role in this development. There are also situations where Wikipedia is the referee, for example when having a discussion with a friend about whether the movie *Memento* was Christopher Nolan's debut, we sometimes turn to Wikipedia (or some other website) for a definitive answer. So, overall, Wikipedia's ranking on these dimensions suggests its position in a region somewhere in the middle of the multidimensional space.

When using IMDB, the flow of information is one-way or two-way, depending on what information is used. Like most Webpages, IMDB is used to browse and look for information, but a user can also make a personal profile, allowing one to grade movies and make a list of movies one wants to see. By doing so, one intentionally offloads information onto IMDB that one can use at some later point. IMDB also personalises in that it monitors and stores one's past browsing behavior in the "Recently Viewed" section. This kind of personalisation is typically unintentional, but when the information in the Recently Viewed section is used for memory purposes, it is a case of two-way information flow. IMDB also provides suggestions of movies related to the entry one is currently reading. So, the entry on *Memento* suggests *American History X*, *Fight Club*, *Shutter Island*, and other films. These are personalised recommendations. In their Frequently Asked Questions, IMDB writes: "First, we take all of the movies and TV shows that you have either rated or added to your Watchlist. Then, we compare your data to ratings made by other users. We can then find movies and TV shows that people with similar tastes to you like."<sup>6</sup>

The information provided by IMDB is typically not evaluated, it concerns facts about films, actors, awards, etc. In case of Wikipedia, all users can alter, delete, or add information, but in case of IMDB, the entries are made by the editors of IMDB. There is no reason to doubt the truth-value of this information. So, the level of (implicit) trust is very high. Both the procedural and informational transparency are very high as well. For those who have a personal profile, some parts of IMDB are individualised but most of it is not. For heavy IMDB-users, it may have transformed their memory about cinema. So, overall, IMDB's ranking on these dimensions suggests a position in a region in the middle of the multidimensional space.

When using Google Search, the flow of information is typically two-way: information is entered into the search-bar (automated suggestions may or may not be used) and the search engine then presents a ranking of the results. The informational accessibility is high as it can be used on smart phones, tablets, laptops, and desktops. The durability of the relation to Google Search depends on the user, but is in general very high. Most students use it daily, in which case a permanent relation to it is established. Head and Eisenberg (2010) show that Google Search is the most used Webpage

---

<sup>6</sup> [http://www.imdb.com/help/show leaf?personalrecommendations&ref=tt\\_rec\\_lm](http://www.imdb.com/help/show_leaf?personalrecommendations&ref=tt_rec_lm)

by students. Trust should be slightly re-interpreted here as Google Search does not contain any information in the way a conventional Webpage does. It merely suggests search queries (which we may or may not use) and ranks Webpages (which we may or may not use). Trust here concerns the students' attitude towards the usefulness and accuracy of both the automated search queries and the page ranking. Whether a student explicitly evaluates the autocomplete suggestions and page ranking depends on the kind of student, but we suspect that most students won't evaluate the page ranking. So, the level of (implicit) trust is high. This is supported by empirical research showing that college students implicitly trust Google's page ranking as they prefer to click on links in higher positions even when the abstracts are less relevant to the task they are doing (Pan et al 2007). Using Google Search is for most students a procedurally transparent process, but the statistics and algorithms behind its page-ranking are for most students opaque. The informational transparency as it concerns the logic behind the ranking is thus not fully transparent. Automated suggestions and page ranking are highly individualised. One can turn off personalisation, but Google's default setting is to personalise. Google Search is the most used Web application for students and has deeply transformed the way they access and consume information. So, overall, Google Search's ranking on these dimensions suggests its position in a region somewhere in the upper middle of the multidimensional space. Below we summarize the outcome of these case studies in a table in which the ranking on the dimensions and overall degree of integration are presented.

<b>Dimension</b>	<b>Wikipedia</b>	<b>IMDB</b>	<b>Google Search</b>
Information flow	One-way	One-way or two-way	Two-way
Accessibility	High	High	High
Durability	Low-High	Medium	High
Trust	Medium-High	High	High
P-transparency	High	High	High
I-transparency	High	High	Low-Medium
Individualization	None	Medium	High
Transformation	Low-Medium	Low-Medium	High
Overall integration	Low-Medium	Medium	Medium-High

## **6 Concluding reflection**

### **6.1 Human-Web systems**

The three case studies demonstrate that most current agent-Web interactions lack the kind of reciprocity and deep integration that characterize the paradigm cases of extended cognitive systems. Given that most of our current interactions with the Web (e.g., Google Maps, YouTube, news sites, TV guides, timetables, Amazon, the weather forecast, etc.) concern one-way information flow structures, it is hard to see those as part of an extended cognitive system. However, emerging Web-applications accessed on one's mobile computing device potentially rank higher on some key dimensions. Particularly relevant here are applications related to memory. We use our mobile devices to store a lot of personally-relevant information in the cloud and on the Web, including appointments, birthdays, shopping-lists, sketches, annotated documents, to-do lists, notes, reminders, photos, and so on. Such applications rank higher on the intensity of information flow between agent and the Web because they concern two-way information flow structures, namely offloading and intake. They potentially rank higher on the accessibility of information, as they are used on mobile devices. Further, because the user has created the information him or herself, it also ranks higher on the amount of trust a user puts into the information, the degree of transparency-in-use, the ease with which the information can be interpreted, and the amount of personalisation. For these reasons, the information is integrated deeper into the cognitive processes that govern our behaviour and it is therefore easier to see it as

part of an extended cognitive system<sup>7</sup> (Smart, Heersmink & Clowes 2017; Smart, Clowes & Heersmink 2017). Moreover, the computerisation of our lifeworld will continue in the future. Research in computer science, machine learning, and artificial intelligence will result in new applications that will be easier to use, are more personalised, and more computationally powerful. The cognitive functions of such computational artifacts and systems will be integrated deeper into a broader spectrum of our cognitive abilities.

One might ask whether some dimensions are more important than others in case of agent-Web interactions. We can single out five dimensions, though other pragmatic explanatory goals may encourage alternative approaches. Information flow is a key aspect of any situated cognitive system. The closer information flow approaches a reciprocal process, either through interactive cueing or through incremental cognitive self-stimulation, the deeper the external information is integrated. Because Wikipedia and many other Web applications have one-way information flow, they rank relatively low on overall integration. Trust is also important, because information on the Web is written by others. If we naively trust external information we find on blogs, forums, or other websites, we run the risk of using false information. Informational transparency is important because online information is not always transparent as it may not be in a language or format a user understands. Finally, a high degree of accessibility is one of the distinguishing features of the Web and therefore transforms, perhaps more so than other technologies, our memory strategies. This is not to say that durability and procedural transparency are not important - they are, and in other explanatory contexts will come to the fore.

Recall Halpin, Wheeler and Clark's suggestion that online information, "when integrated appropriately into the processes that govern an agent's behaviour, may count as part of that agent's cognitive architecture" (2014, p. 24). In this paper, we have provided an initial analysis of what it means for online information to be integrated appropriately into the processes that govern an agent's behaviour. The proposed dimensional analysis provides a new perspective on the conditions for cognitive extension and cognitive distribution. It is not meant to provide a set of necessary and sufficient conditions, but to provide a conceptual and empirical toolbox for investigating the degree and nature of the integration of agent and artifact into "new systemic wholes". The higher an agent-artifact system ranks on the proposed dimensions, the more functional and cognitive integration occurs, and the more tightly coupled is the system (Heersmink 2015).

We have suggested that there are more fine-grained, more interesting, and more empirically productive questions to ask in this domain than simply 'does the Web extend cognition?'. Firstly, the Web is not a single homogeneous thing or artifact, and our cognitive processes and practices stand in different relations to its different applications or features. By identifying a range of dimensions which vary more or less independently, we open space for empirical investigation into different forms and degrees of integration and cognitive transformation. Secondly, there is no need to seek a straightforward identity or even analogy between distributed agent-artifact systems and socially distributed cognitive systems such as those studied in transactive memory theory: technologically and socially distributed cognitive processes have always both been part of the distributed cognition framework, and typically operate together. So, finally, the empirical work with which philosophers interested in cognition and the Web should engage is not only cognitive psychological experiment under controlled and abstracted conditions, but also ethnographic and sociological study of the rich and tangled interactions involved in our changing Web-related cognitive practices.

## 6.2 Epistemic consequences

In this final section, we draw out some of the epistemological implications of the dimensional framework. More specifically, the dimensions of trust, individualization, and transformation are

---

<sup>7</sup> As a reviewer pointed out, mobile devices can potentially be accessed by others, resulting in privacy issues. For some users, this may lead to a reluctance to use their mobile device as an external memory system, at least when it concerns sensitive information (Smart, Clowes & Heersmink 2017).

relevant for one's epistemic success when interacting with the Web. Regarding trust, a distinction can be made between whether a Web-user *ought* to trust online information and whether online information *is* trusted by a user (Smart, Clowes & Heersmink 2017). The first is a normative concern and the second an empirical one. Unreflexively trusting information on the Web may have detrimental epistemic consequences. Consider a striking example provided by Michael Lynch in his recent book *The Internet of Us*. When you ask a question in Google Search, it generates a "featured snippet" at the top of the search page. This snippet is a summary of the answer taken from a top-ranking Webpage. When Lynch (2016, p. 66) searched for the question "What happened to the dinosaurs?", Google Search generated the following snippet:

"The Bible gives us a framework for explaining dinosaurs in terms of thousands of years of history, including the mystery of when they lived and what happened to them. Dinosaurs are used more than almost anything else to indoctrinate children and adults in the idea of millions of years of earth history.<sup>8</sup>"

A naïve user may trust this information and accept it as a fact, but it is clear that one ought not to trust it, as it is false and misleading. Featured snippets and page ranking more generally can be manipulated relatively easily and should therefore be interpreted with a healthy dose of scepticism.

With most cognitive artifacts (e.g., notebooks, shopping lists, and diaries), individualization is done intentionally, but with the Web, individualization is not always an intentional process. Google, for example, uses algorithms to profile users based on their browsing history and to predict what kind of content they may prefer. People may not be aware that page ranking in their online searches is individualized, which might be undesirable because personalized page ranking leaves out certain Webpages and prioritizes others based on one's search-history, resulting in "filter bubbles" and potentially leading to biased views (Simpson 2012; Miller & Record 2013; compare Smart and Shadbolt 2018). Google Search thus generates personalised epistemic choice architectures that are not necessarily epistemically beneficial. Likewise, Google's autocompleted search terms may nudge one onto an epistemically undesirable path (Miller & Record 2017). When one of us (RH) searched for the term 'vaccines', Google suggests the following phrases 'vaccines cause autism', 'vaccines Australia', 'vaccines revealed' and 'vaccines for Vietnam'. Whilst the last three suggestions could potentially be helpful, the first might put a naïve and uninformed enquirer on an epistemically wrong path.

Some have argued that using the Web transforms our biological memory, in that an overreliance on online information may result in not storing that information in biological memory. This phenomenon is referred to as the "Google effect" (Sparrow et al 2011). Nicolas Carr (2011) uses Sparrow et al.'s research to argue that the Web makes us less knowledgeable. Carr borrows an argumentative strategy from Socrates, who argues that written language makes humans forgetful, due to a lack of practice of biological memory. The idea being: why store information in biological memory when it is available in writing? The Web, according to Carr, is relevantly similar to written language. However, Heersmink (2016) has criticised Carr's alarmist claims by questioning the generalizability of Sparrow et al.'s research to the Web. As outlined above in section 4.2, Sparrow et al.'s experiments did not involve using the Web, they only involved storing statements on a desktop computer. Whether storing statements on a desktop computer is relevantly similar to using the Web in ecologically-valid, real-world situations needs more research. It is, of course, likely that the Web transforms human biological memory systems in a variety of ways, but the currently available research does not support strong negative conclusions of the kind voiced by Carr.

---

<sup>8</sup> This snippet is from <https://answersingenesis.org> and was ranked first for the featured snippet due to search engine optimisation.



There are thus a number of possible epistemic issues when using the Web as a scaffolded memory system. To overcome some of these issues, Heersmink (2018) proposes a virtue-epistemic approach to using the Web, particularly search engines. Virtue responsibilism, one of the two main camps in virtue epistemology, emphasises the role of learned cognitive character traits such as open-mindedness, attentiveness, and intellectual autonomy in obtaining knowledge (Baehr 2011). Such truth-conducive character traits are referred to as intellectual virtues. It is important to be intellectually virtuous because such an agent is more likely to obtain true beliefs, knowledge, and understanding than an agent who is less intellectually virtuous (Zagzebski, 1996). Virtue responsibilism aims to provide a theoretical framework for living an intellectually virtuous life and is therefore a promising approach to optimize our epistemic interactions with the Web and other cognitive technologies. This is not to say that virtue reliabilism, the other main camp in virtue epistemology, is not relevant for better understanding the cognitive relations between epistemic agents and technologies. See, for instance, Kourken Michaelian & Santiago Arango-Muñoz (2018) for an interesting analysis of agent-artifact systems from a virtue reliabilist perspective in terms of memory.

Let us briefly illustrate how virtue responsibilism may be helpful as to improve our epistemic relation to search engines by focussing on three intellectual virtues, namely open-mindedness, attentiveness, and intellectual autonomy. Someone who is open-minded does not dogmatically hold on to one's views but considers several alternative views. If these views are better and more convincing, then the agent is willing to change his or her mind. We have seen that personalised page ranking may generate epistemic choice architectures that are not necessarily truth-conducive. One way in which an open-minded agent can deal with this issue is to turn off personalisation, which requires some effort of the agent but is relatively easy to do in Google's settings. Another way to help deal with this issue is to use a search engine that does not personalise such as, for example, DuckDuckGo (Simpson 2012). Someone who is attentive pays attention and has a sustained focus on the cognitive task. The Web, however, is an informational environment that is not particularly conducive to attentiveness. Rather, it promotes skimming of information, cursory reading, and distracted thinking. Howard Rheingold (2012) suggests a number of mindfulness techniques to train oneself to first become aware of distractedness and then to force oneself to stay focused. Another way to prevent online distractedness is to use software programs that block certain webpages such as social media or email at certain times of the day, thereby delegating attentiveness to a software system. Lastly, an intellectually autonomous agent can think for oneself, is cognitively capable, and has a certain degree of scepticism towards information. Regarding the use of search engines, such an agent will interpret Google's epistemic choice architectures, including featured snippets, page ranking, and autocompleted search terms, with a healthy dose of scepticism. If one's information-seeking behaviours are not epistemically virtuous, which we suspect is the case for many people, one can improve one's online epistemic behaviours through practice, training, and education (Heersmink 2018).

Extended and distributed cognition theory but also transactive memory theory are descriptive, that is, they are primarily concerned with *how* we interact with objects and other people, not so much with how we *ought to* interact with objects and people. Virtue responsibilism aims to improve our epistemic skills and can therefore be synthesised with extended, distributed, and transactive approaches, resulting in a richer and more normatively-informed understanding of the cognitive relations between agents and their social and material environment.

## References

Baehr, J. (2011). *The inquiring mind: On intellectual virtues and virtue epistemology*. Oxford: Oxford University Press.

- Berners-Lee, T. et al (2006). A framework for Web science. *Foundations and Trends in Web Science*, 1(1), 1-130.
- Bolter, J.D. & Grusin, R. (1999). *Remediation: Understanding new media*. Cambridge, MA: MIT Press.
- Carr, N. (2011). *The shallows: What the Internet is doing to our brain*. New York: W.W. Norton & Company.
- Clark, A. (1997). *Being there: Putting brain, body and world back together again*. Cambridge, MA: MIT Press.
- Clark, A. (2001). *Mindware: An introduction to the philosophy of cognitive science*. New York: Oxford University Press.
- Clark, A. (2008). *Supersizing the mind: Embodiment, action, and cognitive extension*. New York: Oxford University Press.
- Clark, A. & Chalmers, D. (1998). The extended mind. *Analysis*, 58(1), 10-23.
- Clowes, R.W. (2015). Thinking in the cloud: The cognitive integration of cloud-based technology. *Philosophy and Technology*, 28(2), 262-296.
- van Deursen, A. & van Dijk, J. (2009). Using the Internet: Skill related problems in users' online behavior. *Interacting with Computers*, 21(5-6), 393-402.
- Donald, M. (1991). *Origins of the modern mind: Three stages in the evolution of culture and cognition*. Cambridge, MA: Harvard University Press.
- Erl, A. & Rigney, A. (Eds.) (2009). *Mediation, remediation, and the dynamics of cultural memory*. Berlin: Walter de Gruyter.
- Fallis, D. (2008). Towards an epistemology of Wikipedia. *Journal of the American Society for Information Science and Technology*, (59)10, 1662-1674.
- Halpin, H., Clark, A. & Wheeler, M. (2014). Towards a philosophy of the web: Representation, enaction, collective intelligence. In A. Monnin & H. Halpin (Eds.), *Philosophical engineering: Toward a philosophy of the Web* (pp. 21-30). Malden, MA: Wiley-Blackwell.
- Hargittai, E. (2010). Digital na(t)ives? Variation in Internet skills and uses among members of the "net generation". *Sociological Inquiry*, 80(1), 92-113.
- Harris, C.B, Keil, P.G., Sutton, J. & Barnier, A.J. (2010). Collaborative remembering: When can remembering with others be beneficial? In W. Christensen, E. Schier & J. Sutton (Eds.), *Proceedings of the 9th Conference of the Australasian Society for Cognitive Science* (pp. 131-134).
- Hayles, N.K. (2012). *How we think: Contemporary media and technogenesis*. Chicago: The University of Chicago Press.
- Head, A.J., & Eisenberg, M.B. (2010). How today's college students use Wikipedia for course-related research. *First Monday*, 15(3). DOI: 10.5210/fm.v15i3.2830.
- Heersmink, R. (2015). Dimensions of integration in embedded and extended cognitive systems. *Phenomenology and the Cognitive Sciences*, 14(3), 577-598.
- Heersmink, R. (2016). The Internet, cognitive enhancement, and the values of cognition. *Minds and Machines*, 26(4), 389-407
- Heersmink, R. (2018). A virtue epistemology of the Internet: Search engines, intellectual virtues, and education. *Social Epistemology*, 32(1), 1-12.
- Heidegger, M. (1962). *Being and time*. London: SCM.
- Henrich, J., Heine, S.J., & Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and Brain Sciences*, 33, 61-135.
- Hoskins, A. (2009). The mediatisation of memory. In J. Garde-Hansen, A. Hoskins, & A. Reading (Eds.), *Save as ... digital memories* (pp. 27-43). London: Palgrave.
- Huebner, B. (2016). Transactive memory reconstructed: Rethinking Wegner's research program. *Southern Journal of Philosophy*, 54(1), 48-69.
- Hutchins, E. (1995). *Cognition in the wild*. Cambridge, MA: MIT Press.
- Kirsh, D. (2013). Embodied cognition and the magical future of interaction design. *ACM Transactions on Human-Computer Interaction*, 20(1), 1-30.

- Lewis, K. & Herndon, B. (2011). Transactive memory systems: Current issues and future directions. *Organization Science*, 22(5), 1254-1265.
- Ludwig, D. (2015). Extended cognition and the explosion of knowledge. *Philosophical Psychology*, 28(3), 355-368.
- Lynch, M. (2016). *The Internet of us: Knowing more and understanding less in the age of big data*. New York: W.W. Norton and Company.
- Michaelian, K. & Sutton, J. (2013). Distributed cognition and memory research: History and current directions. *Review of Philosophy and Psychology*, 4(1), 1-24.
- Michaelian, K., & Arango-Muñoz, S. (2018). Collaborative memory knowledge: A distributed reliabilist perspective. In M. Meade, C. Harris, P. van Bergen, J. Sutton, & A. Barnier (Eds.), *Collaborative remembering: Theories, research, applications* (pp. 231-247). Oxford: Oxford University Press.
- Miller, B. & Record, I. (2013). Justified belief in a digital age: On the epistemic implications of secret Internet technologies. *Episteme*, 10(2), 117-134.
- Miller, B. & Record, I. (2017). Responsible epistemic technologies: A social-epistemological analysis of autocompleted web search. *New Media and Society*, 19(12), 1945-1963.
- Menary, R. (2007). *Cognitive integration: Mind and cognition unbounded*. Basingstoke: Palgrave MacMillan.
- Menary, R. (2010). Dimensions of mind. *Phenomenology and the Cognitive Sciences*, 9(4), 561-578.
- Norman, D. (1998). *The invisible computer*. Cambridge, MA: MIT Press.
- Pan, B., et al. (2007). In Google we trust: Users' decisions on rank, position, and relevance. *Journal of Computer Mediated Communication*, 12(3), 801-823.
- Pew Research Center. (2011). Wikipedia, past and present. Retrieved from: [http://www.pewinternet.org/files/old-media//Files/Reports/2011/PIP\\_Wikipedia.pdf](http://www.pewinternet.org/files/old-media//Files/Reports/2011/PIP_Wikipedia.pdf).
- Pew Research Center. (2015). Technology device ownership: 2015. Retrieved from: <http://www.pewinternet.org/2015/10/29/technology-device-ownership-2015/>.
- Rheingold, H. (2012). *Net smart: How to thrive online*. Cambridge, MA: MIT Press.
- Rowlands, M. (2009). Extended cognition and the mark of the cognitive. *Philosophical Psychology*, 22(1), 1-19.
- Simpson, T.W. (2012). Evaluating Google as an epistemic tool. *Metaphilosophy*, 43(4), 426-445.
- Smart, P. (2012). The Web-extended mind. *Metaphilosophy*, 43(4), 446-663.
- Smart, P. (2017). Extended cognition and the Internet: A review of current issues and controversies. *Philosophy & Technology*, 30(3), 357-390.
- Smart, P., Clowes, R. & Heersmink, R. (2017). Minds online: The interface between Web science, cognitive science, and the philosophy of mind. *Foundations and Trends in Web Science*, 6(1-2), 1-234.
- Smart, P., Heersmink, R. & Clowes, R. (2017). The cognitive ecology of the Internet. In S. Cowley & F. Vallée-Tourangeau (Eds.), *Cognition beyond the brain: Computation, interactivity and human artifice* (pp. 251-282). Dordrecht: Springer.
- Smart, P. and Shadbolt, N. (2018). The World Wide Web. In J. Chase and D. Coady (Eds.), *Routledge Handbook of Applied Epistemology*. New York: Routledge.
- Sparrow, B., Liu, J. & Wegner, D.M. (2011). Google effects on memory: Cognitive consequences of having information at our fingertips. *Science*, 333(6043), 776-778.
- Sterelny, K. (2010). Minds: Extended or scaffolded? *Phenomenology and the Cognitive Sciences*, 9(4), 465-481.
- Sutton, J. (2006). Distributed cognition: Domains and dimensions. *Pragmatics and Cognition*, 14(2), 235-247.
- Sutton, J. (2010). Exograms and interdisciplinarity: History, the extended mind and the civilizing process. In R. Menary (Ed.), *The extended mind* (pp. 189-225). Cambridge, MA: MIT Press.

- Sutton, J. (2015a). Scaffolding memory: Themes, taxonomies, puzzles. In L. Bietti and C.B. Stone (Eds.), *Contextualising human memory: An interdisciplinary approach to how individuals and groups remember their pasts*. London: Routledge.
- Sutton, J. (2015b). Remembering as public practice: Wittgenstein, memory, and distributed cognitive ecologies. In D. Moyal-Sharrock, A. Coliva, & V. Munz (Eds.), *Mind, language, and action* (pp. 409-443). Berlin: Walter de Gruyter.
- Sutton, J., Harris, C.B., Keil, P. & Barnier, A.J. (2010). The psychology of memory, extended cognition, and socially distributed remembering. *Phenomenology and the Cognitive Sciences*, 9(4), 521-560.
- Theiner, G. (2013). Transactive memory systems: A mechanistic analysis of emergent group memory. *Review of Philosophy and Psychology*, 4(1), 65-89.
- Tollefson, D., Dale, R. & Paxton, A. (2013). Alignment, transactive memory, and collective cognitive systems. *Review of Philosophy and Psychology*, 4(1), 49-64.
- Ward, A.F. (2013). Supernormal: How the Internet is changing our memories and our minds. *Psychological Inquiry*, 24(4), 341-348.
- Wegner, D.M. (1986). Transactive memory: A contemporary analysis of the group mind. In B. Mullen & G.R. Goethals (Eds.), *Theories of group behaviour* (pp. 185-208). New York: Springer-Verlag.
- Wegner, D.M. (1995). A computer network model of human transactive memory. *Social Cognition*, 13(3), 319-339.
- Wegner, D.M., Giuliano, T. & Hertel, P.T. (1985). Cognitive interdependence in close relationships. In W.J. Ickes (Ed.), *Compatible and incompatible relationships* (pp. 253-276). New York: Springer-Verlag.
- Wegner, D., Erber, R. & Raymond, R. (1991). Transactive memory in close relationships. *Journal of Personality and Social Psychology*, 61(6), 923-929.
- Wilson, R. & Clark, A. (2009). How to situate cognition: Letting nature take its course. In P. Robbins & M. Aydede (Eds.), *The Cambridge handbook of situated cognition* (pp. 55-77). Cambridge: Cambridge University Press.
- Zagzebski, L. (1996). *Virtues of the mind: An inquiry into the nature of virtue and the ethical foundations of knowledge*. Cambridge: Cambridge University Press.