

A Shifting Boundary: the dynamics of internal cognition and the web as external representation

Alan Dix^{1,2}

¹Talis
43 Temple Row
Birmingham, B2 5LS, UK

²InfoLab21
Lancaster University
Lancaster, LA1 4AW, UK

alan@hcibook.com

<http://www.hcibook.com/alan/papers/websci2011-int-ext-cog/>

ABSTRACT

Various disciplines have begun to emphasise the role of the external environment in human cognition. While initially focused on physical interactions, these theories are also of clear relevance to the web, especially in its role as external representation of human knowledge, potentially obviating the need to remember basic facts. Internal representation and cognition remain important, but change in the face of a pervasive digital environment. This paper explores the evolving dynamic between internal and external cognition, in particular the shift from knowledge to meta-knowledge and the way this impacts learning, society and, perhaps, the very nature of our own minds.

Categories and Subject Descriptors

H.5.4 [Information Systems]: Hypertext/Hypermedia – *user issues*; I.2.0 [Artificial Intelligence]: General – *philosophical foundations*; K.4.0 [Computers And Society]: General

General Terms

Human Factors.

Keywords

embodiment, external representation, distributed cognition

1. INTRODUCTION

Over recent years various disciplines have, in different ways, begun to emphasise the role of the external environment in human cognition. The roots of this go back many years: in philosophy [25, 36], in perceptual psychology [19], and in human-computer interaction [43, 28]. While the earliest analysis focused on physical artefacts and interactions, this strand of thinking has been applied extensively to digital interactions and notably the web itself. In particular, Halpin, Clark and Wheeler [23] applied the latest critical thinking from the philosophical side of this area, also known as '4E (embodied, embedded, enactive, extended) cognition', to the web as a form of massive shared external representation.

There may be debate about the nature, accuracy and importance of

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WebSci '11, June 14-17, 2011, Koblenz, Germany.

Copyright 2011 ACM.

internal representations, and indeed the strawman of a veridical internal facsimile is undoubtedly false. However, the existence of some form of internal representation or state is both common sense and an observable phenomenon of agent interactions with the environment beyond trivial stimulus/response.

The crucial theoretical and practical question is how the dynamics of internal and external representation work out on the web: the extent to which it is similar and/or different from previous information technology and how this may transform our very nature as people.

The next section gives a brief introduction to strands of embodiment or enactment thinking from multiple disciplines, especially where these have influence on thinking in human-computer interaction. Those already familiar with this area can skip to section 3, which discusses existing work looking at embodiment and external representation applied to the web. Both enactivist and cognitivist thinking draw heavily on assumptions of optimality assumptions, which are not always borne out in practice. Section 4 examines this 'optimal fallacy' in relation to physical and web embodiment.

Section 5, is the heart of the paper, and explores the kinds of internal representations we can have *of* the web and how these relate to the external representations *in* the web. Most critically, we will consider the growing importance of meta-representations for our understanding of the web and 'post-web cognition' [11]. For example, we see a change from knowledge about things to knowledge about how to find out things, at both a concrete level, "I go to this web page to find X", or a more generic level of knowing, such as the kind of terms to use in search. However, we will also consider the educational impact of this changing internal/external balance: learning often depends on processing, not too readily available material, and reflection is crucial to deeper learning.

The paper ends by briefly exploring some of the broader implications this may have for children growing up in a web connected world, and for our role in society as a whole.

2. A BRIEF HISTORY OF EMBODIMENT

As noted in the introduction, the roots of embodiment go back many years in a number of interlinked disciplines. To some extent the earliest works in each of these areas are reactions to earlier more dualistic positions, notably Descartes' and Kant's focus on internal reasoning as the basis of knowledge, a sort of inside to outside view of the world, that gives primacy to abstract logical argument. As well as being core to science (and academia in general), this found its way into early cognitive science and artificial intelligence, which tried to explain or emulate human action in terms of a pipeline model: sensing – perception –

representation – decision – action. This cognitive model effectively prizes knowledge in the head: perception is the way in which we construct it, and action is what we do with it.

2.1 Philosophy

Most philosophical discussions refer to Heidegger's "*Being and Time*" [25] or Merleau-Ponty's "*Phenomenology of Perception*" [36] and the nature of embodied and enacted being. Central to Heidegger's position is that we cannot view ourselves as in some way sitting back and thinking about the world, but are intimately and continuously part of the world. Whereas earlier (and more recent) conceptions of the mind were focused on the intentional aspects of knowledge (the fact that it is about something in the world) and explicit reasoning, Heidegger focuses on artful interaction and implicit action (the fact that we are someone in the world).

A key concept in Heidegger, that has influenced thought and practice in human–computer interaction, is that of *thrownness* and *breakdown*. Instead of viewing ourselves as mini-logicians, constantly considering each action and making a reasoned decision, Heidegger points out that the normality of human action is (consciously) unconsidered. A key example is the hammer in the hands of a skilful artisan, *ready-to-hand*, not noticed in itself as the focus is on nail and wood. This *thrownness*, when the tool almost becomes an extension of the body, is contrasted with *breakdown*, when the hammer head becomes loose or for some other reason the focus shifts to the tool.

Nowadays the idea that many of our actions are largely unconscious is uncontentious. However, those who have not studied psychology are still often surprised by just how much is going on without conscious consideration. Indeed, even when we think of our own conscious train of thought, what brings something into conscious thought is not itself consciously considered (see [16] for a discussion of this).

Modern philosophers in this area range from the more pragmatic such as Clark [8] to the more mystical such as Varela [44]. Some writings address issues that have cognitive implications and of particular importance is the issue of representation. Traditional views gave central place to internal representation of the world in the head. In contrast, more enactivist views argue that *parsimony* requires that we leave things in the world to represent themselves, and do not construct unnecessary internal representations.

"In general evolved creatures will neither store nor process information in costly ways when they can use the structure of the environment and their operations on it as a convenient stand-in for the information-processing operations concerned." ([7] as quoted in [8])

Clark calls this the "*007 principle*" what we *need to know* in order to act. We will revisit this notion of parsimony later.

To some extent these philosophers of embodiment or enactivism are (or see themselves as) controversial within the discipline. This is partly a matter of vocabulary and definitions. Few doubt the idea that interaction with the world is a crucial part of our being, for example, the person working out a sum with paper and pencil. However, they find more problematic the idea of 'mind' being taken to include not just the grey matter in our heads, but also the paper and pencil in our hands.

2.2 Psychology

In perceptual psychology, Gibson's "*The Ecological Approach to Visual Perception Psychology*" [19], has been influential in establishing a view that regards perception and action not as mere input and output, but intimately connected with one another, and furthermore seeing perception as optimised for this role rather than creating veridical images in our minds. A key term is '*affordance*', the potential for action of an object for a particular actor: a small rock may *afford* picking up or throwing, a larger rock may *afford* sitting on. The traditional (at the time) view of perception was that we see something, work out what we can do with it, and then do it. However, Gibson argued that our vision and other perceptual systems, are fine-tuned for action, and so the affordances of objects are immediately apparent, and directly give rise to action.

The idea of affordance has been influential in HCI, initially through the work of Norman [38] and Gaver [18]. Norman's use of the term led to some confusion (later clarified) being focused on the perceptually apparent potential of an object, such as an onscreen button, which looks as though it protrudes and so '*affords*' pushing. This confusion was incipient in Gibson's own work as his analysis did not distinguish natural objects and environments from the artificial. The immediacy of perception makes sense for natural objects for which our bodies have been fine tuned over millennia, but not for the artificial and especially virtual, except insofar as good design fits the perceptual form of the object (perceived affordances) to its actual potential for action (real affordance). These two concepts, real and perceived affordance, are, by Gibson's arguments, indistinguishable in the natural world, but, as became obvious later, only arise as a consequence of design otherwise [24].

Another critical, but less popularised, part of Gibson's work was that not only is perception part of action, but in a sense action is part of perception. Early computational image processing worried about problem cases. These also turn out *when static* to be problematic for people as evident in 'puzzle' photographs of familiar objects from strange angles. But in real life we simply move our head or move the object that is causing concern; if we are not sure what is round the corner, we look. These are *epistemic actions*, actions intended to render knowledge. Our perceptions are not only fine tuned to deliver action, but also '*expect*' epistemic action. We do not have eyes on the back of our heads because we can turn our neck, and furthermore our peripheral vision is sensitive to movement in order to evoke an instinctive shift of the eyes and head bringing any potentially interesting (or dangerous) object into sharper foveal vision.

While early cognitive science focused on internal representations of the world, in the late 1980's, in reaction to this 'cognitivist' view the concept of '*distributed cognition*' [26, 28], was developed, informed by social anthropology of Peloponnesian navigators and the bridges of large modern military ships. This led to a questioning of cognition and problem solving being things that happened purely in the head. Problems got solved, ships found their way to their destination, but there was often no single person encompassing this within their head, instead the artful connection between people and things (instruments, maps, the wind on sails, the movement of the boat itself), and people and people led to the accomplishment. Seen as a cognitive act, navigation was distributed both environmentally and socially.

Again, the terminology of 'cognition' being distributed is more divisive than the concept. Others took the same extant human behaviours, but saw them in a more traditional cognitive

framework of *'display-based cognition'* that focused on the interactions between sensing and action, but where the internal cognitive models were much as in earlier work [33].

2.3 Human–Computer Interaction

Some of the earliest and most influential models adopted in human–computer interaction were of the Cartesian/cognitivist ilk, notably Card, Moran and Newell's Model Human Processor and GOMS model [5]. The latter focused on the way high-level goals were translated into actions, but in ways that, in most examples, ignored any perception or orientation to the environment. In the view of its later critics, this gave rise to a 'blind' model of human activity driven by pre-conceived plans.

However, almost as early, many of the previously mentioned concepts (breakdown, affordance, distributed cognition) had found their way into mainstream HCI. Of particular note was Suchman's [43] *"Plans and Situated Action"*, which used Garfinkel's ethnomethodology [17], itself influenced by Heidegger's phenomenology. Suchman was looking at engineers repairing photocopying machines (she worked at Xerox PARC!) and found that, rather than following a pre-conceived plan, in contrast their actions were *situated*, derived from the particular circumstances they found. They did not think "here is a check list of things to look for" or "I will follow my diagnostic test schedule", but instead "what's wrong" or "what can I see that is unusual and might be a problem".

Pragmatically and applying layperson's common sense, these views can be seen as alternative and useful perspectives on the same issues. However, the divergent underlying philosophical positions often lead to more radical positioning with most laboratory experiments following a more reductionist and traditional approach and most field studies and ethnographic work adopting a more holistic and situated one.

3. EMBODIMENT AND THE WEB

To some extent embodiment is most clearly seen in relation to physical artefacts and interactions. However, this strand of thinking has been applied extensively to digital interactions and notably the web itself. As already mentioned, breakdown, affordance, distributed cognition and situated action are all common concepts in human–computer interaction applied to the design of purely virtual artefacts such as a GUI interfaces as well as hybrid digital–physical artefacts such as the mobile phone.

Some of the 4E terms (enactive, extended) are easier to see in terms of digital interactions than others (embodied, embedded), however, the basic idea of being creatures of action applies equally to the digital world as to the physical. In WebSci'10, Halpin, et al. [23] explored the latest critical thinking in this area, focusing on three areas: the web as representation, search as enactive and collective intelligence. They note that while ideas of cognitive extension are sometimes seen as controversial, they are not problematic for those involved in the web where concepts such as collective intelligence are already, at a purely practical level, challenge traditional concepts of intelligence and even agency.

In their *"A Manifesto for Web Science?"*, Halford, et al. [22] consider positions from the social sciences that would not normally be seen in the same philosophical vein, and yet focus on the co-constitution of human and non-human actors within the network of the web where *"no entity has existence independent of its relations with other entities"* – surely echoing some of that same rich embedding of human activity within the web.

In distributed cognition the idea of *'offloading'* is critical, the way we use external resources as both a memory and computational aid rather than holding the information, or performing the processing purely in our heads. Sometimes the means of offloading may be additional, such as the paper used to sketch and plan a kitchen. However, sometimes the offloading is intrinsic to the problem being addressed. For example, we do not look at the pieces of a jigsaw puzzle and then attempt to 'solve' it in our heads, but instead manipulate the real pieces on a real surface; similarly only truly exceptional chess players can play blind chess without seeing the board in front of them (although, contrarily, it is against to experiment on the board when planning a move).

For those using digital technology this is hardly new; from phone numbers on your mobile, to the calculator to add up the shopping bill. On the web, we do not need to remember the opening times of the local library, the height of Everest or the birthday of a friend, we simply turn to Google, Wikipedia, or Facebook.

However, external representation in the web is not unproblematic. Search was one of the areas proposed by Halpin et al. as representing external cognition [23], but Carr and Harnard [6] found that Google search functions very differently from human recall. Specifically, they looked at using combination terms closely associated with a topic in order to find Wikipedia pages, but found that typically only the central term was successful. As in other areas, machine intelligence is not so much artificial human intelligence, but alien intelligence. So, shifting from remembering to using search is not simply one of offloading the same methods one would use in the head; either our means of using recall change or the results of our (extended) thinking are different – and probably a combination of the two.

This is not actually so radical. Using a human personal assistant to manage meetings and contacts is not the same as doing it oneself. Socially distributed cognition is inevitably about confronting the alien minds of others and yet, after getting used to one another, we are often able to coordinate and operate effortlessly. Similarly, it can be shocking to see the length of our web history on a day when we were not really aware of extensive web use; web search and browsing are so effortless that we experience Heidegger's thrownness, unaware of the skilled interactions operating moment to moment.

4. INTERNAL REPRESENTATION

4.1 Philosophy vs. common sense

As noted, the various forms of enacted or extended cognition emerged as a reaction to cognitivist or Cartesian views where internal representation and the inner life were seen as primary and the external world as secondary. As part of this reaction some in the area effectively deny any form of internal non-enacted representation. Halpin et al. [23] demure from this issue, instead stating, *"The debate over the existence of internal representations is an empirical debate best left to the neuro-scientific work."* To an outsider, to even question the existence of internal representations, is strange, all we have to do is close our eyes and picture our bedroom, or recall the last ten minutes.

To some extent this apparent clash between philosophy and common sense is again a matter of vocabulary, whereas those in modern computer science are used to ideas of diffuse/distributed 'representations' as found in neural networks, the term in philosophical discourse has come to mean something more akin to an explicit veridical internal facsimile, an exact mental 'scale' model of the world in the head. In Dreyfus' influential commentary [14] on Heidegger's *Being and Time*, he writes of "a

background of everyday practices into which we are socialised but that we do not represent in our minds", referring to the tacit, subconscious and unreflective ways in which we operate routinely, whether walking or making a pot of tea. However, in this case it is obvious, not just from common sense, but also empirically, that there is some level of co-entropy between the contents of our heads and the world. This does not require brain dissection, simply the observation that behaviour is better fitted to a given situation than can be explained by current senses.

What this does emphasise is that, whether we call it 'representation' or invent some other terms such as 'correlated world state', we do have to remember that:

- the large part is not available to conscious reflection – explicit knowledge is the top of an iceberg compared to implicit/tacit
- even where we *can* introspect, most of the time we *do not*; it is available but not salient
- what 'representations' we do have may not be 'accurate' (e.g., much of memory is confabulation, filling in gaps, a real problem in court witnesses)
- they are unlikely to be in a simple isomorphic mapping to the world (with the possible exception of the visual cortex)

4.2 Parsimony cuts two ways

In section 2 we looked at Clark's "007 principle", which suggested that we do not store or process things in our heads that would be better stored or processed in the environment. For example, our short term memory is very limited (around seven 'chunks' of information [37]), so it makes sense to write numbers down on paper during an even moderately complicated sum (e.g. 5139 x 2657), rather than attempt to hold them in your head.

However, parsimony cuts both ways; even moving your eyes or head takes time let alone turning over an object in your hand, writing a number, or walking across a room to take a book from the shelves. This is equally true for digital interactions, where typing or selecting with a mouse, or even taking in what is displayed on the screen takes effort and time. Some very early experiments on GUI interfaces took a well-known application and blanked out the text of menus. When asked to list the menu items the subjects were unable to remember, but when using the application they were able to function despite having no visual indication. That is, while they did not have explicit knowledge of the menu structure they did have tacit/implicit knowledge even though in normal use the menu labels (external representation) was always available.

More recently Gray and Fu [21] modified an experimental application (programming a VCR) in various ways to alter the effort required to find information visually (cost of external representation), and also whether participants had memorised information (cost of internal representation). In all cases they made sure that all information was available in the interface, simply made it more or less easy to see. They based their experiments on Anderson's view of human memory and processing being adaptive and optimising [1]. This led to the hypothesis that when the external representation was more costly and the internal representation less so, the subjects would not bother to seek the information in the interface (epistemic action), but instead rely on recall. The experiments supported exactly this hypothesis, but were not fully consonant with perfect optimisation. The subjects' memory was not perfect and they made mistakes, whilst the information in the interface was perfect. Inaccurate actions led to costly corrections, and yet the subjects

repeatedly (in the words of the paper title) ignored "*perfect-knowledge in the world for imperfect knowledge-in-the-head*". Not only do we rely on internal representations, but, in this case, we do so more than we 'should' based on perfect optimality.

5. THE OPTIMAL FALLACY

Note that we have seen two kinds of appeal to optimality. The first was in Clark's 007 principle and in Gibson's claims about the immediacy of perception, both of which appeal to the eventual optimality of species due to evolution. The second is in Anderson's rational analysis, which is about the eventual optimality of individuals due to learning. While not doubting a general pressure towards optimality in both cases, it is typically not the case that this is always achieved, as Gray and Fu's results [21] indicate. There are various reasons for limited optimality.

5.1 Path effects

In evolutionary development, there is always a given start point; it is hard to make successful radical changes (four arms rather than two), as these need to appear slowly. As we know from algorithmics, hill-climbing leads to local not global optima, and evolutionary change is just that. Similar effects are evident in the web, the success of which was due not least to the way in which it built, in an evolutionary manner, upon earlier *de jure* or *de facto* standards. Twenty years on, the end results, XML, RDF may not be optimal for their purpose, but have grown through a process, a path. Attempts to make non-incremental movements towards 'optimality' may fail, as the XHTML efforts demonstrate.

5.2 Feedback and Self-Reinforcing Structures

Feedback effects, in evolution, notably sexual selection, but potentially other forms of co-evolution, can lead to stable or meta-stable situations that are far from and move further from sensible optimality. The peacock tail is a classic example of this, making it hard to avoid predators, and yet necessary for finding a mate. While sexual selection over aeons is probably not a major issue for web interactions, similar self-reinforcing structures can be seen elsewhere both within algorithms (e.g. see [31, 35]) and also at a large scale where network effects [15, 34] mean that emergent monopolies develop. Once a dominant word processor or social network becomes popular enough it can effectively become the only one, due to mutual reinforcement, whether or not it is the 'best'.

To some extent self-reinforcement may actually be a saving grace, when it comes to making sense of the modern world. Technology must both make sense to its designer to be produced, and make sense to the rest of us to be successful.

5.3 Rapidly changing environment

Evolution is about adaption to a particular niche. Where the niche changes rapidly, then successful species need not be optimal for it. In particular this is true of technological advance, which proceeds at a rate far faster than natural selection. We have already seen this at work in discussing Gibson's concept of immediacy, which while reasonable to apply to the natural world, is not at all clearly applicable to the artificial.

However, whilst a rapidly changing environment acts against evolutionary optimality to a particular niche, Calvin [4] claims it was precisely the rapid (in evolutionary timescale) changes due to the succession of ice ages in Europe that led to the development of modern human intelligence. Looking back to Gibson, we may not have optimal intuition as to the affordances of novel objects, but

we may be naturally good at affordance seeking, coming to know and recognise over time the action potential of things around us: physical or digital. In general, we may be better developed to learn to act than we are simply to act.

It is often argued that the pace of technological change is more rapid now than ever before in human history. There are perhaps other periods, such as the latter half of the 19th century that could claim this also, but we are certainly seeing radical changes within a single lifetime. The crucial question is whether these changes are within our abilities to adapt individually and culturally.

5.4 Resource limits

We can only learn within available sensory and cognitive resources, which have themselves developed (within limitations) through adaption to a natural, not technological environment.

Some years ago the UK National Institute of Agricultural Engineering developed improved shock absorbers for tractors because measurements of low-frequency vibrations had been found to be at dangerous levels. However, when deployed in the field the measured vibrations were no better than before. This was not because the shock absorbers didn't work, but because the tractor drivers drove faster; they drove to the point at which the vibrations caused discomfort, unfortunately well after the point at which they are dangerous. Harmful low-frequency vibrations simply do not exist in nature, and so we do not possess the sensory apparatus to detect them until they have already done damage.

While this is a story of physical sensing, we see similar stories in other areas. For example, as noted earlier, working memory is limited to around seven 'chunks' of information. We often push this to its limit and so forget things when we get interrupted. We have the means to realise when we are trying to remember too much, but no 'advance warning' when we are at near maximum – it is simply not needed in a non-technological world. Within psychology, those that pursue Anderson's rational analysis agenda often do so within a context of cognitive resource limits [27].

Looking at a longer time scale there are more severe resource limitations. Our brains consume significant amounts of energy. An inverse correlation has been noted between species' brain size and stomach size as digestion also requires significant energy and blood supply. There is a three-way trade-off between brain, gut and food, which, has been surmised, is intimately connected to the need for big-brained humans to become meat hunters.

5.5 Optimal optimisers are not optimal

The trade-offs due to resource limitations ultimately lead to the need for imperfect optimisation. If our brains were complex enough to have perfect learning, they would be so big we would die! Similar trade-offs can be found in short-term interactions with the environment. In the wild if you think too long and too hard you are dead, to work out the 'optimal' thing to do leads to a poor outcome, because optimisation is costly. This is even clearer in the case of epistemic action. Doing things to find out is costly in time, energy and maybe risk, as is especially evident in war. And even in design a similar principle holds, where a perfect design is an indicator of a poor design process – if the result is too good, you have spent too much time and money on it.

There is ultimately an infinite regress in optimality. At some point resource limitations mean one should stop attempting to optimise, but, in order to assess whether one has reached this point, one needs information and time to make decisions. That is deciding when you have done enough decision-making is itself

costly. Our human mechanisms for this have cut-offs that need not reflect any pre-conceived idea of 'optimal', not even optimally non-optimal.

6. A DYNAMIC INTERPLAY

So, we are creatures attuned to using the external environment, yet also with internal memory and thought, creatures that manage to use both, but are not necessarily 'optimal' in doing so. The crucial theoretical and practical question is how this dynamic of internal and external representation works out on the web.

6.1 Action and Education

The web is most often seen as a place of information: knowledge needed to do something elsewhere. From an enactment perspective, the web becomes an extension of oneself in one's interactions in the world, and one's actions on the web largely epistemic. This is largely the role that Halpin et al. see when discussing the web as representation and enactive search [23]. Alternatively the web can be viewed as a place of action [13], where one books holidays, buys books, meets friends. In such cases, we see enactment much more directly as the web is the world *within which* we are embodied. Finally, there is that autotelic use of the web for leisure or education: browsing, reading, whether Wikipedia, Peer2Peer university or YouTube.

One of Huizinga's characteristics of play is that it is for itself [29], similarly one of Csikszentmihalyi's criteria for flow is its autotelic nature [9], and while flow was coined to express rich experiences, such as a mountain climber acting at the edge her abilities, the qualities of timelessness and (although not termed as such) thrownness are surely ones familiar to those lost in surfing.

From an educational perspective, there is a tension. Vygotsky's zone of proximal development [45] suggests that we learn best when operating just at the edge of our capabilities, doing things with physical or cognitive aids (scaffolding), which we later cast off as we become proficient and move on to the next level. The web as external representation, or extended knowledge, can be seen as scaffolding. However, the disjointness of multifarious web snippets may not provide the coherence needed for that 'at the edge'-ness, where new knowledge can be linked to old. Furthermore constructive learning theory emphasises the importance of *processing* new information, so that it becomes knitted into our mental fabric. It is not clear that the instant accessibility of Google-augmented activity, albeit powerful for the moment, provides grounding for long-term learning.

However, even if finger-tip information means we fail to learn or even learn not to learn, then surely this is exactly what we would expect from an extended mind. We do not need to learn precisely because the knowledge is (stretching Heidegger's term) '*ready to hand*'. What matters is not what we know, but how we make use of the web; librarian not scholar.

6.2 Meta-representation

The librarian does not know, but knows how to find out. Even some years ago it was evident that '*post-web cognition*' [11] involved a greater reliance on meta-knowledge of various kinds. Table 1 lists various different kinds of knowledge. Some knowledge is first order, knowledge about a real-world domain; this may be concrete knowledge about a particular object, or generic knowledge about a whole class. Some is second order *meta-knowledge* about *what* information is available, and *how* to find it. This again can be specific, knowing that a particular piece of knowledge is on the web, and its URL or appropriate search

terms. Some is more generic, for example knowing that airports usually have associated official web sites.

Table 1. Kinds of Knowledge

	domain	meta-knowledge	
		what	how
concrete	Everest is 2900 ft high	that an Everest web page exists	search "Everest statistics"
generic	airports have runways	airports have web sites	search for airport name, if that fails look at airport Wikipedia page

Going back temporarily to the issues of education; the bottom left cell is potentially problematic. A shift towards retrieving facts (concrete domain knowledge) from the web rather than from our brain, is fine so long as we are able to interpret and understand what we receive – that is have appropriate *generic* domain knowledge. However, generic knowledge is usually obtained through continued exposure and consideration of concrete cases. If this concrete knowledge is increasingly accessed unreflectively, it is unclear how this interpretive framework can develop.

Moving on to meta-knowledge, this is not something new with the web, but something that has been happening continuously as part of our move from oral to written culture. Indeed Hyland-Wood [30] suggests that metadata is as old as data; early promissory tablets were sealed in clay to prevent tampering; on the outside of the clay seal was a description of the contents – metadata. It also appears that even in early times metadata, as ordinary knowledge, was sometimes viewed as dangerous by those in power:

"Metadata can be viewed as an invention of the Near East with substantial refinements by the Greeks. Its spread westward formed the basis for Western libraries. Reasons of culture and governmental authority tended to suppress the use of metadata in the east." [30]

It is important to note that, like our knowledge of other things, meta-knowledge can be tacit (as was evident in the early blank menu experiment) and indeed often needs to be so if it is to be used unreflectively during routine access of external information. The moment we think about what we know or how to find something we are in a 'breakdown' situation thinking about the tool (the web) rather than using it to do something.

However, we should also realise that meta-knowledge, while often skilled and artful, is typically neither 'optimal' nor perfect. Gray and Fu's [21] subjects were effectively applying tacit meta-knowledge as they chose (albeit subconsciously) between using remembered information rather than looking for it in the interface. Arguably the experiment here was short and so it is reasonable that users had not been able to optimise behaviours, but similar effects can be seen even in experts.

Salmoni and Payne performed a series of experiments on users' ability to assess relevance of search results [40]. In one they looked at titles and Google-like snippet, comparing title only, snippet only and title+snippet displays. They found that subjects performed better in the titles-only condition compared with snippets-only. This they attributed to the Google-style snippets bringing together disparate parts of a document making it appear that a multi-term search is more relevant than it really is. One might think that given the combined information in title and snippet, subjects would perform better than both. In fact, it was between the two, better than snippets-only, but worse than titles-only. All the subjects were expert at web search from their day-

to-day work, and yet had not learnt to ignore (or at least use with care) the snippet compared with the title.

Another kind of meta-knowledge concerns trust and authority a crucial issue in any enactivist theory of collective intelligence [23]. In social situations this can be a matter of knowing who to ask, in the words of the music hall song:

*And you can't trust a "Special"
Like the old-time copper
When you can't find your way home.*
<http://ingeb.org/songs/myoldman.html>

The non-optimality of self-reinforcing structures can be a real problem here; there is a human tendency to trust popular opinions, which, in a world of blogs and tweets, then become more cited, so more popular. As Keen points out [32:p.95] this is reinforced further by Google rankings. There is a narrow line between the wisdom of crowds and the madness of the mob.

6.3 Breakdown and Reflection

Accounts of enactment emphasise the normality and power of unconsidered artful interaction with the world. Breakdowns, the moments when we become explicitly aware of the tools and techniques we are using, are considered aberrant interludes.

However, an over-reliance on the tacit can turn the mindfulness of the contemplative to the mindlessness of the brute. This is not to deny the importance of unreflective activity, an athlete who constantly thought about her feet moving whilst in a race would undoubtedly lose, and even Descartes suggested that the level of radical scepticism that led to '*cogito ergo sum*' was a once in a lifetime effort [10]. It is just that breakdown can be valuable as well as problematic as it offers an opportunity for *reflection*,

More practically Schön's observations of expert designers and related professionals [42] emphasises the importance of '*knowing in action*' exercised during the flow of design activity, but also '*reflection in action*', moments during the flow for higher level consideration, and '*reflection on action*', more long term consideration of practices, process and techniques after and between design activity. In my own advice to students I emphasise the importance of *externalisation* [12], making tacit knowledge explicit so as to be amenable to analysis.

Even in *Blink* [20], a book dedicated to the importance of instant snap decisions, multiple examples are given where experts have had to reflect and train themselves to only focus on certain details and ignore others. Once they have done this reflection, snap decisions are possible. Like the athlete studying videos with their trainer between races to improve their physical skills, reflection can be used to train tacit cognitive skills as well as inform more explicit decision-making.

This is critical as it offers a way to break some of the optimisation limits discussed in section 5. For example, once one realises that snippets are misleading one can explicitly remind oneself of this when exploring search results (breakdown), and then hopefully, over time, one learns to do so tacitly (improved artful behaviour).

Furthermore as well as reflection being a potential way to overcome limitations of post-web cognition, computation and the web can themselves be opportunities for reflection. This reflection may be for personal well being (e.g. [41,2]), for health (e.g. fitness web sites such as mapmyrun.com), or for more cognitive ends. One example in *Blink* [20], is the work of cardiologist Lee Goldman, who, through computer analysis, discovered that of the many factors that doctors took into account when deciding on

cardiac care, in fact only three were valuable. The human tendency, just as with the search result titles and snippets, is to try to use all the information available. This is a reasonable approach 'in the wild' when information is maybe sparse, but, in information-rich environments, it can lead to noise reducing the efficacy of decisions; in the case of Goldman's data *unnaturally* ignoring all but three factors led to a 70% improvement in outcome. The web makes more information more easily available, and yet in many situations we appear unable to tacitly learn what is important and what should be ignored. That is the web makes explicit analytic reflection *more* important.

An ongoing design challenge for the web is to create the additional aids that enable people to improve their meta-learning, and maybe even improve ordinary learning, especially of generic and interpretative knowledge. One of the things that helped me personally to learn to spell was a computer spell checker that *did not* suggest corrections. It told me that something was wrong, but I had to think through the spelling or reach for dictionary to correct it, thus forcing processing and hence learning. Are there similar things in other areas?

One candidate is the Mac OS application 'Freedom', which allows you to 'turn off' the internet for a fixed period of time, helping to avoid the inefficient rapid switching between mail, web and what you really want to be doing. Aids would be useful in other areas that we find hard to assess, for example, in provenance tracking and determining authority/reliability of sources, maybe even 'forcing' these to our attention in some way.

7. DEEPER CHANGES

The issues covered by this paper have focused mostly on the changes in cognition experienced by individuals as we adapt to new technology. Mostly these are relatively short term changes that can change again equally easily, for example, those taking an 'internet break' describe themselves being able to focus for longer within weeks of 'kicking the habit'. One of the advantages of our human ability to 'extend' ourselves in the environment, is that we can adopt new things, whether hammers or internet, but equally we can put them down. However, some changes particularly those impacting children and culture have a different dynamic.

Short-term learning can easily be re-learned or unlearned, changing what we know or do, but not who we are. However, understanding of the plasticity of our brains shows that far more dramatic effects are possible and now observable through brain scanning technology; perhaps most well known being the expansion of spatial areas in London cab drivers. So, the changing balance of costs and availability of factual and social information on the web itself may lead to substantive changes in our own cognitive structures. While much has been said about the 'digital native', we are still only just seeing children immersed in digital technology from infancy. It is known that those brought up without exposure to language at critical developmental stages never acquire certain forms of abstract thinking, and Dundee University have noted that, due to gaming and texting, young medical students now have thumb dexterity comparable to older medics' index fingers (cognitive and physical adaptation). What potential deeper cognitive changes are likely?

Whether it is the railroad in the 1800s, universal suffrage in the 1920s, the digital revolution of the 1980s, or the web in the last decade, technological and even social change often moves faster than each individual. Equally clearly, whether in Amazon recommendations or the Arab Spring, digital networks mean that individual actions are increasingly and increasingly rapidly

reflected in the collective, and the collective impinges on individual life. As discussed in section 2, socially distributed cognition is not simply a phenomenon of modern life. Indeed, fragments of cloth from 1300 BC show that several weavers worked simultaneously passing bobbins to one another [3:p.86]. In the past, with pre-digital social interactions, this has tended to be a local phenomenon, with small groups often in physical, aural and visual contact; in contrast, network-distributed cognition, while building upon the same underlying cognitive abilities, may well have very different qualities.

Finally this is not the first technological revolution; can the past inform current, albeit rapid, transformations of human society? Language as perhaps the earliest 'information technology' is crucial here, but also may be so special and deep that it is misleading. The built environment is another comparable innovation, but one that even those from extreme rural environments are able to accommodate with time. Throughout history we have modified our environment and adapted cognitively to make use of it: is the web just more of the same, or the most radical change since the acquisition of language itself?

8. CONCLUSIONS

In the digital world, no less than the physical, we are part of it, continuously and often unconsciously acting on and thinking using the environment around us. However, as evident from both empirical evidence and common experience, we also have an internal life of thoughts and memories brought to bear, whether tacitly or explicitly, as part of our in-dwelling in the world. The resulting dynamics of internal and external representation is not simply one of the web augmenting or assisting existing cognitive behaviours, but potentially changing them. Notable is the shift from knowledge about things to meta-knowledge, knowledge about how information is stored and accessed on the web. As with previous technological changes, it is far from a clear picture of inevitable progress, but understanding these changes, whether good or ill, is a core challenge for technology, for education, for society, and for each one of us.

9. REFERENCES

- [1] Anderson, J. 1990. *The adaptive character of thought*. Erlbaum
- [2] André, P., schraefel, m., Dix, A. and White, R. 2011. Expressing well-being online: towards self-reflection and social awareness. In *Proc. of iConference '11*. ACM Press, 114–121. DOI=10.1145/1940761.1940777
- [3] Barber, E. 1994. *Women's Work: The First 20, 000 Years – Women, Cloth and Society in Early Times*. Norton.
- [4] Calvin, W. 2001. *The Ascent of Mind: Ice Age Climates and the Evolution of Intelligence*. iUniverse.
- [5] Card, S., Moran, T. and Newell, A. 1983. *The Psychology of Human-Computer Interaction*. Lawrence Erlbaum.
- [6] Carr, L. and Harnad, S. 2009. Offloading Cognition onto the Web. *Proc. of WebSci'09*
- [7] Clark, A. 1989. *Microcognition: Philosophy, Cognitive Science and Parallel Processing*. MIT Press.
- [8] Clark, A. 1998. *Being There: Putting Brain, Body and the World Together Again*. MIT Press.
- [9] Csikszentmihalyi, M. 1990. *Flow: The Psychology of Optimal Experience*. New York: Harper and Row.

- [10] Descartes, R. 1644. *Principles of Philosophy*. trans. MacDonald Ross, G., 1998–1999. online version at: <http://www.philosophy.leeds.ac.uk/GMR/hmp/texts/modern/descartes/principles/deprinc.html>
- [11] Dix, A., Howes, A. and Payne, S. 2003. Post-web cognition: evolving knowledge strategies for global information environments. *Int. J. Web Eng. Technol.* 1, 1, 112–126. DOI=10.1504/IJWET.2003.003323
- [12] Dix, A. 2008. Externalisation – how writing changes thinking. *Interfaces*, 76, 18-19. <http://www.hcibook.com/alan/papers/externalisation-2008/>
- [13] Dix, A., Lepouras, G., Katifori, A., Vassilakis, C., Catarci, T., Poggi, A., Ioannidis, Y., Mora, M., Daradimos, I., Md.Akim, N., Humayoun, S. and Terella, F. (2010). From the Web of Data to a World of Action. *Journal of Web Semantics*. 8, 4, 394–408
- [14] Dreyfus, H. 1991, Being-in-the-World: A Commentary on Heidegger's Being and Time, Division I. The MIT Press, Cambridge
- [15] Economides, N. 1996. The Economics of Networks. *Int. J. of Industrial Organization*. 14, 6, 673–699
- [16] Gallagher, S. 2005. *How the Body Shapes the Mind*, Oxford University Press
- [17] Garfinkel, H. 1967. *Studies in Ethnomethodology*. Prentice-Hall.
- [18] Gaver, W. 1991. Technology affordances. In *Proc. of CHI'91*. ACM Press. 79–84.
- [19] Gibson, J. 1979. *The Ecological Approach to Visual Perception*. Lawrence Erlbaum.
- [20] Gladwell, M. 2005. *Blink: The Power of Thinking Without Thinking*. Little, Brown
- [21] Gray, W. and Fu, W. 2001. Ignoring perfect knowledge in-the-world for imperfect knowledge in-the-head. In *Proc. of CHI '01*. ACM Press, 112–119. DOI=10.1145/365024.365061
- [22] Halford, S., Pope, C. and Carr, L. 2010. A Manifesto for Web Science. In *Proc. of the WebSci10*.
- [23] Halpin, H., Clark, A. and Wheeler, M. 2010. Towards a Philosophy of the Web: Representation, Enaction, Collective Intelligence. In *Proc. of the WebSci10*
- [24] Hartson, H. 2003. Cognitive, physical, sensory, and functional affordances in interaction design. *Behav Inform Technol.* 22, 5, 315–338
- [25] Heidegger, M. 1927. *Sein und Zeit*. (English translation: *Being and Time*. Harper, 2008)
- [26] Hollan, J., Hutchins, E. and Kirsh, D. 2002. Distributed Cognition: Towards a New Foundation for Human–Computer Interaction Research. Chapter 4 In: *Human–Computer Interaction in the New Millennium*, Carroll, J. (ed.). Addison-Wesley, 75–94.
- [27] Howes, A., Vera, A. and Lewis, R. 2007. Bounding rational analysis: Constraints on the approach to optimality. In *Integrated Models of Cognitive Systems*, Gray, W. (ed.). Oxford University Press
- [28] Hutchins, E. 1995. *Cognition in the Wild*. MIT Press.
- [29] Huizinga, J. 1944. *Homo Ludens: A study of the play-element in culture*. Routledge.
- [30] Hyland-Wood, D. 2008. *Metadata Foundations for the Life Cycle Management of Software Systems*. Ph.D. Thesis, The University of Queensland, <http://espace.library.uq.edu.au/view/UQ:159852>
- [31] Katifori, A., Vassilakis, C. and Dix, A. 2010. Ontologies and the Brain: Using Spreading Activation through Ontologies to Support Personal Interaction. *Cognitive Systems Research*. 11, 25–41.
- [32] Keen, A. 2007. *The Cult of the Amateur : how today's internet is killing our culture*. Doubleday/Currency.
- [33] Larkin, J. 1989. Display-based problem solving. In *Complex Information Processing: The impact of Herbert A Simon*, Klahr, D. and Kotovsky, K. (eds.). Erlbaum, 319–341.
- [34] Liebowitz, S. and Margolis, S. 1998. Network Externalities (Effects). entry in *The New Palgrave Dictionary of Economics and the Law*. MacMillan.
- [35] Md. Akim, N. Dix, A., Katifori, A., Lepouras, G., Shabir, N. and Vassilakis, C. 2011. Spreading Activation for Web Scale Reasoning: Promise and Problems. In *Proc. of WebSci'11*.
- [36] Merleau-Ponty, M. 1945. *Phénoménologie de la Perception*. (English translation: *Phenomenology of Perception*, Routledge, 1958)
- [37] Miller, G. 1956. The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information. *The Psychological Review*. 63, 81–97
- [38] Norman, D. 1998. *The Design of Everyday Things*. MIT Press.
- [39] Payne, S., Richardson, J. and Howes, A. 2000. Strategic use of familiarity in display-based problem solving. *Journal of Experimental Psychology: Learning, Memory & Cognition*. 26, 6, 1685–1701.
- [40] Salmoni, A. and Payne, S. 2002. Inferences of Content from Search Engine Summaries and Headings. *Proc. HCI2002*, vol 2.
- [41] Sas, S. and Dix, A. (eds.) 2011. Designing for Reflection on Personal Experience (Special Issue). *Int. J. of Human-Computer Studies*. 69, 5, 281–350.
- [42] Schön, D. 1984. *The Reflective Practitioner: How Professionals Think in Action*. Basic Books.
- [43] Suchman, L. 1987. *Plans and Situated Actions : The Problem of Human-Machine Communication*. Cambridge University Press.
- [44] Varela, F., Thompson, E. and Rosch, E. 1991. *The Embodied Mind: Cognitive Science and Human Experience*. MIT Press
- [45] Vygotsky, L. 1978. *Mind and society: The development of higher psychological processes*. Harvard University Press.