# COGWEB—Cognition and the Web: Social Minds, Distributed Brains,

# and Talkative Technologies

Special Session at the 9th International Conference on Advanced Cognitive Technologies and Applications
February 19–23, 2017, Athens, Greece
https://www.iaria.org/conferences2017/COGNITIVE17.html

Paul R. Smart

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

*Abstract*—This special session on Cognition and the Web (COGWEB) features three papers that aim to explore the cognitive significance and relevance of emerging Web technologies. The papers focus on a number of themes and issues. These include the use of conversational interfaces to support human–machine interaction, the use of the Social Web to scaffold the development of intelligent systems, and the development of techniques to support the adaptive organization of distributed cognitive computing cognitive systems.

*Index Terms*—distributed cognition; social intelligence; internet of things; machine intelligence; language.

## I. INTRODUCTION

The growth of the Web and Internet over the past couple of decades has transformed the way a number of familiar human activities are undertaken. The advent of social networking sites, for example, has transformed the way that we establish, maintain and (sometimes) dissolve social relationships. Similarly, the ability to solicit input from large numbers of human individuals has transformed the opportunities and challenges that confront the attempt to solve a smorgasbord of scientific and societal problems [1]. Given the role of social participation in these and other systems, it is easy to see how the scientific study of the Web and Internet could be limited to the realms of social science. However, the Web is a system that is of profound significance and importance to contemporary cognitive science. We see this in work that seeks to understand the effect of the Web on a variety of socio-cognitive phenomena, such as collective intelligence [2], collective problem-solving [3], collective sensemaking [4] and collective creativity [5]. We also see it in work that focuses on cognitive phenomena at the level of individual human agents. This is evidenced by recent work in the areas of memory [6], reading [7], and social cognition [8].

It is, of course, easy to assume that the cognitive significance of the Web is exhausted by the study of *human* cognitive phenomena. In this case, we are perhaps inclined to limit our attention to a set of phenomena that have their origins in the whirrings and grindings of the biological brain, or (in the case of more socio-cognitive phenomena) collections thereof. This, however, is a somewhat impoverished view of the empirical targets and theoretical scope of contemporary cognitive science. For even if we limit our attention to the study of human

cognitive processes, it is far from clear that the machinery of the human mind is restricted solely to the neural realm [9]. This encourages us to take a more situated or ecologically-oriented approach to understanding the effects of the Web and the Internet on human cognitive processes [10]. It is from this new standpoint that we are able to see the Web as supporting the emergence of materially-hybrid cognitive organizations—ones whose cognitive power and potential is tied to the time-variant structural topology of informational circuits that connect biological brains with the resources of the online world.

Aside from its effects on human cognition, the Web can be seen to influence machine-based cognitive capabilities. In support of this, recent years have seen a flurry of interest in systems that attempt to harness the computational and representational resources of the online environment for the purposes of yielding state-of-the-art advances in machine intelligence. Crucially, the advent of the Web can be seen to transform the space of opportunities and challenges that surround the attempt to develop the next generation of intelligent systems. When it comes to the development of Cognitive Computing Systems (CCSs), for instance, Kelly and Hamm [11] note that we need machines that do "much more than calculate and organize and find patterns in data," we also need machines that are able to "sense, learn, reason and interact naturally with people in powerful new ways" (pp. 3–4). This highlights one of the core themes of the special session on Cognition and the Web (COGWEB): the idea that the Web presents us with a relatively new set of challenges, opportunities and incentives when it comes to the attempt to engineer systems that are able to interact and engage with human agents [12][13]. Another theme of the COGWEB special session relates to the extent to which we can view the Web as providing a form of access to the human social environment [12]. Such access is seen to provide a potent opportunity for humanity to participate in the emergence of machine-based cognitive capabilities. Finally, Verma et al. [14] draw attention to issues of adaptive self-organization and distributed processing as part of their discussion of the *distributed brain*—a form of distributed cognitive organization that is "composed [of]...different types of devices...ranging from hand-held devices at the edge of the network to large systems in the cloud." Such claims help to reinforce the idea that the online environment corresponds

to a form of cognitive ecology (see [10])—one that supports the dynamic and *ad hoc* assembly of a variety of cognitive organizations.

## II. TALKATIVE TECHNOLOGIES

The paper by O'Leary et al. [13] tackles a topic of crucial importance in an era characterized by the rapid proliferation of network-enabled devices: how do we support communicative transactions between the entities (both human and machine) that exist within a network environment? The solution proposed by O'Leary et al. capitalizes on the existing familiarity that humans have with natural language. In particular, they propose that a controlled variant of natural language—referred to as a Controlled Natural Language (CNL)—can provide the syntactic foundation for semantically-expressive forms of communication between human and machine agents. Aside from the obvious (human) usability benefits that accompany the use of a natural language format, O'Leary et al. point to a number of other advantages associated with the use of CNLs. These include the ability to represent provenance information and engage in domain-specific reasoning. O'Leary et al. also describe how a CNL can be augmented with a conversational protocol to support dialogic exchanges between CNL-capable agents. Such discursive capabilities are of particular value when it comes to the resolution of semantic ambiguities that arise during the course of communicative exchanges. They also provide an added bonus in terms of the ability to support a range of question/answering capabilities.

One of O'Leary et al.'s key insights concerns the protean nature of the technological environment in which both human and machine agents are situated. The "enactive landscape" (see [15]) associated with a user's home, for example, is subject to constant shifts and deformations as new devices (and information processing capabilities) become available. Here, O'Leary et al. discuss the way in which a CNL can be used to enrich the representational repertoire of a system via the progressive specification and enrichment of semantic models. They also describe how the reasoning capabilities of a computational system can be modified via the addition of linguistically-specified rules. The resulting vision is thus of a complex economy of information processing assets whose cognitive power and potential is grounded in the use of linguistic representations. Relative to this vision, we can begin to see how CNLs are poised to serve as a form of 'linguistic glue' that helps to define the *effective* (see [16]) structure of distributed, task-specific cognitive circuits. Such ideas touch on some of the issues raised by Verma et al. [14] in respect of the adaptive organization of distributed cognitive systems (see Section IV).

## III. SOCIAL MINDS

Language-related issues also surface in the paper by Smart [12]. In particular, Smart suggests that the Web has yielded an online environment that is replete with linguistic content. This, he suggests, has transformed the opportunities (and incentives) that confront the attempt to equip machines with language-related abilities. Such abilities obviously yield a range of communicative benefits, some of which are detailed in the paper by O'Leary et al. [13]. However, while O'Leary et al. focus their attention on the communicative function of language, Smart suggests that language works to instil a range of new cognitive abilities, perhaps resulting from a subtle reorganization of an agent's cognitive economy. It is in this sense that language is seen to constitute a 'gift'—one that has been 'offered' to machines by humanity, and one whose delivery is made possible by the advent of the Social Web.

Such claims form part of a broader argumentative mission, which is intended to highlight the role of the Social Web in supporting the emergence of advanced forms of machine intelligence. In particular, Smart suggests that the Social Web provides an unprecedented form of contact with the human social environment. This is deemed to be important, in a cognitive sense, because it enables machines to benefit from a range of relatively new learning opportunities. In essence, Smart suggests that by virtue of the Social Web, Artificial Intelligence (AI) systems are able to benefit from the sorts of social scaffolding that have long been seen as relevant to the ontogenetic development of human cognitive capabilities [17].

## IV. DISTRIBUTED BRAINS

In discussing the relationship between the Social Web and machine intelligence, Smart [12] suggests that humanity has created an environment that is conducive to the emergence of advanced AI systems. In fleshing out this claim, Smart asks us to consider the way in which the availability of a rich array of online resources (e.g., dictionaries, thesauri, taxonomies, ontologies, encyclopedias, etc.) has supported the development of CCSs, such as the IBM Watson system.

Interestingly, CCSs are the primary target of Verma et al.'s [14] paper. Verma et al. make a distinction between CCSs like IBM Watson, which they gloss as a centralized CCS, and a more distributed or decentralized form of CCS, which they refer to as a distributed brain. A key difference between these categories of CCS is that the cognitive capabilities of the distributed brain are realized by the coordinated activity of materially-heterogeneous elements (e.g., cloud computing sites, robots, sensor systems and human individuals), each of which participates in the network-mediated exchange of cognitively-relevant information. Such a vision, it should be clear, has a natural affinity with work in the area of distributed cognitive science [18]. Indeed, Verma et al. see the distributed brain as a specific kind of distributed cognitive system—one that is poised to press maximal benefit from contemporary network environments as well as a multiplicity of network-enabled computational devices.

Verma et al. point to a number of challenges associated with the development of a distributed brain. They suggest that such systems need to be resilient and self-healing in the sense of exhibiting a degree of fluidity with respect to their structural organization. Distributed brains thus need to support the dynamic, *ad hoc* assembly of information processing circuits in a manner that respects the constraints imposed by the network environment (e.g., connectivity constraints), while simultaneously satisfying the demands of a particular (cognitive) task. Verma et al. also note that such systems should be predictive and proactive, in the sense of anticipating the changing nature of "situations on the ground." Finally, Verma et al. suggest that the distributed brain should be organized so as to make best use of available computational resources. In particular, information processing assets should be strategically

placed within the network so as to avoid the unnecessary transfer of large data sets.

Such capabilities have a distinctly metacognitive feel to them. There is thus a sense in which the properties of the distributed brain are similar to those of other cognitive organizations, especially when it comes to issues of adaptive assembly and dynamic configuration (see, for example, Clark's [9] discussion of the Principle of Ecological Assembly). Beyond this, however, the appeal to neurological metaphors (in the form of the distributed brain concept) helps to establish an important link with recent work in cognitive neuroscience. In particular, there is an interesting parallel between the use of predictive coding regimes in the biological brain [19] and the attempt to achieve economical modes of information transfer via the strategic positioning of computational assets [14]. In both of these cases, the informational exchanges between processing elements are organized in such a way as to minimize the transmission of unnecessary or redundant information. The emphasis on dynamic, *ad hoc* networks in Verma et al. also establishes an interesting point of contact with work that portrays the biological brain as a federated economy of information processing assets. In particular, neural circuits are seen to be subject to a form of dynamic assembly and configuration, such that the neurocomputational capabilities of the biological brain are ones that are delicately tied to the cognitive demands of specific tasks [16][20]. Finally, the emphasis on predictive and proactive capabilities establishes a natural point of contact with predictive processing views of the biological brain [19]. Such views see the biological brain as engaged in the constant attempt to predict the flow of sensory information at a variety of increasingly abstract temporal and spatial scales.

Such parallels point to an interesting form of convergence between the distributed brain and the biological brain. While Verma et al. rely on the concept of the distributed brain to draw attention to the cognitive power and potential of distributed CCSs, their work also helps to reveal a deeper (and more subtle) set of commonalities between the processing dynamics of biological brains and the computational systems that attempt to emulate their cognitive capabilities.

## V. Conclusion

The papers presented at this year's COGWEB special session highlight a range of issues that lie at the intersection of a number of scientific disciplines. Each of the papers makes a distinct contribution to the existing literature in this area, and each embraces a somewhat different view of the cognitive significance and relevance of the Web. In spite of this, there are a number of common threads and themes that run across the papers, and these serve as the focal points for future debate and discussion.

All of the papers in this special session focus on issues of machine intelligence or the cognitive capabilities of hybrid (human–machine) systems. This is perhaps surprising given the interest and concern that is typically expressed in relation to human cognition (e.g., [21]). The current focus on machine intelligence serves as potent reminder that the cognitive significance of the Web is not exhausted by its effects on the human mind. The Web is an environment that provides a dizzying array of opportunities for social scaffolding, distributed computing, data analysis and various forms of socio-technological bonding. Relative to this vision, we should not expect the cognitive effects of the Web to be limited to a particular form of cognitive organization. The Web is a system that may very well be poised to modify multiple kinds of minds: the minds of ourselves, the minds of machines, and the minds of everything in between.

## References

[1] 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.

[2] T. W. Malone and M. S. Bernstein, Eds., *Handbook of Collective Intelligence*. Cambridge, Massachusetts, USA: MIT Press, 2015.

[3] P. Michelucci and J. L. Dickinson, "The power of crowds," *Science*, vol. 351, no. 6268, pp. 32–33, 2016.

[4] P. Krafft, K. Zhou, I. Edwards, K. Starbird, and E. Spiro, "Centralized, parallel, and distributed information processing during collective sense-making," in *ACM Conference on Human Factors in Computing Systems (CHI'17)*, Denver, Colorado, USA, 2017.

[5] L. L. Yu, J. V. Nickerson, and Y. Sakamoto, "Collective creativity: Where we are and where we might go," in *Collective Intelligence*, Cambridge, Massachusetts, USA, 2012.

[6] B. Sparrow, J. Liu, and D. Wegner, "Google effects on memory: Cognitive consequences of having information at our fingertips," *Science*, vol. 333, no. 6043, pp. 776–778, 2011.

[7] G. Fitzsimmons, M. Weal, and D. Drieghe, "Skim reading: An adaptive strategy for reading on the Web," in *6th ACM Conference on Web Science*, Bloomington, Indiana, USA, 2014.

[8] R. I. M. Dunbar, "Do online social media cut through the constraints that limit the size of offline social networks?" *Royal Society Open Science*, vol. 3, no. 1, p. 150292, 2016.

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

[10] 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.

[11] J. E. Kelly and S. Hamm, *Smart Machines: IBM's Watson and the Era of Cognitive Computing*. Chichester, West Sussex, UK: Columbia University Press, 2013.

[12] P. R. Smart, "Machine intelligence and the Social Web: How to get a cognitive upgrade," in *9th International Conference on Advanced Cognitive Technologies and Applications (COGNITIVE'17)*, Athens, Greece, 2017.

[13] N. O'Leary, D. Braines, A. Preece, and W. Webberley, "Conversational homes," in *9th International Conference on Advanced Cognitive Technologies and Applications (COGNITIVE'17)*, Athens, Greece, 2017.

[14] D. Verma, G. Bent, and I. Taylor, "Towards a distributed federated brain architecture using cognitive IoT devices," in *9th International Conference on Advanced Cognitive Technologies and Applications (COGNITIVE'17)*, Athens, Greece, 2017.

[15] D. Kirsh, "Embodied cognition and the magical future of interaction design," *ACM Transactions on Computer-Human Interaction*, vol. 20, no. 1, pp. 1–30, 2013.

[16] K. J. Friston, "Functional and effective connectivity: A review," *Brain Connectivity*, vol. 1, no. 1, pp. 13–36, 2011.

[17] L. S. Vygotsky, *Mind in Society: The Development of Higher Psychological Processes*. Cambridge, Massachusetts, USA: Harvard University Press, 1978.

[18] E. Hutchins, *Cognition in the Wild*. Cambridge, Massachusetts, USA: MIT Press, 1995.

[19] A. Clark, *Surfing Uncertainty: Prediction, Action and the Embodied Mind*. New York, New York, USA: Oxford University Press, 2016.

[20] M. L. Anderson, *After Phrenology: Neural Reuse and the Interactive Brain*. Cambridge, Massachusetts, USA: MIT Press, 2014.

[21] K. K. Loh and R. Kanai, "How has the Internet reshaped human cognition?" *The Neuroscientist*, vol. 22, no. 5, pp. 506–520, 2016.