

Friends at last? Distributed cognition and the cognitive/social divide

Adam Toon

University of Exeter
a.toon@exeter.ac.uk

Abstract

Distributed cognition (d-cog) claims that many cognitive processes are “distributed” across groups and the surrounding material and cultural environment. Recently, Nancy Nersessian, Ronald Giere and others have suggested that a d-cog approach might allow us to bring together cognitive and social theories of science. I explore this idea by focusing on the specific interpretation of d-cog found in Edwin Hutchins’ canonical text *Cognition in the Wild*. First, I examine the scope of a d-cog approach to science, showing that there are important disputes between cognitive and social theorists on which d-cog remains silent. Second, I suggest that, where social explanations can be recast in d-cog terms, this reformulation will not be acceptable to all social theorists. Finally, I ask how we should make sense of the claim that, on a d-cog analysis, social factors *are* cognitive factors.

Keywords

Distributed cognition – Sociology of Scientific Knowledge – Edwin Hutchins –
Ronald Giere – Nancy Nersessian

1 Introduction

There is sometimes thought to be an opposition between cognitive and social theories of science. Perhaps the clearest instance of this opposition is Bruno Latour and Steve Woolgar's infamous "ten-year moratorium on cognitive explanations of science" (1986, p. 280). On one side of the divide, social accounts emphasize scientists' social and political interests or institutional structures. On the other, cognitive accounts refer mainly to scientists' cognitive processes. In philosophical discussions, the term 'cognitive' is often associated with terms such as 'rational' or 'truth-conducive'. In the present context, however, 'cognitive' is used in the sense found in cognitive science, to refer to psychological processes such as perception, reasoning, memory and so on (Giere & Moffatt, 2003, p. 302). Processes that are cognitive in this sense can, of course, fail to be rational or truth-conducive. The divide at issue here is therefore distinct from that concerning the relationship between social factors and the rationality of science (e.g. Longino, 1990; Solomon 2001).

Recently, Nancy Nersessian (2005) has argued that the perceived divide between cognitive and social theories is based on a mistaken "Cartesian" view of cognition, summed up in the tenets of GOFAI ("Good Old Fashioned Artificial Intelligence") (Haugeland, 1985). According to GOFAI, cognition involves computational processes on symbolic representations internal to the individual mind. Within cognitive science, GOFAI is increasingly challenged by a range of different approaches. Particularly important for Nersessian is *distributed cognition* (*d-cog*). In contrast to GOFAI, *d-cog* claims that many cognitive processes are "distributed" across social groups and the wider material and cultural environment. Together with a team of researchers, Nersessian has offered an analysis of the laboratory as an "evolving distributed cognitive system" (e.g., Nersessian, Kurz-Milcke, Newstetter, & Davies, 2003a; Nersessian, Newstetter, Kurz-Milcke, & Davies, 2003b). She

suggests that this approach allows us to “integrate” cognitive and social accounts of science, as well as yielding other important insights. For example, d-cog might allow us to see that the nature of cognitive processes in science has changed over time (2005, p. 52).

Ronald Giere has expressed similar hopes for d-cog (2002a; 2002b; 2002c; 2004; 2006; 2007; 2012; Giere & Moffatt, 2003). As well as applying the approach to a number of aspects of scientific practice, such as the use of diagrams and models, Giere has suggested that we reinterpret well-known studies by sociologists such as Bruno Latour and Karin Knorr-Cetina in d-cog terms (2002b; Giere & Moffatt, 2003). Like Nersessian, Giere believes that d-cog “bridges the often perceived gap between cognitive and social histories of science” (2002a, p. 285) and that the approach offers important historical insights. For example, it allows us to see that “what powered the scientific revolution was an explosion of new forms of distributed cognitive systems” (2002a, p. 298). At times, Giere also seems to go further than Nersessian, arguing that, on a d-cog approach, the cognitive and social “overlap” (2002a, p. 296) or “merge” (Giere & Moffatt 2003, p. 304). For example, referring to Knorr-Cetina’s analysis of experiments in high energy physics, Giere argues that d-cog provides “a complementary, *cognitive*, account of these experiments” (2002b, p. 639). In this account, “we can now say that aspects of the situation that seemed only social are also cognitive” (2006, p. 114).

In this paper I try to assess d-cog’s potential to bring together cognitive and social approaches to science. Exactly what is meant by ‘distributed cognition’ is sometimes unclear. In Section 2, I distinguish between a general and more specific interpretation of the approach. Without further development, the general interpretation of d-cog is difficult to assess. I therefore focus on the more specific interpretation, found in Edwin Hutchins’ *Cognition in the Wild*. Hutchins’ approach is introduced in Section 3. In Section 4, I consider the scope of d-cog: which aspects of science might be understood using this approach and which, if any,

will remain beyond its reach? In doing so, I argue that there are important disputes between cognitive and social theorists on which d-cog remains silent. In Section 5, I suggest that, where social explanations can be recast in d-cog terms, this reformulation will not be acceptable to all social theorists. Finally, in Section 6, I ask how we should make sense of Giere's claim that, on a d-cog analysis, social factors *are* cognitive factors. Can this aspect of d-cog help us to reconcile cognitive and social approaches?

2 What is Distributed Cognition?

The canonical text for distributed cognition is Edwin Hutchins' *Cognition in the Wild* (1995). This is an ethnographic study of navigation on a U.S. Navy ship, which Hutchins calls the *Palau*. Hutchins offers a detailed analysis of how the *Palau*'s crew accomplishes certain tasks, such as determining the ship's position and planning its course. Typically, he finds, these tasks are not performed by any one individual, but by a group of crew members working together. Each of the crew members also makes use of a variety of different tools. The processes that accomplish the tasks are therefore *distributed* across the members of the team and the tools they use. Moreover, the social structure plays a crucial role in the crew's activity: the way in which tasks are performed depends upon the command structures that govern the crew's interaction. On the ship there is a "social distribution of cognitive labor" (Hutchins, 1995, p. 228).

For example, consider the 'fix cycle'. This is the procedure in which the navigation team determines the ship's position by taking visual bearings to landmarks on either side of the ship. The fix cycle involves a number of different crew members. Inside the pilothouse are the navigation plotter and the bearings recorder. Together, they decide on suitable landmarks for obtaining a position fix. They then pass on the names of these landmarks to the

pelorus operators, who stand on either side of the ship. The pelorus operators must identify these landmarks on the horizon and take their bearings. To do so, they use a device called an alidade, which has a hairline sight aligned with a gyrocompass scale. The pelorus operators relay the bearings to the pilothouse, where the recorder notes them in the log and the plotter plots them on the chart to determine the ship's position and project its future course. The plotter uses various tools to carry out this task, such as a hoey, which is a special protractor with a long arm that can be set to the recorded bearing.

Hutchins' analysis thus encompasses both the cognitive acts of individuals (such as reading the bearing of a landmark in the alidade sight) and the social structure guiding their interactions (such as the ship's command hierarchy). In this respect, his analysis is perhaps not so remarkable. After all, as both Giere and Nersessian acknowledge, not all sociological accounts are as hostile to cognitive or psychological explanations as Latour and Woolgar (e.g. Bloor, 1976). Conversely, cognitive theorists sometimes recognize the importance of social context (e.g. Dunbar, 1995). What exactly is distinctive about a d-cog approach? In particular, how does d-cog offer a *reinterpretation* of sociological accounts, such as Knorr-Cetina's?

The distinctive feature of d-cog, it seems, is that it treats social groups, along with tools and parts of the material and cultural environment, as *cognitive systems*. Hutchins' key claim is that the navigation team may be analysed as a "cognitive and computational system" (1995, p. xiv). This passage already points to two different interpretations of the approach, however. According to the first interpretation, to analyze an activity as d-cog is to understand it as a *cognitive* process. According to the second, it is to understand the activity as a *computational* process. While the first approach is not committed to any particular view of the nature of cognition, the second claims that cognition takes a specific form, namely computation.

As we will see in Section 3, it is the computational version of d-cog that underpins

much of Hutchins' analysis in *Cognition in the Wild*. Both Nersessian and Giere sometimes distance themselves from the claim that cognition is computational (Giere, 2006, pp. 107-8; Osbeck & Nersessian, 2006; see also Brown, 2011, p. 28). Nevertheless, it is worthwhile assessing the computational approach, for a number of reasons. First, Hutchins' book remains the *locus classicus* for work on distributed cognition and forms the central inspiration for both Nersessian and Giere. Second, as we shall see in Section 3, the computational version of d-cog has much to recommend it. Third, the computational view has the virtue of being more specific, and therefore easier to assess, than the first interpretation of d-cog. If d-cog claims only that social groups are cognitive systems, without saying more about what cognitive systems are, then it is difficult to see how to evaluate the approach. In what follows, I shall therefore focus on the computational interpretation of d-cog, although some of what I say (especially in Section 6) will apply to the more general interpretation as well.

3 Ships as Computers

At the outset, Hutchins tells us that his study is an attempt “to apply the principal metaphor of cognitive science – cognition as computation – to the operation of [the navigation team]” (1995, p. 49). Although he is skeptical of GOFAI's claim that what goes on inside the head is computational, Hutchins believes that this analysis may be applied to the navigation team as a whole: “The system formed by the navigation team can be thought of as a computational machine” (*ibid.*, p. 228). Moreover, “the computation observed in the activity of [the navigation team] can be described in the way cognition has been traditionally described – that is, as computation realized through the creation, transformation, and propagation of representational states” (*ibid.*, p. 49).

It is important to note that Hutchins does not intend his analysis to be merely

metaphorical. In his view, applying the language of computation to the navigation team is “not a metaphorical extension at all” (1995, p. 364). Hutchins develops his analysis by drawing on David Marr’s (1982) distinction between three different levels on which a cognitive system may be understood. At the *computational* level we have an abstract description of the computation that a system carries out. Thus, fixing a position may be understood as combining two one-dimensional constraints to give a unique position in two-dimensional space. The *representational* level concerns “the choice of representation for the input and output and the algorithm to be used to transform one into the other” (Marr, 1982, pp. 24-25). For example, the crew of the *Palau* uses the standard Western coordinate system and one-dimensional constraints are given by lines of position. Finally, the *implementation* level specifies how the representational system and algorithm are physically implemented. On the *Palau*, for instance, the pelorus operators determine the bearings that give the line of position, while the plotter determines their intersection on the chart.

It will be helpful to see some examples of how Hutchins applies this analysis to social interactions. Within computer science, a *daemon* is an agent that monitors for certain trigger conditions, and takes a specified action when those conditions are met. During the *Palau*’s voyage, Chief Richards instructs another member of the crew, Smith, to watch the fathometer and report when the depth of water falls below 20 fathoms. Hutchins claims that here we have “the social construction of an information-processing mechanism” (1995, p. 192). By giving the order to Smith, Chief Richards reconfigures the system formed by the navigation team to create a daemon which will respond to certain conditions and output a symbolic signal when those conditions are met.

Hutchins offers a similar analysis of “phone talkers”. These are members of the crew posted at each end of the ship’s telephone lines. The job of a phone talker is to receive messages and relay them to the relevant crew members when there is a suitable break in their

activity. In Hutchins' analysis, phone talkers are *information buffers*, which "[permit] communication to take place when the sender and the receiver are not overloaded" (*ibid.*, p. 195). The bearing recorder and log are also information buffers, since they enable the pelorus operators and the plotter to work asynchronously: the plotter need not plot each bearing as it is reported, but can instead refer back to the bearing recorder or log. The log is also a *memory*, as well as a *filter*, which "passes the bearings without passing the temporal characteristics of their production" (*ibid.*).

In Hutchins' original formulation, then, d-cog is a claim about what happens at the implementation level of a computation. Rather than taking place within the mind of any individual, the algorithm determining the ship's position is implemented by a distributed system consisting of the entire navigation crew. This computational version of d-cog has a number of attractions. First, and perhaps most important, it offers a clear rationale for the claim that d-cog integrates the cognitive and the social: according to the computational view, social organization becomes part of the way that the relevant computation is implemented. As Hutchins puts it, we may "treat the social organization as a computational architecture" (*ibid.*, p. 185). Second, the computational approach provides a powerful tool for understanding different social arrangements and assessing their epistemic merits. In this respect, it shares the attraction of Paul Thagard's earlier proposal that we understand science using distributed artificial intelligence (D.A.I.) (1992). Third, d-cog also avoids one of the main objections against Thagard's proposal because, unlike D.A.I., d-cog is not committed to claiming that scientists' *internal* cognitive processes are computational (*ibid.*, p. 58).

4 Uncharted Waters

Neither Nersessian nor Giere claim that d-cog offers a complete theory of science, nor

that it allows us to reconcile all disputes between cognitive and social theorists. In light of this, it is important to ask how far d-cog might be able to take us. Which aspects of science might be analyzed using a d-cog approach and which, if any, will remain out of its reach? These questions are rarely explicitly addressed in the literature. And, at first glance, it might seem difficult to answer them at this stage, since even the computational version of d-cog remains rather general. In itself, the approach claims only that laboratories carry out some form of computation; it does not say what that computation is or exactly how it is implemented. In fact, however, I want to suggest that any computational d-cog analysis will remain silent on some parts of scientific practice.

To see this, recall that d-cog is a claim about the implementation level of a computation: Hutchins shows how the computation determining the *Palau's* position is implemented by the entire navigation team and their social interactions. It is in this sense that d-cog seems to offer a means of integrating cognitive and social factors. Notice, however, that even if we accept such an analysis, a number of issues will remain to be addressed. First, a d-cog analysis will not tell us why a particular computation is being performed or why a specific representational system is being used to carry it out. Second, the analysis will not tell us how a representational system gains its representational status. When applied to science, I suggest, both of these questions receive competing answers from social and cognitive accounts. As a result, they will remain areas where a d-cog analysis of the laboratory cannot help to reconcile debate between the two.

First, consider the choice of representational system. Alongside his study of the *Palau*, Hutchins offers an analysis of traditional navigation techniques in Micronesia. Hutchins argues that the Western and Micronesian systems are, in fact, the same at the computational level: both combine one-dimensional constraints to give a unique position in two-dimensional space. But the two differ radically at the representational level. For example, rather than

taking the boat to be moving across a fixed two-dimensional space, Micronesians think of the canoe as stationary while the water moves past it. Progress during a journey is represented not by a unit of length but by the changing bearing of a reference island. One-dimensional constraints are sometimes provided by sightings of birds, which indicate the canoe's distance from land.

There are many questions that we might ask about these representational systems. Why do Micronesian navigators think of the canoe, rather than the water, as fixed? Why do we employ standard units of distance and time and embody these in charts and other devices, while the Micronesians do not? Are there still other ways in which humans might represent the world to find their way around it? Each of these questions concern the choice of representational scheme used to carry out a computation, not the way that the computation is implemented. As a result, they cannot be answered through a d-cog analysis of a particular navigation practice. Hutchins himself offers many fascinating insights into the origins of the representational scheme used in Western navigation (1995, Ch.2). He does so, however, not by relying on his analysis of the *Palau*, but instead by looking at the history of charts and other navigational tools in the West. In Hutchins' view, the features of the Western representational system are historically contingent, rather than "natural and inevitable or simply the consequences of the interaction of human nature with the demands of a given task" (*ibid.*, p. 114).

Many of the questions that we ask about the sciences also concern representational systems. For example, consider classification schemes. These have been the focus of considerable attention in the sociology of scientific knowledge. In fact, according to David Bloor, "one of the central propositions in the sociology of knowledge" is that "the classification of things reproduces the classification of men" (1982, p. 267). By contrast, some cognitive scientists claim that there are important similarities in the way that different

cultures categorize the world that are due to universal cognitive constraints (e.g. Berlin, 1992). I do not wish to enter into this debate here. All I want to point out is that, like questions concerning the differences between Western and Micronesian navigation, this debate about classification schemes concerns the choice of representational system, rather than its implementation. As a result it seems that, if there is an opposition here between sociological and cognitive accounts, a d-cog analysis of the laboratory cannot help to overcome this opposition.

The second issue that d-cog does not address is that of how representational systems gain their representational status (that is, why they have *meaning*). The reason for this, of course, has to do with a general feature of the computational view. As Tim Crane puts the point, “[a] computational process is, by definition, a rule-governed or systematic relation among representations. To say that some process or state is computational does not explain its representational nature, it presupposes it” (2003, p. 169). In line with Hutchins’ analysis, d-cog would treat the laboratory as a computational system creating, manipulating and destroying representations. This leaves open the question of how those representations come to represent the world. Once again, this is a question that receives very different answers from social and cognitive accounts of science. For example, members of the Edinburgh Strong Programme in the sociology of science propose a theory of meaning known as *finitism*, drawn from the later Wittgenstein (Barnes, Bloor & Henry, 1996). According to finitism, meaning is a fundamentally social phenomenon. By contrast, of course, many cognitive scientists and philosophers reject this view, instead seeking to explain meaning in terms of non-social factors, such as causal relations. It might be thought that d-cog may legitimately defer these disputes about the nature of meaning. And yet finitism is arguably *the* central element of the Strong Programme, which gives rise to many of its key claims, such as epistemic relativism (e.g. Kusch, 2002). Here again, then, we find an important dispute between which d-cog

cannot help to reconcile.

5 A Storm Brewing?

Section 4 pointed to aspects of science that will be omitted from a d-cog analysis. Let us now consider those parts of science that d-cog does seek to analyze, and ask whether this analysis can help to reconcile cognitive and social accounts. Unfortunately, I believe that the prospects for reconciliation here may be limited in an important respect, since the analysis offered by d-cog will not be acceptable to all social theorists. In fact, rather than bridging the gap between cognitive and social theories, d-cog threatens to return us to an old debate between the two, sparked by Peter Slezak's paper, "Scientific Discovery by Computer as Empirical Refutation of the Strong Programme" (1989).

Slezak focused on computer discovery programs such as BACON (Langley, Bradshaw, & Simon, 1983). Given the relevant data, BACON is able to "rediscover" various empirical laws, such as the ideal gas law. According to Slezak, programs like BACON constituted a decisive refutation of the Strong Programme in the sociology of scientific knowledge (SSK), which claims that "social factors are an irreducible component of scientific discovery" (1989, p. 564).¹ In fact, for Slezak, the "very possibility of computer programs making scientific discoveries poses a fundamental challenge to SSK's radical claims" (1989, p. 564).

Slezak's argument proved controversial. Even supporters of cognitive accounts took issue with his claim that BACON was capable of making scientific discoveries. For example, Thagard noted that such programs "are highly simplified and do not constitute full simulations of how discoveries were actually made" (1989, p. 654). Interestingly, however, most commentators agreed that, if computer programs could make discoveries, SSK would be

refuted. For example, in response to Slezak's paper, Giere wrote:

A minimal claim of a sociology of scientific knowledge (SSK) would be that the course of science is necessarily influenced by the human interests of scientists, which are in turn partially shaped by their social relationships. If computers have no human interests and no social relationships, and yet can make scientific discoveries, this minimal claim would be refuted. (1989, p. 639)

Similarly, Harry Collins agreed that "[i]f [Slezak's premise] is taken to mean that BACON reproduces social collectivities then, if true, it would be fatal for 'the strong programme' and its variants." Why? "[B]ecause BACON is not a social collectivity" (1989, p. 614).

The trouble, I think, is that d-cog poses a similar challenge to social theories. Recall that, according to Hutchins, the navigation team is a "computational machine" (1995, p. 185). As Hutchins himself reminds us, computation is independent of the physical medium in which it is implemented (*ibid.*, p. 51). As a result, if the laboratory is a distributed cognitive system, then the computation it performs could in principle be carried out on some other system entirely (Magnus, 2007, p. 299). And this system need not be a social one. In fact, as far as d-cog is concerned, the social is only the "hardware" on which the computation happens to be run. Consider phone talkers, for example. According to Hutchins, they are information buffers, controlling the flow of information between different parts of the system. On the *Palau* this is achieved by assigning crew members a particular role within the social organisation of the ship. But there is nothing essentially social about an information buffer. If Hutchins' analysis is correct, the phone talker might as easily be replaced by a silicon circuit, so long as it implements the same computation.

Of course, d-cog would not support all of the conclusions that Slezak wanted to establish using BACON. For example, Slezak argued that, since BACON implements general

rules of problem solving, its success points to the existence of “context-free or super-cultural norms of rationality” (1989, p. 572). A d-cog approach does nothing to suggest the existence of such norms: the computations carried out in different laboratories might vary widely with different social contexts. Nevertheless, like BACON, d-cog would appear to be at odds with SSK’s claim that science is essentially social. Similar claims are also found in the work of authors outside SSK. For example, Helen Longino writes that science is “necessarily social” (1990, p. 12).

There are a number of ways in which we might try to resolve the apparent conflict here. Recalling our discussion in Section 4, social theorists might argue that social factors are necessary for those aspects of science omitted from d-cog, such as the semantics of representational schemes. Alternatively, we might construe claims regarding the essentially social nature of science in some weaker sense, which is not threatened by the possibility of non-social implementation raised by d-cog. For example, perhaps some social theorists might rest content with the claim that, while a non-social science is possible, actual scientific practice is social, and likely to remain so. Others, however, are keen to argue for a stronger position (e.g. Kusch, 2002). Taking a slightly different line, Thagard argues that D.A.I. is compatible with Longino’s view, since computer networks can capture the critical interactions that she takes to be required for scientific objectivity (1992, p. 61). Once again, these issues cannot be explored fully here. Much will depend upon the details of particular social theories. Nevertheless, I think it is important to note that there *is* an apparent conflict here and that this conflict must be resolved if d-cog is to bridge the gap between cognitive and social accounts of science.

6 Merging the Cognitive and Social

One way in which d-cog promises to bridge the gap between social and cognitive theorists is by revealing that social factors *are* cognitive factors. For example, recall Giere's claim that, on a d-cog analysis, "aspects of the situation that seemed only social are also cognitive" (2006, p. 114). In a similar vein, he writes that

Thinking of science in terms of systems of distributed cognition enlarges the domain of the cognitive in our understanding of science. It is typically assumed that there is a sharp divide between the cognitive and the social. From the perspective of distributed cognition, what many regard as purely social determinants of scientific belief can be seen as part of a cognitive system, and thus within the purview of a cognitive understanding of science. There is no longer a sharp divide. The cognitive and the social overlap. (Giere, 2002a, p. 296)

With d-cog, we are told, the cognitive and social "merge" (Giere & Moffatt 2003, p. 304).

How exactly should we understand these claims? And how might this aspect of d-cog help us to reconcile cognitive and social theories of science? To answer these questions, we first need to understand what is meant by 'cognition' in this context. In everyday contexts, we normally use the term 'cognitive' to refer to internal, psychological processes such as perception, memory, reasoning, and so on. Moreover, it seems that it was broadly this sense of cognition that was at stake in our original dispute between cognitive and social theories of science: social theorists argue for the importance of social factors, such as social status or institutional structures, while cognitive theorists stress internal, psychological processes such as reasoning, perception, memory and so on.

Given our usual notion of cognition, we might wonder how external processes such as

social interactions could possibly count as cognitive. However, Giere is keen to stress that the concept of cognition invoked by d-cog is not our everyday one:

[w]e are developing a science of cognition. In so doing we are free to make *cognition* a technical scientific concept different from everyday notions. (2002b, p. 642; see also Giere 2006, p. 112)

What is this technical notion? What does it mean to say, for example, that the *Palau* or the laboratory is a cognitive system in this technical sense? Giere suggests that “[t]he reason for calling these systems cognitive systems rather than, say, transport systems or agricultural systems, is that they produce a distinctly cognitive product, knowledge” (2002b, p. 642).

Elsewhere, he expands on this idea:

A distributed cognitive system is a system that produces cognitive outputs, just as an agricultural system yields agricultural products. The operation of a cognitive system is a cognitive process. [...] But what makes the output of the system a *cognitive* output? Here I think the only basis we have for judging an output to be cognitive is that it is the kind of output we recognize as the result of human cognition, such as a belief, knowledge, or a representation of something. (Giere, 2006, pp. 112-3)

Thus, a cognitive system is one which produces a cognitive output, such as a belief, knowledge and so on. In this new, technical sense, even distant galaxies used as gravitational lenses for the Hubble Space Telescope may count as part of a cognitive system (Giere, 2012, p. 201).

If we understand ‘cognition’ in this way, then to claim that the laboratory is a cognitive system is simply to claim that its output is a cognitive state, such as a belief, knowledge, and so on. While this claim might be uncontroversial, it is difficult to see how it will help to reconcile debates between cognitive and social theories of science. After all, both sides in this dispute would agree that science produces a cognitive *output*, such as beliefs

(and perhaps also knowledge). What they disagree about, of course, is the nature of the processes that produce that output. Showing that the laboratory is cognitive in the minimal sense that it produces a cognitive output would not appear to resolve any debates over the nature of the processes that lead to that output.

So it seems that, if d-cog's merger of social and cognitive factors is to bridge the gap between cognitive and social theorists, it must employ a more substantial concept of cognition. One obvious strategy would be to stress the computational aspect of d-cog. On this interpretation, to claim that social interactions are cognitive would be to claim that they implement a computation (and perhaps also that they result in a cognitive output, such as a belief). As we saw in Section 5, this claim would be likely to receive more resistance from social theorists. But let us suppose that this resistance could be overcome. Would a computational analysis of social factors then help to reconcile cognitive and social theorists?

Of course, in a sense, it seems likely that disputes will remain. Even if social theorists were to accept that social interactions are computational, they might still disagree over their relative importance compared to internal psychological processes, such as memory or reasoning. Nevertheless, d-cog's proponents might argue that there is now an overarching concept of cognition – namely, cognition as computation – that encompasses both social processes and internal, psychological ones. The difficulty with taking this line, however, is that Hutchins himself does not think that internal cognitive processes are computational. And, as we have seen, both Giere and Nersessian are also wary of the computational view.

Put simply, then, the challenge facing d-cog would seem to be as follows. In the original dispute between social and cognitive theorists, 'cognitive' took something close to its usual meaning, referring to internal, psychological processes such as memory, perception, reasoning, and so on. According to Giere, d-cog invokes a new, and much broader, notion of cognition that applies to systems involving many individuals and their social interactions, as

well as models, diagrams, instruments and sometimes even distant galaxies. At the same time, Giere also argues that d-cog helps to resolve the dispute between social and cognitive theorists since it reveals that social processes *are* cognitive processes. The trouble is that, the further d-cog's technical notion of cognition moves away from the everyday notion at stake in the original dispute, the harder it is to see how d-cog's merger of cognitive and social factors helps to resolve that dispute. Unless d-cog's notion of cognition displays at least some significant similarities with our usual notion, it is difficult to see how it enables us to close the gap between cognitive and social theories of science. Instead, what d-cog would seem to offer is simply a new way of analyzing the social aspects of science (one which understands social interactions as computational).

At this point, one option for d-cog's proponents is to endorse the *extended mind thesis* (Clark & Chalmers, 1998). Like Giere, proponents of the extended mind thesis talk of cognitive systems that exist outside our heads, and even outside our bodies. Unlike Giere, however, they argue that such claims do not involve a merely technical sense of 'cognition'. Consider the well-known case of Otto and Inga. When Inga hears of an exhibition at the Museum of Modern Art (MoMA) she recalls that the museum is on 53rd Street and heads off. Otto is an Alzheimer's patient who carries a notebook with him wherever he goes to record useful information. When Otto hears of the exhibition, he looks up the information in his notebook and heads off. Clark and Chalmers claim that Otto's notebook plays a similar functional role to Inga's biological memory. As a result, they argue, the notebook counts as part of Otto's cognitive processes, not in any merely technical sense, but in precisely the same sense as Inga's biological memory. Before consulting his notebook, Otto literally *believes* that MoMA is on 53rd Street, just as Inga believes this before consulting her memory. If external objects like notebooks can implement recognizably cognitive processes, like memory, then perhaps social interactions can too (Clark & Chalmers, 1998, pp. 17-18).

Giere rejects the extended mind thesis, since he regards it as leading to unnecessary and unanswerable metaphysical puzzles about the mind (e.g. 2006, p. 110-113). It is for this reason that he insists that d-cog involves a technical concept of cognition. If d-cog's proponents were sympathetic to the extended mind thesis, however, then this would seem to offer a way to bridge the gap between cognitive and social theorists. Just like Otto's notebook, scientists' social interactions might be said to implement processes that are cognitive not in any technical sense, but in exactly the same sense as internal, psychological processes, such as memory. Of course, one problem with this approach is that the extended mind thesis remains highly controversial (e.g. Menary, 2010). But another is that supporters of the extended mind use 'cognitive' in a far more restricted manner than Giere. For example, in their original article, Clark and Chalmers argue that, in order to count as part of the cognitive process, an external device must meet conditions of "glue and trust" (Clark, 2010). These conditions require that, like Otto's notebook, an external process must be a constant in the person's life, directly available and automatically endorsed. The glue and trust conditions might perhaps be met in some cases on the *Palau* or in the laboratory. For example, perhaps the plotter always has his slide rule with him and trusts what it says without question. But it seems that many processes, and in particular many social processes, will not meet these conditions. The reports of the pelorus operators are not always directly available to the plotter, for example, and he sometimes questions what they tell him. So it would seem that, even if d-cog's supporters were sympathetic to the extended mind thesis, further work would be needed to show that social processes are also cognitive.

7 Conclusion

Distributed cognition offers a promising approach for understanding scientific practice. My aim in this paper has been to clarify the nature of that approach and explore its

potential for bridging the gap between cognitive and social theories. In particular, I have discussed three apparent limitations on the computational form of d-cog as a means of reconciling cognitive and social theories. First, while d-cog might show how social interactions can implement a computation, it will not tell us why a particular representational scheme is used to carry out that computation or why those representations have meaning. Both of these aspects of science receive competing social and cognitive explanations. Second, the computational form of d-cog will not be acceptable to all social theorists, since it implies that science is not essentially social. Finally, I've suggested that, if d-cog's claim to merge social and cognitive factors involves only a technical concept of cognition, then it is difficult to see how it will help to reconcile debate between cognitive and social theorists. None of these issues should lead us to reject a d-cog approach to science, but they might perhaps help us to understand what such an approach can hope to achieve.

Acknowledgements

I would like to thank Giovanna Colombetti, Ronald Giere, Jeff Kochan, Sabina Leonelli, Nancy Nersessian, Lisa Osbeck, Thomas Sturm and Paul Thagard for extremely helpful discussion and correspondence about distributed cognition. Thanks also to audiences in Aberdeen, Athens, Bielefeld, Exeter and Nancy, and to two anonymous referees, for very insightful comments and criticism.

Notes

1. Following Slezak, I will use 'SSK' to refer to the Strong Programme and related approaches, rather than the sociology of science in general.

References

- Barnes, B., Bloor, D. & Henry, J. (1996). *Scientific knowledge: A sociological analysis*. London: Athlone Press.
- Berlin, B. (1992). *Ethnobiological classification: Principles of classification of plants and animals in traditional societies*. Princeton: Princeton University Press.
- Bloor, D. (1976). *Knowledge and social imagery*. London: Routledge and Kegan Paul.
- Bloor, D. (1982). Durkheim and Mauss revisited: Classification and the sociology of knowledge. *Studies in History and Philosophy of Science*, 13(4), 267-297.
- Brown, M. (2011). Science as socially distributed cognition: Bridging philosophy and sociology of science. In K. François, B. Löwe, T. Müller, & B. van Kerkhove (Eds.) *Foundations of the formal sciences VII* (pp. 17-31). London: College Publications.
- Clark, A. & Chalmers, D. (1998). The extended mind. *Analysis*, 58(1), 7-19.
- Clark, A. (2010). Coupling, constitution and the cognitive kind: A reply to Adams and Aizawa. In Menary, R. (Ed.), *The extended mind* (pp. 81-100). Cambridge, MA: MIT Press.
- Collins, H. M. (1989). Computers and the sociology of scientific knowledge. *Social Studies of Science*, 19, 613-624.
- Crane, T. (2003). *The mechanical mind: A philosophical introduction to minds, machines and mental representation*. London: Routledge. (1st ed. 1995)
- Dunbar, K. (1995). How scientists really reason: Scientific reasoning in real-world

laboratories. In R. J. Sternberg & J. Davidson (Eds.), *Mechanisms of insight* (pp. 363-395). Cambridge, MA: MIT Press.

Giere, R. (1989). Computer discovery and human interests. *Social Studies of Science*, 19, 638-643.

Giere, R. (2002a). Scientific cognition as distributed cognition. In P. Carruthers, S. Stich & M. Siegal, (eds.), *The cognitive basis of science* (pp. 285-299). Cambridge: CUP.

Giere, R. (2002b). Distributed cognition in epistemic cultures. *Philosophy of Science*, 69(4), 637-644.

Giere, R. (2002c). Models as parts of distributed cognitive systems. In N. J. Nersessian and L. Magnani (eds.), *Model-based reasoning: science, technology, values* (pp. 227-241). Dordrecht: Kluwer.

Giere, R. (2004). The problem of agency in scientific distributed cognitive systems. *Journal of Cognition and Culture*, 4(3-4), 759-774.

Giere, R. (2006). *Scientific perspectivism*. Chicago: University of Chicago Press.

Giere, R. (2007). Distributed cognition without distributed knowing. *Social Epistemology*, 21(3), 313-320.

Giere, R. (2012). Scientific cognition: human centered but not human bound. *Philosophical Explorations*, 15(2), 199-206.

Giere, R., & Moffatt, B. (2003). Distributed cognition: Where the cognitive and the social merge. *Social Studies of Science*, 33(2), 301-310.

Haugeland, J. (1985). *Artificial intelligence: The very idea*. Cambridge, MA: MIT Press.

Hutchins, E. (1995). *Cognition in the wild*. Cambridge, MA: MIT Press.

Kusch, M. (2002). *Knowledge by agreement*. Oxford: OUP Press.

Langley, P., Bradshaw, G. L., & Simon, H. A. (1983). Rediscovering chemistry with the BACON system. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach* (pp. 307 – 330). San Mateo, CA: Morgan Kaufmann.

Latour, B. & Woolgar, S. (1986). *Laboratory life: The construction of scientific facts*. Princeton: Princeton University Press. (1st ed. 1979)

Longino, H. (1990). *Science as social knowledge: Values and objectivity in scientific inquiry*. Princeton: Princeton University Press.

Magnus, P. D. (2007). Distributed cognition and the task of science. *Social Studies of Science*, 37(2), 297-310.

Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco: Freeman.

Menary, R. (2010). (Ed.) *The extended mind*. Cambridge, MA: MIT Press.

Nersessian, N. J. (2005). Interpreting scientific and engineering practices: Integrating the cognitive, social, and cultural dimensions. In M. Gorman, R. Tweney, D. Gooding, & A. Kincannon, (Eds.) *Scientific and technological thinking* (pp. 17-56). New Jersey: Erlbaum.

Nersessian, N. J., Kurz-Milcke, E., Newstetter, W. C., & Davies, J. (2003a). Research laboratories as evolving distributed cognitive systems. *Proceedings of the 25th Annual Conference of the Cognitive Science Society*: 857-862.

Nersessian, N. J., Newstetter, W. C., Kurz-Milcke, E. & Davies, J. (2003b). A mixed-method approach to studying distributed cognition in evolving environments. *Proceedings of the International Conference on Learning Sciences*: 307-314.

Osbeck, L. M. & Nersessian, N. (2006). The distribution of representation. *Journal for the Theory of Social Behaviour*, 36(2), 141-160.

Slezak, P. (1989). Scientific discovery by computer as empirical refutation of the strong programme. *Social Studies of Science*, 19, 563-600.

Solomon, M. (2001). *Social empiricism*. Cambridge, MA: MIT Press.

Thagard, P. (1989). Welcome to the cognitive revolution. *Social Studies of Science*, 19, 653-657.

Thagard, P. (1992). Societies of minds: Science as distributed computing. *Studies in History and Philosophy of Science*, 24(1), 49-67.