This is a pre-publication version of a paper that will appear as *Smart, P. R., Engelbrecht, P. C., Braines, D., Strub, M., & Giammanco, C. (2010). The Network-Extended Mind. In D. Verma (Ed.), Network Science for Military Coalition Operations: Information Extraction and Interaction. Hershey, Pennsylvania, USA.: IGI Global*. The text of this pre-publication version differs from that which will appear in the aforementioned book. Copies of the book can be purchased from Amazon.co.uk (http://www.amazon.co.uk/Network-Science-Military-Coalition-Operations/dp/1615208550).

The Network-Extended Mind

Paul R. Smart[1], Paula C. Engelbrecht[2], Dave Braines[3], Mike Strub[4] and Cheryl Giammanco[5]

[1]School of Electronics and Computer Science, University of Southampton, Southampton, SO17 1BJ, UK.
[2]School of Psychology, University of Southampton, Southampton, SO17 1BJ, UK
[3]Emerging Technology Services, IBM United Kingdom Ltd, Hursley Park, Winchester, SO21 2JN, UK
[4]Dstl Porton Down, Salisbury, Wiltshire, SP4 0JQ, UK (on secondment from U. S. Army Research Laboratory).
[5]Human Research & Engineering Directorate, U. S. Army Research Laboratory, Aberdeen Proving Ground, Maryland, 21005-5425, USA.

Abstract
Whereas the traditional view in cognitive science has been to view mind and cognition as something that is the result of essentially inner, neural processes, the extended cognition perspective claims that at least some human mental states and processes stem from complex webs of causal influence involving extra-neural resources, most notably the resources of our social and technological environments. In this chapter, we explore the possibility that contemporary and near-future network systems are poised to extend and perhaps transform our human cognitive potential. We also examine the extent to which the information and network sciences are relevant to our understanding of various forms of cognitive extension, particularly with respect to the formation, maintenance and functioning of extended cognitive systems in network-enabled environments. Our claim is that the information and network sciences are relevant on two counts: firstly, they support an understanding of the mechanisms underpinning socially- and technologically-mediated forms of cognitive extension; secondly, they serve to guide and inform engineering efforts that strive to enhance and expand our cognitive capabilities. We discuss the relevance and applicability of these conclusions to current and future research exploring the contribution of network technologies to military coalition operations.

Introduction
The traditional view in the sciences of the mind sees the human brain as occupying a rather special place in the material fabric associated with the realization of human mental states and processes. One only has to flick through the pages of any contemporary text on cognitive neuroscience to appreciate the considerable dominance of what one might call the 'neurocentric view'. And it is a view that is reinforced by (and reflected in) a steady stream of brain imaging studies, many of which claim to have isolated the neuroanatomical basis of some aspect of our everyday psycho-cognitive functioning. The traditional view thus sees human mental states and processes as the direct product

of what the brain does. It claims that the machinery of the mind is housed largely within the head, and that to understand more about our cognitive profile we need to understand more about how the brain works. Eventually, it is claimed, we will have a complete theory of human cognition, and within this theory the human brain will occupy centre-stage.

The validity of this neurocentric, or intra-cranial, perspective has recently been challenged by those who embrace situated, embodied or distributed approaches to cognition (Clark, 1999; Haugeland, 1998; Hutchins, 1995a; Pfeifer & Bongard, 2007; Robbins & Ayded, 2009). Such approaches challenge the notion that mind and cognition are solely internal (neural) phenomena by emphasizing the role played by extra-neural and extra-bodily factors in shaping the profile of much real-world cognitive processing. One view that is perhaps maximally opposed to the internalist or individualistic conception of the human mind (the notion that the mind is the result of purely internal processes) is the thesis of the extended mind (Clark & Chalmers, 1998). This view explicitly endorses the idea that the human mind is not solely the product of what the brain does, and that the boundaries of the human mind are not necessarily co-extensive with the biological boundaries of the brain. Instead, the claim is that much of the machinery of the human mind extends beyond the brain to encompass a much larger nexus of extra-neural (and sometimes extra-organismic) resources. According to the extended mind perspective, human mental states and processes are not always in the head; they can sometimes extend beyond the brain to encompass aspects of the external technological and social environment.

Claims about the distributed or extended nature of human cognition are commonplace in the scientific and philosophical literature (Clark, 1997, 2003, 2008; Clark & Chalmers, 1998; Dennett, 1996; Haugeland, 1998; Hollan, Hutchins, & Kirsh, 2000; Hurley, 1998; Hutchins, 1995a; Kirsh, 1996, 2006; Norman, 1993; Wilson, 1994; Wilson & Clark, 2009). But what do such claims really amount to when we consider the potential impact of network systems and technologies on our current cognitive profiles? And what role do the information and network sciences play when it comes to understanding socially- and technologically-mediated forms of cognitive extension? One thing is relatively clear: it is that as we move into an era of pervasive computing and ubiquitous network access, much of our material world is becoming infused with greater computational potential, both for ourselves and the social collectives of which we are a part. If we want to understand the opportunities (as well as the hazards[1]) for cognitive transformation in this new era, we need to have theories and approaches that are capable of operating at the interfaces of the engineering, cognitive and social sciences. It is our claim, in this chapter, that the information and network sciences are a vital source of such theories and approaches; they are suitably poised to advance our understanding of the mechanisms underpinning socially- and technologically-mediated forms of cognitive extension.

Recognizing the contribution of the wider social and technological environment to cognitive processing (at both the individual and collective level) is of particular relevance in military coalition environments. Such environments are often conceptualized in terms of multiple interconnected networks (i.e. networks of networks) that subtend the human, technological and informational domains. Such networks interact in complex, non-linear ways throughout the course of coalition

[1] Not all forms of cognitive extension are necessarily guaranteed to impact cognition in positive ways; some forms of cognitive extension may prove deleterious to the cognitive capabilities of the larger system. This issue is taken up in the 'Human-Centered Cognitive Extension' section.

operations, and the challenge for military coalitions is often to coordinate the structure and activity of these networks in ways that meliorate cognitive performance. The ability of a specific coalition element to respond in an adaptive and intelligent manner is, for example, often based on the broader ability of the coalition formation to properly create, encode, select, retrieve, transform and communicate information-bearing structures (representations), and such information manipulation processes often need to be sensitive to the structure of existing communication and social networks. Notions of distributed and extended cognition have a special relevance here because they focus attention on the fundamental interdependencies between specific cognitive performances and the wider webs of social and technological scaffolding in which such performances take place. In this chapter, we aim to show why distributed and extended approaches to human cognition are relevant to our understanding of the inter-relationships between coalition networks and cognitive processing at both the individual and collective levels.

The structure of this chapter is as follows. The section entitled 'Externalism and Extension: A Brief History' provides an introduction to externalist approaches to the human mind. It reviews the key arguments associated with two forms of externalism, namely content externalism and vehicle externalism. Both of these forms of externalism raise doubts about the philosophical and scientific integrity of, what might be called, internalism (the idea that mind and cognition can be understood solely by focusing on internal, intra-cranial states-of-affairs). This leads on to a discussion about notions of cognitive extension in the section entitled 'Cognitive Extension'. Cognitive extension has been introduced using a number of real-world examples in the literature. These include long multiplication (see Wilson & Clark, 2009), ship navigation (Hutchins, 1995a), academic paper writing (Clark, 1997), puzzle solving (Kirsh, 2009; Kirsh & Maglio, 1994; Maglio, Matlock, Raphaely, Chernicky, & Kirsh, 1999), and the process of artistic creation (see Clark, 2001). In this chapter we introduce the notion of cognitive extension using a 'simple' non-cognitive example, namely the process of spider web building behaviour. This example is intended to show how a collection of capabilities that is ostensibly the product of a centralized neurological resource (the spider's nervous system), actually turns out, on closer inspection, to involve a variety of more far-flung forces and factors. Spider web weaving thus emerges as an example of what has, in the literature, been dubbed 'non-trivial causal spread' (Wheeler, 2005; Wheeler & Clark, 1999), a feature that characterizes many cases of environmentally-extended cognition. After presenting the case for network scientific approaches to extended cognitive systems in the section on 'Extended Cognitive Systems', we then present the extended mind thesis in the section entitled 'The Extended Mind'. An extended mind can be thought of as a particular kind of extended cognitive system, namely one that relies on the more or less permanent coupling of a human agent with cognitively-potent technological add-ons. The section entitled 'The Web-Extended Mind: A Thought (Provoking) Experiment' extends the discussion about the extended mind thesis and applies it to putative cases of cognitive extension involving the World Wide Web. In this case, we engage in a thought experiment regarding the close coupling of a human agent with near-future Web-based technologies. The thought experiment gives rise to a number of issues regarding potential shifts in our conception of ourselves as cognitively- and epistemically-bounded agents. Issues relating to socially-extended cognition (i.e. cases of cognitive extension involving other human agents) are reviewed in the section entitled 'Socially-Extended Cognition', and this is followed by a discussion of the kinds of cognitive extension that are likely to be encountered in military coalition contexts in the section entitled 'Extended Cognitive Systems and Military Coalitions'. A number of defence-related research programs, including the new

Network Science Collaborative Technology Alliance (CTA), feature research that is highly relevant to some of the issues raised in this chapter, and an overview of such programs is provided in the section on 'Relevant Defence-Related Research Programs'. Within the same section we briefly present work within the joint U.S./U.K. International Technology Alliance (ITA) research program, which specifically seeks to explore a number of issues related to cognitive extension in network environments. Such work, we suggest, can be seen as the intellectual lynchpin that connects work in many other research programs, such as the Network Science and Cognitive Neuroergonomics CTAs and the Tactical Human Integration with Networked Knowledge Army Technology Objective (THINK ATO). In the 'Conclusion' section we summarize the main arguments motivating a consideration of the information and network sciences to our understanding of network-mediated forms of cognitive extension. The section also reiterates the main points of relevance regarding cognitive extension research and military coalition operations.

Externalism and Extension: A Brief History

Historically, cognitive science has embraced a particular view of the mind, one which sees human mental states and processes as largely the product of inner, neural mechanisms. Human mental states and processes, the view maintains, are essentially realized by physical mechanisms inside the head of human subjects, and thus the mechanistic boundaries of the human mind are roughly co-extensive with those of the biological brain. This particular view of the mind (which, following Wilson and Clark (2009), we will refer to as individualism) maintains that the human mind can be studied and understood independently of any reference to the external environment. It essentially advocates what Jerry Fodor (1980) once referred to as 'methodological solipsism', the idea that the cognitive sciences can limit their study to the individual, effectively bracketing off the world in which the individual is embedded. On the individualist view, cognition is something that is wedged between perception (on the input side) and action (on the output side), constituting the filling of what Susan Hurley (1998) refers to as a 'cognitive sandwich'. The individualist conception recognizes the role of the wider environment as an input/output space for cognitive processes, but it does not afford any constitutive role for extra-organismic elements in those processes. Despite a recent emphasis on situated (Robbins & Ayded, 2009), distributed (Hutchins, 1995a) and embodied (Clark, 1999; Pfeifer & Bongard, 2007) approaches to cognition, internalism is still very much apparent in the sciences of the mind[2]. In addition, it is likely that some form of neurocentric individualism best captures our contemporary 'common-sense' notions about the material origins of the human mind. As Noë (2009) points out:

> *"We live in a time of growing excitement about the brain…Perception, memory, our likes and dislikes, intelligence, morality, whatever – the brain is supposed to be the organ responsible for all of it."* (Noë, 2009; pg xi)

Doubts about the integrity of individualism first arose during the 1970s in the work of Hilary Putnam (Putnam, 1975) and Tyler Burge (Burge, 1979). The predominant concern was that individualism failed to adequately account for the content or *meaning* of mental representations. Putnam (1975) thus argued that mental states could not be individuated in accord with the constraint of

[2] Harnad and Dror (2006) thus state "…cognition takes place entirely within the brains of cognizers…The causes and effects stretch more distally, but not the cognition; cognition begins and ends at the cognizer's sensor and effector surfaces."

individualism because the content of a mental representation might be determined by facts external to the individual, specifically the environmental or historical location of an individual. It is only by referencing these external facts that the meaning of an inner representational item can be discerned. This form of externalism, which has been referred to as taxonomic externalism (Wilson, 2000, 2004) or content externalism (Rowlands, 2006), is clearly contrary to the main thrust of the individualist thesis. It asserts that although intentional mental states exist as internal (e.g. neural) states of an individual, they are not (in virtue of their content) supervenient[3] on purely internal (intra-individual) factors.

Despite its appeal to physical, social and historical factors, the form of externalism just described (content externalism) is still largely committed to an internalist perspective about the location of the physical structures associated with mental states and processes. Even though the content of mental representations is deemed to depend on the external environment, the physical vehicles of cognition[4], it is claimed, are still likely to be situated within the head of the individual. An alternative, and more radical, claim is that even the physical vehicles of cognition need not be restricted to the internal realm. Instead, so the claim goes, the vehicles of both mental states and mental processes are perfectly able to *extend* beyond the head into the external world. Mind, and the cognitive processes that constitute it, sometimes extend into the physical and social environment of the individual human agent.

This second form of externalism (which we will refer to as vehicle externalism) goes by a variety of names, including locational externalism (Wilson, 2000, 2004), active externalism (Clark & Chalmers, 1998), vehicle externalism (Hurley, 1998; Rowlands, 2006), environmentalism (Rowlands, 1999), and the extended mind (Clark & Chalmers, 1998). What unites all these terms is a theoretical commitment to the idea that the physical boundaries of a cognitive system should not be assumed to coincide with the traditional biological boundaries of skin and skull. Instead, cognition is seen as something that is often an environmentally-extended process. To fully understand human cognition, it is not enough to focus solely on the inner states of the individual; for such a focus reveals only a partial picture of cognitive processing. Instead, we need to look beyond the individual, to understand the way in which cognition is fundamentally situated and embedded within a larger nexus of physical and social influences.

The notion of vehicle externalism, as just described, is something that will occupy us for the remainder of this chapter. However, before we embark on that discussion, it is important to point out that claims about vehicle externalism are largely orthogonal to those of content externalism. Content externalism is a theory about how the content of (inner) mental representations supervene

---

[3] The notion of supervenience represents a kind of dependency relationship between sets of properties. A set of properties (X) is said to supervene on another set of properties (Y) if objects that are indistinguishable from the perspective of Y properties are also indistinguishable from the perspective of X properties (see Braddon-Mitchell & Jackson, 2007). Thus mental states (x) supervene on brain states (y) if brain states that are physically indistinguishable are associated with mental states that are also indistinguishable. The claim of content externalism is that this is not the case: the content of mental states supervenes on facts that are external to the neurophysiological details.

[4] The vehicles of cognition are the physical states and processes associated with mental states and processes. The distinction between contents (as in content externalism) and the vehicles of contents (vehicle externalism) is a distinction between the content (or meaning) and the thing that has the content (or meaning). For example, the content of a written sentence is the meaning of the sentence, while the thing that has the content (or is the bearer of the content) is the sequence of written words.

on facts about the historical and environmental location of an individual. Vehicle externalism, in contrast, does not concern itself with how representational vehicles acquire the meaning or content they do; it is primarily a theory about the nature and interaction of the vehicles themselves. In advocating vehicle externalism, we are essentially committing ourselves to an understanding about how the physical, social and technological environment contributes to the material realization of specific states and cognitive performances; we are relatively less concerned with how those states and performances acquire their specific contents.

Cognitive Extension

The claims of vehicle externalism have a somewhat radical sounding flavour to them[5], but the notion that the physical vehicles of cognition are not restricted to the inner, neural realm is a notion that is perfectly compatible with the claims of both physicalism and functionalism (see Braddon-Mitchell & Jackson, 2007). And although the dominant view in artificial intelligence research (at least in the last century) was guided by predominantly individualistic and internalist conceptions of the mind, there is nothing in the bedrock claims of classical cognitivist theory (Newell, 1980; Newell & Simon, 1976; Pylyshyn, 1984) that necessarily binds intelligence to internally-situated mechanisms[6]. In spite of this, the claims of vehicle externalism are largely counter to our common-sense intuitions about the human mind, and this often results in a deep scepticism about the tenability of the core claims. To make both the claims of the vehicle externalist thesis clearer, and to invite a principled consideration of the relevant ideas, it helps to start with the simple (and in our case non-cognitive[7]). Therefore, this section begins with a simple, but powerful, demonstration of how intelligent behaviour can emerge from the delicate interplay of forces and factors that extend beyond the neural realm. In

[5] Indeed, the notion of cognitive extension has been the subject of a lively debate in the philosophical and cognitive scientific literature. Criticisms of the extended mind thesis centre on issues of cognitive and computational control (Butler, 1998), the distinction between intrinsic and derived contents (Adams & Aizawa, 2001, 2008, 2009, in press), and worries about the vulnerability of external resources to damage and social manipulation (Sterelny, 2004). All of these concerns have been addressed by Andy Clark in a series of recent publications (Clark, 2005, 2007a; Clark, in press-a, in press-b; Wilson & Clark, 2009). Clark (2008) provides a good summary of the criticisms and associated responses.

[6] Indeed, Edwin Hutchins (1995a) depicts the symbol-manipulating vision of classical cognitive science as, in fact, a vision of environmentally-situated problem-solving. According to this vision, the human agent implements a serial, symbol manipulating processing economy by virtue of his or her interaction with a variety of external props, aids and artefacts.

[7] The extent to which this behaviour is, in fact, non-cognitive depends very much on one's view of what constitutes cognition. The problem is that what is and what is not a cognitive process is often determined by ostensive definition. We can therefore point to examples of cognitive processing (e.g. perceiving, reasoning, thinking and so on), but establishing precisely what it is that makes something a cognitive process is much harder. Adams and Aizawa (2001) favour a view of cognition that highlights the role of representations with 'intrinsic' as opposed to 'derived' intentionality. Unfortunately, however, it is not entirely clear what is meant by the notion of intrinsic intentionality, or when we confront representations whose content is intrinsically given. Rowlands (2006) defines a cognitive process as "one that: (i) is required for the accomplishing of a cognitive task, (ii) involves information processing, and (iii) is of the sort that is capable a yielding a cognitive state" (pg. 32). In this definition, the notion of a 'cognitive task' is defined by ostension, and the notion of a cognitive state is construed as a genuinely representational state; i.e. a state that can be seen as representational in virtue of its satisfaction of a host of additional criteria. The main problem here, of course, concerns the fact that we are still relying on ostensive definitions for the notion of a cognitive task. We also encounter problems regarding the precise conditions under which a physical state should count as one that is genuinely representational.

subsequent sections, we expand on this initial case study and show how vehicle externalist views can be applied to more complex forms of 'human-level' problem-solving.

Web Construction

The web of the garden cross spider (*Araneus diadematus*), like that of most orb web spiders, is a compelling example of how our initial intuitions about the problem-solving potential of a seemingly simple bio-computational system (in this case an insect nervous system) can founder in the face of real-world performance. The spider's central nervous system is composed of a number of ganglia (collections of neural tissue), of which the most prominent are the supraesophagal and subesophagal ganglia. These ganglia serve to implement and coordinate the majority of the spider's sensorimotor functions. The total number of neurons in the central nervous system is small, about 30,000 neurons in the case of the orb-web spider *Argiope* and 100,000 neurons in the case of the larger wandering spider *Cupiennius* (Foelix, 1996). This compares with somewhere in the region of 100 billion neurons for the average human brain. Given the scale of the spider's nervous system, we might expect its behavioural capacities to be somewhat limited. And yet spiders are capable of surprisingly complex behaviours[8], of which the most well known is probably web construction. The spider's web is architecturally complex, composed as it is of multiple types of silk thread, each laid down in a specific sequence and geometric pattern. Specific types of thread need to be produced at just the right time, and the overall design of the web has to be sensitive to a number of factors including the size of the prey to be caught and the shape of the local environment (the shape made available by local branches or other supporting structures). The problem might be easier if it was possible to use visual information to guide action selection processes; however, *Araneus* is practically blind and does not rely on visual information to complete the web construction process (Witt, Reed, & Peakall, 1968). The average human being, blindfolded and presented with the task of creating a complex geometric structure from multiple types of building material, might be hard pressed to match the spider's feat of engineering, and this is despite the fact that our own neural systems far outstrip the size and complexity of those possessed by the average orb web spider. The feat of web construction seems to require a capacity for judgement, decision-making and planning that is profoundly out of kilter with our expectations and intuitions about what the spider should be capable of. So just how does the spider do it?

The answer seems to lie in the spider's exploitation of bodily contingencies and the power of the local environment to structure and guide action choice. A detailed ethological examination of web spinning behaviour suggests that spiders are sensitive to certain bodily contingencies involving the relative positioning of their legs on certain types of silk thread (Krink & Vollrath, 1997, 1998, 1999). As the web develops, the positioning of the legs becomes a reliable cue as to what type of action needs to be executed next, as well as what type of silk needs to be produced. In essence, the web serves as "its own best model" (Brooks, 1991) of what needs to be accomplished, and the spider need only be responsive to local information concerning the structural organization of threads in the immediate vicinity of its body. At each stage of the web construction process, each of the spider's legs need only perform a local (spatial) search for the nearest thread, and, once located, the relative positioning of the legs (as well as the type of thread they are in contact with) 'represents' the web's

[8] The araneophagic spiders, in particular, have been shown to engage in a variety of complex behaviours, ranging from optimal route selection (Tarsitano & Jackson, 1997) to deceptive signalling (Wilcox & Jackson, 2002).

structural status. In response to this rich body of local information, the spider need only implement locally-effective rules concerning which action to perform. And it turns out that aspects of spider web weaving behaviour can be modelled using a relatively simple (and minimal) set of rules (Krink & Vollrath, 1997, 1999). Importantly, each rule exploits facts about the spider's bodily design, and its outputs specify actions that are geared to structuring the problem space in ways that guide, constrain and simplify subsequent behaviour. The spider, it seems, distributes the computational burden associated with web spinning behaviour across a complex system that comprises its brain, body and aspects of the (self-structured) external environment[9].

So perhaps the reason we find the spider's web spinning behaviour both remarkable and mysterious (relative to its rather meagre neuro-computational resources) is because we fail to appreciate the behaviour for what it really is: a compelling example of environmentally-extended bio-morphological computation[10], one in which neural, bodily and environmental factors play representationally and computationally-significant roles. The central nervous system of the spider no doubt plays a very important coordinative role in the process of web construction, but it is only one element of a complex, environmentally extended system, and its representational resources and computational capabilities are geared not towards to the manipulation and transformation of abstract disembodied symbolic representations that occupy some inner, neural realm, but rather to the generation of temporally extended action sequences, actions that themselves serve to progressively structure and restructure the target problem-space in computationally- and representationally-potent ways.

The moral of this story, then, is that it is easy to be misled into thinking that intelligent action is always the sole product of neural mechanisms – that the point source of intelligent behaviour is always something that must reside in the 'head' of an agent. For what the case of web construction teaches us is that agents may often co-opt a variety of far flung forces and factors into a problem-solving routine, and not all of these forces and factors need to be biological in nature. We should not necessarily be surprised by this outcome. Evolution does not care about the material nature of problem-solving resources; it only cares about how those resources can be exploited to meet adaptive behavioural ends. Artificial evolutionary processes attest to the variety of ways in which seemingly irrelevant forces and factors may be co-opted into a design solution. Thus, in using genetic algorithms to evolve real electronic circuits, Bird and Layzell (2002) managed to create an 'oscillator circuit' whose systemic oscillatory behaviour was parasitic on the radio signals being generated from a nearby computer. In essence, the evolving circuit had generated the correct oscillatory behaviour, but had done so not by creating a genuine oscillator circuit; it had solved the problem by evolving radio reception capabilities and relaying the oscillations created by nearby circuits. Such phenomena are a common feature of many evolutionary processes. Thompson, Harvey and Husbands (1996) thus argue that during the evolution of electronic circuitry:

> *"…it can be expected that all the detailed physics of the hardware will be brought to bear on the problem at hand: time delays, parasitic capacitances cross-talk, meta-*

[9] Its body is (perhaps non-accidentally) designed so as to best exploit this state of affairs – you can represent quite a lot of information when your representational repertoire is sensitive to the spatial dynamics of a system comprising eight articulated appendages!

[10] Morphological computation concerns the way in which the physical body of a robot or organism can be used to perform computationally-significant functions (Paul, 2004, 2006).

> *stability constraints and other low-level characteristics might all be used in generating the evolved behaviour." (pg. 21)*

What we begin to see, therefore, is that for any given problem-solving process, evolution may often assemble solutions that pay scant regard to the manner in which the problem is solved. In many cases, the nature of the solution yielded by an evolutionary process will draw on whatever resources are available to meet the representational and computational demands of the problem at hand. And the responsibility for yielding adaptive behavioural success will, in many cases, be distributed across a broad coalition of neural, bodily and environmental resources.

We thus approach the main take home message of this section. It is that when seen in a certain light, the external environment emerges as more than just a space for sensory inputs and motor outputs; it is also poised to play an important (explanatorily-potent) role in the mechanisms by which that behaviour is realized. Intelligent behaviour, we might say, is at least sometimes realized by processing loops that *extend* beyond the neural realm and productively incorporate a variety of extra-neural resources. Some forms of behavioural intelligence are, we might say, environmentally-extended with regard to their mechanistic realization.

There is a parallel here – one that follows on nicely from the account of spiders and evolutionary processes – with Richard Dawkins' (1982) account of the extended phenotype. As part of his introduction to *The Extended Phenotype*, Dawkins (1982) encourages us to ignore the traditional biological boundary of the body and instead focus on the way in which external structures can form part of an extended system, one that is both created and maintained by specific genetic influences. From this 'extended' viewpoint, we can, he suggests, regard the spider's web as part of the spider's phenotype; it is a system that, just like the spider's body, determines the extent to which the spider's genes will be transmitted to future generations. The spider's web, when viewed through the special lens of the extended phenotype, thus emerges as a more-or-less equal partner in a complex matrix of phenotypic structures (some biological and others not) all of which are subject to evolutionary selection pressures.

But there is a deeper analogy here, one that goes beyond the level of extended phenotypes and extended behavioural mechanisms. It is the role that genes themselves play with regard to the generation of phenotypic structures. For in many ways, we suggest, the mechanisms by which genes control, regulate and contribute to the emergence of ontogenetic and cellular processes via their participation in genetic regulatory networks is directly analogous to the role played by the spider's nervous system in architecting its web. Just as the spider's web-spinning performances can seem remarkable relative to its available neuro-computational resources, so the morphological and physiological complexity of organisms can often seem surprising relative to the number of genes encoding their development (Claverie, 2001). Studies in functional genomics, for example, reveal that the number of protein-coding genes in the case of the human genome is about 20000-25000 genes (International Human Genome Sequencing Consortium, 2004), while that for the rather unsophisticated nematode worm, *Caenorhabditis elegans*, is a surprising 20,000 (C. elegans Sequencing Consortium, 1998). These results are surprising because, inasmuch as one sees genes as coding directly for specific aspects of physical form and function, one would have expected relative differences in large-scale phenotypic complexity to be reflected in large-scale differences in gene number. So how do we reconcile the apparent similarity of gene numbers in the case of *C. elegans* and *H. sapiens* with the apparent differences in physiological and structural complexity manifested by the two species?

One approach to answering this question is to emphasize the complex relationships that exist between an organism's physical structure and the genetic substrates that supposedly encode aspects of that structure. Thus, we now recognize that genes participate in complex regulatory networks that, in addition to producing structural proteins, also serve to constrain and control the expression of specific genes via protein-based feedback mechanisms (see Kauffman, 1995). Genes do not, therefore, seem to encode directly for specific aspects of physical structure; instead, they participate in the creation of complex networks of feedback and feedforward influence that, in conjunction with other factors, contribute to much of the biological complexity that we ultimately observe. Commenting on the surprising similarity of gene numbers between species, Buchanan (2002) points out that genes encode for proteins, and it is these proteins, interacting in complex webs of causal influence, that determines the differences between species. In order to understand the real role and function of genes, therefore, one needs to adopt a perspective that is specifically geared to understanding the complexity of network systems:

> *"To comprehend what makes us alive, and especially what distinguishes us from plants, will require insight into the architecture of this vast network; our sophistication is not due to one or another protein, but to the delicate design of the entire network." (Buchanan, 2002; pg. 16)*

The analogy with the spider's web building behaviour is thus revealed. In both cases what we seem to confront is the presence of a core biological resource (neuronal or genetic) whose function it is to create networks of causal influence (some of which operate in the manner of a closed-loop feedback control system). Such networks, in conjunction with the core biological resource, realize functions whose complexity far outstrips that made possible by the initial encodings or (in the case of the brain) computational processes. To see the core biological resource as causally-relevant to the final outcome (i.e. behaviour or phenotype) of the network in question is not, of course, incorrect, but it is important to give proper explanatory weight to the role played by the networks that extend beyond the boundaries of the core resource. And it is important, in both cases, to recognize the functional contributions of the neural and genetic resources for what they really are: mechanisms to create, maintain and exploit networks of causal influence that subtend a variety of organismic and extra-organismic resources. It is not possible, we suggest, to understand the proper function and significance of the core resource (genome and brain) in the absence of this network-oriented perspective, and we certainly cannot afford to restrict our scientific attention to these resources if we ever hope to understand how higher-level phenomena (such as biological structure and intelligent behaviour) are produced. For to divest these resources of their inter-relationships with the complex networks in which they participate (and often create) is to lose sight of something explanatorily vital in our quest to understand the contribution of those resources to the target phenomena of interest. It is lose sight of the fact that the functional significance of neural and genetic resources is often determined by networks that extend far beyond the neural and genetic realms.

Puzzles, Papers and Human-Level Problems

The critic will, of course, have identified a particular problem associated with the foregoing discussion, namely that, at least in the case of arachnid behaviour, we have focused on a form of intelligent behaviour that is far removed from the traditional targets of cognitive scientific enquiry (e.g. the realm of deliberative thought, planning, complex problem-solving and so on). This we accept, although it is not always clear to what extent ostensibly simple forms of adaptive behaviour should always be regarded as essentially non-cognitive in nature (see note 7). In spite of this, it is important to show how the notion of vehicle externalism can be applied to behaviours that are less controversially construed as cognitive. In the current section, therefore, we introduce a few more examples of intelligent behaviour in which human-level cognitive capabilities seem to draw on a variety of causal influences distributed across brain, body and world.

Consider first the case of multiplying two three digit numbers. A purely internalist account of how we are able to multiply the two numbers might emphasize how we first derive some symbolic encoding of the visual (or auditory) input corresponding to the two numbers. It would then invoke a computational account according to which the inner symbols are manipulated in some way so as to achieve the correct mathematical outcome. Now contrast this with what is surely a more accurate (ecologically-realistic) picture of how we implement long multiplication in the real-world. This alternative picture involves the active manipulation of external symbols in such a way that the kind of problem confronting the biological brain is profoundly simplified. In place of purely inner computational operations we see a pattern of perception-action cycles in which single digit numbers are compared and intermediate computational results are stored in an external medium using (e.g.) pen and paper. This example, described in Wilson and Clark (2009), is a case of what we might call environmentally-extended computation or 'wide computationalism' (Wilson, 1994). It takes what is, ostensibly, an inner cognitive capability (an ability to do long multiplication) and shows how crucial aspects of the problem-solving process can be (and usually are) delegated to aspects of the external environment. Importantly, the human agent in this situation emerges as a cognitive agent that (by virtue of culturally-scaffolded educational regimes) is able to make best use of a number of external props, aids and artefacts in order to meliorate problem-solving. Such melioration often occurs as a result of the way in which physical actions are used to structure and restructure aspects of the local external environment. In most cases, the result of the environmental restructuring is to radically simplify or transform the kind of problem-solving process in which the biological brain must engage.

Moving beyond the case of long multiplication, we encounter a number of cases where real-world action has been accorded an important role in enabling human subjects to navigate complex (and perhaps otherwise intractable) problem domains (Kirsh, 2009; Kirsh & Maglio, 1994; Maglio et al., 1999). David Kirsh (1995), for example, suggests a mechanism by which we are able to achieve success in the game of Scrabble[11]. Cast as a purely internal process, the cognitive demands of Scrabble seem considerable, but our problem-solving performances in the real-world often circumvent these overheads by relying on physical actions that simplify the kind of problem we are confronted with. Thus, in playing Scrabble, we typically engage in a process of active manipulation of the Scrabble tiles so as to construct spatial orderings and configurations that work in concert with the pattern matching and pattern completing capabilities of the human brain. Some initial (perhaps random) spatial orderings serve to prompt the recall of specific word candidates, and these can then

[11] See Maglio et al (1999) for some empirical research on this issue.

be evaluated and extended by further letter juxtapositions and spatial configurations. What is important here, as elsewhere, is to recognize the important and powerful role that physical action and environmental structure plays with regard to the larger problem-solving process. In the case of Scrabble and other problem domains (see Kirsh, 2009), Kirsh and Maglio (1994) suggest that certain types of action play key roles in enabling us to solve the problem in question. They refer to such actions as *epistemic actions*. These are actions that enable us to make information available[12] in ways that meliorate some aspect of our problem-solving performances. And it is epistemic actions, Clark (2008) suggests, that occupy centre-stage in discussions about how extended cognitive systems are brought into existence on the back of our active physical engagement with the external world:

> *"...epistemic actions, I want to suggest, are paramount among the ways in which bodily activity yields transient but cognitively crucial extended functional organizations." (Clark, 2008; pg. 70)*

As a final example of extended cognition in action (!), consider the process of writing an academic paper or report, such as the one that confronts you now. One view as to how we generate such artefacts might emphasize the role of purely inner resources in contributing to fully-formed thoughts, which are then serialized as words on paper. But this, of course, is seldom, if ever, how real academic texts get written. For better or worse, what generally tends to happen is that we start by writing down a few fragmentary thoughts and ideas, and these then prompt further thoughts and ideas. As the paper emerges, a variety of external resources, such as text and papers, often themselves heavily annotated with notes and marginalia, are continually consulted. As Clark (1997) argues:

> *"[the text] does not spring fully formed from inner cogitations. Instead, it is the product of a sustained and iterated sequence of interactions between my brain and a variety of external props. In these cases, I am willing to say, a good deal of actual thinking involves loops and circuits that run outside the head and through the local environment. Extended intellectual arguments and theses are almost always the products of brains acting in concert with multiple external resources. These resources enable us to pursue manipulations and juxtapositions of ideas and data that would quickly baffle the un-augmented brain." (pg. 207)*

Note that what is important here is the way in which some of the environmentally-extended processing loops are deemed to be *constitutive* of the thought processes giving rise to the finished article. Thinking, on this view, is not something that occurs solely within the head; it is also something that can be spread across a variety of extra-neural and extra-corporeal resources. Thinking, as with other types of cognitive processing, is sometimes literally extended into the world outside the head.

Extended Cognitive Systems

Our aim in this section has been to highlight the way in which some forms of intelligent behaviour seem to depend on the interaction of a variety of resources, including body morphology, environmental structure, and neural processing. In fact all of the examples presented in this section are examples of what has been called 'non-trivial causal spread' (Wheeler, 2005; Wheeler & Clark,

---

[12] For an extended discussion of the notion of 'making information available', particularly with respect to Gibson's (1966) theory of visual perception, see Rowlands (2006; pg. 34-40).

1999). This is something which occurs whenever we encounter a phenomenon that has the initial appearance of being the product of a well-demarcated system, but which, on closer inspection, turns out to involve the exploitation of a variety of more far-flung forces and factors. Whenever we have a case of non-trivial causal spread, we also have a case of explanatory spread; i.e. a relative expansion of our explanatory frameworks to account for the phenomenon in question. Such spread seeks to give explanatory weight to factors that we initially supposed were causally-irrelevant with respect to some target phenomenon. In cases where the target phenomenon is a cognitive process, then it makes sense to see the causally-active physical vehicles of the process as extending beyond the inner, neural realm. And, inasmuch as we equate the boundaries of a cognitive system with the physical limits of the mechanisms that comprise that system's cognitive processing routines, then cognition is, at least sometimes, not bounded by the traditional borders of skin and skull; it emerges as something that is perfectly able to extend beyond the head and seep into the world.

Of course, in order to make this radical-sounding claim stick, we need to do adequate justice to the notion that patterns of causal influence and dependence are sufficient to warrant a readjustment of cognitive system boundaries. It is not enough to claim that an external resource becomes part of the system simply because it exerts a causal influence on some aspect of system processing. What is needed is a clear understanding of *when* environmentally-situated cognition becomes a case of genuine cognitive extension. When, in other words, does some external tool or resource become incorporated into an agent's cognitive processing routines?

There are a number of ways to approach this problem (see Haugeland (1998) and Clark (2007b) for two related, but subtly different, accounts), but much clearly rests on the extent of functional integration between the candidate component and the larger system. We tend to recognize a functionally-unified system, we suggest, when the various components of that system participate in the realization of some goal or purpose it is the system's job to achieve. What seems to be important then in the case of cognitive extension is that we confront a set of distinct components (brain, body and worldly elements) that are connected together in such a way that their functional inter-operation makes them part of a functionally-integrated (yet internally differentiated) whole. In other words, what seems to be important is the specific way in which the components cooperate in the processing and exchange of information for the purposes of accomplishing some specific task or objective, a task that we typically identify as the responsibility of a specific agent (in most cases, an individual human agent). What makes something a part of an extended cognitive system, we claim, relates to the details of the functional connectivity and patterns of information flow and influence that characterize the inter-operation of the various system components. It is in precisely this sense that we conceive of an extended cognitive system as consisting of a network of heterogeneous elements, each of which makes a specific functional contribution to the shape and profile of the cognitive performances manifest by the larger system[13].

[13] This is not to say that such contributions are always indispensable – take away the physical rotation of Tetris zoids (see Kirsh & Maglio, 1994) and the subject may still be able to make do with purely internal rotational strategies. This does not, however, detract from the fact that when external resources are available, and productively coupled into ongoing sequences of neural operations and world-involving actions, they can still become incorporated into transient systemic wholes whose purpose is the efficient realization of a cognitive task.

Given this characterization of an extended cognitive system as a coordinated pattern of information flow and influence between networked components, it should be obvious why we see a role for the network and information sciences as contributing to our understanding of extended cognition. For such sciences are ideally poised to inform our understanding of how various heterogeneous components can interact in highly complex, nested and non-linear ways in order to realize cognitive functions. In addition, such a role is perfectly commensurate with the role to which such sciences are already being applied in the areas of neuroscience, economics, ecology, cellular biology, organizational analysis and epidemiology (Barabasi, 2002; Buchanan, 2002; Watts, 2003). Few would dispute the claim, we suspect, that network science is relevant to the project of understanding how large-scale neuronal ensembles are able to give rise to cognitively-interesting phenomena[14]; our claim is simply that the analytic targets of network science will often have to encompass a much broader range of resources when it comes to understanding the profile of much (but not necessarily all) real-world cognition. In this respect, the application of network science to extended cognitive science is perfectly compatible with existing research efforts in the information and network sciences; it simply extends the traditional focus of analysis to a much broader range of material resources.

One might, of course, be inclined to point out that the extended networks we see in the case of extended cognitive systems are not like those we encounter in conventional forms of network scientific analyses, especially those focused on the neural domain. The networks associated with an extended cognitive system seem to include a broad range of disparate elements (brains, bodies, and external artefacts), and this makes such networks unlike those that are the typical focus of neuroscientific enquiry. Doesn't the heterogeneity of elements within such networks mitigate against network-based analysis, and shouldn't we perhaps try to understand the capabilities of the neural sub-systems independently of the other, bio-external, elements?

We reject this claim for a number of reasons, not least because it is unclear whether the capabilities and performance profile of an extended cognitive system can be understood by a strategy of piecemeal decomposition and componential analysis (see the discussion on emergent capabilities in the section entitled 'The Web-Extended Mind: A Thought (Provoking) Experiment'). Moreover, the heterogeneity of extended cognitive networks is, in our view, a reason why we should embrace network- and information-based scientific approaches. The components that may comprise an extended cognitive system are indeed wildly disparate and various. They may include simple textual cues and prompts, or they may involve specific cognitive artefacts, such as slide rules, compasses, and so on (see Hutchins, 1995a). In some cases, the external technological resource may participate in computational processes independent of the human agent (e.g. mobile devices or decision support systems), or the resource may not even be technological in nature (it may, for example, be a another human agent – see the section entitled 'Socially-Extended Cognition'). Such heterogeneity merits and perhaps even necessitates the analytic techniques and conceptual theorizing of disciplines whose empirical targets are those of patterns of information-based flow and influence in materially-abstract functional organizations. The information and network sciences are ideally poised to provide this kind of abstract, functional analysis of extended cognitive systems.

[14] See Sporns (2002) and Sporns, Chialvo, Kaiser and Hilgetag (2004) for some applications of network science to systems-level neuroscience.

Another reason why we suggest the information and network sciences are relevant to the study of extended cognitive systems relates to the fact that we are not always solely interested in analysis. Part of our interest in understanding extended cognitive systems is to be able to engineer new systems, or at least engineer environments and resources in which cognitively-relevant mergers, interactions and alliances can be established. What we need to understand, as engineers, are the kinds of technologies that are apt for integration and incorporation into existing and sometimes novel cognitive routines. Some of this is, of course, the focus of existing and well-established scientific disciplines, such as the disciplines of Human-Centered Technologies and Human-Centered Computing (Norman, 1993, 1998). But in the case of our current profile of technological innovation and development, the sciences that deal with patterns of network-mediated interaction and influence have a special relevance. This is precisely because ours is an era in which information and communication networks, as well as a host of networked multimedia devices, are both pervasive and increasingly intertwined with our daily problem-solving activities and routines. If we are to exploit the power and potential of these new network-enabled environments, then we need tools, techniques and ways of thinking that are inherently sensitive to the features of network systems. It is precisely for this reason that the information and network sciences are relevant to our effort to understand and engineer network-mediated forms of bio-technological intelligence.

The Extended Mind

The previous section highlighted the way in which certain types of intelligent behaviour and cognitive processing seem to include (as wholes do their proper parts) mechanisms that extend beyond the traditional biological borders of skin and skull. The specific claim was that, under at least some conditions, we are warranted in seeing cognition as, quite literally, extending into the extra-organismic environment. The argument as currently presented, however, might be seen as applying to a narrow subset of mental states and processes, relative to those that we typically associate with a human mind. In accounting for much of the behaviour of both ourselves and others we typically make reference to a set of common-sense, mentalistic terms (such as belief, desire, hope, fear, and so on), and these are seen as playing a genuine explanatory role in psychologically-interesting patterns of behaviour. Thus my action to retrieve a beer from the fridge is explained in terms of my 'desire' to drink a beer and my 'belief' that a beer could be found in the fridge. It is this kind of intentional characterization (the ascription of intentional mental states) that helps us make sense of (to understand) patterns of human behaviour – it enables us to gain a predictively and explanatorily potent toehold on patterns of behaviour that would otherwise be psychologically unintelligible to us. So the question that arises in the case of cognitively-extended systems is whether the notion of cognitive extension gains any purchase in the more ethereal domain of folk-psychological discourse (the strategy of explaining human behaviour with respect to mental states, such as belief and desire). Can the notion of cognitive extension, as currently presented, be extended to account for the mental states that are posited as causally-relevant to the psychological understanding of our everyday patterns of behaviour? Can we, in other words, extend the case of an environmentally-extended cognitive system to the more general case of an environmentally-extended mind?

It is here (perhaps not surprisingly) that the philosophical waters begin to run deep. Perhaps the most lucid and influential account of why we should take notions such as extended belief states seriously is provided by Clark and Chalmers (1998) in their classic paper, 'The Extended Mind'. Clark and Chalmers (1998) ask us to imagine two individuals: Inga and Otto, both of whom are situated in

New York City. Inga is a normal human agent with all the usual cognitive competences, but Otto suffers from a mild form of dementia and is thus impaired when it comes to certain acts of information storage and recall. To attenuate the impact of his impairment on his daily behaviour, Otto relies on a conventional notebook which he uses to store important pieces of information. Otto is so reliant on the notebook and so accustomed to using it that he carries the notebook with him wherever he goes and accesses the notebook fluently and automatically whenever he needs to do so. Having thus set the stage, Clark and Chalmers (1998) ask us to imagine a case where both Otto and Inga wish to visit the Museum of Modern Art to see a particular exhibition. Inga thinks for a moment, recalls that the museum is on 53rd street, and then walks to the museum. It is clear that in making this episode of behaviour intelligible (or psychologically transparent) to us Inga must have *desired* to enter the museum, and it is clear that she walked to 53rd street because she *believed* that that was where the museum was located. Obviously, Inga did not believe that the museum was on 53rd street in an occurrent sense (i.e. she has not spent her entire life consciously thinking about the museum's location); rather, she entertained the belief in a dispositional sense. Inga's belief, like perhaps many of her beliefs, was sitting in memory, waiting to be accessed as and when needed.

Now consider the case of Otto. Otto hears about the exhibition, decides to visit the museum, and then consults his notebook to retrieve the museum's location. The notebook says the museum is on 53rd street, and so that is where Otto goes. Now, in accounting for Otto's actions we conclude, pretty much as we did for Inga, that Otto *desired* to go to the museum and that he walked to 53rd street because that is where he *believed* the museum was located. Obviously, Otto did not believe that the museum was on 53rd street in an occurrent sense (Otto has not spent much of his life constantly looking at the particular page in his notebook containing museum-related facts); rather, he entertained the belief in a dispositional sense. Otto's belief, like perhaps many of his beliefs, was sitting in the notebook, waiting to be accessed as and when needed.

Clark and Chalmers (1998) thus argue that the case of Otto establishes the case for a form of externalism about Otto's states of dispositional believing. The notebook, they argue, plays a role that is functionally akin to the role played by Inga's onboard bio-memory. If this is indeed the case, then it seems to make sense to see the notebook as part of the material supervenience base for some of Otto's mental states, specifically his states of dispositional belief (such as those involving museum locations). The main point of the argument is to establish a (potential) role for external artefacts in constituting the physical machinery of at least some of our mental states and processes. If, as Clark and Chalmers (1998) argue, the functional contribution of an external device is the same as that provided by some inner resource, then it seems unreasonable to restrict the material mechanisms of the mind to the inner, neural realm. It seems possible, at least in principle, for the human mind to occasionally extend beyond the head and into the external world.

Such claims are, understandably, disconcerting, and it is important that we understand the precise nature of the claim that is being made. One immediate cause for concern relates to the notion of functional equivalence between the inner (e.g. bio-memory) and outer (e.g. notebook) contributions. If we allow any form of externally-derived influence to count as part of the mechanistic substrate of the mind, then doesn't this cast the mechanistic net too widely? Don't we end up confronting cases that are so blatantly counter-intuitive that they undermine the very notion of the mind as a proper focus of scientific and philosophical enquiry? Consider, for example, the case where two people have a conversation on the bus. Does this mean that their respective minds have merged into one integrated whole? And what about cases where we have some very loose coupling with an external information source, say the kind of access we have to information in a conventional textbook? Clearly, not all of the technologies or external resources that we encounter are apt to engage in the kind of bio-technological hybridization envisioned by the extended mind hypothesis. As Clark (1997) argues:

> *"There would be little value in an analysis that credited me with knowing all the facts in the Encyclopaedia Britannica just because I paid the monthly installments and found space for it my garage" (pg. 217).*

Similarly, we suggest, it would be foolish to equate my personal body of knowledge and beliefs as co-extensive with the informational contents of the internet simply because I have an internet-enabled mobile phone. What, then, are the conditions under which we count a set of external resources as constituting part of an environmentally-extended mind? In answering this question, Clark and Chalmers (1998) embrace a particular set of criteria, ones that appeal to the accessibility, portability, reliability and trustworthiness of the external resource. The criteria are that:

1. "…the resource must be available and typically invoked" (Clark, in press-b). **[Availability Criterion]**
2. "…any information…retrieved from [the non-biological resource must] be more-or-less automatically endorsed. It should not usually be subject to critical scrutiny (unlike the opinions of other people, for example). It should be deemed about as trustworthy as something retrieved clearly from biological memory" (Clark, in press-b). **[Trust Criterion]**
3. "…information contained in the resource should be easily accessible as and when required". (Clark, in press-b) **[Accessibility Criterion]**

Clearly, such criteria serve to guide and constrain our intuitions about the kind of bio-artifactual and bio-technological couplings that are relevant to the formation of an extended mind. And they do so precisely because they delimit the range of situations under which we recognize the capabilities engendered by an external resource as being (most plausibly) that of a specific individual (or agent). In other words, what is important about the various criteria Clark and Chalmers (1998) propose is that they ensure that the capacities of an environmentally-extended, bio-technologically hybrid system are most plausibly seen by external observers (and perhaps by the agents themselves – see below) as the capacities and features of a particular agent. As Wilson and Clark (2009) suggest:

> *"We properly expect our individual agents to be mobile, more or less reliable, bundles of stored knowledge and computational, emotional and inferential capacities. So we need to be persuaded that the new capacities enabled by the addition of the notebook are likewise sufficiently robust and enduring as to contribute to the persisting cognitive profile we identify as Otto the agent. The bulk of Clark and Chalmers' (1998) work was an attempt to isolate and defend a specific account of the conditions under which we would be justified in identifying such an extended mind." (pg. 67).*

What Wilson and Clark (2009) are suggesting here, we think, is not that the conditions cited in Clark and Chalmers (1998) (the conditions of trust, reliability, portability and so on) are necessary for *all* forms of cognitive extension. Instead, they are suggesting that the conditions apply in the specific case of the extended mind, and perhaps even here – although Wilson and Clark (2009) do not explicitly state it – they are really only relevant to the specific case of dispositional beliefs. What we seem to confront then is a set of what might generally be referred to as *coupling conditions*, conditions that determine when we are and when we are not justified in identifying cases of cognitive extension that apply to the realm of folk-psychological theorizing. In all cases of cognitive extension, we claim, what is important is a particular pattern of temporally fine-tuned information flow and influence within a networked ensemble of diverse resources. This network constitutes the mechanistic substrate of an extended *cognitive* system whenever the objective of that system, or the task in which it is engaged, is recognizably cognitive in nature (see note 7). However, this networked ensemble need not be permanent in nature. It can be a one-off organization that is assembled for the purposes of a specific cognitive task, or it can be a temporary but repeatable organization that is assembled to deal with an intermittent or periodically-occurring task (see Wilson & Clark, 2009). When the organization is more permanent, we approach the kind of conditions under which we count the external resource as constituting part of the material supervenience base associated with an agent's daily patterns of psychologically-interesting behaviour. These are precisely the kind of conditions under which we are justified in seeing the emergence of an environmentally-extended mind.

The Web-Extended Mind: A Thought (Provoking) Experiment

Clark and Chalmers' (1998) original presentation of the extended mind thesis relies on a thought experiment involving a simple augmentative resource – a conventional notebook. It is perfectly correct and appropriate to ask whether this notebook is actually the kind of resource that could (in virtue of the kinds of human-artefact interaction it supports) fulfil the conditions for an extended mind. And, in fact, it is not clear that any actual notebook currently carried by a human agent could fulfil the criteria of portability, accessibility, reliability and so on, to the extent required. For all that, however, the main point of the notebook case was to highlight the mere possibility of an extended mind; it was not meant to suggest that most cases of notebook use actually result in genuine cases of cognitive extension. But now that the notions of cognitive extension and the extended mind have been fleshed out, we can dispense with such technologically low-grade examples and focus our attention on the role played by the rich variety of emerging technologies and resources that we see in today's hi-tech environment, most of them relying, in one form or another, on complex networks of information exchange, distribution and transformation. To what extent do ubiquitous modes of network-mediated information access, as well as portable devices and wearable computers, contribute to the technological realization of extended cognitive systems and the possibility of environmentally-extended minds?

To pursue this notion in the context of our own research we have posited an extension to the original thesis of the extended mind. The thesis is called the thesis of network-enabled cognition (Smart, Engelbrecht, Braines, Hendler, & Shadbolt, 2008) (or more recently the thesis of the network-extended mind), and it makes a specific claim about the role of network systems in constituting some parts of an extended computational system, one that is capable of implementing cognitive operations and contributing to the realization of certain contentful mental states. The thesis is as follows:

> *Thesis of the Network-Extended Mind: The technological and informational elements of large-scale information and communication networks can, under certain circumstances,*

constitute part of the material supervenience base for (at least some of) an agent's mental states and processes.

Clearly, one of the things to be assessed in evaluating this thesis is whether the kinds of technologies and resources that are being made available as a result of recent research and development in the electronics and computer science domains are sufficiently well-suited to meet the kind of criteria that Clark and Chalmers (1998) insist are important to the emergence of an extended mind. In some of our recent work we have examined this claim with regard to our (currently) best example of a large-scale networked information environment, namely the World Wide Web (Smart, Engelbrecht, Braines, Strub, & Hendler, 2009). What emerges from this analysis (see also Smart et al., 2008) is that, in many cases, the general trend of technological evolution is suitably well-aligned with the kind of criteria proposed by Clark and Chalmers (1998). Thus, in terms of concerns about portability we highlighted the fact that the current state-of-the-art in mobile computing devices has already given us devices that are at least as portable as the conventional notebook in Clark and Chalmers' (1998) discussion. Moreover, in terms of the accessibility of information content, it is significant that the focus of many research and technology efforts, particularly in the context of the World Wide Web, are geared towards improving user access to online information. Work of particular note here includes the development of natural language question-answering systems (Lopez, Pasin, & Motta, 2005; Tablan, Damljanovic, & Bontcheva, 2008), user-friendly semantic information browsers (schraefel et al., 2005), the use of sub-vocalization techniques to support Web navigation (Jorgensen & Binsted, 2005), and the use of intelligent forward caching and data charging mechanisms to mitigate download delays and the effects of intermittent network connectivity (Cherniack, Franklin, & Zdonik, 2001). New technologies in the field of wearable computing are also likely to enhance our access to information. Mobile device eyewear systems[15], for example, display information directly to a user's visual field using conventional eyewear equipment (e.g. spectacles). Some of the applications envisioned for this new technology include location-aware social network services, real-world visual overlays for environment navigation, battlefield situation awareness displays, and immersive virtual reality systems for education and entertainment. Such systems tend to reduce the cost of information access[16], and, we argue, they introduce new ways in which network-accessible information content can be co-opted into the information processing loops of cognitively-extended agents.

It is also important to note (and this is where our philosophical interests start to converge with our own scientific research programs) that as we move forward into an era of next-generation Web technologies, we are witnessing a move away from document-centric modes of information encoding to more data-centric modes. Document-centric modes of information encoding are those typically encountered on the conventional Web, where task-relevant information is often embedded in resources such as Web pages, often surrounded by (in many cases) irrelevant or redundant information. Think about the problem of accessing factual information from a web-accessible resource, such as Wikipedia. Even if the delays associated with document retrieval (i.e. downloading) and presentation are resolved, the user is still confronted with the onerous task of surveying the

[15] http://www.microvision.com/wearable_displays/index.html

[16] The notion of cost is important here because empirical studies suggest that the cost of accessing information from external resources has a significant impact on whether the resource is actually used (Gray & Fu, 2004). Information access cost is typically quantified in terms of temporal considerations, but it is possible that other types of consideration (such as physical effort) may also be important.

document for relevant information content. In most cases, this requires the user to scroll through the web page and process large amounts of largely irrelevant content in order to identify the small amount of information that is actually needed. This is a very inefficient means of information access. Even if the user tries to isolate specific information items for use on multiple occasions, they cannot do this without reliably fixing the physical location of the information (perhaps by copying the required information to a local resource[17]).

What is important for the emergence of network-extended minds, we suggest, are flexible modes of data integration, aggregation and presentation, in conjunction with an ability to gear information retrieval operations to suit the task-specific needs and requirements of particular problem-solving contexts. Such capabilities are being progressively unleashed by new approaches to data representation and information access on what is (presently) the core technological infrastructure of the conventional World Wide Web. Thus notions such as the Semantic Web and Linked Data[18] initiatives (Berners-Lee *et al.*, 2006; Berners-Lee, Hendler, & Lassila, 2001; Shadbolt, Hall, & Berners-Lee, 2006) countenance an approach to data modelling and representation that is largely independent of specific presentational formats or usage contexts. Commenting on the relationship of the Semantic Web to the conventional Web, Berners-Lee et al (2006) write:

> *"The SW [Semantic Web] tries to get people to make their data available to others, and to add links to make them accessible by link following. So the vision of the SW is an extension of Web principles from documents to data." (pg. 18)*

This shift of emphasis (from linked documents to linked data) is, we suggest, an important milestone in enabling the kind of selective data integration, aggregation, and filtering that undergirds the emergence of cognitively-extended systems and the mechanistic realization of extended mental states.

To make this vision a little more concrete, we present a thought experiment involving a near-future case of Web-mediated information access in the context of a fully interactive (in the sense of extended Web 2.0 capabilities), linked data web environment. Imagine that our future (in our case cognitively unimpaired) human agent is equipped with a mobile networked device (a mobile phone will do), an information presentation device (such as the aforementioned mobile eyewear devices, or the memory aids being developed by the Memory Glasses project at the Massachusetts Institute of Technology[19]), and a means of controlling information access and navigation in a simple and effective manner (for the sake of argument imagine an advanced form of the electromyographic, electroencephalographic and electrooculographic interfaces being developed by a variety of academic and commercial organizations (Mason, Bashashati, Fatourechi, Navarro, & Birch, 2007; Nicolelis, 2001; Pfurtscheller, Scherer, & Neuper, 2007; Stix, 2008))[20]. Thus equipped, our future agent is able to retrieve information from the Web, on demand, in a manner that is delicately

[17] Links to sections within the page will not work because Wikipedia, like most Web 2.0 applications, features dynamic content, and the physical location of specific information items is liable to change across multiple usage contexts.

[18] http://linkeddata.org/

[19] http://www.media.mit.edu/wearables/mithril/memory-glasses.html

[20] This is probably the most problematic aspect of our discussion: how to afford access to and interaction with network systems in low-cost ways. We revisit this later on in the discussion about context-aware information retrieval (see 'Human-Centered Cognitive Extension' section).

geared[21] to shaping, influencing and constraining ongoing sequences of thought and action. Our subject could, for example, be guided as to the location of interesting spatial targets by the use of simple geo-registered directional indicators overlaid onto the visual field. Our subject would not, therefore, have to rely on bio-memory to recall facts, such as the location of the Museum of Modern Art, because location-aware services would retrieve and present this information in a way that would serve to guide ongoing behaviour. Similarly, imagine that our subject has an interest in baseball and that baseball facts and figures are continually posted on the Web in a form that permits flexible forms of retrieval, combination, aggregation and inference (e.g. using the Resource Description Framework[22] or Web Ontology Language[23]). In this situation, our subject would be able to retrieve any piece of baseball-related information, on demand, in a manner that is robustly and continuously available. What, we might wonder, would our scientific, social and (indeed) subjective intuitions be in such a situation? Would it be appropriate for us to say that the subject pretty much 'knows' everything there is to know about baseball, at least in terms of the information that is posted on the Web? If this claim seems profoundly implausible or inappropriate to you, think for a moment about what it is that determines what you think you already know. What seems to determine whether we know or do not know something is not the fact that we are continuously, consciously aware of relevant facts and figures. What seems to count is more the kind of access we have to the relevant information, the fact that when we need to recall the information it is there, easily (and sometimes not so easily) made available to us by our bio-memory systems. But need our bodies of personal knowledge be so reliant on biologically-based modes of information storage? What if our access to externally-located information was just as reliably, easily and continuously available as the kind of access afforded by our own bio-memories? It this case, it seems, there is no principled reason to suggest that the external information would not count as part of your own personal body of knowledge and dispositional beliefs. As Clark (2003) argues:

> *"..it sometimes makes both social and scientific sense to think of your individual knowledge as quite simply whatever body of information and understanding is at your fingertips; whatever body of information and understanding is right there, cheaply and easily available, as and when needed." (pg. 42).*

If this is indeed what it means to know something, then the epistemic implications of our future contact with network systems and resources could be significant. For in such situations the boundaries of what we know seems to be limited only by the accessibility we have to various sources of environmental information, and if that information consists in the sum total of human knowledge, as stored in some large-scale networked space, then the epistemic limits of the network-extended mind are of a scale and potential that surpasses anything we have yet seen in the course of human history. What might be the long-term effect of such a cognitively-extended system on our familiar notions of knowledge-guided competence? And what might be the effect of such forms of epistemic contact on our core notions of who and what we are?

---

[21] In the sense that only relevant information gets presented. The mode of information presentation is also important here. In particular, it is important to avoid concerns about information overload (see the section on 'Human-Centered Cognitive Extension'). Ideally, information should be presented in the form of simple, perhaps subliminal (see DeVaul, Pentland, & Corey, 2005), cues and prompts that serve to guide thought and action in cognitively productive ways.

[22] http://www.w3.org/RDF/

[23] http://www.w3.org/2004/OWL/

One thing that is worth considering at this point is that a potential shift in our notions of knowledge-guided competence, as applied to other agents, might also be accompanied by a correlative shift in our own subjective impressions of ourselves. In order to make this idea intelligible, think for a moment about the light in a refrigerator. If we did not know better we might be inclined to say that the light in the refrigerator is always on. Indeed, whenever we open the door to check whether the light is on, the light is, in fact, on. It seems to us as though the light is continuously lit because it is lit whenever we choose to look at it. In a similar vein, our sense of the detail in a visual scene may be attributable to the fact that the details of the scene are always made available to us whenever we try to look for them (see Myin & O'Regan, 2009). Our sense of 'seeing all the detail' in a visual scene is not necessarily because all aspects of the scene are explicitly represented[24]; rather, our conscious experience of seeing all the detail stems from the fact that we can continually visit and revisit all aspects of the scene (by moving our head and eyes) whenever we feel the need to do so. Arguments such as this form part of an influential theory of our conscious experience (Noë, 2004; Noë, 2009; O'Regan & Noë, 2001), which emphasizes how our subjective perceptual experiences are dependent on an implicit knowledge of sensorimotor dependencies (knowledge or expectations concerning the effect of movement or change on sensory stimulation). The claim that we want to make here is that this approach to accounting for our conscious experience may also be relevant in accounting for what we 'feel' or 'sense' we do and do not know. In this case, our sense of what we know would be guided by our ability to make knowledge and information available whenever we choose to do so, or are required to do so. In a way that is similar to our sense of the detail in a visual scene, we sense that we know something because the thing that is known can be easily accessed and co-opted into ongoing problem-solving sequences whenever it needs to be so. The claim is that if, by virtue of our experiences, we come to learn that certain bodies of information and knowledge (perhaps past experiences) can be easily accessed at will, then we will genuinely *feel* as if those bodies of knowledge and information are part of us, that they are part of *our* personal body of knowledge and experience.

This touches on an issue that we address in other work (Smart, O'Hara, Engelbrecht, Giammanco, & Braines, in press), namely the extent to which 'our' memories can be externally-located and perhaps even externally-manipulated. What we suggest is that our memories need not always be in the head, and if they are not in the head then they can be manipulated in a variety of ways. This notion of manipulation touches, of course, on the classic studies in false memory research (e.g. Loftus, 1997), but the implications are somewhat broader here. What we suggest is that if our personal memories are partially constituted by what is outside the head, then we open up opportunities for radical forms of re-personalization, experiential reprogramming and memory configuration. In the extreme case, imagine if your sense of what your memories are is partially constituted by your access to various sources of external information. Now imagine that if, after some head trauma (or perhaps deliberate neurological intervention), you lose all your bio-based memories. Now your memories are entirely constituted by what is made available to you by your external cognitive aids. If you wished, you could have someone manipulate the information contents of those aids and give you, what is in effect, a new set of memories!

Whether such claims can be substantiated or not is something that only future research and engineering efforts can address. For now, the main point of our argument is to highlight the mere

[24] See Noë (2004) for further discussion on this issue.

logical possibility of some forms of cognitive extension and to explore their implications in terms of our cognitive capabilities and potential. Such forms of cognitive extension are not necessarily contingent on any radically new forms of science fiction style neural rewiring or neural interfacing technology, as has been proposed in other parts of the scientific literature (e.g. Stix, 2008). They are, potentially at least, part and parcel of the current trend towards increasingly intimate forms of bio-technological merger with our best network-enabled devices and network-accessible information stores. They are, we might say, intermediate stopping points *en route* to our network-enabled cognitive destiny.

What this section has intended to show is that the notion of the extended mind is an important and powerful thesis when it comes to understanding the potential impact of new network-enabled technologies on our core notions of cognitive capability and knowledge-guided competence. The point is not lost on those who embrace an extended mind perspective. Thus Clark (2008) argues:

> *"…as we move toward an era of wearable computing and ubiquitous information access, the robust, reliable information fields to which our brains delicately adapt their inner cognitive routines will surely become increasingly dense and powerful, perhaps further blurring the boundaries between the cognitive agent and her best tools, props and artifacts." (pg. 41)*

As engineers, interested in technology-mediated modes of cognitive augmentation, we can and should strive to support the emergence of systems that meet the criteria for cognitive and mental extension. The philosophically-derived coupling conditions, in this case, provide a rough set of criteria as to the required performance characteristics of putative mind-extending technologies. Such criteria obviously need to be supported by future empirical studies regarding the specific kinds of information access that are required to motivate a shift in our social, scientific and (perhaps) subjective tendencies regarding the intentional characterization of behaviour. But the conditions clearly do provide a set of useful targets for future requirements analysis and requirements-driven technology development.

And what of the role of network scientific analyses and network-theoretic approaches in supporting the emergence of network-extended minds? We saw, in the case of extended cognitive systems, that network scientific analyses were merited by virtue of their potential to shed new light on the emergence, maintenance and operation of circuits supporting cognitive extension. Such merits are equally applicable when it comes to understanding the contribution of information and communication networks to network-extended minds. This is so, even though the nature of the external resource (e.g. network-enabled device or network-accessible information resource) may be somewhat more dynamic and invested with greater computational potential compared to the kind of cognitive artefacts featuring in traditional extended cognition/extended mind accounts. Another reason to embrace the network and information sciences in relation to the thesis of the network-extended mind concerns the contribution such approaches make with respect to the development and configuration of new network-enabled artefacts and networked environments. Given the potential for our minds to become partially constituted by external technological resources, it is surely important that we seek to design those technologies so as to deliver the best profile of cognitive performance capabilities and, in the case of our adversaries, limitations.

There is a one particular sense in which network-theoretic approaches are perhaps crucial to our understanding of the role of new technologies in building network-extended minds. It concerns the way in which the integration of external resources into a hybrid cognitive processing routine (one that straddles the biological and technological realm) sometimes results in the emergence of capabilities and competences that are not reducible to those of the constituent parts. In some of our most compelling cases of cognitive extension, the incorporation of an external resource does not merely result in the augmentation or enhancement of some well-established ability; it engenders entirely new forms of cognitive processing capability. One has only to think of the impact that written and spoken forms of language have had on our cognitive profile (see Clark, 2008; chapter 3) to appreciate the extent to which our cognitive potential can be transformed following certain forms of cognitive merger and integration. And it is here that network science plays another potentially significant role. For network science, as a specialized branch of complexity science, is concerned with themes of self-organization, emergence and systems-level thinking. This makes it ideally poised to deal with cases in which we cannot understand the abilities of an extended cognitive system by a process of piecemeal decomposition and additive reassembly. As Wilson and Clark (2009) point out:

> *"To understand the integrated operation of the extended-thinking system, created, for example, by combining pen, paper, graphics program, and a trained mathematical brain, it may be quite insufficient to attempt to understand and then combine the properties of pens, graphics programs, paper, and brains." (pg 73).*

The reason why this is inappropriate is suggested by areas of scientific study like cognitive neuroscience. Cognitive neuroscience does not try to understand the cognitive capabilities of the human brain by exclusively focusing its analysis on the processing potential of individual neurons. Instead, in order to understand the contribution of neurons to cognitive phenomena, cognitive neuroscientists embrace the principles of systems-level analysis, thinking and modelling. They do this precisely because the capabilities of large neuronal ensembles are not those of the individual elements (i.e. neurons) that comprise the ensemble. We properly recognize, in this case, that the phenomena of interest – the ones concerning cognitive processing capabilities – emerge at the systems level. So too when it comes to cases of cognitive extension. We should not necessarily assume that we can study the elements of an extended cognitive system in isolation from the complex webs of causal influence and informational exchange that effectively couple those elements into functionally integrated systemic wholes. For, in many of our most compelling cases of cognitive extension, the capabilities of the whole cannot be understood by a simple strategy of componential analysis. Network science, as a specialized branch of complexity science, should be at the heart of our effort to understand the actual and potential capabilities of network-based bio-technological organizations.

Socially-Extended Cognition

The discussion so far has focused on how external, technological resources may become integrated into extended cognitive systems centered on individual human agents. However, this discussion overlooks an important aspect of human cognition – the fact that it is often embedded in complex networks of social influence and interaction. What is the relationship between technologically-mediated forms of cognitive extension and forms of cognitive extension in which the external resources consist of other human agents?

In unravelling the notion of socially-extended cognition, there are a number of different perspectives that might be taken[25]:

1. The first is that within a large-scale information and communication network environment we might see a variety of socially-derived information resources as contributing to individual forms of cognitive extension. This notion of extended cognition is, at best, a weak form of socially-extended cognition. It emphasizes the role that social interactions and collaborations play with respect to the development and maintenance of external, shared resources. However, the resources in question are little more than virtual surrogates, or stand-ins, for more direct forms of social contact and communication.
2. A stronger form of socially-extended cognition sees other agents as directly constituting the supervenience base for individual forms of cognitive extension. In this case, an individual human agent (X) would become so tightly coupled with another human agent (Y), from the perspective of some cognitive processing routine, that Y would come to constitute part of the machinery associated with X's cognitive profile. Both would, essentially, become integrated into a single cognitive system. Whether the right kind of coupling relationship between the agents could ever, in practice, be established is unclear, but some theorists, such as Tollefsen (2006), seem favourably inclined to such a view.
3. The strongest from of socially-extended cognition is what might be called the group mind or collective mind thesis. The idea here is that a group of human agents is so organized (with regard to the flow of information and influence between them) that the group itself becomes the bearer of genuine mental states.

Clearly, the third of these is the most contentious option, and few theorists seem inclined to support it[26]. Rather than try to review or progress the philosophical debate surrounding this issue, our aim, in the current section, will be to highlight a number of issues and observations that we see as most relevant to the future study of network-mediated cognitive processing involving multiple human agents.

Firstly, we suspect that the best way of thinking about socially-extended cognition is in terms of the role that contemporary and near-future network systems might have in coordinating the thoughts and actions of a group of problem-solving agents. One way of thinking about this is to consider our

[25] There is also, potentially at least, another option here. This is the idea that cognitive processes are distributed across the members of a group in such a way that neither individual forms of cognitive extension nor collective minds need emerge. Much of the work in distributed cognition (Hutchins, 1995a; Tribble, 2005) can perhaps be seen in this way.

[26] The key problem in this debate may centre on our ability to ascribe mental constructs to groups of people – groups of people just aren't the kind of things that behave in ways that warrant thought ascription. Something similar may confront individual forms of the extended mind thesis. When we say, for example, that an agent, in conjunction with external technologies, solves a particular problem, there is a potential mismatch between the system that is causally implicated in the expression of behaviour (the extended cognitive system) and the thing to which mentalistic constructs get ascribed (the patterns of behaviour of a specific component of the larger system, namely the human agent). Given that beliefs are what gets ascribed to patterns of behaviour, and the behaviour that is produced (at least in the case of human agents) always results from proximate mechanisms that reside in the biological realm, does this mean that we have an inherent tendency (or bias) to always perceive the biological agent as the proper target of mental state ascriptions? Do we have an inherent tendency to discount the wider nexus of extra-biological causal influences that ultimately contributes to the profile of behaviour warranting thought ascription?

earlier claim that many episodes of externally-directed cognitive processing or intelligence are at least partly constructive in nature (recall the role of genes or a spider's 'brain'[27] in actively creating structures that subsequently contribute to much of the complexity we observe at the phenotypic or behavioural level). Can something like this vision be applied to the socio-cognitive realm, the realm where cognitive processes are distributed across a network of interacting human agents?

One way in which the notion might be unpacked is by drawing attention to the way in which many cognitively-potent external resources are the creative result of the collective actions of multiple individuals. Thus consider the mechanisms that lie at the heart of termites' abilities to construct termite mounds. Much of this ability seems to rely on the use of stigmergic processes (Bonabeau, 1999; Bonabeau, Dorigo, & Theraulaz, 1999), processes that serve to progressively structure and coordinate collective action via the presence of simple external cues. As one termite drops a mud-ball, it leaves a pheromone marker that encourages other termites to deposit mud-balls nearby. As the collection of mud-balls increases in size, so specific architectural structures begin to emerge as the result of collective, pheromonally-mediated behaviours (see Camazine et al., 2001; chapter 18). The key point about such examples of collective intelligence and self-organization is that they show how the collective actions of multiple individuals can serve to progressively structure the environment (or at least a key problem-solving resource) in ways that meliorate some aspect of individual or collective problem-solving[28].

Perhaps, in the World Wide Web context, systems like Wikipedia are already good examples of this. Such systems highlight the role that networks (in this case physical, communication networks) play in enabling individual contributions to assemble complex resources that subsequently constrain, influence and shape the profile of individual (and perhaps collective) thought and action. Sometimes, when we are confronted with such resources, we are enabled to pursue cognitive goals that would be difficult, if not impossible, to accomplish by ourselves. An illustrative example of this may be the way in which scientific open access initiatives[29], in conjunction with global information networks, serve to facilitate creative insight and intellectual progress in the domain of scientific endeavour. As Stevan Harnad (1999) rightly notes, the Web allows us to accomplish something akin to 'scholarly skywriting' – scientific theories, thoughts, ideas, experimental results, and sometimes data, are made available in ways that are increasingly accessible to fellow academics and scientific colleagues. It is almost as if the outputs of scientific and intellectual enquiry were written in the sky for all to see.

One idea that we want to canvass here is that the key virtue of this mode of information distribution and dissemination is that it effectively establishes linkages between ideas, thoughts and concepts that would otherwise have been too widely separated to be linked by agents engaged in individual forms of reason-constrained thought and inference. Imagine, for example, that many of the scientifically-interesting ideas which we are capable of entertaining are the nodes in a complex network whose linkages correspond to the individual transitions in a reason-respecting chain of thought. Paths through this network of ideas would then correspond to the intellectual arguments

[27] Technically, the spider does not have a brain; its central nervous system is composed of a number of ganglia, of which the most prominent are the supraesophagal and subesophagal ganglia.

[28] Within the domain of computer science, the notion of stigmergy informs many approaches to complex problems concerning optimization, coordination and self-organization (see Ajith, Crina, & Vitorino, 2006).

[29] http://www.eprints.org/openaccess/

or theses that flow from some set of initial ideas, assumptions or observations. Such models, while perhaps faithful depictions of the inference chains of classical expert systems, seem congenitally ill-suited to capturing much of our human potential for creativity and insight. Perhaps this is because such models overlook the fact that our intellectual excursions are not limited to logically-constrained trajectories through a space of scientific ideas (an idea space); instead, at least in some cases, we are able to jump around in this space by virtue of our exposure to the thoughts and ideas of others. This, in conjunction with our ability to combine slow, deliberative forms of rational thought with a capacity for analogical reasoning and abstract pattern matching, enables us to effectively form conduits or shortcuts[30] to distant parts of the idea space, parts that would have been too distant or disconnected to be linked by individual (and socially-unaided) modes of exploration and search. The vision, then, is one of networks enabling individuals to exploit and benefit from mechanisms of collective search, establishing new trails through a space of ideas, some of which may, on occasion, result in discontinuous steps forward in scientific thinking, innovation and discovery.

In addition to the role of networks in supporting the collective creation of cognitively-potent artefacts and resources, there is a body of empirical research that draws close attention to the role of network structures in influencing collective problem-solving abilities. This research seeks to illuminate the specific role that factors like network topology play in enabling groups of problem-solving agents to make effective decisions and discover optimal solutions to problems. In one study involving human subjects, Mason, Jones and Goldstone (2005) explored the effect of different network topologies (e.g. small-world, random, full-connected, etc.) on the ability of groups of people to correctly guess a randomly selected number between 1 and 100. On each trial of the experiment, subjects attempted to guess the target number and were provided with feedback about the correctness/accuracy of their own response, as well as the responses of their immediate neighbours in the network (i.e. the human subjects to which they were directly connected). Mason et al (2005) found that when subjects were given simple problems involving a single target number, the fully connected networks were most effective in enabling groups to collectively settle on the correct solution. However, when the problem was more complex and involved a three-peaked payout function (one optimal solution and two sub-optimal solutions), the networks with the longest average path lengths were the most effective in enabling groups to find the optimal solution. Summarizing these results, Goldstone, Roberts and Gureckis (2008) conclude:

> *"Problem spaces requiring substantial exploration may benefit from networks with mostly locally connected individuals. The problem with fully connected networks is that everybody ends up knowing the same information, and they thereby become too like-minded, acting like a single explorer rather than like a federation of independent explorers."*

Similar results to these have been reported by Lazer and Friedman (2007), who conducted studies with synthetic agents, again using different network topologies. Their results suggest that when agents are dealing with complex problems, the more efficient the network is at disseminating information, the better the short-run performance of the system (relative to network structures that are less efficient at disseminating information). However, as the performance of the system is

[30] Such conduits may, in some sense, be akin to the short-cuts in small-world networks (Watts & Strogatz, 1998). The shortcuts reduce the path length between our current thoughts and ideas and those that are remote, or even impossible, to reach by virtue of reason-constrained forms of inference.

monitored across time, those network structures that are less efficient at disseminating information are able to deliver better performance outcomes. In essence, the more efficient networks are better at solving problems under heavy time-constraints; however, when temporal considerations are not so important, the less efficient networks are able to deliver better long-term performance outcomes.

We thus encounter strong support for the claim that networks supporting rapid information dissemination (small-world and fully-connected networks) are more suitable for what might be called simple or 'high-tempo' problems. This contrasts with the case where the problem to be solved is more difficult and can be tackled at a more leisurely rate. In this case, more locally-connected network structures may be preferable. The reason for sub-optimal performance (at least on difficult problems) in groups connected by low average path length networks (e.g. small-world networks) seems to centre on the group's tendency to prematurely settle on sub-optimal solutions – to be drawn into sub-optimal solution outcomes on the basis of initial shared information. Such results are of potential relevance to a number of findings in the social psychological literature. They include the phenomena of groupthink (Janis, 1982), production blocking (Diehl & Stroebe, 1987) and the common knowledge effect (Stasser & Titus, 1985)[31], all of which seem to be characterized by a group's inability to find optimal solutions based on some form of precipitant interaction or early information sharing.

The empirical results of Mason et al (2005) and Lazer and Friedman (2007) are important because they highlight two things about the role of networks in socio-cognitive processing. Firstly, the suitability of a particular network structure to enable a group of problem-solving agents to reach an optimal solution outcome may depend on both the nature of the task in which the group is engaged as well as the structure of the solution landscape. Secondly, the differential effectiveness of the network structure in supporting certain group-level outcomes may be accounted for by variables, such as the rate of information dissemination, that depend as much on the dynamic, time-variant functional connectivity of the network, as they do its static, structural characteristics. In respect of this latter issue, note that just because a network structure supports rapid information dissemination this does not mean that the actual flow of information through the network must be necessarily rapid. Agents or nodes within the network can effectively modulate the speed with which information is transmitted by selectively ignoring information, or by only intermittently processing information (in fact this was precisely one of the manipulations employed by Lazer & Friedman (2007)). In human networks, there are clearly a variety of factors that might contribute to the rate of information distribution. These include things such as the tendency to hoard information, willingness to cooperate, vulnerability to copying/transmission errors[32], and trust. Also, of course, in

---

[31] Hinsz, Tindale and Vollrath (1997) have also highlighted some of the dangers associated with a group's over-reliance on shared information. Such insights, in combination with the results reported here, should give us pause for thought when it comes to notions of shared situation awareness (Nofi, 2000) and shared understanding (Smart, Huynh et al., 2009). Inasmuch as the interventions used to enhance shared situation awareness and shared understanding depend on the sharing of common sets of information, it is important that we do not create a situation in which group-level problem-solving abilities are undermined as a result of trying to achieve some other human factors objective.

[32] Lazer & Friedman (2007) evaluated the impact of copying errors in their computer simulation studies. They report that, in the long-run, systems with high error rates in the copying process outperformed those in which copying errors were minimized. The explanation for these results seems to be the same as that proposed for the effect of network structure on performance, namely that "Error rates in copying….alter the balance

situations involving mobile ad hoc networks many nodes may be expected to have only occasional and intermittent connectivity, and this may effectively impede the spread of information between all network nodes.

The studies on socio-cognitive processing in group situations highlights the relevance of network scientific approaches to our understanding about how to analyze and engineer network environments so as to best support collaborative problem-solving and decision-making. There is clearly much more work to be done here, but one thing does seem relatively clear at this early stage: it is that the kinds of information and network-theoretic approaches we advocated in the case of individual forms of cognitive extension (i.e. extended webs of information flow and influence spun around individual human agents), are readily applicable to the study of systems in which the webs of information flow and influence subtend multiple agents. Whether one wants to refer to such systems as socially-extended cognitive systems, or group minds, is, to our mind at least, largely irrelevant (although much may depend on whether we recognize some higher-level agency to which cognitive states and processes can be readily ascribed – see note 26). What seems important is that cognitive processing can take place in group situations, and that much of it can be supported by features of the network structure that acts to mediate group interactions. In such situations, the tools, principles and techniques of information and network science are just as relevant to our ultimate understanding of the cognitive capabilities of social organizations as they are to our understanding of extended cognitive systems involving individual human agents.

Extended Cognitive Systems and Military Coalitions

In considering the possibility of cognitive extension in military coalition environments, we can discern two distinct ways in which cognition may be extended beyond the bounds of individual human agents. One of these forms of cognitive extension is centered on the individual human agent. It sees the cognitive capabilities of the human agent as, in part, realized by complex webs of information flow and influence between a variety of inner (biological) and outer (social, technological and informational) resources. This is the form of cognitive extension that is most commonly encountered in the philosophical and scientific literature, and it is the form of cognitive extension that has occupied us for most of the current chapter. There is, however, a second way in which cognitive processes may be extended beyond the biological borders of specific individuals. This is the form of cognitive extension that we encounter in cases of distributed cognition research (Hutchins, 1995a, 1995b; Tribble, 2005). It emphasizes the way in which cognitive processes inhere in the complex webs of information processing that connect multiple human agents with a variety of non-biological props, aids and artefacts. This form of multi-agent cognitive extension can be discriminated from individual forms by virtue of the emphasis placed on the larger socio-technical system in which much of the relevant cognitive processing is deemed to occur. Thus, while individual forms of cognitive extension focus on the individual human agent as the target system of interest, the distributed cognition movement tends to see the larger socio-technical system as the relevant unit of cognitive analysis. Relative to this larger system, the activities of individual human agents form part of a complex web of coordinated computational activity, one that serves to propagate and

between exploration and exploitation in the system, increasing the amount of experimentation but reducing the rate with which successful strategies spread" (Lazer & Friedman, 2007).

transform representations in ways that ultimately lead to coherent patterns of system-level behaviour.

These two forms of cognitive extension are, we suggest, highly relevant to our understanding of coalition-based cognitive capabilities. Although the multi-agent form of cognitive extension might, at first glance, seem more interesting and relevant from the perspective of military coalitions (not least because it affords an opportunity to see *entire* military coalitions as functionally integrated cognitive systems), we suggest that both forms of cognitive extension are, in fact, important foci of philosophical and scientific attention. The reason for this is that we see the global effectiveness of a military coalition as dependent (at least in part) on the cognitive capabilities of both individual soldiers and the wider socio-technical systems in which these soldiers are embedded. Increases in cognitive productivity at the individual level, as produced by cognitive extension, may be magnified many times once such human-centered extended systems are combined and integrated into larger webs of collective cognitive processing. In both cases (of individual and collective cognitive processing), our understanding of how to create, configure and maintain multiple types of networks in ways that best serve the information processing objectives of the larger coalition organization is of paramount importance.

In this section we review the opportunities and challenges for cognitive extension in military coalitions, focusing exclusively on the two forms of cognitive extension identified above. The section on 'Human-Centered Cognitive Extension' reviews issues and research associated with cognitive extension at the level of the individual soldier or warfighter; the section on 'Coalitions as Extended Cognitive Systems' explores the opportunities for cognitive extension at the level of entire military coalitions (or at least significant elements thereof).

Human-Centered Cognitive Extension

The notion of human-centered cognitive extension is simply the notion of cognitive extension that has occupied us for much of the current chapter. It is the idea that the physical machinery underpinning at least some of the cognitive capabilities of an individual human agent need not necessarily reside in the head of the human agent. In understanding how to support cognitive extension at the individual human level, we have argued that we should focus on the nature of the relationships between the human agent and network-accessible information resources. Thus, in order for human-centered extended cognitive systems to emerge, we need to ensure that the appropriate channels of information flow and influence are established between the human agent and the surrounding nexus of cognitively-relevant social, technological and informational scaffolding. One of the most important issues here concerns the *bi-directional* exchange of information between individual soldiers and other (non-biological) elements of the extended cognitive system. In particular, we need to ensure that the information provided by some external resource is sufficiently poised to guide response selection and response execution processes in adaptive and intelligent ways. Furthermore, the biological elements of the soldier-centered cognitive hybrid need to be appropriately interfaced with the non-biological elements such that the hybrid system can function as a single functionally-integrated cognitive whole. What this means, in practice, is the deployment of technologies that work in concert with the human agent – technologies that are sensitive and responsive to aspects of human psycho-biological functioning, and which are capable of adapting their functional profile to meet the problem-solving goals and objectives of the larger hybrid system. Research programs such as the DARPA-funded Augmented Cognition program and the recently announced Cognition and Neuroergonomics CTA both boast scientific and technology development goals that are directly aligned with these requirements.

One problem that seems particularly pertinent to the possibility of human-centered forms of cognitive extension concerns the way in which human agents are enabled to play an active role in the retrieval, structuring and transformation of information from non-biological sources. Thus, recall that in many cases of cognitive extension (see the earlier section on 'Cognitive Extension') what we seem to encounter are feedback-loops that involve the active manipulation of an external resource by a core biological agent. Recall, also, the thought experiment discussed in the section on 'The Web-Extended Mind: A Thought (Provoking) Experiment'. This thought experiment was intended to provide a vision of the impact of near-future technologies on our traditional notions of knowledge-guided competence at the scientific, social and (perhaps) subjective levels. But of all the technological elements described as part of that thought experiment, one element emerges as particularly problematic with respect to the current state-of-the-art. This is the way in which information retrieval operations (from bio-external media) are initiated and controlled by the human agent. In the thought experiment, we discussed the use of complex sensor devices that were sensitive to minute patterns of muscular or neural activity. However, the current functionality of such devices is limited, and it is not always clear that they can be used to good effect across different situations. In the military context, for example, soldiers are typically engaged in highly intense physical activity, and such activity interferes with the controlled and deliberate expression of both muscular and neural response profiles. Ongoing work within the DARPA-funded Augmented Cognition program, as well as the forthcoming Cognition and Neuroergonomics CTA, may help to address some of these issues (see Relevant Defence-Related Research Programs'), but, in the meantime, what other strategies might we pursue in order to support the retrieval and presentation of information in ways conducive to the emergence of network-extended minds?

One potentially relevant line of research here concerns the attempt to support context-aware modes of just-in-time information retrieval (Bahrami, Yuan, Smart, & Shadbolt, 2007; Rhodes & Maes, 2000). This research seeks to proactively present relevant information by monitoring specific aspects of the task or environmental context. Complementing this research effort is work in the ITA program that seeks to monitor and infer mission status information on the basis of both physical and contextual cues (Poltrock, Handel, Bowyer, & Waggett, 2008). Importantly, once we are able to detect features of the problem-solving context, we are able to proactively disseminate information to individual agents in ways that supports the effective realization of individual and collective problem-solving goals. Clearly, our ability to dynamically configure the physical network in a way that supports this mode of context-sensitive information distribution is of vital importance (see 'Coalitions as Extended Cognitive Systems'), as is our ability to create and exploit representations of (e.g.) coalition plans (Mott & Hendler, 2007) that could be used to control information distribution and adapt communication network topologies.

One concern in relation to network-mediated forms of information retrieval and presentation involves the notorious problem of information overload. In this sense, network access is both a boon and a burden. It is a boon inasmuch as it creates new opportunities for situation awareness and improved decision-making, but it is a burden inasmuch as it runs the risk of overwhelming the capacity of the individual human agent to adaptively exploit available information in the context of ongoing decision-making processes. There are a number of lines of research that might be pursued here. One strategy is to rely on the aforementioned mechanism of context-sensitive information retrieval to limit the amount of information that is presented to a user (Bahrami et al., 2007). Another is to rely on alert and notification systems that can be tailored to a user's specific goals, interests and concerns (Smart, Russell et al., 2009). There is also an important body of research that concerns the use of subliminal[33] cuing techniques to influence behavioural output (DeVaul et al., 2005). Such techniques are important because they provide a route to behavioural influence that does not involve conscious processing.

One question that we should ask in light of these ongoing research efforts is the extent to which the various technological add-ons, changes in information accessibility and so on, are genuinely enhancing or augmenting the cognitive capabilities of a particular human agent. The answer to this question is perhaps not quite as straightforward as it might initially seem, especially since there is nothing in the bedrock claims of the extended mind thesis to suggest that all cases of cognitive extension need to be uniformly beneficial from a performance perspective. Indeed, some commentators have suggested that network technologies may have a somewhat negative impact on human cognitive processing (e.g. Carr, 2008; Greenfield, 2003). Carr (2008), for example, bemoans the apparent impact the Web is having on his cognitive capabilities:

> *"As the media theorist Marshall McLuhan pointed out in the 1960s, media are not just passive channels of information. They supply the stuff of thought, but they also shape the process of thought. And what the Net seems to be doing is chipping away at my capacity for concentration and contemplation. My mind now expects to take in information the way the Net distributes it: in a swiftly moving stream of particles. Once I was a scuba diver in the sea of words. Now I zip along the surface like a guy on a jet ski" (Carr, 2008; pg 57).*

Clearly, we should not assume that such anecdotal reports provide any insight into the Web's true effects on human cognitive functioning[34]. Nevertheless, the cautionary flavour of Carr's (2008)

[33] Subliminal in this context means a perceptual cue that is presented at a level of intensity or duration below that necessary for it to become part of conscious awareness.

[34] One reason to be cautious of such claims is that the extended mind thesis obliges us to take a systems-level perspective when thinking about the capabilities of network-extended cognitive systems. Thus, just because some aspect of the psycho-cognitive functioning of an individual seems to have been altered as a result of a specific biotechnological merger, this does not mean that those capabilities (or extensions of those capabilities) are not manifest at the system level. To put this into context, think about the role that language plays in augmenting our cognitive capabilities (see Clark, 2008; chapter 3). It may well be that human agents are increasingly delegating many of their cognitive burdens to the Web, but is this really any different from the role that written and spoken forms of language already play for us? No one, we suspect, would be comfortable with the claim that we should abandon written forms of language because they undermine the (pure) cognitive integrity of the 'real' environmentally-decoupled human agent. And this is not just because individually and collectively we are better off, in a cognitive sense, for the development of writing systems. It is because such innovations are now so deeply integrated into our everyday problem-solving routines that the very notion of establishing a neat separation between the true capabilities of the human agent and their language-infected capabilities seems untenable. Many of us, we suspect, feel that linguistically-enabled capabilities are an intrinsic part of our own personal cognitive repertoire. We see language as less a form of technological enhancement and more an aspect of our own idiosyncratic cognitive profile. The long-term vision of the network-extended mind theorist is no different in this respect. The vision is that as network technologies become more permanent, reliable and accessible, so they will become increasingly integrated into our cognitive self image – our image of who *we* are and what *we* are capable of.

commentary is well taken, and, pending further research, we should perhaps be somewhat cautious of the kind of bio-technological unions we make ourselves susceptible to.

Aside from the potential negative effects of cognitive extension on human cognitive performance, it is not always clear that technologically-mediated forms of cognitive extension should always be considered augmentative, even when the presence of such technologies seems to bolster cognitive performance. The reason for this, we suggest, is that the boundaries of the extended cognitive system are not the same as the boundaries of the individual human cognitive agent, and, inasmuch as the cognitive performances in question are attributed to the extended cognitive system (rather than the cognitive agent), it may be inappropriate to regard the capabilities of the human agent as significantly altered by the emergence of environmentally-extended cognitive circuits. Here we see a potential tension with regard to notions of cognitive agency and the physical machinery that supports cognitive processing. When the mechanisms supporting a particular cognitive performance extend beyond the biological boundaries of an individual human agent, then we arguably confront a genuine case of cognitive extension. However, it is not always clear, in such cases, that the cognitive capabilities of the larger, mechanistically-extended cognitive system should always be equated with those of the individual, biologically-bounded, human agent. Something along these lines may underlie the apparent confusion in the philosophical and cognitive scientific literature concerning the augmentative status of a number of cognitive technologies. Thus, while many commentators talk of external resources acting to augment or enhance human memory, Hutchins (1995b) suggests that we should not see such resources as enhancing the memory of individual human agents *per se*; rather, we should see the augmented capabilities as those of a new human-technology hybrid system. For example, in discussing the way speed bugs[35] contribute to memory functions in an airplane cockpit, Hutchins (1995b) argues:

> *"Individual pilot memory has not been enhanced; rather, the memory function has now become a property of a larger system in which the individual engages in a different sort of cognitive behavior...To call speed bugs a 'memory aide' for the pilots is to mistake the cognitive properties of the reorganized functional system for the cognitive properties of one of its human components. Speed bugs do not help pilots remember speeds; rather, they are part of the process by which the cockpit system remembers speeds." (pg. 282-283)*

Such views serve to remind us that issues of cognitive extension cannot necessarily be divorced from ones of cognitive agency. In attempting to understand the extent to which the cognitive capabilities of agents are enhanced (or undermined) as a result of particular bio-technological mergers, we may need to account for how the boundaries of specific cognitive agents are identified and how the cognitive capabilities of those agents get ascribed.

Coalitions as Extended Cognitive Systems

In addition to seeing cognitive extension as something that can take place at the individual, human agent level, it is also possible (on occasion) to see much larger systems, comprising multiple agents and material artefacts, as extended cognitive systems. The distributed cognition movement, for example, seeks to account for the performance of large-scale socio-technical systems in terms of the interactions of multiple human agents with a surrounding nexus of social, technological and

[35] A speed bug is an indicator that highlights specific speeds on a airspeed instrument panel.

informational resources (Hutchins, 1995a, 1995b; Tribble, 2005). Within such systems, the cognitive capabilities of the individual human agent are important but, by themselves, they are often inadequate in terms of accounting for the real targets of scientific enquiry: the systemic cognitive properties of the larger system.

The notion that large-scale socio-technical systems may be seen (and analyzed) as cognitive systems in their own right has been championed by the cognitive anthropologist Edwin Hutchins. In his studies of both ship navigation (Hutchins, 1995a) and airplane piloting (Hutchins, 1995b), Hutchins makes it clear that the proper focus of cognitive scientific attention is not always the individual human agent; often it is the larger system of social and technological resources in which the human agent is embedded. In cases where we seek to understand the cognitive capabilities of a single human agent, of course, this perspective reduces to the case of individual forms of cognitive extension (cases where the cognitive capabilities of an individual human agent are of primary interest). But in a range of cases the cognitive system in question is composed of a larger aggregation of human agents working together in support of some common or shared goal, and it is in these cases that a consideration of entire military coalitions as extended cognitive systems seems most appropriate.

Of course, just because we confront a system in which cognitive processing is distributed across a much broader nexus of resources than is the case for individual forms of cognitive extension, this does not mean that the kind of coupling conditions we saw as relevant in the case of individual forms of cognitive extension are not equally important in the case of more distributed cognitive systems. Recall our core claims about the importance of functional integration in the case of human-centered extended cognitive systems:

> *What seems to be important then in the case of cognitive extension is that we confront a set of distinct components (brain, body and worldly elements) that are connected together in such a way that their functional inter-operation makes them part of a functionally-integrated (yet internally differentiated) whole. In other words, what seems to be important is the specific way in which the components cooperate in the processing and exchange of information for the purposes of accomplishing some specific task or objective, a task that we typically identify as the responsibility of a specific agent (in most cases, an individual human agent). What makes something a part of an extended cognitive system, we claim, relates to the details of the functional connectivity and patterns of information flow and influence that characterize the inter-operation of the various system components. It is in precisely this sense that we conceive of an extended cognitive system as consisting of a network of heterogeneous elements, each of which makes a specific functional contribution to the shape and profile of the cognitive performances manifest by the larger system.*

Much the same can be said when we confront a large-scale distributed cognitive system. Although it is not always clear that we can talk of such distributed systems as cognitive agents (at least in the same kind of way that we talk of human beings as cognitive agents), it does seem that the same kind of functional integration and coordination of the various system components is important to the cognitive outputs of the system. And just as we advocated the use of network scientific approaches in the case of human-centered cognitive extension, so it is important, we argue, to apply the tools and techniques of the information and network sciences to the case of large-scale distributed cognitive systems, ones in which a variety of biological and technological elements cooperate in producing globally-coherent patterns of systemic behaviour.

One reason why network-based analyses are important when it comes to military coalition systems is that such systems are typically seen as composites of multiple interacting and interconnected networks (i.e. a network of networks). The relationships between human agents, for example, may be seen as forming one kind of network (e.g. a social network), while the relationships between elements of the physical communication infrastructure may be seen as forming a different kind of network (i.e. a communication network). Importantly, the various types of network one sees in a coalition environment interact in highly complex ways, and we may expect the systemic cognitive capabilities of a coalition-based cognitive system to depend greatly on the adaptive alignment between the disparate networks. Inasmuch as systemic cognitive performances are influenced by the structure and dynamics of the various networks comprising a military coalition system, then such networks can, we suggest, be seen as candidate elements of the cognitive machinery for coalition-based distributed cognitive systems.

Given the interdependencies between (e.g.) information, communication and human networks in coalition environments, we expect to see an important role for studies that shed light on the adaptive configuration of such networks throughout the course of coalition deployments. One, relatively simple, example of this is the case where we seek to change the topology of physical communication networks in order to promote the appropriate exchange and transfer of information between spatially-distributed coalition elements. Another example is the case where we seek to modify the human social network in order to bring people with different (albeit related) bodies of knowledge and expertise together in order to solve some specific problem (see Huang, Contractor, & Yao, 2008). In all cases, what seems to be important is an ability to dynamically configure the structure and activity of multiple networks so as to best support the realization of organization-level or system-level goals.

The importance of dynamic configuration and functional coordination is recognized by those working in the area of distributed cognitive systems:

> *"In distributed cognition, one expects to find a system that can dynamically configure itself to bring subsystems into coordination to accomplish various functions." (Hollan et al., 2000; pg. 175)*

Interestingly, this is a view that is echoed by the cognitive scientist, Marvin Minsky, in his book *The Society of Mind* (Minsky, 1986). Minksy argues that the human mind can be seen as a large system of experts or agencies that are dynamically assembled together in various ways in order to accomplish specific cognitive tasks. Of course, one issue that is particularly difficult to resolve here concerns the mechanisms that support the dynamic configuration of networks, or network elements, in ways that best support the realization of cognitive goals. Ideally, network structures within military coalition environments should be capable of *automatic* modes of adaptation in order to ensure that the various elements of the coalition network are functionally aligned with respect to force-level objectives. What kind of system could support such automatic modes of adaptive configuration in coalition networks?

Perhaps one way of answering this question is to turn to biology and examine the kind of solutions that nature has derived for managing information flows in large-scale network systems. Van Essen, Anderson and Olshausen (1994), for example, posit the existence of 'control neurons' in the brain of mammalian nervous systems whose function is to regulate information flows between various neural processing resources. Such regulation constitutes, they argue, a key mechanism by which we are able to adaptively focus attention on specific subsets of environmental information and efficiently organize action output in the face of competing motor commands. Their analogy is with the division of labour in a large-scale commercial organization in which the key focus of organizational activity is not the actual generation of the final product *per se*, but rather the internal trafficking of information and materials. Perhaps, therefore, some of the insights gleaned from this, and other nature-inspired solutions to the problem of automatic network configuration, could be applied to case of network configuration in military coalition systems.

Other challenges posed by a consideration of coalition-level forms of cognitive extension arise from the cultural, linguistic and technological differences between coalition force elements. One challenge, for example, relates to the need to ensure a common (or shared) understanding of informational cues against a backdrop of community-specific interpretational biases and linguistic conventions (see Smart, Huynh et al., 2009). A lack of shared understanding may compromise the functional integration of coalition system components and thereby contribute to a breakdown in collective cognitive processing. Indeed, a sufficient level of shared understanding may be deemed to be one of the factors that determines whether the coalition formation can operate as an extended cognitive system. Similar threats to functional integration stem from problems associated with information exchange, trust and technological compatibility.

Relevant Defence-Related Research Programs

A number of defence-related research programs feature research that is relevant to cognitive extension in network-enabled contexts. Of course, military research has often been at the forefront of efforts to develop technologically-mediated forms of cognitive augmentation and enhancement. This is reflected in the efforts to develop sophisticated wearable computing devices as part of the Land Warrior and Future Force Warrior programs (see Ashok & Agrawal, 2003). More recently, one sees efforts to develop and exploit the kind of portable projection systems that present information directly to an individual's field of view. This work, being undertaken as part of the DARPA-funded ULTRA-Vis program[36], promises to deliver the kind of portable display systems that make network-accessible information better placed to influence ongoing sequences of thought and action.

Another strand of defence-related research concerns the attempt to augment human cognition by adaptively controlling the way in which information is presented to a user based on cognitive state information. The DARPA Augmented Cognition program, for example, seeks to monitor cognitive state using a variety of physiological sensors for the purposes of adapting display configurations and task commitments. Similar research goals are expressed by the recently proposed Cognition and Neuroergonomics CTA. This program seeks to:

> *"…enhance Soldier-system performance in complex operational settings by optimizing information transfer between the system and the Soldier, identifying mental processes*

[36] http://www.darpa.mil/ipto/Programs/uvis/uvis.asp

> *and individual differences that impact mission-relevant decision-making, and developing technologies for individualized analyses of neurally-based processing in operational environments."* (United States Army Research Laboratory, 2008)

Such efforts promise to improve our understanding of the opportunities for cognitive mergers that involve various forms of bio-artifactual coupling. By focusing on ways to optimize the presentation of task-relevant information and adapt systems to work in concert with biological functioning, these programs promise to yield valuable insights into the opportunities for cognitive extension at the level of individual human agents.

Research programs such as those mentioned above are most obviously concerned with the functioning of individual soldiers, and this makes them relevant to what we have referred to as human-centered forms of cognitive extension. There are, however, a number of research programs that seek to tackle issues related to cognitive performance in more 'distributed' contexts. One such program is the joint U.S./U.K. ITA program, which focuses on military coalition contexts and includes a specific research task devoted to issues of network-mediated cognitive extension. One strand of research in this task is to explore the relationship between network-level variables and collective cognitive performances. The studies reviewed in the section on 'Socially-Extended Cognition', specifically those of Mason et al (2005) and Lazer and Friedman (2007), attest to the importance of network-level considerations in understanding collective problem solving, but such studies possess a number of limitations that undermines their applicability to the military coalition environment. One of these limitations concerns the effect of different network structures (small-world, fully-connected, etc.) and patterns of information flow in different types of task context. The studies of both Mason et al (2005) and Lazer and Friedman (2007) focus on a particular type of task involving parallel search in a fixed solution space. But the military domain features many different types of task (e.g. hierarchical planning), and some of these tasks mandate that specific agents (or agent teams) are assigned to different aspects of a more general problem. What profile of network connectivity and information flow best supports collaborative problem-solving in these more systematically-structured task contexts? A second limitation of the studies of Mason et al (2005) and Lazer and Friedman (2007) concerns the use of fixed solution spaces. In this case, the assumption is that the correct solution is the same throughout the entire course of the experiment, but in real-world environments the most appropriate solution outcome may be linked to a dynamic and evolving situational context. As the situation changes, the nature of the correct solution may also change. In light of this, it is important to study the effects of variables like network topology in problem-solving domains that feature dynamic solution spaces. Thirdly, the studies of Mason et al (2005) and Lazer and Friedman (2007) employ static network structures whose topology does not change throughout the course of the problem-solving exercise. Such static networks are unlike those typically encountered in military coalition operations where there is an increasing move towards mobile ad hoc networks with dynamically changing topologies. In order to extend the results of Mason et al (2005) and Lazer and Friedman (2007) with respect to dynamic network environments, research in the ITA program is currently exploring the effect of dynamic changes in network topology on collective problem-solving performance in simulated agent communities (see Smart, Sycara, & Huynh, in press).

Another research program where we see a potential emphasis on coalition-level cognitive outcomes is the Network Science CTA. This CTA aims to undertake fundamental research in network science in order to understand the complex interplay between a variety of types of networks (e.g. social/cognitive, information and communication networks). Such research goals are clearly relevant to our understanding of how complex networks of informational, technological and social resources can be dynamically adapted in order to support cognitive processing at the collective, multi-agent, level.

Finally, the THINK ATO is a research program that aims to enhance warfighter cognitive performance in complex dynamic network environments. It was established as a collaborative venture between the U.S. Army Research Laboratory, the U.S. Army Research Institute, and the Communications and Electronics Research, Development and Engineering Center (CERDEC). It is geared towards making best use of available networked knowledge and information in order to support individual warfighter cognitive performance, as well as collaborative decision-making and distributed problem-solving. Clearly, such objectives coincide nicely with those of research into network-mediated forms of cognitive extension at both the individual and collective levels.

Conclusion

The traditional view in cognitive science is that cognition is inside the head of individual human agents. In contrast to this view, the notion of cognitive extension maintains that, at least in some situations, cognition is extended beyond the traditional biological borders of skin and skull. This latter view draws on an emerging wealth of empirical data concerning in way in which the facts of material embodiment and environmental embedding contribute to the emergence of cognitive processing routines that are distributed across the brain, body and world. In this chapter, we have suggested that notions of cognitive extension can be used to understand the transformative potential of a variety of network-enabled devices and network-accessible information resources on human cognitive processing. We have also proposed an extension to the original extended mind thesis, one that specifically caters for the potential role of network systems in extending the bounds of human cognition.

Our review of the literature relating to cognitive extension and the extended mind has highlighted a number of ways in which the information and network sciences are relevant to our understanding of extended cognitive systems. These include, but are not necessarily limited to, the following:

1. Extended cognitive systems consist of networks of information flow and influence between a variety of heterogeneous resources. Network science is well suited to assist us with the project of understanding how extended cognitive circuits operate with respect to the cognitive capabilities of the larger systemic organization.
2. Network-based approaches are merited in the specific case of network-extended minds because physical networks are at the heart of contemporary technology-mediated forms of cognitive extension. Our future attempts at engineering network-extended cognitive systems, or at least enabling them to emerge, will be dictated by our ability to develop and configure network technologies in ways that expand our human cognitive potential.
3. As a specialized branch of complexity science, network scientific approaches can help us understand the emergent capabilities of extended cognitive systems. Given the complex, nested and non-linear interactions between the components of an extended cognitive system, the capabilities of the larger system are not always guaranteed to be mere

augmentations or enhancements of some existing capability; they can sometimes be entirely new forms of cognitive capability and competence.

The application of network scientific approaches to both the analysis and engineering of extended cognitive systems is relevant to military coalition operations because such approaches help us to understand the factors that contribute to both the efficiency and quality of problem-solving processes in collaborative, network-enabled, distributed teams. By developing a better understanding of the cognitive impact of network systems on both individual and collective problem-solving, we are in a much better position to engage in interventions that enhance the cognitive power and potential of military coalition formations.

Of course, the possibility for network-mediated forms of cognitive extension is not something that is relevant just to military coalitions; the increasing ubiquity and pervasiveness of network systems motivates a more general interest in the effect of network technologies on our human cognitive potential. As Hollan et al (2000) comment:

> *"As we build richer, more all-encompassing, computational environments it becomes more important than ever to understand the ways human agents and their local environments are tightly coupled in the processing loops that result in intelligent action." (pg. 186)*

The advent of new computing technologies and network-enabled capabilities highlights a potential milestone in our human cognitive evolution. Just as the ability to use and exploit linguistic encodings marked a seachange in our individual and collective cognitive abilities, so the development of ubiquitous network systems, wearable computing devices and pervasive computing, presents us with unparalleled opportunities for cognitive extension at both the individual and collective levels. Ours, we suggest, are 'fishnet' minds, ones that are increasingly enmeshed in complex networks of technological, linguistic and social influence. As we learn to exploit those networks for our cognitive good or ill, so too we must cast our philosophical and scientific explanatory nets ever wider. In this way we may, at last, come to see the human mind for what it really is: not as some immaterial spirit stuff that emerges solely from the machinations of the human brain, but as a set of physical processes that occasionally escape their cranial confines and extend out into the world.

Additional Reading

Clark, A. (2001). Reasons, Robots and the Extended Mind. Mind & Language, 16(2), 121-145.

Clark, A. (2003). Natural-Born Cyborgs: Minds, Technologies and the Future of Human Intelligence. Oxford, UK: Oxford University Press.

Clark, A. (2007b). Re-Inventing Ourselves: The Plasticity of Embodiment, Sensing, and Mind. Journal of Medicine and Philosophy, 32(3), 263-282.

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

Clark, A. (in press-b). Memento's revenge: the extended mind, extended. In R. Menary (Ed.), The Extended Mind. Aldershot: Ashgate Publishers.

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

Goldstone, R. L., Roberts, M. E., & Gureckis, T. M. (2008). Emergent Processes in Group Behavior. Current Directions in Psychological Science, 17(1), 10-15.

Hutchins, E. (1995). Cognition in the Wild. Cambridge, Massachusetts: MIT Press.

Kirsh, D. (2009). Problem solving and situated cognition. In M. Ayded & P. Robbins (Eds.), Cambridge Handbook of Situated Cognition. Cambridge, UK: Cambridge University Press.

Noë, A. (2009). Out of Our Heads: Why You Are Not Your Brain, and Other Lessons from the Biology of Consciousness. New York, USA: Hill & Wang.

Pfeifer, R., & Bongard, J. (2007). How the Body Shapes the Way We Think: A New View of Intelligence. Cambridge, Massachusetts, USA: MIT Press.

Wilson, R. A., & Clark, A. (2009). Situated Cognition: Letting Nature Take its Course. In M. Ayded & P. Robbins (Eds.), Cambridge Handbook of Situated Cognition. Cambridge, UK: Cambridge University Press.

Acknowledgements

Research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the author(s) and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

References

Adams, F., & Aizawa, K. (2001). The bounds of cognition. *Philosophical Psychology, 14*(1), 43-64.

Adams, F., & Aizawa, K. (2008). *The Bounds of Cognition*. Oxford, UK: Blackwell.

Adams, F., & Aizawa, K. (2009). Why the mind is still in the head. In M. Ayded & P. Robbins (Eds.), *Cambridge Handbook of Situated Cognition*. Cambridge, UK: Cambridge University Press.

Adams, F., & Aizawa, K. (in press). Defending the bounds of cognition. In R. Menary (Ed.), *The Extended Mind*. Aldershot, UK: Ashgate Publishers.

Ajith, A., Crina, G., & Vitorino, R. (Eds.). (2006). *Stigmergic Optimization*. Berlin, Germany: Springer-Verlag.

Ashok, R., & Agrawal, D. (2003). Next-generation wearable networks. *Computer, 36*(11), 31-39.

Bahrami, A., Yuan, J., Smart, P. R., & Shadbolt, N. R. (2007). *Context-Aware Information Retrieval for Enhanced Situation Awareness.* Paper presented at the Military Communications Conference (MILCOM), Orlando, Florida, USA.

Barabasi, A.-L. (2002). *Linked: The New Science of Networks*. Cambridge, Massachusetts, USA: Perseus Publishing.

Berners-Lee, T., Hall, W., Hendler, J. A., O'Hara, K., Shadbolt, N., & Weitzner, D. J. (2006). A framework for web science. *Foundations and Trends in Web Science, 1*(1), 1-130.

Berners-Lee, T., Hendler, J., & Lassila, O. (2001). The Semantic Web. *Scientific American, 284*(4), 34-43.

Bird, J., & Layzell, P. (2002). *The evolved radio and its implications for modelling the evolution of novel sensors.* Paper presented at the Congress on Evolutionary Computation (CEC'02), Washington D.C., USA.

Bonabeau, E. (1999). Editor's introduction: stigmergy. *Artificial Life, 5*(2), 95-96.

Bonabeau, E., Dorigo, M., & Theraulaz, G. (1999). *Swarm Intelligence: From Natural to Artificial Systems*. New York, USA: Oxford University Press.

Braddon-Mitchell, D., & Jackson, F. (2007). *Philosophy of Mind and Cognition: An Introduction* (2nd ed.). Oxford, UK: Blackwell Publishing.

Brooks, R. A. (1991). Intelligence without representation. *Artificial Intelligence, 47*, 139-160.

Buchanan, M. (2002). *Nexus: Small Worlds and the Groundbreaking Science of Networks*. New York, USA: W. W. Norton & Company.

Burge, T. (1979). Individualism and the mental. In P. French, T. Uehling & H. Wettstein (Eds.), *Midwest Studies in Philosophy* (Vol. 4, pp. 73-121). Minnesota, Minneapolis, USA: University of Minnesota Press.

Butler, K. (1998). *Internal Affairs: Making Room for Psychosemantic Internalism*. Dordrecht, The Netherlands: Kluwer Academic Publishers.

C. elegans Sequencing Consortium. (1998). Genome sequence of the nematode C. elegans: a platform for investigating biology. *Science, 282*(5396), 2012-2018.

Camazine, S., Deneubourg, J.-L., Franks, N. R., Sneyd, J., Theraulaz, G., & Bonabeau, E. (2001). *Self-Organization in Biological Systems*. Princeton, New Jersey, USA: Princeton University Press.

Carr, N. (2008). Is Google Making Us Stupid? *The Atlantic, 302*(1), 56-63.

Cherniack, M., Franklin, M. J., & Zdonik, S. (2001). Expressing user profiles for data recharging. *IEEE Personal Communications: Special Issue on Pervasive Computing, 8*(4), 32-38.

Clark, A. (1997). *Being There: Putting Brain, Body and World Together Again*. Cambridge, Massachusetts: MIT Press.

Clark, A. (1999). An embodied cognitive science. *Trends in Cognitive Science, 3*(9), 345-351.

Clark, A. (2001). Reasons, Robots and the Extended Mind. *Mind & Language, 16*(2), 121-145.

Clark, A. (2003). *Natural-Born Cyborgs: Minds, Technologies and the Future of Human Intelligence*. Oxford, UK: Oxford University Press.

Clark, A. (2005). Intrinsic content, active memory and the extended mind. *Analysis, 65*(285), 1-11.

Clark, A. (2007a). Curing Cognitive Hiccups: A Defense of The Extended Mind. *Journal of Philosophy, 104*(4), 163-192.

Clark, A. (2007b). Re-Inventing Ourselves: The Plasticity of Embodiment, Sensing, and Mind. *Journal of Medicine and Philosophy, 32*(3), 263-282.

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

Clark, A. (in press-a). Coupling, Constitution and the Cognitive Kind: A Reply to Adam and Aizawa. In R. Menary (Ed.), *The Extended Mind*. Aldershot, UK: Ashgate Publishers.

Clark, A. (in press-b). Memento's revenge: the extended mind, extended. In R. Menary (Ed.), *The Extended Mind*. Aldershot: Ashgate Publishers.

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

Claverie, J.-M. (2001). Gene number: What if there are only 30,000 human genes? *Science, 291*(5507), 1255-1257.

Dawkins, R. (1982). *The Extended Phenotype: The Long Reach of the Gene*. Oxford, UK: Oxford University Press.

Dennett, D. C. (1996). *Kinds of Minds: Toward an Understanding of Consciousness*. New York, USA: Basic Books.

DeVaul, R., Pentland, A., & Corey, V. (2005). *The memory glasses: subliminal vs. overt memory support with imperfect information.* Paper presented at the 7th IEEE International Symposium on Wearable Computers, White Plains, New York, USA.

Diehl, M., & Stroebe, W. (1987). Productivity loss in brainstorming groups: toward the solution of a riddle. *Journal of Personality and Social Psychology, 53*, 497-509.

Fodor, J. (1980). Methodological solipsism considered as a research strategy in cognitive psychology. *Behavioral and Brain Sciences, 3*(1), 63–73.

Foelix, R. F. (1996). *Biology of Spiders*. New York, USA: Oxford University Press.

Gibson, J. J. (1966). *The Senses Considered as Perceptual Systems*. New York, USA: Houghton-Mifflin.

Goldstone, R. L., Roberts, M. E., & Gureckis, T. M. (2008). Emergent Processes in Group Behavior. *Current Directions in Psychological Science, 17*(1), 10-15.

Gray, W. D., & Fu, W. T. (2004). Soft constraints in interactive behavior: The case of ignoring perfect knowledge in-the-world for imperfect knowledge in-the-head. *Cognitive Science, 28*(3), 359-382.

Greenfield, S. (2003). *Tomorrow's People: How 21st Century Technology is Changing the Way We Think and Feel*. London, England: Penguin Books Ltd.

Harnad, S. (1999). The future of scholarly skywriting. In A. Scammell (Ed.), *I in the Sky: Visions of the Information Future*. London, UK: Aslib.

Harnad, S., & Dror, I. E. (2006). Distributed cognition: Cognizing, autonomy and the Turing Test. *Pragmatics & Cognition, 14*(2), 209-213.

Haugeland, J. (1998). Mind embodied and embedded. In J. Haugeland (Ed.), *Having thought: Essays in the Metaphysics of Mind*. Cambridge, Massachusetts, USA: Harvard University Press.

Hinsz, V., Tindale, R., & Vollrath, D. (1997). The emerging conceptualization of groups as information processors. *Psychological Bulletin, 121*, 43-64.

Hollan, J., Hutchins, E., & Kirsh, D. (2000). Distributed cognition: toward a new foundation for human-computer interaction research. *ACM Transactions on Computer-Human Interaction (TOCHI), 7*(2), 174-196.

Huang, Y., Contractor, N., & Yao, Y. (2008). *CI-KNOW: recommendation based on social networks.* Paper presented at the 9th Annual International Digital Government Research Conference, Montreal, Canada.

Hurley, S. (1998). *Consciousness in Action*. Massachusetts, USA: Harvard University Press.

Hutchins, E. (1995a). *Cognition in the Wild*. Cambridge, Massachusetts: MIT Press.

Hutchins, E. (1995b). How a cockpit remembers its speeds. *Cognitive Science, 19*(3), 265-288.

International Human Genome Sequencing Consortium. (2004). Finishing the euchromatic sequence of the human genome. *Nature, 431*(7011), 931-945.

Janis, I. L. (1982). *Victims of Groupthink* (2nd ed.). Boston, USA: Houghton Mifflin.

Jorgensen, C., & Binsted, K. (2005). *Web Browser Control Using EMG Based Sub Vocal Speech Recognition.* Paper presented at the 38th Annual Hawaii International Conference on System Sciences (HICSS'05), Big Island, Hawaii.

Kauffman, S. (1995). *At Home in the Universe: The Search for Laws of Self-Organization and Complexity*. London, England: Penguin Books Ltd.

Kirsh, D. (1995). The intelligent use of space. *Artificial Intelligence, 73*(1-2), 31-68.

Kirsh, D. (1996). Adapting the environment instead of oneself. *Adaptive Behavior, 4*(3/4), 415-452.

Kirsh, D. (2006). Distributed cognition: A methodological note. *Pragmatics & Cognition, 14*(2), 249-262.

Kirsh, D. (2009). Problem solving and situated cognition. In M. Ayded & P. Robbins (Eds.), *Cambridge Handbook of Situated Cognition*. Cambridge, UK: Cambridge University Press.

Kirsh, D., & Maglio, P. (1994). On distinguishing epistemic from pragmatic action. *Cognitive Science, 18*, 513-549.

Krink, T., & Vollrath, F. (1997). Analysing Spider Web-building Behaviour with Rule-based Simulations and Genetic Algorithms. *Journal of Theoretical Biology, 185*(3), 321-331.

Krink, T., & Vollrath, F. (1998). Emergent properties in the behaviour of a virtual spider robot. *Proceedings of the Royal Society of London B: Biological Sciences, 265*(1410), 2051-2055.

Krink, T., & Vollrath, F. (1999). Virtual Spiders Guide Robotic Control Design. *IEEE Intelligent Systems, 14*(5), 77-84.

Lazer, D., & Friedman, A. (2007). The Network Structure of Exploration and Exploitation. *Administrative Science Quarterly, 52*(4), 667-694.

Loftus, E. (1997). Creating false memories. *Scientific American, 277*(3), 70-75.

Lopez, V., Pasin, M., & Motta, E. (2005). *AquaLog: An Ontology-portable Question Answering System for the Semantic Web.* Paper presented at the 2nd European Semantic Web Conference (ESWC 2005), Heraklion, Greece.

Maglio, P., Matlock, T., Raphaely, D., Chernicky, B., & Kirsh, D. (1999). *Interactive skill in Scrabble.* Paper presented at the 21st Annual Conference of the Cognitive Science Society, Vancouver, British Columbia, Canada.

Mason, S., Bashashati, A., Fatourechi, M., Navarro, K., & Birch, G. (2007). A comprehensive survey of brain interface technology designs. *Annals of Biomedical Engineering, 35*(2), 137-169.

Mason, W. A., Jones, A., & Goldstone, R. L. (2005). *Propagation of innovations in networked groups.* Paper presented at the 27th Annual Conference of the Cognitive Science Society, Stresa, Italy.

Minsky, M. (1986). *The Society of Mind*. New York, USA: Simon & Schuster, Inc.

Mott, D., & Hendler, J. (2007). *Progress on the Collaborative Planning Model.* Paper presented at the 1st Annual Conference of the International Technology Alliance (ACITA), Maryland, USA.

Myin, E., & O'Regan, J. K. (2009). Situated perception and sensation in vision and other modalities: a sensorimotor approach. In M. Ayded & P. Robbins (Eds.), *Cambridge Handbook of Situated Cognition*. Cambridge, UK: Cambridge University Press.

Newell, A. (1980). Physical symbol systems. *Cognitive Science, 4*, 135-183.

Newell, A., & Simon, H. (1976). Computer science as empirical enquiry: Symbols and search. *Communications of the Association of Computing Machinery, 19*, 113-126.

Nicolelis, M. (2001). Actions from thoughts. *Nature, 409*(6818), 403-407.

Noë, A. (2004). *Action In Perception*. Massachusetts, USA: MIT Press.

Noë, A. (2009). *Out of Our Heads: Why You Are Not Your Brain, and Other Lessons from the Biology of Consciousness*. New York, USA: Hill & Wang.

Nofi, A. A. (2000). *Defining and Measuring Shared Situational Awareness* (No. CRM D0002895.A1/Final). Alexandria, Virginia, USA: Center for Naval Analyses.

Norman, D. A. (1993). *Things That Makes Us Smart*. Reading, Massachusetts, USA: Addison-Wesley.

Norman, D. A. (1998). *The Invisible Computer*. Cambridge, Massachusetts, USA: MIT Press.

O'Regan, J. K., & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences, 24*(5), 939-973.

Paul, C. (2004). *Morphology and Computation.* Paper presented at the From Animals to Animats 8: Proceedings of the Eighth International Conference on Simulation of Adaptive Behavior, Los Angeles, California, USA

Paul, C. (2006). Morphological Computation: A Basis for the Analysis of Morphology and Control Requirements. *Robotics and Autonomous Systems, 54*, 619-630.

Pfeifer, R., & Bongard, J. (2007). *How the Body Shapes the Way We Think: A New View of Intelligence*. Cambridge, Massachusetts, USA: MIT Press.

Pfurtscheller, G., Scherer, R., & Neuper, C. (2007). EEG-based brain-computer interface. In R. Parasuraman & M. Rizzo (Eds.), *Neuroergonomics: The Brain at Work*. New York, USA: Oxford University Press.

Poltrock, S., Handel, M., Bowyer, H., & Waggett, P. (2008). *A dynamic model of mission context.* Paper presented at the 2nd Annual Conference of the International Technology Alliance (ACITA'08), London, UK.

Putnam, H. (1975). The meaning of "meaning". In K. Gunderson (Ed.), *Language, Mind and Knowledge* (pp. 131-193). Minneapolis, Minnesota, USA: University of Minnesota Press.

Pylyshyn, Z. (1984). *Computation and Cognition*. Cambridge, Massachusetts, USA: MIT Press.

Rhodes, B., & Maes, P. (2000). Just-in-time information retrieval agents. *IBM Systems Journal, 39*(3-4), 685-704.

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

Rowlands, M. (2006). *Body Language: Representation in Action*. Cambridge, Massachusetts, USA: MIT Press.

schraefel, m. c., Smith, D. A., Owens, A., Russell, A., Harris, C., & Wilson, M. L. (2005). *The evolving mSpace platform: leveraging the Semantic Web on the trail of the memex.* Paper presented at the Hypertext 2005 Conference, Salzburg, Austria.

Shadbolt, N., Hall, W., & Berners-Lee, T. (2006). The semantic web revisited. *IEEE Intelligent Systems, 21*(3), 96-101.

Smart, P. R., Engelbrecht, P., Braines, D., Strub, M., & Hendler, J. (2009). *Cognitive Extension and the Web.* Paper presented at the Web Science Conference: Society On-Line, Athens, Greece.

Smart, P. R., Engelbrecht, P. C., Braines, D., Hendler, J. A., & Shadbolt, N. R. (2008). *The Extended Mind and Network-Enabled Cognition* (No. ITA/P12/NEC). Southampton, UK: School of Electronics and Computer Science, University of Southampton.

Smart, P. R., Huynh, T. D., Mott, D., Sycara, K., Braines, D., Strub, M., et al. (2009). *Towards an Understanding of Shared Understanding in Military Coalition Contexts.* Paper presented at the 3rd Annual Conference of the International Technology Alliance (ACITA'09), Maryland, USA.

Smart, P. R., O'Hara, K., Engelbrecht, P., Giammanco, C., & Braines, D. (in press). Memory, Networks and the Extended Mind: Individual Memory and Technology-Mediated Memory Enhancement. In T. D. Huynh, P. R. Smart, D. Braines & K. Sycara (Eds.), *Network-Enabled Cognition: The Contribution of Social and Technological Networks to Human Cognition*. Raleigh, North Carolina, USA: Lulu Press.

Smart, P. R., Russell, A., Liang, S. F., Shadbolt, N. R., Booth, C., Briscombe, N., et al. (2009). *Using Semantic Technologies to Improve Information Exploitation in Military and Civilian Application Contexts.* Paper presented at the Knowledge Systems for Coalition Operations Conference (KSCO'09), Southampton, UK.

Smart, P. R., Sycara, K., & Huynh, T. D. (in press). Exploring the Dynamics of Belief Propagation and Collective Problem Solving in Simulated Agent Networks. In T. D. Huynh, P. R. Smart, D. Braines & K. Sycara (Eds.), *Network-Enabled Cognition: The Contribution of Social and Technological Networks to Human Cognition*. Raleigh, North Carolina, USA: Lulu Press.

Sporns, O. (2002). Network analysis, complexity, and brain function. *Complexity, 8*(1), 56-60.

Sporns, O., Chialvo, D., Kaiser, M., & Hilgetag, C. (2004). Organization, development and function of complex brain networks. *Trends In Cognitive Sciences, 8*(9), 418-425.

Stasser, G., & Titus, W. (1985). Pooling of unshared information in group decision making: biased information sampling during group discussion. *Journal of Personality and Social Psychology, 48*, 1467-1478.

Sterelny, K. (2004). Externalism, Epistemic Artefacts and the Extended Mind. In R. Schantz (Ed.), *The Externalist Challenge: New Studies on Cognition and Intentionality* (pp. 239-254). Berlin, Germany: de Gruyter.

Stix, G. (2008). Jacking into the brain - is the brain the ultimate computer interface? *Scientific American, 299*(5), 56-61.

Tablan, V., Damljanovic, D., & Bontcheva, K. (2008). *A Natural Language Query Interface to Structured Information.* Paper presented at the 5th European Semantic Web Conference (ESWC'08), Tenerife, Spain.

Tarsitano, M., & Jackson, R. (1997). Araneophagic jumping spiders discriminate between detour routes that do and do not lead to prey. *Animal Behaviour, 53*(2), 257-266.

Thompson, A., Harvey, I., & Husbands, P. (1996). Unconstrained Evolution and Hard Consequences. In E. Sanchez & M. Tomassini (Eds.), *Toward Evolvable Hardware: The Evolutionary Engineering Approach* (pp. 136-165). Berlin, Germany: Springer-Verlag.

Tollefsen, D. P. (2006). From extended mind to collective mind. *Cognitive Systems Research, 7*(2-3), 140-150.

Tribble, E. (2005). Distributing Cognition in the Globe. *Shakespeare Quarterly, 56*(2), 135-155.

United States Army Research Laboratory. (2008). *Cognition and Neuroergonomics Collaborative Technology Alliance: Draft Program Announcement* (No. W911NF-08-R-0014). USA: United States Army Research Laboratory.

Van Essen, D., Anderson, C., & Olshausen, B. (1994). Dynamic routing strategies in sensory, motor, and cognitive processing. In C. Koch & J. L. Davis (Eds.), *Large-Scale Neuronal Theories of the Brain* (pp. 271–299). Massachusetts, USA: MIT Press.

Watts, D. J. (2003). *Six Degrees: The Science of a Connected Age*. London, UK: William Heinemann.

Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of "small word" networks. *Nature, 393*, 440-442.

Wheeler, M. (2005). *Reconstructing the Cognitive World*. Cambridge, Massachusetts, USA: MIT Press.

Wheeler, M., & Clark, A. (1999). Genic representation: reconciling content and causal complexity. *British Journal For The Philosophy Of Science, 50*(1), 103-135.

Wilcox, S., & Jackson, R. (2002). Jumping Spider Tricksters: Deceit, Predation, and Cognition. In M. Bekoff, C. Allen & G. M. Burghardt (Eds.), *The Cognitive Animal*. Massachusetts, USA: MIT Press.

Wilson, R. A. (1994). Wide computationalism. *Mind, 103*(411), 351-372.

Wilson, R. A. (2000). The mind beyond itself. In D. Sperber (Ed.), *Misrepresentations: A Multidisciplinary Perspective* (pp. 31-52). New York, USA: Oxford University Press.

Wilson, R. A. (2004). *Boundaries of the Mind: The Individual in the Fragile Sciences: Cognition*. New York, USA: Cambridge University Press.

Wilson, R. A., & Clark, A. (2009). Situated Cognition: Letting Nature Take its Course. In M. Ayded & P. Robbins (Eds.), *Cambridge Handbook of Situated Cognition*. Cambridge, UK: Cambridge University Press.

Witt, P. N., Reed, C. F., & Peakall, D. B. (1968). *A Spider's Web: Problems in Regulatory Biology*. Heidelberg, Germany: Springer.