Machine Intelligence and the Social Web: How to Get a Cognitive Upgrade

Paul R. Smart

Electronics & Computer Science
University of Southampton
Southampton, UK
email: ps02v@ecs.soton.ac.uk

***Abstract*—The World Wide Web (Web) provides access to a global space of information assets and computational services. It also, however, serves as a platform for social interaction (e.g., Facebook) and participatory involvement in all manner of online tasks and activities (e.g., Wikipedia). There is a sense, therefore, that the advent of the Social Web has transformed our understanding of the Web. In addition to viewing the Web as a form of information repository, we are now able to view the Web in more social terms. In particular, it has become possible to see the Web as providing access to the human social environment. This is important, because issues of social embedding and social interaction have been seen to contribute to the emergence of human cognitive capabilities. The ability of the Web to provide access to the human social environment thus raises an important question: Can humanity play a productive role in the emergence of advanced forms of machine intelligence by virtue of their interactions and engagements with the online realm? The present paper attempts to show why this question is worth asking. It also attempts to highlight some of the ways in which Web-based forms of contact with the human social environment may be relevant to research into machine intelligence.**

***Index Terms*—social web; social intelligence; language; machine intelligence; machine learning.**

I. INTRODUCTION

The World Wide Web (Web) is a technology that was created by humanity, and its implications for humanity—e.g., its effects on human cognitive and social processes—are, to a large extent, the primary focus of our current interest and concern [1]. Such a preoccupation is, of course, perfectly understandable. It is natural for us to wonder (and sometimes worry) about the implications of the Web for our species, especially when it comes to the effect of the Web on our cognitive capabilities. For such capabilities are the hallmark of our species: it is our cognitive profile that sets us apart from other forms of terrestrial life, and it is such capabilities that enable us (and only us) to actively shape the course of our cognitive destiny—to engineer something like the Web, and then worry about its cognitive consequences.

The present paper attempts to approach the Web from a somewhat different perspective. Rather than attempt to consider the implications of the Web for future forms of *human* intelligence, it aims to consider the implications of the Web for future forms of *machine* intelligence. In order to lay the foundation for this appraisal, it is important that we first recognize the status of the Web as a social environment, or at least as an environment that provides an important form of informational contact with the human social environment. Such a claim is unlikely to require much in the way of a detailed defence. There can be little doubt that the Web has come to play a significant role in supporting all manner of social activities and processes. One need only reflect on the popularity of systems such as Facebook, Instagram, Twitter and Snapchat to appreciate the status of the Web as a form of social technology, i.e., as a technology that can be used to support and enable various forms of social interaction and engagement. This does not mean, however, that the social significance of the Web is exhausted by systems that support human–human communication. Beyond the social networking and instant messaging systems, we thus encounter a rich array of systems where issues of social participation are of paramount importance. These include systems that rely on large-scale social participation to yield substantive bodies of online content (e.g., Wikipedia). It also includes so-called human computation systems [2], which rely on social participation for the purposes of completing computationally difficult or intractable tasks. In general, the contemporary Web is dominated by an array of systems that enable human users to generate, edit and organize online content. Such systems are, of course, the familiar socio-technical denizens of what we now refer to as the Social Web.

Why should the social properties of the Web, however we choose to define them, be of any interest or relevance to the emergence of advanced forms of machine intelligence? To answer this question, it will help to introduce two substantive strands of empirical and theoretical research: one focused on the evolution of human intelligence, the other, on the ontogenetic development of human cognitive capabilities.

Evolutionary matters first. Humans, it should be clear, are prodigious cognizers, capable of traversing cognitive terrains that other organisms seem congenitally ill-equipped to navigate. What is it that explains this remarkable difference between ourselves and every other species that has inhabited Planet Earth? Surely not the properties of the physical environment in which we are embedded; for we are not alone in having to cope with the problems the physical world throws at us. An alternative possibility is that our particular form of intelligence is tied to the properties of the social environment in which we are embedded. It is thus the peculiar features of the *human social environment* that best explains the evolutionary emergence of the human mind. This idea actually comes in a variety of flavours, including the social brain

ISBN: 978-1-61208-531-9

hypothesis [3], the machiavellian intelligence hypothesis [4], the cultural intelligence hypothesis [5], and the social intelligence hypothesis [6]. The specific details of these hypotheses need not concern us here; what is important, for present purposes, is simply the idea that a considerable body of work associates the evolutionary emergence of human intelligence with the peculiar properties of the human social environment.

The importance of the human social environment has also been highlighted by work that seeks to explain the development of human cognitive capabilities. "[S]ociality," it has been suggested, "lies at the heart of cognitive development" [7, p. 7]. This is a view whose lineage can be traced to the work of the Soviet psychologist, Lev Vygotsky. Vygotsky argued that it is the nature of our interaction with socially-significant others that holds the key to understanding human cognition (see [8]). In the absence of our ability to interact and engage with other human agents, we would, Vygotsky suggests, be unable to acquire the sorts of abilities that are the hallmark of human cognizing. Similar sentiments are expressed by those who emphasize the importance of enculturation to human cognitive development. Tomasello [9], for example, suggests that:

> ...if a human child were raised from birth outside of human contact and culture, without exposure to human artifacts or communal activities of any kind, that child would develop few of the cognitive skills that make human cognition so distinctive... [9, p. 359]

The general idea, then, is that a capacity for social interaction, and an ability to exploit the products of human culture, are of crucial importance when it comes to understanding the distinctive shape of the human cognitive economy. Once we combine this idea with the earlier claim regarding the role of the human social environment in human cognitive evolution, it is easy to see how the putative social properties of the Web might pique the interests of the machine intelligence community. For inasmuch as we accept the claim that the Web affords an unprecedented form of access to the human social environment, then it seems that we are provided with a novel opportunity to expand the cognitive repertoire of systems that inhabit the online realm. The argumentative basis for this claim is as follows:

P1: The Web provides an important form of contact with the human social environment.
P2: The development of human intelligence is tied to the fact that we are socially-situated agents that are embedded within the human social environment.
C1: The Web is poised to play a productive role in the emergence of advanced forms of machine intelligence by virtue of the kind of contact it provides with the human social environment.

The present paper aims to marshal support for the claim that the Web provides an important form of contact with the human social environment. It also attempts to highlight some of the ways in which such forms of contact may be relevant to research into machine intelligence. The paper makes no attempt to evaluate the extent to which the (contemporary) Web is able to yield actual advances in machine intelligence—such efforts must, of course, await the results of empirical research. Instead, the aim of the present paper is more to identify a number of avenues for future research and explain why such avenues are worth pursuing.

II. THE HUMAN CLOUD

Elizabeth Shaw: *How do you...? How do you know that?*
David: *I watched your dreams.*

—Prometheus, 20th Century Fox, 2012

As the Web has developed, it has yielded an unprecedented form of access to the human social environment. The Social Web has, of course, played a particularly important role in this transition. With the advent of social networking sites, microblogging services, and media sharing systems, the online environment seems to afford ever-deeper insights into the dynamics of human social behavior [10]. Beyond this, however, the Web is a technology that has become deeply embedded in society, playing a crucial (and sometimes indispensable) role in all manner of social activities and processes. In the extreme case, such forms of socio-entanglement may lead to the conclusion that the Web has become an intrinsic part of society—part of the physical machinery that (at least in part) realizes a rich array of social processes (see [11]).

When we look at the Web through human eyes, what we see is a global space of networked information assets and computational resources. It is a space of almost unimaginable scale and complexity. But now imagine that you are a machine agent that is embedded in this space. From this new perspective, you can see yourself, perhaps, as having access to the social environment in which we—us humans agents—live. It is, perhaps, an imperfect form of access—you only see the digital shadows of the human agents that inhabit the offline world. But it is arguably a form of access, nevertheless.

As a means of helping to effect this shift of perspective, consider the claim that our current arsenal of Internet-enabled devices—our phones, watches, tablets and so on—serve as bidirectional 'plug points' [12]. In one sense, such devices enable us to 'plug into' the online environment, typically for the purposes of pursuing some cognitive, social or epistemic objective. In another sense, however, these devices also enable the machines of the online world to plug into us! We are thus the resources that exist at the edge of the Web, just as (from our human perspective) it is the machines that are available at the end of an HTTP request.

Such forms of bidirectional contact lie at the heart of what has been referred to as the "human cloud" [13]. The human cloud, in this case, is the human counterpart of the conventional 'cloud', i.e., the suite of online resources and services that are the typical targets of cloud computing initiatives [14]. In essence, the notion of the human cloud encourages us to see the human social environment as a kind of computational resource—one that can be used to assist with certain kinds of information processing activity and (perhaps) the storage of certain kinds of information. A number of engineering efforts are relevant to this cloud-based view of the human social environment. Such efforts include the use of service-oriented protocols to discover and access human agents [15], the extension of traditional Web service description languages to accommodate the possibility of human involvement [15], and the emergence of programming frameworks that are specifically geared to deliver "complex computation systems incorporating large crowds of networked humans and machines" [16, p. 124].

The concept of the human cloud is important because it helps to highlight some of the ways in which machine-based systems can draw on the human social environment as a means of solving certain kinds of problem. Thus just as we humans rely on the online environment to support our problem-solving efforts, so too, perhaps, we can see machine-based systems as 'tapping' into the human cloud as a means of bolstering their 'cognitive' performance profiles. In some cases, such forms of socio-technical coupling can be seen to yield hybrid problem-solving organizations whose computational capabilities surpass those of their constituent (human and machine) elements. The forms of socio-technical entanglement that are enabled by the Web thus provide a range of novel opportunities to support large-scale problem-solving efforts that combine the distinctive strengths (and perhaps weaknesses) of both human agents and conventional computational systems [2][17][18].

There is, however, another point that is worth mentioning here. It relates to the way in which various forms of human–machine interaction can play a productive role in extending the reach of machine-based cognitive capabilities. By reaching out to the human cloud, for example, resources that were previously too ill-structured to support machine learning can sometimes be transformed into something that is much better aligned with the requirements of machine learning algorithms. Consider, for example, the way in which the addition of descriptive tags and annotations to a set of image resources can assist with the development of automated image classification (machine vision) systems [19]. Such possibilities are explicitly recognized by those who seek to engage human subjects in computationally-difficult tasks. With respect to citizen science systems, for example, Lintott and Reed [20] note that one of the limiting factors in the development of automated processing solutions is the availability of sufficiently well-structured training data sets, and that one of the key advantages of citizen science projects is the provision of such data sets. Similarly, when it comes to a class of systems known as Games With A Purpose (GWAPs), von Ahn and Dabbish [21] are keen to stress the role of human contributions in giving rise to evermore intelligent (and human-like) forms of machine-based processing:

> By leveraging the human time spent playing games online, GWAP game developers are able to capture large sets of training data that express uniquely human perceptual capabilities. This data can contribute to the goal of developing computer programs and automated systems with advanced perceptual or intelligence skills. [21, p. 67]

III. LEARNING FROM EXPERIENCE

Roy Batty: *If only you could see what I've seen with your eyes!*

—Blade Runner, Warner Bros., 1982

The concept of the human cloud helps to highlight the status of the human social environment as a form of complementary computational resource—one that can be used to circumvent the limitations of more conventional forms of (silicon-based) computational processing. This is clearly important when it comes to our ability to support novel kinds of problem-solving organization. However, it also serves as a reminder that human agents are the locus of particular forms of skill and expertise that are often grounded in the extensive experience we have with particular domains. It is here that the Web provides us with an opportunity to extend the reach of machine-based capabilities. The basic idea is that the Web can be used as a form of *social observatory*—one which enables machines to observe the human social environment and acquire information about various forms of human competence. From this perspective, the Web can be seen to support a particular form of *social learning*: it enables us to treat the human social environment as a source of information and knowledge that can be mined and monitored in order to reduce the 'experiential gap' that otherwise limits the cognitive and epistemic reach of intelligent systems.

To help us understand this claim in a little more detail, consider the effort to develop self-driving cars. Such efforts clearly depend on advances in our ability to engineer sophisticated forms of information processing, especially in the visual domain. However, they also rely on advanced control systems that are able to respond in an intelligent manner to a multitude of road-relevant situations. In order to emulate the behavior of human road users, it thus seems important to capture some of the knowledge that human drivers have acquired as a result of their experience behind the wheel. Such experience underlies our ability to anticipate the likely behavior of other road users, our ability to behave appropriately at an intersection, our ability to adjust our driving behavior given specific meteorological conditions, and so on. An experienced human driver thus embodies a wealth of knowledge and experience that is clearly pertinent to the design of autonomous vehicles, and this is especially so if (as seems likely) we encounter a transitional era in which self-driving vehicles must share the road with human-driven vehicles.

How do we go about building cars that possess the behavioral competence and road-related *savoir faire* exhibited by the typical human driver? One option is to enlist the use of conventional knowledge elicitation techniques [22] in order to create formal models of the relevant body of human knowledge. The problem with this approach is that it is likely to require substantial time and effort, especially when one considers the complexity of the target domain and the diversity of driving practices exhibited by both individuals and cultural groups (consider the norms and conventions that appear to characterize the behavior of British and Italian drivers!).

Here is another approach: track the behavior of human-driven vehicles as they move around the road network and attempt to extract and formalize interesting regularities from the resultant body of 'experiential data'. Such data sets are likely to be particularly valuable in cases where it is possible to track the precise behavior of vehicles at particular locations, such as at an intersection, a roundabout or a notorious black spot. Additional value comes from the ability to track other kinds of information, such as the use of driver signalling mechanisms (e.g., the use of indicators and headlights) and information about prevailing meteorological conditions (e.g., the presence of fog).

The main point of this example is that it helps us see how a particular form of access to the human social environment can provide insight into bodies of experientially-grounded knowledge, some of which may be relevant to the attempt to engineer intelligent systems. The vision is thus one in which advanced forms of machine intelligence come about as the result of a deliberate attempt to monitor and learn from the human social environment. According to this vision, machine intelligence is, in some sense, parasitic on human experience: it relies on the experience that humans have in order to short-circuit the acquisition of particular forms of cognitive and behavioral competence.

It is here that we encounter an interesting point of contact with research in the Web Science community. For inasmuch as we accept the idea that the Web provides access to the human social environment, then it seems that we should be able to mine the Web for certain kinds of knowledge. This is the general idea behind a body of work that goes under the heading of *experiential knowledge mining* [23]. Experiential knowledge mining is defined as "the process of acquiring experiential knowledge, as opposed to a priori knowledge, from a variety of multimedia sources that describe human experiences of various sorts" [23, p. 33]. In a Web context, such forms of knowledge acquisition aim to shed light on (among other things) the norms and conventions that surround particular patterns of human behavior. A similar sort of claim is sometimes encountered in the nascent sub-discipline of computational social science [10]. In this case, the Web is seen to provide an unprecedented opportunity to learn about the human social environment, enabling us to acquire the sorts of insights that other approaches may be unable to provide.

IV. ACTIVE LEARNING

Caleb: *Did you program her to flirt with me?*
Nathan: *If I did, would that be cheating?*

—Ex Machina, Universal Pictures, 2015

We have seen that the Web provides us with an unprecedented opportunity to observe the human social environment (or at least aspects thereof). But the notion of the Web as a form of social observatory comes with an attendant risk. The risk is that we lose sight of the way in which online systems can play an active role in shaping the course of their own cognitive development. When we view the Web as a form of social observatory, there is a danger that we see machines as the purely passive observers of some distant and perhaps impervious social realm. This is a highly impoverished view of social learning, and it is one that seems at odds with the profile of human learning.

There is, however, no reason why we should restrict ourselves to this purely passive view of machine learning. There are, in fact, a number of ways in which we can view machines as playing a more active role in the learning process. One example of this comes from work into what is conveniently called *active learning* [24]. Active learning is a form of machine learning in which the machine attempts to actively intervene in the learning process, structuring its training experiences in a manner that yields the best learning outcome. Such forms of active intervention have been shown to yield a number of pedagogical pay-offs. For example, active learning has been shown to improve the efficiency of the learning process by reducing the number of training examples that are required to reach near-optimal levels of performance [25].

A good example of active learning in a Web context is provided by Barrington et al. [26]. Barrington et al. describe the use of an online game, called Herd It, in which groups of human individuals annotate a musical resource with descriptive tags. These annotations are used to train a supervised machine learning system that ultimately aims to perform the annotation task independently of the human agents. All this, of course, is broadly in line with the general shape of machine learning methods. But what makes Barrington et al.'s system of particular interest is the way in which the machine actively *directs* the course of its own learning. It does this by actively selecting the musical resources that will be the focus of future tagging efforts by the human game players. This is important, because it gives the machine an opportunity to select those forms of feedback that are likely to be of greatest value relative to its subsequent cognitive development. In the words of Barrington et al. [26], "the machine learning system actively directs the annotation games to collect new data that will most benefit future model iterations" (p. 6411).

A consideration of active learning thus expands our understanding of the forms of contact that the Web provides with the human social environment. Rather than seeing the Web simply as a form of social observatory—one that permits a largely passive form of observational contact with humanity—we can now entertain a more active (and interactive) view of the Web. On this view, the Web provides machines with an opportunity to structure their contact with the human social environment, enabling machines to influence human behavior in a manner that befits the demands of a particular cognitive task.

V. THE GIFT OF LANGUAGE

Louise: *It's not a weapon, it's a gift—their language.*

—Arrival, Paramount Pictures, 2016

Given that much of the content of the Web is expressed in the form of natural language, it is perhaps unsurprising that the advent of the Web has led to something of a renaissance in language-related computational research. Such interest is evidenced by research into Natural Language Processing (NLP) (e.g., [27]), information extraction [28], and sentiment analysis [29]. It is also evidenced by the effort to develop various forms of language-enabled agents, i.e., computational agents that are able to exhibit proficiency in the use of natural language expressions. This includes work relating to so-called social bots [30], chatbots [31] and conversational agents [32]. The reason for this renewed interest in language-related technologies is, at least in part, due to the wealth of linguistic content that is available in the online realm. Such content provides us with a substantive body of empirical data that can be used to inform large-scale analytic efforts. It should also be clear that the Web has transformed the incentives that drive research and development in this area—consider, for example, the potential use of Twitter feeds as a means of predicting the outcome of political elections [33].

How does this renewed interest in linguistic analysis impact the present discussion on machine intelligence? The most obvious answer to this question is that machines will become increasingly proficient in understanding human language, and as a result of this understanding, they will be better placed to exploit our orthographic contributions to the online realm (e.g., they will have an improved ability to distil information and knowledge from resources such as Wikipedia, Twitter, Facebook and so on). It should also be clear that enhancements in linguistic proficiency often go hand-in-hand with improvements in communicative ability. There can be little doubt that such communicative abilities play an important role in extending the cognitive reach of an agent community. Indeed, we can view communication as a form of networking capability that enables agents to 'connect' with a range of cognitively-relevant resources. This applies as much to human agents as it does to their synthetic counterparts. As noted by Merlin Donald [34], when it comes to human language, "Individuals in possession of reading, writing, and other visuographic skills...become somewhat like computers with networking capabilities; they are equipped to interface, to plug into whatever network becomes available" (p. 311).

The communicative function of language is no doubt important when it comes to future forms of machine intelligence. But there is another view of language that is of potential significance here. This view is sometimes referred to as the supracommunicative view of language [35][36]. The general idea, in this case, is that language plays a role in transforming the cognitive capabilities of the language-wielding agent. Echoes of this view are apparent in the work of the philosopher, Daniel Dennett [37]. He suggests that our ontogenetic immersion in a linguistic environment contributes to an effective reorganization of the human cognitive economy, such that the parallel-processing dynamics of the biological brain are transformed into something that more closely resembles a conventional (symbol-manipulating) computational machine. Strikingly, Dennett proposes that some of the most distinctive features of human cognition (including the phenomenon of human consciousness) emerge as a result of our attempts to get to grips with the linguistic domain. Inasmuch as we accept these claims, it should be clear that a simple *communicative* view of language is unlikely to do justice to the potential impact of the Web on future forms of machine intelligence: by immersing intelligent systems in a linguistically-rich environment, and by forcing such systems to assimilate linguaform representations deep into their cognitive processing routines, we potentially endow machines with the sorts of abilities and insights that only us language-using human agents are able to grasp.

VI. CHILD MACHINES

David: *Big things have small beginnings.*

—Prometheus, 20th Century Fox, 2012

In the Introduction to this paper (see Section I) we encountered the idea that issues of social embedding and social interaction play a crucial role in the development of human cognitive capabilities. Such claims are relevant to the present discussion in the sense that we can think of the Web as enabling human agents to participate in the socially-mediated development of machine-based cognitive capabilities. However, in considering the ways in which the Web can be used to scaffold and support the cognitive development of intelligent systems, it is easy to overlook the fact that human infants are not born with the same sorts of cognitive resources as their adult counterparts. It is here that we encounter a productive point of contact with work that shows how maturational shifts in cognitive, sensory and motor capabilities may be of crucial relevance to the emergence of advanced forms of cognitive competence [38]–[41]. Such ideas are typically encountered in the context of human language learning. Bjorklund [40], for example, suggests that by imposing constraints on the kinds of information that can be processed, maturational mechanisms can be seen as supporting the progressive reshaping of the 'effective' structure of a human infant's linguistic environment, transforming what might seem like an impossible language learning task into something a little more congenial. Similar ideas can be found in more recent work in developmental robotics [38][41]. Gómez et al. [38], for example, describe an intriguing set of results concerning the development of sensorimotor capabilities in a real-world robotic system. They report that a developmental profile characterized by progressive increments in the complexity of sensory, motor and neurocomputational subsystems results in a profile of task performance that is superior to that of a robot in which the relevant maturational processes are disabled. Commenting on this developmentally-grounded dissociation in 'adult' performance profiles, they suggest that:

> ...rather than being a problem, early morphological and cognitive limitations effectively decrease the amount of information that infants have to deal with, and may lead to an increase in the overall adaptivity of the organism. [38, p. 119]

Such findings intimate at the potential relevance of maturational parameters in the acquisition of advanced forms of cognitive competence. In particular, they suggest that various forms of cognitive immaturity may be of adaptive value in terms of a system's ability to achieve the sorts of cognitive success that mark the end of the developmental process (see [42]).

What implications do these insights have for our understanding of machine intelligence? In answering this question, it helps to consider the notion of *incremental learning* [43]. Incremental learning, as defined by Lungarella and Berthouze [41], is the "idea of some learning-related resource (e.g., memory, or attention span), starting at a low value, which then gradually increases while (but not necessarily because) the organism matures" (p. 1). In essence, the claim is that in our attempts to yield advanced forms of machine intelligence, we should attempt to emulate the developmental profile of human infants. We should, in other words, seek to create what Alan Turing [44] once referred to as "child machines"—machines whose cognitively-relevant processing capabilities emerge as the result of a particular form of artificial ontogenesis.

In the context of the present discussion, the notion of incremental learning helps to reveal an important research opportunity. This relates to the extent to which maturational shifts in computational parameters can assist with the task of pressing maximal cognitive benefit from Web-based forms of informational contact with the human social environment. There are a number of ways in which we might seek to explore this dynamic dovetailing of intrinsic information processing capabilities with the structure of the relevant learning environment. One possibility is for a machine learning system to exert some degree of control over the sorts of scaffolding that are supplied by the human social environment, in the manner, perhaps, of the active learning system described by Barrington et al. [26]. Alternatively, a system could be configured so as to process inputs of increasing complexity as learning progresses. One way of accomplishing this is to manipulate the computational and representational resources that are available to the system as it attempts to learn about a particular domain. Elman [39] provides a nice demonstration of this sort of intervention. By altering the configuration of a neural network, Elman was able to introduce a processing constraint that limited the network's ability to process complex natural language sentences. As a result of this limitation, the network was able to achieve a level of 'linguistic' performance that was otherwise difficult to attain.

What is important, here, is not the details of how incremental learning could be could be implemented by a machine learning system. Instead, what is important is that we appreciate the role of maturational processes in shaping the course of cognitive development. In the absence of such an awareness, it might be all too easy to think that only advanced forms of machine intelligence are able to benefit from the sorts of contact the Web provides with the human social environment, and this is especially so once we consider the complexity of the digital traces that mark our occasional forays into the online world. There is thus a danger that our reasoning becomes somewhat circular: only advanced forms of machine intelligence are able to press maximal cognitive benefit from the Web, and only machines that press maximal cognitive benefit from the Web are able to exhibit advanced intelligence. In this sense, the notion of incremental learning serves a useful prophylactic purpose: it helps to remind us that big things often have small beginnings. Indeed, it is perhaps by virtue of being small that certain kinds of big thing are ever able to exist.

VII. THE ENGINEERS

Elizabeth Shaw: *We call them Engineers.*
Fifield: *Engineers? Do you mind, um, telling us what they engineered?*

—Prometheus, 20th Century Fox, 2012

One of the characteristic features of the Social Web is that everyone can contribute to it: every time we edit a Wikipedia article, write a blog, or post a tweet, we are, in some sense, contributing to the totality of information that is available on the Web. This feature is, of course, so well-established as to be hardly worth mentioning. And yet the idea that the online environment emerges as a result of human contributions is an important one. In particular, inasmuch as we see the Web as a form of environment or ecology (see [1]), then our creative contributions to the online realm can perhaps be glossed as a form of *ecological engineering*. This helps to establish an interesting point of contact with work that goes under the heading of *niche construction*. Niche construction is a term used by evolutionary biologists to refer to the ways in which animals actively engineer their local environments (niches) so as to alter the sorts of selective pressure that apply to future generations [45]. The idea is relevant to the present discussion, because the notion of niche construction has been implicated in the evolutionary emergence of human intelligence [46]. The full details of this proposal need not detain us here; the main point, for present purposes, is simply the idea that humans have the potential to engineer their environments in ways that alter the evolutionary trajectory of future generations.

Issues of ecological engineering and niche construction are important when we consider the potential role of the Web in supporting the emergence of future forms of machine intelligence. Thus just as progressive alterations in the structure of the physical and social environment may help to establish the conditions that favor the evolutionary emergence of human intelligence (see [46]), so too, perhaps, our current forms of interaction and engagement with the Web can be seen to yield an ecological niche that is conducive to the emergence of advanced (perhaps human-like) forms of machine intelligence.

As a means of helping us get a better grip on what is being proposed here, it may help to focus our attention on the way in which the Global Positioning System (GPS) has contributed to the emergence of spatially-aware autonomous systems. The constant stream of data provided by GPS satellites has obviously impacted the way in which we humans navigate and locate ourselves in space. But it should also be clear that the *very same system* has had a significant impact on the development of a rich array of intelligent systems. The availability of GPS signals has thus transformed the kinds of approach that can be adopted with respect to the implementation of (e.g.) driverless cars and pilotless drones, enabling forms of navigational competence that may have been difficult, if not impossible, to achieve in the absence of such a suitably structured environment. Crucially, once we focus our attention on the sorts of capabilities that are exhibited by driverless cars and pilotless drones, the 'cognitive' relevance of the environment in which such systems are embedded starts to come into clearer view. Absent the constant stream of data provided by GPS satellites and the complexity of the navigational problem may become so severe as to stymie our best attempts at automation. The moral to emerge from all this is that we should not underestimate the transformative potential of an appropriately configured and suitably enriched environment. When it comes to issues of machine intelligence, the ability to engage in various forms of intelligent response may have as much to do with the environment in which a system is embedded as it does with the sophistication of the system's inner information processing mechanisms.

We can clearly think of GPS signals as establishing a particular kind of digital data ecology—one that has proved to be of value when it comes to the implementation of certain kinds of intelligent system. But much the same could be said of the digital data ecology that we encounter in the case of the Web. Consider, for example, one of the key exemplars of the emerging cognitive computing paradigm: IBM Watson. Watson's virtuoso performance in answering an array of difficult questions is undoubtedly an important demonstration of the growing sophistication of Artificial Intelligence (AI) algorithms, especially in the areas of natural language processing, machine learning and domain-specific reasoning. But, in the course of being awestruck by (at least some of) Watson's outputs, it is easy to overlook the simple fact that many of the inputs to the system—especially the information resources that Watson exploits in supplying its human interlocutors with answers—are ones that are, in general, developed by large numbers of human individuals. Such resources include online encyclopedias, dictionaries, thesauri, taxonomies, ontologies and so on [47]. None of these resources, it should be clear, were specifically intended to support Watson, or indeed any of the other cognitive computing systems that are the focus of current research efforts. Nevertheless, such resources play a crucial role in enabling Watson to exhibit capabilities that surpass those of the average human individual. Perhaps if Watson had emerged in an earlier era, we would have been astounded by its question/answering capabilities. And yet Watson is, at root, exactly the same kind of symbol manipulating machine that has long been the focus of philosophical theorizing and the primary instrument of cognitive scientific practice. Arguably, what has changed in recent times is not so much the underlying technologies as the nature of the environment in which such technologies are situated. The emergence of the Web has thus yielded an environment that has altered the 'evolutionary' landscape for intelligent systems. It is an environment that we humans have created, and it is one that we will bequeath to our biological progeny. But perhaps this particular form of niche construction is unlike those that preceded it. Perhaps it is no longer just the cognitive destiny of our own species that is affected by our attempts to create, configure and construct the informational ecology of the online digital world.

VIII. CONCLUSION

On the basis of the foregoing analysis, we are in a position to identify a number of ways in which the Web can be seen to provide access to the human social environment. One kind of access is captured by the notion of the human cloud (see Section II). In this case, we saw how the Web could be used to 'recruit' human agents into complex information processing tasks. A similar idea emerged in the discussion of active learning (see Section IV). Here, we saw how machine learning systems could actively shape their learning trajectories by soliciting particular kinds of input from the human social environment. Finally, in Section III, we saw how human contributions to the online environment could support the processes of social learning and experiential knowledge mining.

In addition to outlining the ways in which the Web provides access to the human social environment, the present paper also highlights a number of areas for future research. For example, in Section V, we touched on the idea that language has a supracommunicative function, helping to promote productive shifts in the cognitive capabilities of language-wielding agents. We also encountered the idea that maturational shifts in cognitively-relevant parameters may play a crucial role in enabling an intelligent system to acquire certain kinds of cognitive and behavioral competence (see Section VI). Such insights are relevant to the development of intelligent systems that are able to press maximal cognitive benefit from their contact with the human social environment.

The advent of the (Social) Web marks a potentially significant milestone in the development of advanced forms of machine intelligence. Traditionally, issues of social embedding, social interaction and enculturation have been discussed in relation to human intelligence. It is thus the nature of our contact with the human social environment that has been seen to underlie the kinds of cognitive capabilities that are the hallmark of our species. Prior to the advent of the Web, the opportunities for machines to be embedded within the human social environment were somewhat limited. Now, in an era where a significant proportion of humanity engages with the Web on a more-or-less daily basis, it is increasingly difficult for online forms of machine intelligence to ignore the vagaries of the social world. Inasmuch as we accept that human intelligence emerges as the result of our attempt to navigate the complexities of the social realm, then perhaps the online environment is the perfect place to look for machines whose capabilities enable them to reach those parts of the cognitive terrain that only our (human) minds are able to call their home.

ACKNOWLEDGMENT

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

[1] P. R. Smart, R. Heersmink, and R. W. Clowes, "The cognitive ecology of the Internet," in *Cognition Beyond the Brain: Computation, Interactivity and Human Artifice*, 2nd ed., S. J. Cowley and F. Vallée-Tourangeau, Eds. London, England, UK: Springer-Verlag, in press.

[2] E. Law and L. von Ahn, "Human computation," *Synthesis Lectures on Artificial Intelligence and Machine Learning*, vol. 5, no. 3, pp. 1–121, 2011.

[3] R. I. M. Dunbar, "The social brain hypothesis," *Evolutionary Anthropology*, vol. 6, no. 5, pp. 178–190, 1998.

[4] R. W. Byrne and A. Whiten, *Machiavellian Intelligence: Social Expertise and the Evolution of Intellect in Monkeys, Apes, and Humans*. Oxford, UK: Oxford University Press, 1988.

[5] E. Herrmann, J. Call, M. V. Hernández-Lloreda, B. Hare, and M. Tomasello, "Humans have evolved specialized skills of social cognition: The cultural intelligence hypothesis," *Science*, vol. 317, no. 5843, pp. 1360–1366, 2007.

[6] H. Kummer, L. Daston, G. Gigerenzer, and J. B. Silk, "The social intelligence hypothesis," in *Human By Nature: Between Biology and the Social Sciences*, P. Weingart, S. D. Mitchell, P. J. Richerson, and S. Maasen, Eds. Mahwah, New Jersey, USA: Lawrence Erlbaum Associates, 1997.

[7] K. Dautenhahn and A. Billard, "Studying robot social cognition within a developmental psychology framework," in *Third European Workshop on Advanced Mobile Robots*, Zürich, Switzerland, 1999.

[8] J. Lindblom and T. Ziemke, "Social situatedness of natural and artificial intelligence: Vygotsky and beyond," *Adaptive Behavior*, vol. 11, no. 2, pp. 79–96, 2003.

[9] M. Tomasello, "Primate cognition: Introduction to the issue," *Cognitive Science*, vol. 24, no. 3, pp. 351–361, 2000.

[10] M. Strohmaier and C. Wagner, "Computational social science for the World Wide Web," *IEEE Intelligent Systems*, vol. 29, no. 5, pp. 84–88, 2014.

[11] P. R. Smart and N. R. Shadbolt, "Social machines," in *Encyclopedia of Information Science and Technology*, M. Khosrow-Pour, Ed. Hershey, Pennsylvania, USA: IGI Global, 2014.

[12] P. R. Smart, "Extended cognition and the Internet: A review of current issues and controversies," *Philosophy & Technology*, in press.

[13] E. Kaganer, E. Carmel, R. Hirschheim, and T. Olsen, "Managing the human cloud," *MIT Sloan Management Review*, vol. 54, no. 2, pp. 23–32, 2013.

[14] B. Hayes, "Cloud computing," *Communications of the ACM*, vol. 51, no. 7, pp. 9–10, 2008.

[15] D. Schall, "Service oriented protocols for human compuation," in *Handbook of Human Computation*, P. Michelucci, Ed. New York, New York, USA: Springer, 2013.

[16] P. Minder and A. Bernstein, "CrowdLang: A programming language for the systematic exploration of human computation systems," in *4th International Conference on Social Informatics*, Lausanne, Switzerland, 2012.

[17] J. Hendler and T. Berners-Lee, "From the Semantic Web to social machines: A research challenge for AI on the World Wide Web," *Artificial Intelligence*, vol. 174, pp. 156–161, 2010.

[18] R. J. Crouser, A. Ottley, and R. Chang, "Balancing human and machine contributions in human computation systems," in *Handbook of Human Computation*, P. Michelucci, Ed. New York, New York, USA: Springer, 2013.

[19] S. Dieleman, K. W. Willett, and J. Dambre, "Rotation-invariant convolutional neural networks for galaxy morphology prediction," *Monthly Notices of the Royal Astronomical Society*, vol. 450, no. 2, pp. 1441–1459, 2015.

[20] C. J. Lintott and J. Reed, "Human computation in citizen science," in *Handbook of Human Computation*, P. Michelucci, Ed. New York, New York, USA: Springer, 2013.

[21] L. von Ahn and L. Dabbish, "Designing games with a purpose," *Communications of the ACM*, vol. 51, no. 8, pp. 58–67, 2008.

[22] N. R. Shadbolt and P. R. Smart, "Knowledge elicitation: Methods, tools and techniques," in *Evaluation of Human Work*, 4th ed., J. R. Wilson and S. Sharples, Eds. Boca Raton, Florida, USA: CRC Press, 2015.

[23] S.-H. Myaeng, Y. Jeong, and Y. Jung, "Experiential knowledge mining," *Foundations and Trends in Web Science*, vol. 4, no. 1, pp. 1–102, 2012.

[24] B. Settles, "Active learning," *Synthesis Lectures on Artificial Intelligence and Machine Learning*, vol. 6, no. 1, pp. 1–114, 2012.

[25] A. Holub, P. Perona, and M. C. Burl, "Entropy-based active learning for object recognition," in *IEEE Online Learning for Classification Workshop*, Anchorage, Alaska, USA, 2008.

[26] L. Barrington, D. Turnbull, and G. Lanckriet, "Game-powered machine learning," *Proceedings of the National Academy of Sciences*, vol. 109, no. 17, pp. 6411–6416, 2012.

[27] F. Ciravegna, S. Chapman, A. Dingli, and Y. Wilks, "Learning to harvest information for the Semantic Web," in *First European Semantic Web Symposium*, Heraklion, Crete, Greece, 2004.

[28] S. Sarawagi, "Information extraction," *Foundations and Trends in Databases*, vol. 1, no. 3, pp. 261–377, 2008.

[29] B. Pang and L. Lee, "Opinion mining and sentiment analysis," *Foundations and Trends in Information Retrieval*, vol. 2, no. 1–2, pp. 1–135, 2008.

[30] E. Ferrara, O. Varol, C. Davis, F. Menczer, and A. Flammini, "The rise of social bots," *Communications of the ACM*, vol. 59, no. 7, pp. 96–104, 2016.

[31] R. Dale, "The return of the chatbots," *Natural Language Engineering*, vol. 22, no. 5, pp. 811–817, 2016.

[32] J. Lester, K. Branting, and B. Mott, "Conversational agents," in *The Practical Handbook of Internet Computing*, M. P. Singh, Ed. Boca Raton, Florida, USA: Chapman & Hall/CRC, 2004.

[33] A. Tumasjan, T. O. Sprenger, P. G. Sandner, and I. M. Welpe, "Predicting elections with Twitter: What 140 characters reveal about political sentiment," in *Fourth International AAAI Conference on Weblogs and Social Media*, Washington D.C., USA, 2010.

[34] M. Donald, *Origins of the Modern Mind: Three Stages in the Evolution of Culture and Cognition*. Cambridge, Massachussets, USA: Harvard University Press, 1991.

[35] A. Clark, "Magic words: How language augments human computation," in *Language and Thought: Interdisciplinary Themes*, P. Carruthers and J. Boucher, Eds. Cambridge, UK: Cambridge University Press, 1998.

[36] ——, "How to qualify for a cognitive upgrade: Executive control, glass ceilings and the limits of simian success," in *The Complex Mind: An Interdisciplinary Approach*, D. McFarland, K. Stenning, and M. McGonigle-Chalmers, Eds. Basingstoke, England, UK: Palgrave Macmillan, 2012.

[37] D. Dennett, *Conciousness Explained*. Boston, Massachusetts, USA: Little, Brown & Company, 1991.

[38] G. Gómez, M. Lungarella, P. Eggenberger Hotz, K. Matsushita, and R. Pfeifer, "Simulating development in a real robot: On the concurrent increase of sensory, motor, and neural complexity," in *Fourth International Workshop on Epigenetic Robotics*, Genoa, Italy, 2004.

[39] J. L. Elman, "Learning and development in neural networks: The importance of starting small," *Cognition*, vol. 48, no. 1, pp. 71–99, 1993.

[40] D. F. Bjorklund, "The role of immaturity in human development," *Psychological Bulletin*, vol. 122, no. 2, pp. 153–169, 1997.

[41] M. Lungarella and L. Berthouze, "Adaptivity through physical immaturity," in *2nd International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*, Edinburgh, Scotland, 2002.

[42] D. F. Bjorklund and B. L. Green, "The adaptive nature of cognitive immaturity," *American Psychologist*, vol. 47, no. 1, pp. 46–54, 1992.

[43] S. K. Chalup, "Incremental learning in biological and machine learning systems," *International Journal of Neural Systems*, vol. 12, no. 6, pp. 447–465, 2002.

[44] A. M. Turing, "Computing machinery and intelligence," *Mind*, vol. 59, no. 236, pp. 433–460, 1950.

[45] K. N. Laland, J. Odling-Smee, and M. W. Feldman, "Niche construction, biological evolution, and cultural change," *Behavioral and Brain Sciences*, vol. 23, no. 1, pp. 131–146, 2000.

[46] K. Sterelny, "Social intelligence, human intelligence and niche construction," *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 362, no. 1480, pp. 719–730, 2007.

[47] D. Ferrucci *et al.*, "Building Watson: An overview of the DeepQA project," *AI Magazine*, vol. 31, no. 3, pp. 59–79, 2010.