The World Wide Web

Paul Smart[1] and Nigel Shadbolt[2]
[1]Electronics and Computer Science, University of Southampton, Southampton, SO17 1BJ, UK
[2]Department of Computer Science, University of Oxford, Oxford, OX1 3QD, UK

Introduction

The World Wide Web (henceforth the "Web") is a large-scale digital compendium of information that covers practically every sphere of human interest and endeavour. For this reason, it should come as no surprise to learn that the Web is a prominent target for epistemological analysis. To date, search engines (Heintz, 2006; Miller & Record, 2013; Simpson, 2012), Wikipedia (Coady, 2012; Fallis, 2008, 2011) and the blogosphere (Coady, 2012; Goldman, 2008) have all been the focus of epistemological attention. Other systems, while relevant to epistemology, have attracted somewhat less scrutiny. These include microblogging platforms, (e.g., Twitter), social networking sites (e.g., Facebook), citizen science projects (e.g. Galaxy Zoo), and human computation systems (e.g., Foldit). One of the aims of this chapter is to introduce the reader to these systems and highlight their relevance to applied epistemology. A second aim is to review existing epistemological analyses of the Web, and, where necessary, point out problems with the philosophical narrative. A third and final objective is to highlight areas where the interests of epistemologists (both theoretical and applied) overlap with the interests of those who seek to understand and engineer the Web. One of the outcomes of this analysis is a better understanding of the ways in which contemporary epistemology can contribute to the nascent discipline of Web science (see Smart et al., 2017).

Personalized Search: Epistemic Boon or Burden

One of the major areas of epistemological enquiry into the Web concerns the epistemic impact of search engines, such as Google Search (Miller & Record, 2013; Simpson, 2012). A particular focus of attention relates to the effect of personalized search mechanisms, which filter search results based on a user's prior search activity. Such mechanisms, it is claimed, can result in so-called "filter bubbles" (see Pariser, 2011), which have the effect of limiting a user's awareness of important bodies of epistemically-relevant information. Epistemologists are largely in agreement concerning the negative effects of personalized search. Simpson (2012), for example, argues that filter bubbles accentuate the problem of confirmation bias and undermine users' access to objective information. Similar views are expressed by Miller and Record (2013). They claim that the justificatory status of an agent's beliefs are undermined as a result of exposure to personalized search results.

Concerns about the epistemic sequelae of personalized search have led epistemologists to make a number of practical suggestions as to how to avoid filter bubbles, or at least minimize their epistemic effects. Simpson (2012) thus suggests that users should turn off personalization, or resort to search engines that do not use personalization mechanisms (he cites DuckDuckGo[1] as a prime example). Simpson also suggests that there is a *prima facie* case to be made for government regulation of search engine providers. Echoing the views of Introna and Nissenbaum (2000), he argues that search engines are in the business of providing an important public service and that regulation is required to ensure they operate in an objective and impartial manner.

[1] See https://duckduckgo.com/.

Other proposals to address the problem of personalized search centre on the epistemic responsibilities of Web users. Miller and Record (2013) thus suggest that search engine users "can use existing competencies for gaining information from traditional media such as newspapers to supplement internet-filtered information and therefore at least partly satisfy the responsibility to determine whether it is biased or incomplete" (p. 130).

Finally, Knight (2014) draws attention to the efforts of computer scientists in developing "diversity-aware search" techniques. These are deemed to enable users to break out of their filter bubbles via the active inclusion of 'diverse' information in search results.

We thus have a range of proposals concerning the practical steps that could be (and perhaps should be) taken by users to obviate the negative effects of personalized search. But before we accept such proposals, we should at least question the (largely implicit) assumption upon which all these proposals are based. Do personalized search engines really undermine the epistemic status of their users? And, if so, are we justified in condemning personalized search engines on account of their poor veritistic value? In responding to these questions, we suggest it helps to be aware of a range of issues that, to our knowledge, have not been the focus of previous epistemological attention. While these issues do not exclude the possibility that personalized search may, on occasion, harm the epistemic standing of individual Web users, they do at least provide reasons to question the epistemological consensus that has emerged in this area.

The first issue to consider relates to the way in which search engines are actually used. Waller (2011), for example, discovered that almost half (i.e., 48%) of the queries entered by search engine users appeared to be directed towards the retrieval of information about a specific website. In other words, it seemed that users were relying on a search engine, at least in part, as a means of providing quick and easy access to familiar sources of information. These findings are important, for they suggest that the discovery of new information is not the sole purpose of search engines; instead, it seems that search engines may also be used quickly access sources of information that a Web user is already aware of. When seen in this light, it is far from clear that personalized search mechanisms should be seen as *always* working against the epistemic interests of the human individual. In fact, there is perhaps a risk that by interfering with personalization mechanisms, we will disrupt a set of well-honed techniques for quickly and efficiently accessing familiar bodies of task-relevant information.

This is not to say that Waller's (2011) findings eliminate concerns about the epistemic implications of personalized search. The use of search engines as a convenient shortcut to familiar sources of information may, from an epistemic perspective, be more-or-less hazardous depending on the kind of information that is being accessed, and clearly nothing about this particular way of using search engines guarantees the objectivity or impartiality of the actual information source.[2] In spite of these caveats, Waller's (2011) findings are important because they draw attention to the *different* ways in which Web users may exploit the functionality of personalized search engines. One question for future research is to ascertain whether all these modes of use are equally injurious to an individual's epistemic health and standing, and whether some modes may actually be of productive value in enhancing an individual's epistemic functioning.

[2] We are grateful to an anonymous reviewer for highlighting this particular point.

A second issue to consider relates to the broader ecological setting in which search engines are used. Here we suggest that epistemological analyses can benefit from the sort of perspectives that have long been embraced by the cognitive science community, especially those that emphasize the situated and environmentally-embedded nature of cognitive processing (Robbins & Aydede, 2009). In particular, we suggest that it is helpful to think of Web users as embedded in multiple networks of information flow and influence, each of which presents the user with a diverse (even if filtered) stream of facts, ideas and opinions. This broader informational ecology, we suggest, might work to mitigate the negative epistemic effects of personalized search (if indeed there are any). The sociological concept of *network individualism* (Rainie & Wellman, 2012) may be potential value here. Networked individualism refers to the way in which society is changing as a result of the introduction of new media technology. In particular, it emphasizes the manner in which people connect, communicate and exchange information following the advent of the Web and the growth of mobile communications technology. According to Rainie and Wellman (2012), society is increasingly organized along the lines of multiple, overlapping social networks, each of which is characterized by fluid and dynamic membership. As a result of these shifts in social structure, individuals are likely to be exposed to multiple sources of heterogeneous information, and this may help to allay concerns about the selective exposure effects that filter bubbles are deemed to produce. A consideration of network individualism thus reminds us that a user's informational ecology is not necessarily exhausted by the nature of their interaction with a particular search engine. Once this broader informational ecology is taken into consideration, concerns about the epistemological impact of personalized search may appear somewhat overblown.

Finally, a more positive perspective on personalized search is provided by the notion of *mandevillian intelligence* (Smart, in press-b, in press-c). Mandevillian intelligence is a specific form of collective intelligence in which the cognitive shortcomings and epistemic vices of the individual agent are seen to yield cognitive benefits and epistemic virtues at the collective or social level of analysis; e.g., at the of level of collective doxastic agents (see Palermos, 2015) or socio-epistemic systems (see Goldman, 2011). According to this idea, personalized search systems may play a productive role in serving the *collective* cognitive good, providing a means by which an individual vices (e.g., a tendency for confirmation bias) are translated into something that more closely resembles an epistemic virtue (e.g., greater cognitive coverage of a complex space of thoughts, ideas, opinions, and so on). Consider, for example, the way in which personalized search may help to focus individual attention on particular bodies of information, thereby restricting access to a larger space of ideas, opinions and other information. While such forms of 'restricted access' or 'selective information exposure' are unlikely to yield much in the way of an epistemic benefit for the individual agent, it is possible that by exploiting (and, indeed, accentuating!) an existing cognitive bias (e.g., confirmation bias), personalized search technologies play an important role in maintain cognitive diversity, helping to prevent precipitant forms of cognitive convergence (see Zollman, 2010) and supporting an effective division of cognitive labour (see Muldoon, 2013). This possibility reveals something of a tension in how we interpret or evaluate the veritistic value of a particular technology or epistemic practice. In particular, it seems that assessments of veritistic value vary according to whether our epistemological attention is directed to the level of individual epistemic agents or the collective ensembles (e.g., socio-epistemic systems) in which those individuals are situated.

Needless to say, much more work needs to be done to evaluate these claims about the potential epistemic benefits of personalized search (as well as perhaps other forms of online information

filtering). Note, however, that in the absence of the notion of mandevillian intelligence, the epistemic consequences of personalized search might have seemed self-evident and thus unworthy of further scientific and philosophical scrutiny. This helps to highlight one of the ways in which the notion of mandevillian intelligence is relevant to applied epistemology: it helps to provide the conceptual basis for novel investigative efforts that seek to explore the epistemic consequences of (e.g.) technological interventions at *both* the individual and collective (social) levels.

The of mandevillian intelligence is also relevant to a number of existing epistemological debates and discussions, many of which are of an applied nature. Consider, for example, the way in which major technology vendors, such as Facebook and Google, have been the subject of recent criticism by prominent political leaders for failing to provide details of their information filtering algorithms.[3] From the standpoint of mandevillian intelligence, such critiques are revealed as epistemologically inadequate, in the sense that they fail to consider the way in which the collective epistemic good may be served by the technological accentuation of individual forms of epistemic vice — in essence, mainstream views appear to trade the epistemic standing of the one (i.e., the individual agent) against the epistemic capabilities of the many (e.g., a particular society). In this respect, the notion of mandevillian intelligence may be of practical significance in terms of its ability to shape Web-related legislative programs and social policy agendas.

Web-Extended Knowledge

One of the ways in which epistemologists have sought to understand the epistemic effects of the Web is by drawing on externalist approaches to mind and cognition (see Clark, 2008). According to the notion of *active externalism*, for example, the causally-active physical vehicles that realize mental states and processes can sometimes extend beyond the traditional biological borders of the brain (and body) to include a variety of non-biological (i.e., extra-organismic) resources (Clark & Chalmers, 1998). This idea is sometimes presented as a thesis about the explanatory kinds of interest to cognitive science (in which case it is commonly referred to as the Hypothesis of Extended Cognition or HEC), and sometimes it is presented as a thesis about mentalistic folk categories, such as states of belief (in which case it is commonly referred to as the Extended Mind Thesis or EMT).

With respect to the EMT, it has been suggested that the nature of our interaction with the Web supports the emergence of *Web-extended minds*; i.e., forms of bio-technological merger in which the Web serves as part of the realization base for some of our folk psychological mental states, most notably states of dispositional belief (Smart, 2012). This possibility has been discussed in relation to a number of criteria that have been discussed in relation to the EMT (see Clark & Chalmers, 1998). Smart (2012) thus talks about the Web in terms of the opportunities it provides for quick and easy access to online information and the way in which these opportunities speak to at least one of the criteria for cognitive extension discussed by Clark and Chalmers (1998), namely the accessibility criterion.

Recently, the notion of Web-extended minds has led epistemologists to make a number of claims about the impact of the Web on our epistemic profiles. One implication of the Web-extended mind concept, Ludwig (2015) argues, is that we are able to envisage a profound transformation of our doxastic potential. In particular, Ludwig anticipates “an explosion of dispositional beliefs and

[3] See http://www.bbc.co.uk/news/technology-37798762.

knowledge that is caused by digital information resources such as Wikipedia or Google" (p. 355). Similar views are expressed by Bjerring and Pedersen (2014). They argue that the Web enables us to enjoy various forms of "restricted omniscience," wherein we have more-or-less "complete knowledge about a particular, fairly specific subject matter" (p. 25). We thus arrive at a claim that seems to follow quite naturally from the possibility of Web-extended minds — a claim that is nicely captured by the following hypothesis:

> **Web-Extended Knowledge Hypothesis**
>
> Cognitively-potent forms of bio-technological union between human agents and the Web serve as the basis *Web-extended knowledge*, i.e., epistemically-relevant doxastic states that supervene on material elements forming part of the technological and informational fabric of the Web.

Unfortunately, there are a number of problems confronting this hypothesis. One of the most pressing problems relates to the way in which the criteria for cognitive extension (e.g., those proposed by Clark and Chalmers) work against the epistemic interests of the technologically-extended agent (see Smart, in press-a). In order to help us understand this, consider the accessibility criterion, as discussed by Clark and Chalmers (1998). The general idea behind the accessibility criterion is that external information should be quickly and easily accessible — it should be possible for agents to draw on external information whenever it is required and easily incorporate this information in their cognitive processing routines. Accessibility thus seems to demand a degree of *fluency* with respect to the kind of interaction an agent has with bio-external resources, where the notion of fluency can be understood (at least in part) as the "subjective experience of ease or difficulty with which we are able to process information" (Oppenheimer, 2008, p. 237).

Now, the problem with claims regarding easy access and fluent interaction is that these properties seem to be in some tension with the possibility of Web-extended forms of knowledge. One of the key insights to emerge from research on fluency, for example, is that fluent processing is often associated with a "truth bias," in which the 'truth' of some body of external information is judged relative to the subjective ease with which it is processed (Alter & Oppenheimer, 2009). In the context of the Web, where information is of variable reliability, this particular kind of cognitive bias looks set to undermine the epistemic integrity of the Web-extended cognizer. Indeed, it seems reasonable to think that, in the interests of preserving positive epistemic standing, Web users should be somewhat circumspect about online information. At the very least, it seems important for the epistemically responsible agent to subject online information to critical evaluation (Heersmink, in press; Record & Miller, in press). But now note how this seemingly sensible demand for critical evaluation conflicts with the putative role that fluency plays in extending the epistemic reach of the Web-extended cognizer. Fluency thus seems to speak in favour of the possibility of Web-extended minds, but it seems to work against the interests of Web-extended knowers (i.e., agents whose epistemic credentials are *enhanced* as a result of Web-based cognitive extension). We thus encounter the following problem (see Smart, in press-a):

> **Extended Cognizer vs. Extended Knower Problem**
>
> The properties that work to ensure that an external resource can be treated as a candidate for cognitive incorporation are also, at least in some cases, the very same properties that work to undermine or endanger the positive epistemic standing of the technologically-extended agent.

Somewhat surprisingly, this problem highlights a potential tension between our notions of extended cognition and extended knowledge. Contrary to the idea that the Web-extended minds are the natural harbingers of Web-extended knowledge, cognitive extension may lead to a form of epistemic diminishment, undermining the extent to which extended cognizers are the proper targets of knowledge attribution.

Epistemic Feelings

In addition to ideas concerning Web-extended knowledge and Web-extended knowers, there is an additional way in which active externalism is relevant to epistemological analyses of the Web. This is revealed by the results of recent empirical studies investigating the effect of Web access on subjective, epistemically-relevant experiences, such as the *feeling of knowing* (Fisher et al., 2015; Ward, 2013). The feeling of knowing is one of a range of epistemic feelings that have been studied by epistemologists (Michaelian & Arango-Muñoz, 2014). It refers to the experience of being able to retrieve or access some piece of information (e.g., the answer to a specific question), typically from bio-memory. In situations where people use the Web to search for online information, however, it seems that this feeling of knowing 'extends' to include the informational contents of the online realm. Searching for information, Fisher et al. (2015) suggest, "leads people to conflate information that can be found online with knowledge in the head" (p. 675). Similarly, Ward (2013) notes that as people turn to the "cloud mind of the internet, they seem to lose sight of where their own minds end and the mind of the internet begins. They become one with the cloud, believing that they themselves are spectacularly adept at thinking about, remembering, and locating information" (p. 88).

These findings are of interest, because they have long been anticipated by those working in the active externalist camp. In 2007, for example, Clark proposed that our subjective sense of what we know is informed by the kind of access we have to bio-external information:

> Easy access to specific bodies of information, as and when such access is normally required, is all it takes for us to factor such knowledge in as part of the bundle of skills and abilities that we take for granted in our day to day life. And it is this bundle of taken-for-granted skills, knowledge, and abilities that...quite properly structures and informs our sense of who we are, what we know, and what we can do. (Clark, 2007, p. 106)

Such comments seem particularly prescient in view of the findings by Ward (2013) and Fisher et al. (2015). Indeed, from the standpoint of active externalism, it might be thought that the results of Ward (2013) and Fisher et al. (2015) are largely consistent with the idea of online information being incorporated into an individual's body of personal beliefs and (perhaps) knowledge. Inasmuch as we accept this to be the case, then it potentially alters our views about the significance of Web-induced changes in the feeling of knowing. The aforementioned quotes from Ward (2013) and Fisher et al. (2015) both sound something of a cautionary note regarding the extent to which changes in the feeling of knowing should be seen as marking a genuine shift in an individual's epistemic and cognitive capabilities — at the very least, such comments appeal to a distinction between what might be called 'knowledge-in-the-head' and 'knowledge-on-the-Web'. Extended approaches to cognition and knowledge encourage us to question this distinction. From an active externalist perspective, it is entirely possible that changes in epistemic feelings are merely the subjective

corollary of a particular form of cognitive extension, one that accompanies particular kinds of interaction with the informational contents of the online realm.

It goes without saying, of course, that feelings of knowing are not sufficient for genuine knowledge attribution — we may feel we know lots of things without actually knowing anything! It is thus important to note that while the work of Ward (2013) and Fisher et al. (2015) might be seen to support claims about the Web-extended mind, this does not necessarily tell us anything about Web-based forms of extended knowledge. From the perspective of applied epistemology, it will be important, in future work, to consider the extent to which Web-based shifts in epistemic feelings provide a reliable indication of what we do and do not know. It will also be important to consider the extent to which changes in subjective experience alter our tendency to engage in epistemically-relevant processes and practices (e.g., those that help to ensure the modal stability of our beliefs across close possible worlds). Interestingly, research in social psychology suggests that changes in self-related perceptions of expertise contribute to a more closed-minded or dogmatic cognitive style (Ottati et al., 2015). In view of such results, it is natural to wonder whether changes in feelings of knowing (such as those accompanying the use of Web technologies) might lead individuals to become more dogmatic and thus diminish their epistemic standing under a virtue-theoretic (especially, a virtue responsibilist) conception of knowledge (see Baehr, 2012).

Social Machines

Despite the fact that the Web is a relatively recent phenomenon, it plays a crucial role in an ever-expanding array of social processes. Indeed, the sudden disappearance of the Web would, in all likelihood, result in a severe disruption of society, on a par perhaps with that resulting from a coordinated nuclear strike. (This is somewhat ironic given that the Web emerged on the back of research efforts to support the continued functioning of society in the face of a nuclear attack!) For this reason, it is appropriate to think of the Web as a form of critical infrastructure for society, resembling, perhaps, the more traditional elements of national infrastructure, such as the road, rail and electricity distribution networks. Arguably, the reason why the Web has emerged to occupy this role is because of its ever-more intimate integration into practically every aspect of social life. For better or worse, the Web has now become an integral part of the structures and processes that make our society what it is — part of the integrated physical fabric that makes our society materially possible.

This vision of socio-technical integration lies at the heart of an important concept that has emerged in the context of the Web science literature. This is the concept of *social machines* (Palermos, in press; Smart & Shadbolt, 2014). Social machines are systems in which human and (Web-based) machine elements are jointly involved in the mechanistic realization of phenomena that subtend the computational, cognitive and social domains (Smart & Shadbolt, 2014). From an epistemological perspective, a particular category of social machines are of particular interest. These are known as *knowledge machines* (Smart et al., 2017). A knowledge machine is a social machines that participates in some form of knowledge-relevant process, such as the process of knowledge acquisition, knowledge discovery, knowledge representation, and so on. Citizen science systems, such as Galaxy Zoo (Lintott et al., 2008), are one kind of knowledge machine that has been the focus of considerable attention. These have grown in prominence over recent years, to the point where they play an important role in many forms of scientific practice (see Meyer & Schroeder, 2015). Such characterizations are sufficient to make citizen science systems worthy of applied epistemological

analysis, and this is especially so given the interest in applying epistemological theory to the understanding and analysis of scientific processes (e.g., Palermos, 2015).

Another important class of knowledge machines are *human computation systems* (Law & von Ahn, 2011), which seek to incorporate human agents into some of form computational processing. One example of such a system is the online protein folding game, Foldit (Cooper et al., 2010). This system incorporates the pattern matching and spatial reasoning abilities of human participants into a hybrid computational process that aims to predict the structural properties of protein molecules. The role of the human participants, in these sorts of systems, should not be underestimated. In many cases, the task being performed by the larger socio-technical ensemble — the one involving both human and machine elements — is not one that could be (easily) performed in the absence of the (often large-scale) socio-technical infrastructure that social machines make available. This is something that is often explicitly recognized by those who seek to harness the epistemic potential of social machines. In one of the early papers published in respect of the Foldit system, for example, the authors of the paper explicitly acknowledge the contributions made by more than 57,000 users of the Foldit system (Cooper et al., 2010).

One of the things that is revealed by a consideration of citizen science and human computation systems is the extent to which social machines draw on the complementary contributions of both human and machine elements. Human agents are thus the locus of particular kinds of capability that subtend the epistemic, cognitive, perceptual, behavioural, social, moral, emotional, affective and aesthetic domains; computing technologies, in contrast, are renowned for their speed of processing, their ability to engage in repetitive symbolic manipulation, their capacity for digital data storage, and so on. By bringing these diverse capabilities together in the context of a complex task, social machines are potentially poised to tackle problems that currently lie beyond the cognitive and epistemic reach of our species (see Hendler & Berners-Lee, 2010).

Network Epistemology

One of the goals of the social machine research effort is to gain a better understanding of the forces and factors that influence the performance profile of social machines relative to the kinds of tasks in which they are involved. In the case of knowledge machines, for example, scientists are interested in understanding how different organizational schemes (characterized as the pattern of information flow and influence between human and technological elements) affect the quality of specific epistemic products, such as the reliability of propositional statements.

It is here that we encounter a potentially productive point of contact between the scientific goals of social machine researchers and the philosophical concerns of the epistemological community. Goldman (2011), for example, identifies a specific form of social epistemology, called *systems-oriented social epistemology*, whose primary objective is to understand the veritistic value of different kinds of socially-distributed epistemic practice and social organization. This, it should be clear, is very much in accord with the goals of social machine researchers. It is also something that is well-aligned with a body of work in network science that seeks to illuminate the ways in which the topological structure of social networks influences the dynamics of belief formation and collective cognitive processing within a community of interacting cognitive agents (Glinton et al., 2010; Kearns, 2012). Such forms of *network epistemology* (see Zollman, 2013) (or, more generically, *computational social epistemology*) promise to inform our understanding of the complex interactions that occur

between forces and factors at a variety of levels (e.g., the cognitive, the social and the technological), as well as the ways in which these interactions influence the epistemic properties (e.g., truth tracking capabilities) of individual agents and socio-epistemic organizations.

An Epistemically Safe Environment?

One issue that typically arises in debates about the epistemic impact of the Web concerns the extent to which the Web can serve as a source of reliable information. At first sight, it would seem that the open and democratic nature of the Web (i.e., the fact that pretty much anyone can participate in the creation of online content) poses a problem for claims about the reliability of online content. The problem, of course, is that by enabling every Tom, Dick or Harry to add or edit content we run the risk of contaminating the online environment with misleading and inaccurate information. In the face of such epistemic risks and hazards, is there any reason to think that the Web is apt to serve the epistemic interests of our doxastic systems?

One response to this question involves an appeal to sorts of social participation that are enabled by the Web. Of particular interest is the scale of Web-based social participation — the fact that many hundreds or thousands (and sometimes millions) of individuals are involved in the creation and curation of specific bodies of online information (consider, for example, the number of people who have contributed to the Wikipedia system). Interestingly, large-scale social participation may help to be relevant to some of the concerns that have raised in respect of the reliability of online content. To help us see this, consider Google's PageRank algorithm (Brin & Page, 1998), which is used to support the ranking of Web search results. Part of the reason the PageRank algorithm works is because it exploits the linking behaviour of human users on a *global* scale, and this helps to ensure that the efforts of a 'few' malign individuals will be swamped by the efforts of the more virtuously-minded masses. Similar kinds of approach are used by a variety of Web-based systems. When it comes to human computation or citizen science systems, for example, contributions are typically solicited from multiple independent users as a means of improving the reliability of specific epistemic outputs.

One thing that is worth noting here is that these sorts of *reliability mechanism* are, in many cases, very difficult to sabotage. When it comes to Google Search, for example, any attempt to 'artificially' elevate the ranking assigned to specific contributions (e.g., a user's website) is offset by the globally-distributed nature of the linking effort, coupled with the fact that links to a specific resource are themselves weighted by the ranking of the resource from which the link originates. In view of such safeguards, it is difficult for individual agents to 'artificially' recreate the sort of endorsement that is reflected in the results of the PageRank algorithm. At the very least, it is difficult to see how such endorsement could be manufactured in the absence of a large-scale, socially-coordinated effort.

Reliability Indicators and Trust

Trust is topic that lies at the intersection of both Web science (Golbeck, 2006) and contemporary epistemology. From an epistemological perspective, issues of trust are typically discussed in relation to what is dubbed *testimonial knowledge* (Lackey, 2011), i.e., the knowledge communicated by other individuals. The fact that so much of our knowledge is based on the testimony of others raises questions about the extent to which we are justified in believing what others tells us. In the context of face-to-face encounters, of course, there are a variety of cues — or, in the terminology of Craig (1990), "indicator properties" — that are apt influence our judgements as to who is a trustworthy

informant (see Sperber et al., 2010). Such cues are likely to play an important role in influencing decisions as to who we select as a source of information, as well as the extent to which we endorse the information provided by a particular source (e.g., information may be rejected if an informant shows signs of dishonesty or incompetence while communicating information). By being responsive to such cues, it seems that we are able to exercise considerable ability with respect to the epistemically-virtuous selection and endorsement of sources of testimonial information, a claim that is broadly consistent with the tenets of virtue reliabilistic approaches to knowledge (Greco, 2007).[4]

It is easy to see how the judicious exploitation of reliability-relevant cues could serve as the foundation for testimonial knowledge in face-to-face encounters. But are such strategies relevant to the processing of information derived from the online realm. Do we, in other words, encounter cues on the Web that could be used to assess the reliability of information? And are these cues, if they exist, actually used to judge the trustworthiness or credibility of particular information sources?

There are a number of strands of Web science research that speak to these issues. In terms of the information that is presented on a typical website, studies have revealed that user credibility judgements tend to be influenced by relatively superficial features, such as a website's appearance, structure and navigability (Fogg et al., 2003; Metzger, 2007; Wathen & Burkell, 2002). Other studies have explored the relevance of social cues to credibility assessments. Westerman et al. (2012), for example, investigated the relationship between perceptions of source credibility in the context of the Twitter microblogging system. Their results revealed the presence of a curvilinear relationship, with too many or too few Twitter followers having a negative impact on perceptions of expertise and trustworthiness.

The problem with the sorts of cues investigated by these studies (e.g., site design features, number of Twitter followers, etc.) is that they are relatively easy to 'fake'. Site designs are easily modified, and fake Twitter accounts (and thus non-existent followers) are relatively easy to manufacture. Ideally, what is required is a set of cues that provide something akin to an *honest signal* in evolutionary theory (see Pentland, 2008). In other words, an important property of an online reliability indicator is that it *reliably* indicates the reliability of online content. A crucial question, therefore, is whether the Web provides access to these particular kinds of ('honest') reliability indicators.

In fact, such indicators are available, and as with almost everything on the Web, they rely heavily on the fact that the Web is as much a social environment as it is a technological one. Some examples of such indicators, as identified by Taraborelli (2008), include the following:

1. implicit indicators of individual endorsement (such as indicators that a specific user selected/visited/purchased an item);
2. explicit indicators of individual endorsement (such as explicit ratings produced by specific users);
3. implicit indicators of socially aggregated endorsement (such as density of bookmarks or comments per item in social bookmarking systems);

[4] By being the reliable receivers of testimony, for example, our cognitive abilities play an important role in explaining why it is that we believe the truth in testimonial exchanges.

4. explicit indicators of socially aggregated endorsement (such as average ratings extracted from a user community);
5. algorithmic endorsement indicators (such as PageRank and similar usage independent ranking algorithms);
6. hybrid endorsement indicators (such as interestingness indicators in Flickr, taking into account both explicit user endorsement and usage-independent metrics).

All of these indicators, it should be clear, are ones that rely, to a greater or lesser extent, on the behaviour of other users. They are, as such, reminiscent of work that seeks to investigate the phenomenon of *social proof* (Cialdini, 2007). As work in this area suggests, a tendency to rely on the actions of others does not always yield positive results; sometimes it can lead to herd behaviour, and in an epistemic context, there is a risk that users will erroneously equate popularity with reliability. Nevertheless, there are, it seems, a rich variety of cues that users can exploit as part of the epistemically-virtuous selection and endorsement of online content, and these cues do seem to play an important role in guiding users' actual judgements as to the credibility of online content (see Metzger et al., 2010). It thus seems that rather than being an environment that is deficient or impoverished with respect to the availability of reliability-indicating cues, the Web may, in fact, afford access to cues that are both more varied and perhaps more reliable than those encountered in face-to-face testimonial contexts. Key issues for future research in this area concern the extent to which features of the online socio-technical environment can be used to support the construction, evaluation and validation of epistemically-relevant indicator properties. It will also be important to assess the extent to which socially-constructed indicator properties are immune to the various forms of epistemic injustice that have been discussed in the epistemological literature (see Fricker, 2003).

Conclusion

The Web provides access to a digital compendium of information that is unprecedented in terms of its scale, scope and accessibility. It is, in addition, a resource that plays an ever-greater role in shaping our epistemic capabilities at both an individual and collective level. The Web is, as such, a valuable form of epistemic infrastructure for our species, influencing the kinds of beliefs we form and providing a platform for us to discover, manage and exploit epistemic resources. As a discipline whose primary focus is to understand the factors that influence our epistemic capabilities, applied epistemology establishes a natural point of contact with contemporary Web science, helping to reveal the Web's epistemic properties and informing the search for interventions that maximize its epistemic power and potential.

Acknowledgements

This work is supported under SOCIAM: The Theory and Practice of Social Machines. The SOCIAM Project is funded by the UK Engineering and Physical Sciences Research Council (EPSRC) under grant number EP/J017728/1 and comprises the Universities of Southampton, Oxford and Edinburgh.

References

Alter, A. L., & Oppenheimer, D. M. (2009) Uniting the tribes of fluency to form a metacognitive nation. *Personality and Social Psychology Review*, 13(3), 219-235.

Baehr, J. (2012) *The Inquiring Mind: On Intellectual Virtues and Virtue Epistemology*. Oxford University Press, Oxford, UK.

Bjerring, J. C., & Pedersen, N. J. L. L. (2014) All the (many, many) things we know: Extended knowledge. *Philosophical Issues*, 24(1), 24-38.

Brin, S., & Page, L. (1998) *The anatomy of a large-scale hypertextual search engine*. 7th Annual Conference of the World Wide Web, Brisbane, Australia.

Cialdini, R. B. (2007) *Influence: The Psychology of Persuasion*. HarperCollins, New York, New York, USA.

Clark, A. (2007) Soft Selves and Ecological Control. In D. Ross, D. Spurrett, H. Kincaid & G. L. Stephens (Eds.), *Distributed Cognition and the Will: Individual Volition and Social Context*. MIT Press, Cambridge, Massachusetts, USA.

Clark, A. (2008) *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford University Press, New York, New York, USA.

Clark, A., & Chalmers, D. (1998) The Extended Mind. *Analysis*, 58(1), 7-19.

Coady, D. (2012) *What to Believe Now: Applying Epistemology to Contemporary Issues*. Blackwell, Oxford, UK.

Cooper, S., Khatib, F., Treuille, A., Barbero, J., Lee, J., Beenen, M., Leaver-Fay, A., Baker, D., Popović, Z., & Foldit Players. (2010) Predicting protein structures with a multiplayer online game. *Nature*, 466(7307), 756-760.

Craig, E. (1990) *Knowledge and the State of Nature*. Oxford University Press, Oxford, UK.

Fallis, D. (2008) Toward an epistemology of Wikipedia. *Journal of the American Society for Information Science and Technology*, 59(10), 1662-1674.

Fallis, D. (2011) Wikipistemology. In A. I. Goldman & D. Whitcomb (Eds.), *Social Epistemology: Essential Readings*. Oxford University Press, New York, New York, USA.

Fisher, M., Goddu, M. K., & Keil, F. C. (2015) Searching for Explanations: How the Internet Inflates Estimates of Internal Knowledge. *Journal of Experimental Psychology: General*, 144(3), 674-687.

Fogg, B., Soohoo, C., Danielson, D. R., Marable, L., Stanford, J., & Tauber, E. R. (2003) *How do users evaluate the credibility of Web sites? A study with over 2,500 participants*. Conference on Designing for User Experiences, San Francisco, California, USA.

Fricker, M. (2003) Epistemic justice and a role for virtue in the politics of knowing. *Metaphilosophy*, 34(1-2), 154-173.

Glinton, R., Paruchuri, P., Scerri, P., & Sycara, K. (2010) Self-Organized Criticality of Belief Propagation in Large Heterogeneous Teams. In M. J. Hirsch, P. M. Pardalos & R. Murphey (Eds.), *Dynamics of Information Systems: Theory and Applications*. Springer, Berlin, Germany.

Golbeck, J. (2006) Trust on the World Wide Web: A survey. *Foundation and Trends in Web Science*, 1(2), 131-197.

Goldman, A. I. (2008) The Social Epistemology of Blogging. In J. van den Hoven & J. Weckert (Eds.), *Information Technology and Moral Philosophy*. Cambridge University Press, New York, New York, USA.

Goldman, A. I. (2011) A Guide to Social Epistemology. In A. I. Goldman & D. Whitcomb (Eds.), *Social Epistemology: Essential Readings*. Oxford University Press, New York, New York, USA.

Greco, J. (2007) The nature of ability and the purpose of knowledge. *Philosophical Issues*, 17(1), 57-69.

Heersmink, R. (in press) A virtue epistemology of the Internet: Search engines, intellectual virtues, and education. *Social Epistemology*.

Heintz, C. (2006) Web search engines and distributed assessment systems. *Pragmatics & Cognition*, 14(2), 387-409.

Hendler, J., & Berners-Lee, T. (2010) From the Semantic Web to social machines: A research challenge for AI on the World Wide Web. *Artificial Intelligence*, 174, 156-161.

Introna, L., & Nissenbaum, H. (2000) Defining the Web: The politics of search engines. *Computer*, 33(1), 54-62.

Kearns, M. (2012) Experiments in social computation. *Communications of the ACM*, 55(10), 56-67.

Knight, S. (2014) Finding knowledge - what is it to "know" when we search? In R. König & M. Rasch (Eds.), *Society of the Query Reader: Reflections on Web Search*. Institute of Network Cultures, Amsterdam, The Netherlands.

Lackey, J. (2011) Testimony: Acquiring Knowledge from Others. In A. I. Goldman & D. Whitcomb (Eds.), *Social Epistemology: Essential Readings*. Oxford University Press, New York, New York, USA.

Law, E., & von Ahn, L. (2011) Human Computation. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 5(3), 1-121.

Lintott, C. J., Schawinski, K., Slosar, A., Land, K., Bamford, S., Thomas, D., Raddick, M. J., Nichol, R. C., Szalay, A., Andreescu, D., Murray, P., & van den Berg, J. (2008) Galaxy Zoo: Morphologies derived from visual inspection of galaxies from the Sloan Digital Sky Survey. *Monthly Notices of the Royal Astronomical Society*, 389(3), 1179-1189.

Ludwig, D. (2015) Extended cognition and the explosion of knowledge. *Philosophical Psychology*, 28(3), 355-368.

Metzger, M. J. (2007) Making sense of credibility on the Web: Models for evaluating online information and recommendations for future research. *Journal of the American Society for Information Science and Technology*, 58(13), 2078-2091.

Metzger, M. J., Flanagin, A. J., & Medders, R. B. (2010) Social and heuristic approaches to credibility evaluation online. *Journal of Communication*, 60(3), 413-439.

Meyer, E. T., & Schroeder, R. (2015) *Knowledge Machines: Digital Transformations of the Sciences and Humanities*. MIT Press, Cambridge, Massachusetts, USA.

Michaelian, K., & Arango-Muñoz, S. (2014) Epistemic feelings, epistemic emotions: Review and introduction to the focus section. *Philosophical Inquiries*, 2(1), 97-122.

Miller, B., & Record, I. (2013) Justified belief in a digital age: On the epistemic implications of secret Internet technologies. *Episteme*, 10(02), 117-134.

Muldoon, R. (2013) Diversity and the division of cognitive labor. *Philosophy Compass*, 8(2), 117-125.

Oppenheimer, D. M. (2008) The secret life of fluency. *Trends In Cognitive Sciences*, 12(6), 237-241.

Ottati, V., Price, E. D., Wilson, C., & Sumaktoyo, N. (2015) When self-perceptions of expertise increase closed-minded cognition: The earned dogmatism effect. *Journal of Experimental Social Psychology*, 61, 131-138.

Palermos, S. O. (2015) Active externalism, virtue reliabilism and scientific knowledge. *Synthese*, 192(9), 2955-2986.

Palermos, S. O. (in press) Social Machines: A Philosophical Engineering. *Phenomenology and the Cognitive Sciences*.

Pariser, E. (2011) *The Filter Bubble: What the Internet is Hiding from You*. Penguin, London, England, UK.

Pentland, A. (2008) *Honest Signals: How They Shape Our World*. MIT Press, Cambridge, Massachusetts, USA.

Rainie, L., & Wellman, B. (2012) *Networked: The New Social Operating System*. MIT Press, Cambridge, Massachusetts, USA.

Record, I., & Miller, B. (in press) Taking iPhone Seriously: Epistemic Technologies and the Extended Mind. In A. J. Carter, A. Clark, J. Kallestrup, O. S. Palermos & D. Pritchard (Eds.), *Extended Epistemology*. Oxford University Press, Oxford, UK.

Robbins, P., & Aydede, M. (Eds.) (2009) *The Cambridge Handbook of Situated Cognition*. Cambridge University Press, New York, New York, USA.

Simpson, T. W. (2012) Evaluating Google as an epistemic tool. *Metaphilosophy*, 43(4), 426-445.

Smart, P. R. (2012) The Web-Extended Mind. *Metaphilosophy*, 43(4), 426-445.

Smart, P. R. (in press-a) Emerging Digital Technologies: Implications for Extended Conceptions of Cognition and Knowledge. In A. J. Carter, A. Clark, J. Kallestrup, O. S. Palermos & D. Pritchard (Eds.), *Extended Epistemology*. Oxford University Press, Oxford, UK.

Smart, P. R. (in press-b) Mandevillian intelligence. *Synthese*.

Smart, P. R. (in press-c) Mandevillian Intelligence: From Individual Vice to Collective Virtue. In A. J. Carter, A. Clark, J. Kallestrup, O. S. Palermos & D. Pritchard (Eds.), *Socially-Extended Epistemology*. Oxford University Press, Oxford, UK.

Smart, P. R., 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. R., & Shadbolt, N. R. (2014) Social Machines. In M. Khosrow-Pour (Ed.), *Encyclopedia of Information Science and Technology*. IGI Global, Hershey, Pennsylvania, USA.

Sperber, D., Clément, F., Heintz, C., Mascaro, O., Mercier, H., Origgi, G., & Wilson, D. (2010) Epistemic vigilance. *Mind & Language*, 25(4), 359-393.

Taraborelli, D. (2008) How the Web is changing the way we trust. In K. Waelbers, A. Briggle & P. Brey (Eds.), *Current Issues in Computing and Philosophy*. IOS Press, Amsterdam, The Netherlands.

Waller, V. (2011) Not just information: Who searches for what on the search engine Google? *Journal of the American Society for Information Science and Technology*, 62(4), 761-775.

Ward, A. F. (2013) *One with the Cloud: Why People Mistake the Internet's Knowledge for Their Own.* Harvard University, Cambridge, Massachusetts, USA.

Wathen, C. N., & Burkell, J. (2002) Believe it or not: Factors influencing credibility on the Web. *Journal of the American Society for Information Science and Technology*, 53(2), 134-144.

Westerman, D., Spence, P. R., & Van Der Heide, B. (2012) A social network as information: The effect of system generated reports of connectedness on credibility on Twitter. *Computers in Human Behavior*, 28, 199-206.

Zollman, K. J. S. (2010) The epistemic benefit of transient diversity. *Erkenntnis*, 72(1), 17-35.

Zollman, K. J. S. (2013) Network Epistemology: Communication in Epistemic Communities. *Philosophy Compass*, 8(1), 15-27.