Against `Distributed Cognition'

Graham Button

Theory Culture Society 2008; 25; 87
DOI: 10.1177/0263276407086792

Against 'Distributed Cognition'

Graham Button

WITHIN DISCIPLINES that draw upon cognitive science, such as anthropology, questions are being asked about the status of some of its prime contentions (cf. Lave, 1988). Some have become hostile to the idea that 'thinking only goes on inside the heads of individuals' (Hutchins, 1996), and, though they might not necessarily put it as such, they are gesturing towards a social conception of mind. Their critique originates in their appreciation that, among other things, individuals engage in conduct with others; that they live and work within a culture, and that they cooperate with one another. Taking account of at least these facts about the social character of human life, they have become uneasy with cognitive science's preoccupation with an inner world as the basis for the explanation of human doings.

Some of their uneasiness has been prompted by Suchman's book *Plans and Situated Actions* (1987). Suchman mounted an attack upon iconic figures in cognitive science such as Herbert Simon for their failure to understand the implication of the fact that human action is situated in the social and cultural world for the explanation of action. In particular, she undermined the idea that human action can be accounted for in terms of 'cognitive science' uses of such otherwise mundane terms as 'plans', which has become an important concept in the development of Artificial Intelligence (AI). Developing the notion of 'script' proposed by Shank and Ableson (1977), researchers in AI such as Wilensky (1981) have proposed that before engaging in a course of action people first, mentally, develop a plan of action, and that their subsequent action is then the fulfilment of that plan. Consequently, human action can be explained in a quasi-causal vocabulary in terms of the putatively 'mental' antecedent of the plan. Suchman's rebuttal of such an explanatory framework made reference to the fact that all action takes place within a swarm of socio-cultural contingencies that cannot be covered in full, and in advance, by a plan. A plan may be a guide for some actions, but it will have to be applied within the context of a range of

■ *Theory, Culture & Society* 2008 (SAGE, Los Angeles, London, New Delhi, and Singapore), Vol. 25(2): 87–104
DOI: 10.1177/0263276407086792

contingencies that have not been pre-formulated, and, in that respect, it cannot be a causal antecedent of action. Indeed, many actions can be undertaken without someone having formulated a plan at all.

Suchman's work has been taken up in many quarters that contend with a cognitive presence in their disciplines, such as anthropology and the field of Computer Supported Collaborative Work (CSCW). In this respect, the development of 'distributed cognition' may be appealing to those who have found Suchman's work on situated action of interest. Thus, Hutchins, the founder of 'distributed cognition', writes: 'The emphasis on finding and describing "knowledge structures" that are somewhere "inside" the individual encourages us to overlook the fact that human cognition is always situated in a complex socioculural world and cannot be unaffected by it' (Hutchins, 1996: xiii). Those who regard 'distributed cognition' as a development of the critique of cognitive science are probably drawn to a particular part of this quotation that emphasizes 'that human cognition is always *situated in a complex socioculural world* and cannot be unaffected by it' (emphasis added). Thus, Hutchins is widely regarded as a champion of a socio-cultural view of human activity rather than as a cognitivist.

However, emphasizing this aspect of the above quotation overlooks another important part of what Hutchins is proposing: '*that human cognition* is always situated in a complex socio-cultural world and cannot be unaffected by it' (emphasis added). Hutchins' supposed critique of cognitive science is not proposing that the very idea of cognition is itself a mistake, the idea that is articulated in Suchman's arguments, and most forcefully promulgated by other critiques of cognitive science which, drawing from the work of Wittgenstein, explicitly argue the case for the social status of mind (cf. Coulter, 1979; Williams, 1999). Instead, Hutchins' position is that it is possible to provide a cognitive science treatment of the socio-cultural world; thus extending the remit of cognitive science to the socio-cultural world: 'Culture is a . . . process. It is a human cognitive process that takes place both inside and outside the minds of people' (Hutchins, 1996: 354). Thus, the appeal to the socio-cultural world is not done to critique cognitive science; Hutchins is rather chiding cognitive scientists for having allowed the skull to define the boundaries of their discipline. Thus, in his book *Cognition in the Wild* (1996), Hutchins locates his examinations of the activities involved in piloting and navigating large vessels within Marr's (1982) cognitivist consideration of information-processing and Simon's (1981) representational description of problem-solving. Consequently, rather than bringing into a consideration of mind the social character of human thought in order to confront an inappropriate computational model of mind articulated by cognitive science, Hutchins wants to actually extend its provenance: 'I will move the boundaries of the cognitive unit of analysis out beyond the skin of the individual person . . .' (Hutchins, 1996: xiv).

In contrast to treating 'distributed cognition' as a welcome corrective to cognitive science that aligns with other sociologically derived criticisms, it will be argued here that 'distributed cognition' promulgates mistaken

ideas about mind and meaning that originate in cognitive science. The main thrust of the argument is quite simply that 'distributed cognition' *is* cognitive science. 'Distributed cognition' does not renounce cognitive science; it actually extends it by describing social phenomena in a redundant cognitive vocabulary. Just as simply, it is argued here that it should be dismissed for the very same reasons that cognitive science should be dismissed: because it argues the plausibility of the dichotomy between an inner and outer world. To make this argument, the thesis of 'distributed cognition' will be considered in terms of Wittgenstein's arguments against the idea of a private language, and through a consideration of the way in which 'distributed cognition' theorizing considers 'thinking in context'. Following a brief consideration of some of its main arguments we will start our critique by asking what actually is explained by the theory of 'distributed cognition'?[1]

'Distributed Cognition'

'Distributed cognition' asserts that so-called cognitive phenomena such as 'intelligent processing', are not just the province of individual minds, but are distributed across individuals, and that what are said to be the representational states of individuals are mediated through cognitive artefacts such as technological phenomena. Thus, for instance, 'distributed cognition' would address the work setting by arguing that it is not possible, as traditional cognitive science would argue, to account for the activities going on in terms of individual minds processing their environment and then reacting in certain ways on the basis of that processing. Rather, 'distributed cognition' would view that processing as going on across the boundaries of the individuals involved. Thus, what to the sociologist or economist might be a division of labour in a work setting, is to the distributed cognitive scientist a cognitive system in which individuals are coordinating their individual processing through shared representations.

For example, at the beginning of his book *Cognition in the Wild*, Hutchins provides a dramatic account of what happened when the steering gear of USS *Palau* failed when entering San Diego harbour and how the ship was guided to a safe anchorage. He writes of the work this involved:

> . . . no single individual on the bridge acting alone – neither the captain nor the navigator nor the quartermaster chief supervising the navigation team – could have kept control of the ship and brought it safely to anchor. Many kinds of thinking were required to perform this task. Some of them were happening in parallel, some in co-ordination with others, some inside the heads of an individual, and some quite clearly both inside and outside the heads of the participants. (1996: 5–6)

'Distributed Cognition' then, views the cultural world and the activities and interactions between people that make it up as 'cognitive systems'. Its particular contribution is not so much to take observations about that world as the basis for the reformulation of cognitive science, but to use the

conceptual apparatus of cognitive science to say things about those observations. Instead of, for example, describing the activities that went on aboard the *Palau* when its steering failed in socio-cultural terms, for example through a system of reciprocally shared rules (cf. Parsons), Hutchins describes their activities in cognitive terms, as the result of brain processing, inside an individual head and also brain processing being coordinated between heads. 'Distributed cognition' is clearly an attempt to let the cognitive genie out of the bottle and loose it on the social world.

One result of so doing is that everyday objects in and for a workplace become re-described in the arcane language of cognitive science and a computational model of mind. For instance, in his description of the maturation of the ship's log from a rope with knots in it, to a structure built into the ship, Hutchins, writes:

> Putting the calibrating nails into the deck is a way of creating a memory for the lengths between knots in the log line in a medium that has psychical properties that match the computational needs of the task. In this case, the marks on the deck are a *memory* for distance. (1996: 106)

By way of another example, the chart the navigator uses in plotting the position of a vessel or the heading the ship should be placed on is not just a chart that documents certain information about the sea and the seabed, which can be used in navigating, it too can be considered as a 'memory'. Within cognitive science, memory is described in computational terms as 'a store' of information awaiting retrieval. In these terms, a chart can be viewed as containing a store of information about the position of objects, what those objects are, the topography of the seabed, the composition of the seabed, and so on. Therefore, if we take Admiralty chart No. 742, *Mahe, Praslin and Adjacent Islands*, we would find the 'memory' that Capucin Rock is 0.4 of a mile SSW of Pte Police on the south west of the Island of Mahe, and that water breaks over it. However, unlike an individual's memory, which is understood in cognitive science as an information store that an individual retrieves information from, an object such as a chart is a *shared* memory. Charts, as memories, are stores of information that can be accessed by anyone. The chart, for Hutchins, thus embodies a memory of where the 5-meter contour line in Anze Latio bay is that can be accessed by all those who read it, and, consequently, is something people can use to do things in common with one another, for example, anchoring up together on the 5-meter contour mark in Anze Latio bay on the island of Praslin. Within the theory of 'distributed cognition', the use of technology externalizes certain cognitive processes, and, if the technology is being used publicly in a division of labour, then some parts of that processing can be observed by the others. It is through such objects that it is said that processing can go on between people. Consequently, the use of the technology can be both a part of a person's private cognitive processes and also an element of communication.

For Hutchins, then, the work of manoeuvring such a complex of technology that constitutes the *Palau* is a piece of what he calls cultural cognition. Persons are able to accomplish the overall operation through the coordination of their thinking through the socially constituted objects relevant for their tasks, but which are reinterpreted in the terms of the computational model of mind that is paramount in cognitive science.

What Is Explained by Cognitive Systems?

'Distributed cognition' views the socio-cultural world as a cognitive system within which the processing that is said to be going on in one person's head is coordinated with the processing that is said to be going on inside another person's head through cognitive artefacts, often technological ones. Following a cognitivist tradition, 'distributed cognition' understands 'mind' computationally, running processes on inner representations, and it is by reference to these inner representations that human action is understood. A first question that needs to be asked of 'distributed cognition' in this respect is: 'What is actually described by stating that the socio-cultural world is a cognitive system?' Specifically: 'What is described by stating that a work setting is a cognitive system?' And, even more specifically: 'What is described in stating that a device or a tool used in that setting is a cognitive artefact?'

Starting with the last question; if it is necessary to have an inner representation of something in order to orient to it through some mental process, does merely saying that having an inner representation of something actually explain what something is or how it is being used by someone in the accomplishment of their work, or how it is used between people who are collaborating in some work task? Hutchins gives numerous examples of tools being used by people in their work of navigation: logs, charts, dividers and the like. So we could ask with respect to one of these, for example a chart, 'How does saying that in order to use a chart we have a mental representation of it in our heads actually explain what a chart is or how we use a chart?' If we were to explain what a chart was to someone who has not seen one we would have to make recourse to its public character, not some inner representation. We would point out its role in the organized world of the practices of navigation where maritime bodies have laid down conventions as to what objects depicted on the chart mean, and how to depict objects that are important to the community of seafarers. Thus, for instance, charts are not very detailed, when compared to ordnance survey maps, about the features of the land. Very few objects are depicted, and there are certainly no contour lines and the other cartographic notations. What is depicted are only features along a coastal strip and these tend to be features that could be used by navigators in order to make accurate fixes; for example, tall buildings or chimneys that can be seen out to sea. Such features and their depiction on charts is made relevant by, and embedded, within, the fabric of navigating practices.

There has also been a public agreement within the seafaring community about how features such as chimneys are depicted, so that any navigator knows through having learnt how to navigate what the symbols on a chart represent, and any navigator displays that they know these things in their very practices of navigating. Thus, the IALA (International Association of Lighthouse Authorities) A buoyage system has been agreed across a range of the seafaring nations and lays down that red buoys are port-hand markers and green buoys are starboard-hand markers, and the UK hydrographic office, which produces Admiralty charts, uses standard symbols to depict, for example, the presence of rocks and make these symbols publicly available. Thus, in order to understand what a chart is, and why certain sorts of objects are shown and others are not, it is necessary to make reference to its public character, and its role in public practices, in this case the practices of navigating.

However, not only would an inner representation not explain the meaning that charts have for the navigating community in this respect, nor how they are used by them, there would also presumably have to be multiple representations of the same objects, for there is not only one buoyage system but two: IALA A and IALA B. Those seafaring nations that have agreed to IALA A have agreed that a green marker marks the starboard side of a navigable stretch of water such as a channel in a river, while those seafaring nations that have agreed to the IALA B system have agreed that a green marker marks the port side of the channel. So, in the UK, green means starboard and red means port, but, in the USA, green means port and red means starboard. Presumably, then, we would have to have two inner representations of a marker, a UK one and an American one. Also, presumably, if we made a mistake and ran aground because we had assumed that, because we were in the *British* Virgin Islands, they operated IALA A when in fact they have opted for IALA B, we would have mixed up our mental representations. However, what is it to say we mixed up our mental representations? Is it to say that I mixed up one mental picture of a green cone with another mental picture of a green cone? The way we would explain our mistakes to any passer-by as we languished waiting for the tide is not in such mentalistic terms, rather it would be in just the terms we have used above. We would make recourse to the fact that we confused the buoyage systems, and if pushed because the person to whom we were giving the explanation did not understand, we would go on to explain the different meanings the markers have in the different chart conventions and explain that we confused one system with the other. We would have recourse to the public and agreed meanings things have, not to mental representations. Describing a chart as part of a cognitive system does not, then, explain what a chart *is* or how it is used within the communal practices of navigation – it is just an assertion that does not add to our understanding of something.

'Distributed cognition' also views cognitive artefacts as the mechanisms through which individual processing is coordinated. The reason for this is that they are described as a memory that can be accessed by the

different people involved. Obviously, a chart may be used by two people to coordinate their actions with one another. Thus, a watch navigator (at least in the navigation of small vessels) may lay down courses in pencil on a chart and make pencilled notations as to the position of the vessel, the time of the fix, the distance run, and even the heading. They will also enter information lifted from the chart and other instruments into the logbook. When the watch changes, the new watch navigator will use all this information on the chart and quickly orient to the current circumstances as depicted by his/her predecessor. However, what is added to this description if we say that the chart is a cognitive artefact within a cognitive system and which works to coordinate individual processing? If we were to explain how the two navigators coordinated their activities by using a chart we would have to explain how they use the chart within a community of navigational practitioners to perform, according to agreed conventions, certain calculations that they and their colleagues can use in navigating a vessel. Thus, as part of the community, one navigator can use the calculations made from the chart by another navigator to take over when the watch changes. Again, the ability to use a chart in navigational calculation and the ability for one navigator to use the calculations of another navigator resides in there being a public standard regarding the doing of navigational calculations which people who use the chart and do the calculations express in what they do, in their practices of navigating. Does saying that a chart is a memory, then, actually explain what the navigators are concertedly doing? 'Distributed cognition' is trying to align the world to its cognitivist model by reclassifying things found in the world according to cognitivistic categories; the point being made in this article is that this is a redundant exercise because the reclassification does not help us to understand how these things play out in the social world where their meaning is publicly agreed. And that if we want to explain something in that world, we would not be able to do so through those cognitive theories, we would use publicly agreed meanings.

The everyday tools of the trade of navigation are not only construed as memories, Hutchins is also fond of describing some of them as 'analogue computers' as in his description of the ancient Greek 'astrolabe', a mechanical device that modelled the movement of heavenly bodies: 'the astrolabe is not just a memory for the structure of the sky; it is also an analogue computer' (1996: 98).

When astrolabes were in use by astrologers, it was an everyday term for them, just as radar, sextant and global positioning system, are everyday terms for those who use and work with these things today. Such technical vocabulary may not be familiar to those outside of the particular group that frequently uses them. However, although some terms that are used frequently by some groups of people may be less familiar to the rest of us, nevertheless, they are also concepts in our language, the meaning of which are available to us through an *explanation* of their meaning by someone who knows what they mean. Not knowing what something means does not mean that we *cannot* know what it means, and when we are given an explanation

or a description of something we are being given its meaning. When someone explains a word or an idea to us, what is explained is just simply what that thing means. An explanation of the meaning of some technical object may involve describing what they are used for, and how they are used, and it may be that we are only given a simple understanding if they are particularly technical matters. However, even then, how the description is done is in everyday language and what we learn is what they mean as they are described in everyday language use. We have, quite simply, ways of describing things in our socio-cultural world, even if those things are technical, and even if some things are difficult to understand. The socio-cultural world is, then, one that comes with its own description, so to speak, and that description is done in everyday language.

Like us today, people outside of the particular community that used astrolabes might not understand what they were. However, if at the time of their use someone asked of someone who used them what an astrolabe was, what they would receive as an answer, if their question were respected and treated seriously, would be a description couched in ordinary language. They would not have been told, however, that it is an analogue computer. Even people who at the time of their use were knowledgeable about astrolabes would not describe them to one another as an analogue computer because astrolabes were used before the advent of analogue computers. An astrolabe was not at the time of its introduction an analogue computer simply because analogue computers had not been invented.

Following the invention of the analogue computer, even if an astrolabe does exactly what the analogue computer does, it is still not correct to call it an analogue computer. We could say, if it were true, that it was the precursor of the analogue computer, or that our forebears had invented something that does what an analogue computer does, if indeed it does do that.

Calling astrolabes analogue computers is then to *re-describe* them, just as to describe a chart as a memory is to re-describe what a chart is. For like charts, astrolabes would not be ordinarily described in the way suggested by Hutchins, they would be described in terms of their use in and relevance for the community who used them. 'Distributed cognition' is then in the intellectual business of re-describing the socio-cultural world, and as such it follows in the footsteps of the human sciences which have, in the main, taken as their own intellectual task the re-description, done in unfamiliar terms, of everyday activities that people already have descriptions of and which they use as common currency in their daily lives. When this re-description goes on it does so in the terms of a particular theoretical orientation. Thus 'distributed cognition' is re-describing the socio-cultural world in the terms of the theory of cognitive science, and in this respect the everyday descriptions of things is being replaced by descriptions that draw on a particular theory of mind, a computational model of mind in which the mind is rendered in terms of computational processing. Objects in the world, activities, interactions, perfectly well describable by using everyday descriptions, are being reinterpreted in computational and, hence,

unfamiliar ways. The question we then have to ask is why re-describe them, what is gained in re-describing them?

In order to understand why 'distributed cognition' theorists may wish to engage in this re-description of the socio-cultural world, we have to remember that the objective of 'distributed cognition' is, as has been argued here, to move cognitive science onto new ground. By re-describing aspects of that world in a cognitive science vocabulary, 'distributed cognition' theorizing is then attempting to show that the cultural world can indeed be handled in the very terms of cognitive science. Thus, showing that it is possible to re-describe the world in cognitive terms is to demonstrate the extendibility of cognitive science from descriptions of a supposed inner world, to descriptions of the outer world.

We could grant, however, that calling a chart a memory makes a certain point, in the same way as using a metaphor may illuminate something about an object we had not appreciated before. Saying that someone's hat looks like the Taj Mahal, or that the world is a stage, is a way of drawing attention to the hat's overbuilt qualities or that people may be acting out their lives. Calling a chart a memory may draw attention to the ways in which charts document a fact, which means that we do not have to work out for ourselves and on every occasion we might pass over a sea area what the depth is. However, as Louch (1966) points out, when we say these sorts of things metaphorically, we are not producing a *literal* description of something. Therefore, we do not mean that the hat is, literally, the Taj Mahal, or, literally, that the world is a theatre.

However, the re-description of the socio-cultural world in disciplinary terms is intended to be a literal description, because it is intended to provide us with knowledge about objects in that world that we previously did not have. The point being argued here is that this wholesale replacement of our everyday terms with seemingly scientific terms is untenable, for calling a chart a memory does not then tell us anything about charts that we did not know, or calling an astrolabe an analogue computer does not tell us anything about astrolabes that we could not know from reading about the use of navigational instrument; if anything, it can be actually confusing. If we want to know about things in our world that we do not know about them then we come to know about them through the explanation of their use, not through the application of a scientific term. Even here, we have to be careful, because the so-called scientific terms are drawn from our everyday language use, so the terms of cognitive science are familiar to us, for we all use terms such as 'memory', 'plans', 'intentions' and the like. The point is, then, not so much that we are applying scientific terms, the point is that we are *misusing everyday concepts*. The confusion that can arise from calling a chart a memory originates in the inappropriate use of the word 'memory'.

The description of the socio-cultural world in the terms of cognitive science is then a re-description of a world already known to us through our everyday descriptions. Re-describing is then quite analytically circular and idle.[2]

The Unsustainability of an Inner/Outer Dichotomy

Harold Garfinkel, the founder of ethnomethodology, once remarked that there is nothing inside our heads but brains, and this succinctly characterizes the arguments levelled against cognitive science by those who, drawing on Wittgenstein, have articulated a social understanding of mind. A major plank in their argument is that it is not possible to sustain the idea of an inner world within which the outer world is rendered in terms of representations. Within cognitive science there are, as Ryle (1973) put it, the parallel worlds of the inner and the outer, and the uneasy shuffle between the two. It is the claim of 'distributed cognition' that it is moving beyond the inner world of cognitive science. However, this movement does not entail the abandonment of the world of the inner for the public world; 'distributed cognition' holds tight to the idea of an inner world and its movement into the public world is to reinterpret the public world in the coinage of the inner world. Simply, the problematics associated with cognitive science and the objections that are levelled against it apply equally to 'distributed cognition' as they do to cognitive science in general.

The principal objection by many to the idea that it is in an inner world of sensations and representations that people fix the meaning and their understanding of objects in the world, an inner ostensive definition process as Williams (1999) refers to it, is drawn from Wittgenstein's articulation of an argument against a private language. This is not the place to reproduce his arguments but a short examination of them through his example of the beetle in the box can serve our purposes.

Wittgenstein writes:

> Suppose everyone had a box with something in it: we call it a 'beetle'. No one can look into anyone else's box. – Hence, it would be quite possible for everyone to have something different in his box. One might even imagine such a thing constantly changing. – But suppose the word 'beetle' had a use in these people's language – If so it would not be used as the name of a thing. The thing in the box has no place in the language-game at all; not even as a *something*; for the box might even be empty. – No, one can 'divide through' by the thing in the box, it cancels out, whatever it is.
>
> That is to say: if we construe the grammar of the expression of sensation on the model of 'object and designation' the object drops out of consideration as irrelevant. (Wittgenstein, 1958: para. 293)

This example addresses a major fallacy with the notion of an inner world. Wittgenstein is showing us that if the meaning of something is fixed internally then a question that could be asked, for example, is 'How can I know that my pain is the same pain as you experience?' If my pain is in my head, and you cannot feel it and I cannot feel yours how can we know we are feeling the same thing? Wittgenstein's point is that this question is a meaningless one to ask and that the experience is irrelevant to the question of how we *use* the word 'pain'.

The reason for this is because there is just no way that I can feel your pain, nor is there any way in which you can feel mine. In this respect, it is just not sensible to ask the question because it is not possible to answer it, and it is irrelevant for our understanding and descriptions of pain. This is the argument that is being articulated in Wittgenstein's beetle in the box example. Because I cannot see into your box or you into mine I just cannot know what is in your box or you what is in mine. They may or may not be the same thing but we have no way of telling. The only thing we could say of the beetle is that it is what we are calling 'the thing in the box', but whether or not that thing is the same we will never know, and even then, there may not be *anything* inside the box. As with the beetle in the box, my inner sensations of pain are just not available to you, so become irrelevant in our descriptions of pain.

In this respect, we can know that someone is in pain without recourse to their sensations; we know that someone is in pain because of the circumstances in which they are in, or because they are acting as if they are in pain. And – as suggested by Wittgenstein's example that there may be nothing in the box, and the thing in the box may change – whether or not they are actually experiencing pain, or experiencing the pain they felt when they were last in pain, is irrelevant for our description of them being in pain. They may or may not be in pain, but that need not affect our description of them being in pain. Equally, someone may be in a great deal of pain and they may conceal this from us by not doing the things that people do when they are in pain. If they did not do the things that people do when they are in pain then we would never know if they were in pain or not; their actual experiences of pain are irrelevant for our understanding that they are in pain, for if they do not let on then how are we to know? It is, then, public practice not the sensation that is at issue in recognizing that someone is in pain or describing someone as being in pain.

Someone may act as if they are in pain and may tell us that they are in pain but they may be lying to us. In various circumstances, we simply will not know, and, in that respect, their non-experience of pain is not the issue. We may decide that they are lying to us if we covertly see them acting differently when they think we are out of sight, or they may tell us later that they had fooled us because they were not really in pain when they said they were. Again, how are we to know? Someone may say they are in pain but not do the things people do when they are in pain; we may say they are being stoic, or we may say we do not believe them. Again, their experiences of pain are not the relevant issue; the relevant issue is the public display.[3]

Turning these arguments back to a consideration of distributed cognition, you may say you know how to navigate, and that you know the meaning that charts have. If meanings were private matters, how could I know if that meaning was the same meaning that I have? How could I know if the meaning you give to a black over yellow marker is the same meaning I give? The only way we would know is if we made these meanings public, for example, when we started to work together on a navigational problem. If it

turns out that you mean a south cardinal and I mean a north cardinal, then it is not just the case that we give different meanings to the marker. Rather, you are wrong because the symbol has been agreed to as a north cardinal; its meaning has been constituted through public agreement, not through private assignment. If you insisted on calling it a southern cardinal and in your navigating practices displayed that then you could not be said to be properly navigating, you would not be trusted or deemed to be competent by your peers.

In this sense, then, the chart does not act to mediate and coordinate between individual cognitive processing units that could assign different meanings to objects. That there is a common meaning to the symbols, and that subsequently two people can use the chart to coordinate their activities with one another is just not happenstance between those people, it is because of the public character of the symbols, the public character that the meaning 'north cardinal' has in our language. A chart exists in the public world, it is a communally agreed to object. In this respect, a chart is not a means for making the private meanings of individuals public, or a means for coordinating or negotiating two people's activities or meanings. A chart is by its very character a public object the meaning of which is agreed to by those who properly use it.

You may use the chart and privately assign the term southern cardinal to a black and yellow marker. However, you may in working with me translate it as a northern cardinal and in your activities with me you may engage in navigating practices that display the appropriate understanding of the cardinal. However, like the beetle, and like pain, what you might privately call it or experience it as is irrelevant for our interactions over the chart table if you are using the symbol correctly in your navigation practices.

We might say that the chart allows two individuals to work out in their heads a solution to a navigational problem. They take information from the chart which they are able to process individually, and they do the calculations in their heads. We may therefore conclude that the chart is a means for coordinating individual and separate processing. We do say that people can work things out in their heads, as we can say that people can say things silently. But, as Wittgenstein has pointed out, there is not anything different in kind going on when someone says something silently or says something aloud; they are in both instances using language according to its conventional character. And the same can be said of calculating something in your head; for it to be calculation it has to be done according to the rules and conventions of mathematics, there is nothing different in kind going on when I work out something in my head to when I work it out on paper; both activities are subject to mathematical convention. In providing information that has been publicly agreed to, the chart does not fuel an individual processing unit that allows processing to migrate from one individual to another. It just provides information that I and you can use in working something out in our heads or on paper. Our ability to do this, and our ability to compare what we have done and to check each other's workings out, are

entirely to do with the public and conventional character of the charts and of mathematics.

'Distributed cognition' holds firm to the idea of an inner realm. Its attempts to move beyond what it calls individual processing are attempts to align supposed inner realms, not to do away with the concept of an inner realm. Its stated contribution is the extension of cognitive science to areas that cognitive science has not otherwise touched. In this respect, 'distributed cognition' is about the cognitive science understanding of social doings, and consequently its mistake is the mistake of cognitive science. To criticize cognitive science understandings of an inner/outer dichotomy is to criticize 'distributed cognition' for the two are indivisible in this regard.

Context

One of the problems elaborated on by Suchman (1987) with respect to cognitive science was that it failed to take account of the way in which people use the setting within which they conduct their activities to condition them. All action is 'situated' in some setting and the participants orient to the details of that situation in shaping what they do. People's actions can be inspected by others to see what details of the setting are relevant for what they may be currently doing. Suchman's critique of plans as adequate causal explanatory devices for human action is that they do not take account of the details of the setting, and the way in which these can occasion the unfolding of activities in ways that cannot, necessarily, be determined in advance. Therefore, since by its very nature a plan has to be pre-formulated prior to engaging in activity, it cannot take into account the details of settings within which a course of action and interaction will be embedded. This is not to say that people do not make plans, but it is to say that a plan is not a causal antecedent for action because by their nature implementations can fail and resulting activity not go according to plan. Indeed, plans can also be revised in the light of circumstances. Cognitive science is asking of a concept, perfectly well used in social life, to do too much explanatory work. In as much as it is possible to see that people in these settings *do* take account of these details, then the idea of a plan as a metal prerequisite for action is inadequate as an explanatory device for action.

One of the reasons that 'distributed cognition' may have been embraced by those who may be themselves critical of traditional cognitive science is that it might seem to be attempting to take this type of argument into account. 'Distributed cognition' views the socio-cultural world as part of a cognitive system, and it places emphasis upon objects to be found within particular settings. Thus, for 'distributed cognition' so-called artefacts found in settings and the multiplicities of people in settings are constitutive elements of a cognitive system. We have been mainly concentrating upon Hutchins' work and, to illustrate our arguments, have been drawing from his account of the navigation and pilotage of large vessels. However, practitioners of 'distributed cognition' have examined a range of settings: Rogers (1992, 1993) has studied engineer practice; Halverson (1995) has

examined air traffic control; Hutchins and Klausen (1996) have studied cognition in the cockpit, to name some of the significant work in 'distributed cognition'. Common to all of these examinations is that they use fieldwork techniques that would be familiar to ethnographers; their articles also include data that would be familiar to the ethnographer such as transcriptions, quotations from the field and ethnographic descriptions of what people are doing. Consequently, in as much as 'distributed cognition' makes reference to the setting or the situation of action it may be viewed to be a way of rising above the criticisms of cognitive science, even an ally to those critics.

However, in this enthusiasm it is important not to lose sight of what 'distributed cognition' is attempting to do with the observations that are being made. Certainly, we are presented with observations about what the observer has seen people doing in actual situations, and we have in this respect publicly available data. However, 'distributed cognition' is not just concerned with this type of publicly observable data, it is also concerned with brain processing and what that entails. Thus, unlike many who use ethnographic techniques for studying human action in context where all that is on hand is the observed data, 'distributed cognition' is attempting to relate that data to computational processing in the brain. We should thus have two sets of data: data drawn from the field and the data about processing in the brain, and we could then relate the two within the supposed cognitive system. However, while we are certainly presented with observational data we are most certainly *not* presented with data on this inner computational processing. And how could we be? We cannot observe brain processing and we cannot measure it.

Therefore, it is just not possible to relate the data observed in a setting to the data that is said to be processed in the brain. Champions of 'distributed cognition' argue that people are engaged in doing computational processing as they undertake the doing of their actions and interactions, and that this processing and conduct goes on as part of a cognitive system. However, while they can say this they cannot actually show this to be the case. All they can do is to *infer* that processing is going on as conduct is engaged in. They can only *read out* of the data observed in the setting what is going on in the brain, they cannot actually show that processing is going on. In this respect, all they have – at the very best – is just a hypothesis that there is data being processed in the brain. But it is a hypothesis that can never be checked out for they just do not have access to this posited brain processing. In this respect, if we re-examine Hutchins' description of the cognitive system involved in the navigation of the USS *Palau*, we find that he produces deep descriptions of what he saw people in the situation doing but then he is only able to *assert* that brain processing was going on as well. He does not, and cannot describe that processing. Neither does he say, nor is he able to say, what relationship the asserted brain processing has to the actions and interactions he has witnessed. Consequently, we might say that nothing was being added to our understanding in making reference to brain processing.

It could be supposed that a response to this argument in the cognitive science tradition could be that, because it is possible to measure the electrical activity going on in the brain, we do indeed have evidence of processing going on. Because we can wire a person up, present them with different circumstances and see that the electrical activity in their brains changes it is possible, in fact, to gather data on brain processing. Therefore, we could imagine that 'distributed cognition' practitioners might wire up actors in social settings and record the electrical patterns that go on as they do their work. That, for example, Hutchins could prove his statement that in manoeuvring the USS *Palau* to safety there was individual processing going on because he had wired the people up and could now point to the print-outs as data, thus having data from the setting and data on brain processing.

Not unsurprisingly, 'distributed cognition' does not do this. This is not, however, because of the practicalities involved, but because it would not actually answer the problem. There is a raft of problems in cognitive science with respect to the transformation of neurological data into claims about cognitive processing. Certainly there are fluctuations in levels of neural activity that can be seen when people are faced with certain conditions, but what is the correlation between the neural activity and the posited cognitive process? It is not even possible to lay this correlation out in broad terms, let alone do the detailed mapping between the neurological activity and the cognitive phenomena, and there would indeed have to be a one-to-one mapping if the point was to be made. Thus, it is not possible to correlate a neural activity with a mental concept such as a motive so that we could say that when this pattern of electrical activity was taking place someone had a motive. Even if we granted the plausibility of that idea, the idea it subsumes that mental events are discreet is, however, hard to sustain. For instance, we might have many motives for acting; how do we determine which motive it is that is mapped onto the observed neural read-out?

Within cognitive science, there are only *programmatic claims* that there is a relationship between so-called cognitive events and electrical activity in the brain. Cognitive science is beset with a range of unresolved problems, this being merely one. What 'distributed cognition' is doing, then, is to bring these problems into considerations of social doings; but the waters of the human sciences are cloudy enough and beset by the raging of their own tempestuous problems. What, if anything, is gained in our understanding of human doings by adding the conceptual confusions of cognitive science to the conceptual confusions of the human sciences?

However, let us say that in arguing this we are short-sighted and have lost our wits, for we have got the conceptual problem wrong; indeed, that the march of science has somehow proved us wrong and that it does become possible to do a one-to-one mapping. But even if it is possible to conceive of this, then are we actually any the wiser about how the people on the USS *Palau* did their work? Do we know anything more about how that work is put together? Does it actually tell us anything about the navigating of large

vessels? Correlating data about brain processing with data about the social world does not tell us anything about how that social world is made up and organized. It is merely to say that, when engaged in an activity such as greeting someone, there is a pattern of electrical activity going on in the brain. It does not tell us anything about the organization of greetings in our culture.

The problem with 'distributed cognition' is that, like many other explanatory frameworks in the human sciences, it looks outside of the setting in order to explain how the setting is being organized. It thus looks to the brain rather than to the setting itself. In doing this, it is not alone. Garfinkel has chided many social science perspectives for looking outside of the scene by looking to value systems for example. Garfinkel's point is that settings are self-organizing, that the social action is ongoingly ordered as it is conducted. The virtue of this is that the data for the descriptions of how that ordering is being done is what was observably done in the setting. At least it is possible to see what is being talked about.

Conclusion

Our argument has been twofold. First, it would be a mistake to view 'distributed cognition' as a corrective to traditional cognitive science. Its approaches to the socio-cultural world are not for the purposes of challenging major precepts in cognitive science about mind. Rather, 'distributed cognition' is attempting to make the socio-cultural world a further topic for cognitive science.

Second, 'distributed cognition' does not bring anything to the table. It is perfectly possible to take out statements about cognitive systems, and brain processing from 'distributed cognition' accounts and still be left with the descriptions of the settings examined. Thus, we could take out all references to cognitive systems and processing in Hutchins' *Cognition in the Wild* and we would still be left with a description of what Hutchins had seen: people working together to navigate a large vessel. That description does not depend upon the assertions Hutchins makes about brain processing or cognitive systems. Though whether or not Hutchins' actual descriptions amount to an interesting consideration of the work of navigating is another matter.

Notes

Thanks to Wes Sharrock, Jeff Coulter and Mike Lynch for their reflections and comments on this article, and also to four anonymous reviewers, one of whose reflections on the description of the sensation of pain was the most useful and helpful set of comments I have ever had from a reviewer.

1. One reviewer of this article argued that I had missed the point of 'distributed cognition' because, the reviewer asserts, distributed cognition is an attempt to acknowledge the importance of a cultural world for the description of human action, and is therefore a corrective of 'cognitive theory'. This reviewer, however, misses my point. His point is just what Hutchins and those who have taken up his argument

would say of 'distributed cognition'. The point that I lay out here, and attempt to explicate rather than assert, is that this is in fact a misplaced understanding of 'distributed cognition' and that, despite the conventional assertion (there is, sorry to say, no explication associated with this conventional reading) that Hutchins is articulating a cultural corrective, he ends up merely articulating 'cognitive theory'. The whole point of this article is to show this, not just accept what Hutchins and others claim for his work. The reviewer, who just reiterates Hutchins as if he were Hutchins himself, asked if I had actually read the work in its entirety; I likewise wonder if the reviewer had actually read this current article.

2. This is not just a problem with 'distributed cognition'; it is a problem that could be levelled against a whole range of human science undertakings. Making this charge points to an alternative task for human science to the one of re-description. As opposed to reinterpreting the social world in the terms of particular social theories the alternative is, as provided for in Garfinkel's ethnomethodology (Garfinkel (1967), to describe the knowledge that people have of their own doings and how they use that knowledge in the ordering of their organized social affairs.

I offended another reviewer when I argued that it was not useful to re-describe an astrolabe as an analogue computer. He/she argued that this would mean that we could not say that people died of the bubonic plague because at the time this understanding of their death was not known. The point I am trying to make is simply this: what is gained in re-describing the astrolabe as an analogue computer? Do we understand anything more about it or how it was used? No, we do not; in fact we can start to confuse ourselves because we can say of a computer that 'I used the computer to calculate the distance of one star from another' and I can say that 'the computer calculated the distance of one star from another', and I can say that 'I used the astrolabe to calculate the distance of one star from another' but I *cannot* say that the 'the astrolabe calculated the distance of one star from another'. By conflating the astrolabe and the computer it is possible to start to misunderstand one or the other. So in this respect I cannot say that people were using an instrument that they did not know – but we now know – was a computer simply because astrolabes could not, then or now, compute, unlike a computer, which can compute. (Of course this should not be misinterpreted as conflating computer processing with human reasoning.) But we can say that people died in London in 1665 of, what we now know – though at the time people did not know – was the bubonic plague. The issue is that in the human sciences we start to confuse ourselves by compounding misuses of our language, and 'cognitive theory' exists in that confusion of misuse which re-naming artefacts of the past in terms of artefacts of the present compounds.

3. Wittgenstein, and indeed Ryle, have been misunderstood as advocating a behaviourist theory of mind and I do not want to give fuel to that misinterpretation in this discussion of public display. The point I want to hold on to is that the inner sensation of pain is not at issue in the description of pain.

References

Coulter, J. (1979) *The Social Construction of Mind: Studies in Ethnomethodology and Linguistic Philosophy*. London: Macmillan.

Garfinkel, H. (1967) *Studies in Ethnomethodology*. Englewood Cliffs, NJ: Prentice Hall.

Halverson, C.A. (1995) 'Inside the Workplace: New Technology and Air Traffic Control', PhD thesis, Dept of Cognitive Science, University of California San Diego.

Hutchins, E. (1996) *Cognition in the Wild*. Cambridge, MA: MIT Press.

Hutchins, E. and T. Klausen (1996) 'Distributed Cognition in an Airline Cockpit', in D. Middleton and Y. Engestrom (eds) *Communication and Cognition at Work*. Cambridge: Cambridge University Press.

Lave, J. (1988) *Cognition in Practice*. Cambridge: Cambridge University Press.

Louch, A.R. (1966) *Explanation and Human Action*. Oxford: Basil Blackwell.

Marr, D. (1982) *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. Cranbury, NJ: W.H. Freeman.

Rogers, Y. (1992) 'Ghosts in the Network: Distributed Trouble Shooting in a Shared Working Environment', pp. 295–315 in *Proceedings of the Conference on Computer-Supported Cooperative Work*. New York: ACM.

Rogers, Y. (1993) Coordinating Computer-Mediated Work, *Conference on Computer-Supported Cooperative Work* CSCW 1: 295–315.

Ryle, G. (1973) *The Concept of Mind*. Harmondsworth: Penguin.

Shank, R.C. and R.F. Abelson (1977) *Scripts, Plans, Goals and Understanding*. Hillsdale, NJ: Lawrence Erlbaum.

Simon, H.A. (1981) *The Sciences of Artificial Intelligence*. Cambridge, MA: MIT Press.

Suchman, L. (1987) *Plans and Situated Actions: The Problem of Human–Machine Communication*. Cambridge: Cambridge University Press.

Wilensky, R. (1981) 'PAM', in R.C. Shank and C.K. Riesbeck (eds) *Inside Computer Understanding*. Hillsdale, NJ: Lawrence Erlbaum.

Williams, M. (1999) *Wittgenstein, Mind and Meaning: Towards a Social Conception of Mind*. London: Routledge.

Wittgenstein, L. (1958) *Philosophical Investigations*, 2nd edn. Oxford: Basil Blackwell.

Graham Button is Executive Dean of the Faculty of Arts, Computing, Engineering and Sciences at Sheffield Hallam University. Prior to this he was Laboratory Director at the European Research Centre in Grenoble, France.