

Friends with benefits! Distributed cognition hooks up cognitive and social conceptions of science

P.D. Magnus and Ron McClamrock
Philosophical Psychology
DOI 10.1080/09515089.2014.964857

Authors' final draft, July 2014.
<http://www.fecundity.com/job>

ABSTRACT: One approach to science treats science as a cognitive accomplishment of individuals and so defines a scientific community as an aggregate of individual enquirers. Another treats science as a fundamentally collective endeavor and so defines a scientist as a member of a scientific community. Distributed cognition has been offered as a framework that could be used to reconcile these two approaches. Adam Toon has recently asked if the cognitive and the social can be friends at last. He answers that they probably cannot, posing objections to the would-be rapprochement. We clarify both the animosity and the tonic proposed to resolve it, ultimately arguing that that worries raised by Toon and others are unconvincing.

KEYWORDS: science studies, distributed cognition, d-cog, SSK, scientific knowledge, division of cognitive labor

Treating science as a cognitive accomplishment of individuals, we might define a scientific community as an aggregate of individual enquirers. Contrariwise, treating science as a fundamentally collective endeavor, we might define a scientist as a member of a scientific community. The framework of *distributed cognition* has been heralded as a way of doing both at once (e.g., by Ronald Giere in a series of publications; 2002a, 2002b, 2003a, 2003b, 2003c, 2006, 2007, 2013; Giere&Moffatt, 2003). Adam Toon (2014) has recently posed objections to this would-be rapprochement between the cognitive and the social. Our aim is to clarify the issues and ultimately to show that distributed cognition can deliver on its promise of hooking up cognitive and social conceptions of science.

In what follows, we use the phrase ‘d-cog’ instead of ‘distributed cognition’ as a reminder that this is a technical term of art rather than just the composition of the usual words ‘distributed’ and ‘cognition’. There is still a question of how the term of art should be understood, of course, and we adopt and argue for a specific criterion below.

1 Framing the Debate

Toon does not give a clear characterization of either ‘d-cog’ or of the ‘cognitive’ and ‘social’ extremes which it is meant to reconcile. Instead, he identifies them just by offering exemplars. Latour and Woolgar (1986), for example, stand in for the social extreme. The description of navigation teams given by Ed Hutchins (1995) stands in for d-cog. Before considering Toon directly, we want to offer a more detailed characterization both of d-cog and the tension it is meant to resolve.

1.1 The cognitive and the social

Put modestly, Giere’s suggestion is that distributed cognition provides “a way of overcoming some of the opposition that many have felt between social and cognitive understandings of science” (Giere, 2006, p. 114). Put more boldly, it is that d-cog can provide “a productive theoretical framework that all contributors to science studies, including historians, philosophers, psychologists, sociologists and anthropologists, can share” (Giere, 2007, p. 319).

The cognitive approach is individualist and presumes that the primary unit of analysis is the individual organism or thinker. This presumption is sometimes just a visceral one; John Searle comments in a recent interview, “it upsets me when I read the nonsense written by my contemporaries, the theory of extended mind makes me want to throw up” (Boag, 2013). It also stems from some common methodologies. In philosophy, it can be traced to the tradition which treats the task of epistemology as analyzing ‘S knows that P’ for an individual knower S and some proposition P. Communities, which typically go unmentioned, are made up of atomic individual Ss who must each meet the condition for knowledge. In cognitive science, it can come from the Good Old-Fashioned AI (GOF AI) approach which attempts to model cognition entirely as software which runs inside an individual organism, connected to the world by the inputs and outputs of biological hardware.

The social approach presumes that the real scientific action happens in public. SSK, the sociology of scientific knowledge, is typically offered as exemplary of the social approach. Some SSK scholars express explicit hostility to epistemology. Latour and Woolgar proposed setting it aside entirely and called for a “ten-year moratorium on cognitive explanations of science” (1986, p. 280). Giere invokes Latour explicitly, and Toon calls the proposed moratorium “[p]erhaps the clearest instance” of the cognitive/social divide (Toon, 2014, p. 112).

Latour and Woolgar’s ten-year moratorium ended almost twenty years ago, so one might suspect that invocations of it are just dying echoes of the 1990s Science Wars. However, although the war is over, recent discussions of science sometimes fall along the same battle lines. For example, Philip Mirowski writes in discussing the commercialization of science:

...I want to set out some relatively tractable notions of “good science,” starting from some aggregate measures, and then ask what has happened to them. Philosophers have turned out not very helpful in this regard, primarily acting as though there abided generic benchmark characterizations of reliable knowledge, meanwhile drawing the bulk of attention away from the social conditions and structures that might improve or degrade the process of validating or augmenting the quality of knowledge. Perhaps even more disturbing, the profession of science studies has become absorbed with the prospect of “democratizing” science and diminishing the power of expertise, to the neglect of considerations of the degradation of the knowledge base that might possibly ensue as a consequence of “opening up” science to various constituencies. (Mirowski, 2011, p. 288)

Mirowski wants a way of thinking about what it would mean for the scientific community to be operating well. He complains that philosophers have talked too much about individual knowers, and social theorists have nothing to say about cognitive accomplishments of science. We might point to writers both in philosophy and science studies who have done some of the work that Mirowski wants, but our point is simply that he describes the two approaches which Giere suggests d-cog can reconcile.

1.2 Characterizing d-cog

For the sake of concreteness, we will start from a specific criterion for d-cog given by Magnus (2007).¹

It is standard, following David Marr (1982), to distinguish levels of analysis. We distinguish the *task* of an activity from the *process* that implements it. The task is given as “an abstract and idealized specification of the behavior to be achieved” (McClamrock, 1995, p. 20), situated in a context in which the behavior is specified. The process, by contrast, is given by the actual local mechanisms and procedures which accomplish the behaviors in that context. Applied to one of Marr’s examples (1982, pp. 22–33): Tallying the total cost of your purchases is the task accomplished by a cash register. The way numbers are represented in the cash register and the physical circuitry are both part of the process.²

Now we can say that an activity is ‘d-cog’ if (a) the task is one that we would consider cognitive if it were performed within a single individual and if (b) the process extends beyond the individual. (Magnus, 2007, p. 300) This has the virtue of applying clearly to exemplars of d-cog and generalizing well to other cases. For example, adding sums in your head is a paradigm cognitive activity, so using a cash register to tally large sums counts as d-cog. Before moving on, we take up several objections that have been raised against this criterion.

First, a single system can be characterized in different ways relative to different task specifications. Hyundeuk Cheon objects to this account that “task-specification *relativity* is hardly justified” and “task-specification relativity leads to the relativity of cognition, thus making cognition so superficial” (Cheon, 2013). Cheon’s worry only makes sense if we think that ‘System S is d-cog’ must be a well-defined monadic predicate. We suggest instead that it should be understood instead as a relation ‘System S is d-cog for task T’.

Second, as we readily acknowledge, there is no single task which is performed by cognition simpliciter (Magnus, 2007, p. 307). So Cheon objects, “We may therefore ask why [Magnus] demands the specification of the task only for distributed cognition, but does not for undistributed cognition” (Cheon, 2013). This objection misunderstands

¹ The criterion given by Magnus (2007), although advocated there by just one of us, is posed in terms of a distinction that had earlier been articulated by the other of us (McClamrock, 1991; 1995).

² Note that this takes both Marr’s *algorithm* and *implementation* to be aspects of the *process*.

the suggestion. There may be no general way of characterizing a task which is performed by all cognition simpliciter, but particular instances of cognition — distributed or not — must have particular tasks. For example, when a person calculates a sum in her head, the task is to perform addition. So cognition simpliciter is also a relation, something of the form ‘System S implements the cognitive task T’. Undistributed cognition is the special case where the system is entirely within a single thinking individual.

Note that task-relativity is not an ad hoc peculiarity of d-cog explanation. For all kinds of systems, structures often have multiple functions with respect to different tasks. To take a non-cognitive example, bird feathers contribute both to flight and to temperature regulation. Understanding and assessing birds requires recognizing the separate tasks that their feathers perform. To take a homey cognitive example, imagine someone who is panicking because they are afraid of the dark and who attempts to calm themselves by reciting multiplication tables. Treating their recitation as a performance of the task *calculating products* misses important aspects of what is going on, namely that it is also directed at the task *controlling attention and heart rate*.

Third, Cheon suggests that the account conflates notions of ‘distributed cognition’ and ‘extended cognition’ which have different motivations and different consequences. As Cheon has it, “extended cognition performs some tasks that one individual, in principle can perform” whereas “distributed cognition” pursues “goals that one human agent cannot solely achieve” (Cheon, 2013).³ If this is supposed to mean that d-cog must have goals which could not even *in principle* be carried out by an individual: We do not know what tasks Cheon has in mind, but we worry that they would only count as ‘cognitive’ in some contentious sense. A collective accomplishment which is nothing like what an individual thinker could achieve is *ipso facto* not something akin to thought. If Cheon means instead to say that d-cog

³ Hutchins (2014) uses the phrase “distributed cognition” to mean a methodological approach which analyzes cognition in terms of the interaction between separate functional elements, even if all of the functional components are contained inside a single organism. Our characterization of *d-cog* explicitly requires distribution beyond the boundary of an organism, because we are especially interested in understanding science. For our purposes, what is crucial is precisely the involvement of instruments and communities of people.

must *as a practical matter* be something that cannot be done by one individual: We readily admit that a great many cognitive tasks are impossible for any human to do in their head because of basic limits like brain size, but nevertheless we can recognize the tasks as ones that are cognitive. For example, there is a limit to the size of arithmetic calculations that a given material brain can perform, but performing larger calculations still counts as a cognitive task. If my task is adding a long series of numbers, the pad of paper or calculator I use is a critical part of my task solution, such that I could not do the task in the same way (and maybe not even at all) without such a tool. If Cheon allows that a large enough paper-and-pencil calculation counts as distributed cognition, then it seems natural to also count a smaller calculation which I could perhaps have done in my head when I actually opt to carry it out longhand.⁴

Fourth, one might worry d-cog entails the existence of extended minds in a metaphysically thick sense of *mind*. In Clark's way of putting it, extended cognition is the view that "the actual local operations that realize certain forms of human cognizing include inextricable tangles... of brain, body, and world. The local mechanisms of mind, if this is correct, are not all in the head" (Clark, 2008, p. xxviii). This can be taken to mean that parts of the world outside the organism not only contribute to the content of cognition but also bear mental properties. There has been prominent and healthy debate over this thesis. (Clark, 2008; Adams & Aizawa, 2008) For our current purposes, however, we need not take a stance in that debate. Much ballyhooed questions about *mind* are not our target. For something to be d-cog, it suffices that the mechanisms underlying someone's ability to accomplish cognitive tasks are in part outside that person's body. Extended mind, read in a strong way, is the additional claim that at

⁴ One might wonder about the relation between d-cog and *embedded cognition* — the view that, as Rupert puts it, "cognitive processes depend heavily, and in hitherto unexpected ways, on organismically external properties and devices, and on the structure of the external environment in which the cognition takes place" (2004, p. 393). The central claim of both is that mechanisms, objects, and structures instantiated outside the body are often critical supporting networks to the task accomplished in cognition. If one reads embedded cognition as something that happens inside the organism but requires the external structure as props, then it is a somewhat more conservative view. Alternately, one might read embedded cognition as d-cog.

least some of those mechanisms are also constitutive of the person having a mind or at least a mind of a certain sort.

Fifth, some authors think it is important to distinguish distributed cognition which distributes the process across props and tools from distributed cognition which distributes the process across several individuals. If one begins by wondering about the boundaries of one's own mind, one might be interested in the way that instruments extend one's abilities but not in the way other people do. The other people, after all, have their own minds. So one could characterize distribution in a way which only allows there to be one thinker at the center of it. Contrariwise, Giere allows distributed cognition to involve more than one person. Yet he requires it to involve tools and instruments, counting cognition which includes multiple people but no tools as mere 'collective cognition' (Giere, 2007). We do not make either restriction, but rather intend d-cog as the broader and more inclusive category including both variants. Although tools may be crucial to some d-cog explanations, others are predominantly social. In order to understand how science as a social endeavor can do cognitive work, it is precisely the involvement of so many people — where none of them has the whole scheme in their own mind — that is crucial. So we neither require instruments nor preclude other people being involved in d-cog. We insist that the process which implements d-cog will include something outside the individual organism, but those outside contributors might be other people, things, or both. Just as we are agnostic on the question of whether a d-cog system of individual with a smartphone has an *extended mind*, we are agnostic as to whether a scientific team has some kind of *hive mind*.

Sixth, and finally, some authors have argued that a definition of d-cog should be broader than this criterion. For example, Matthew Brown (2011) suggests moving beyond the task/process distinction to a tri-fold distinction between operations, actions, and activities. So a system S might be d-cog *qua* activity system, without specifying any further terms. He concludes provocatively, "presumably, *I* am some kind of cognitive system, even though I am not built to carry out one specific and well-bounded task, even though my cognitive activities evolve, and aren't always as well-bounded as certain cognitive theories might presuppose" (Brown, 2011, p. 28). Note, however, that Brown's alternative has not been developed in any detail. We worry about whether it can be applied in a clear way to substantive cases.⁵

⁵ We have the same worry about Cheon's (2013) positive definition, which includes many unexplicated jargon terms: "cognitive

Regardless, the objection is only that the conception is too narrow — not that it is too lax. The worry misses its mark, since the condition is not a definition but only a sufficient condition (Magnus, 2007, p. 300). A sufficient condition, even if it does not exhaust the phenomena, is enough to positively diagnose cases of d-cog. If some more expansive definition could be elaborated and fruitfully employed, it would encompass both the cases we discuss below and more besides.

Having cleared this ground, let's turn to Toon's objections.

2 Toon's challenge

Adam Toon (2014) raises three kinds of problems for d-cog. Here are the three problems, briefly. We consider each at greater length in subsequent sections.

1. A d-cog description may explain *how* a task is performed but not *why* it is performed.
2. A d-cog description is only contingently social, whereas many social theorists insist that science is essentially social.
3. A d-cog description is cognitive only in a technical sense of 'cognitive', not in the pretheoretic sense that frames the opposition between the social and the cognitive.

2.1 Why the task is performed

Toon points out that a d-cog description of some cognitive activity does not explain why the activity was performed. Hutchins' analysis of navigation describes the activities of a ship's crew and can explain why particular elements of the navigation cycle happen as they do (for example, why the plotter draws a particular line on a map). Yet it cannot explain why the whole cycle occurs at all or even, in Toon's view, help to account for such matters.

We argue that Toon's reasons for this depend on misconstruing the three-level analysis that he inherits from Marr and Hutchins. As originally posed by Marr, the analysis separates the *computational theory* (concerned with "the goal of the computation, why is it appropriate, and what is the logic of the strategy by which it can be

outputs" which result from "coordinated information processing" and the propagation of "semantic information". Although the phrases use familiar English words, we do not have a precise enough understanding of them to apply the definition to cases.

carried out”), the account of *representation and algorithm* (which asks “How can this computational theory be implemented? In particular, what is the representation for the input and output, and what is the algorithm for the transformation?”), and *hardware implementation* (concerned with “How can the representation and algorithm be realized physically?”) (Marr, 1982, p. 25).⁶

Toon applies this distinction in a way that inappropriately seems to undermine d-cog. In his characterization, “d-cog is a claim about what happens at the implementation level of a computation”, by describing how its physical implementation might be distributed in the physical world, and “it is in this sense that d-cog offers a means of integrating cognitive and social factors”. Because he takes d-cog analysis to just be about hardware implementation, Toon thinks it “will not tell us why a particular computation is being performed or why a specific representational system is being used to carry it out” and that “d-cog analysis... will not tell us how a representational system gains its representational status.” So he claims that d-cog “cannot help to reconcile debate between” social and cognitive answers to such questions (Toon, 2014, p.117).

However, showing how activities are d-cog is a matter of showing facts about not only implementation but about representation and algorithm as well. The analysis of the division of labor in navigation is not only about what physical things play various roles, but is even more centrally about how the process that accomplishes the task is divided into a set of cooperating sub-procedures and exchanges of information — that is, the representations and algorithms for accomplishing the task of navigation. Such analyses of information processing within the system, rather than exclusively hardware implementation descriptions, are just the sort of things that in turn inform and illuminate the content-level questions that should be addressed at the level of task. (McClamrock, 1995) Contra Toon, much of a d-cog analysis *will* “concern the choice of representational scheme used to carry out a computation” rather than merely “the way that the computation is implemented”; cf. (Toon, 2014, p.117). So the broader questions may well be at least partially answered within the d-cog framework.

⁶ Ultimately, the three ‘level’ analysis must be adjusted to account for different perspectives on various levels of organization of complex systems; see McClamrock (1991). As David Danks put it, Marr’s three levels conflate “a multidimensional space of theoretical commitments” (2013, p. 2129).

Of course there are limitations, but d-cog explanations are no different than traditional cognitive explanations in this regard. Consider David Danks' criticism of rational analyses which explain an agent's input or output in terms of its solving a particular task in a specified environment (2008). Such analyses do not ultimately explain the agents' input and output, Danks notes, because they do not explain how the agent developed so as to employ the optimal strategy. Showing how the task is functional does not show how it came to be undertaken. This is only a problem for the d-cog approach if broader context would involve incompatible descriptions in social or in cognitive terms, so that the tension between the two returns. A broader analysis will necessarily be something *further*, but it may also be d-cog. The navigation cycle determines the position of the ship, and position information is used in various ways in the broader operation of the ship. Using different task analyses of the activity makes it possible to situate it in these broader contexts. The navigation cycle occurs to determine the present position of the ship, which is used to plot the movement of the ship over time for purposes of navigation and planning. Those tasks figure in the life of the ship, which is a unit in the navy, which in turn is a strategic instrument of the nation. The regress of explanation must come to a stop somewhere, because we only follow it as far as our explanatory interests require and our abilities allow, but that alone will not drive a wedge between the cognitive against the social. If we persist in asking *why* questions about each broader element, we will eventually reach a point where d-cog explanations break down — but we will also reach a point where neither social nor cognitive explanations are available.

2.2 Whether science is essentially social

Toon points out that a d-cog description makes social processes logically contingent. For scholars who prefer social explanations, Toon suggests, science being social is essential and logically necessary. So a d-cog description would not satisfy them.

Stated in this way, it is not much of an objection. Admittedly, d-cog alone does not and should not absolutely necessitate *social* distribution of the tasks among multiple agents. Yet the centrality of social distribution to the activity and task of gaining scientific knowledge is what we should expect, given our cognitive limits and social natures — facts about us which are in principle contingent, but deeply general and entrenched. Because of the kinds of creatures we are, we cannot really do science without working together. So science's

being a social enterprise is practically indispensable for us. The mere logical possibility of an epistemic angel who does science alone in its own mind is too abstruse to matter for understanding science as a thing that finite creatures do. Science as socially distributed cognition is a variety of d-cog, not something in conflict with it.

Toon presses the point by observing that d-cog explanation does not explain how d-cog systems have representational content (that is, how they come to have meaning). The d-cog explanation can describe the operation of social and material representations, but it “leaves open the question of how those representations come to represent the world.” He adds that “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” (Toon, 2014, p. 118).

In reply, it is important to consider why the social nature of science is important to thinkers like those in the Edinburgh school. On closer inspection, their reasons to insist that science is social are not reasons to quarrel with d-cog. David Bloor, one of the architects of the Edinburgh school and the strong program, explicitly rejects the division between the cognitive and the social. The division reflects what Bloor calls “the ‘distortion’ model”, i.e. “the traditional rationalist dichotomy between the rational and the social” (1991, p. 169). The distortion model only allows for a weak sociology of knowledge, because social factors would only be factors distorting rather than contributing to knowledge. Such a weak program is just what the strong program is defined against.⁷

Still, Toon is right that the strong program takes the social nature of science to be essential. Bloor writes that “our best and most cherished scientific achievements could not exist as they do without having the character of social institutions” (1991, p. 164). Bloor argues for this by imagining several separate people who look at the same measurements and infer “their own personal version of Boyle’s Law.” The solipsist-scientists would not yet have “Boyle’s Law as we know it”, however, because they would not have “a scientific community with a shared body of knowledge” (Bloor, 1991, p. 168). The community having a publicly available theory as a shared object is different than each individual having private representations which just happen to coincide. However, this strikes us as entirely compatible

⁷ Bloor himself expresses the opposition in these terms, e.g. (Bloor, 1991, p. 165).

with d-cog explanations. There would be some commonality between the two cases (insofar as the imagined personal versions of Boyle's Law and actual public versions of Boyle's Law can both count as versions of Boyle's Law in some sense) but there is the deeply important difference between them (insofar as one is the imagined feat of an almost inhuman private cognition and the other is the actual product of human d-cog).

Finally, some of the apparent disagreement is just terminological, because parties define 'knowledge' in different ways. The Edinburgh school defines 'knowledge', as distinct from individual belief, in collective terms. Bloor writes: "Of course knowledge must be distinguished from mere belief. This can be done by reserving the word 'knowledge' for what is collectively endorsed" (1991, p. 5). Steven Shapin writes similarly that "the very idea of *knowledge* implicates a public and shared commodity, to be contrasted with the individual's state of *belief*" (1996, p. 106). If knowledge must be social simply as a matter of definition, then so too for scientific knowledge *qua* knowledge.⁸ Epistemologists who define 'knowledge' differently are disagreeing about how to use words rather than about substance.⁹ It seems to us that d-cog is independent of these matters. It is compatible with both collectivist and individualist definitions of 'knowledge', and so with the collective nature of 'scientific knowledge' being necessary

⁸ In addition to SSK, Toon mentions the work of Helen Longino. He writes, "Similar claims are also found in the work of authors outside SSK. For example, Longino writes that science is 'necessarily social' " (2014, p. 120). In the passage he cites, Longino argues that "the development of knowledge is a necessarily social rather than individual activity" (Longino, 1990, p. 12). Here the terminological difference seems especially acute. Longino offers several definitions of 'knowledge', all of which include the community as an explicit term (Longino, 2002, pp. 135–140).

⁹ Note that the issue can be made substantive by combining an individualist definition of knowledge with the requirement that cognitive activity yield knowledge. For example, Giere (2013) insists that science must yield human knowledge for an individual. Krist Vaesen (2011) argues that this risks allowing any d-cog analysis to be refigured in individual terms. Although our sympathies lie with Vaesen, d-cog as we understand it is compatible with either Giere or Vaesen being right.

by definition or merely necessary given deep facts about the kinds of creatures we are.

2.3 What sense of ‘cognitive’ is in play

Toon objects that d-cog is only ‘cognitive’ in a technical sense. It does not mean, for example, the thinkings of a mind. Indeed, its status as a term of art was our motivation for adopting the abbreviation ‘d-cog’ in the first place. Yet, Toon suggests, scholars who prefer cognitive explanations do so for some ordinary sense of ‘cognitive’. So there is a possible gap between the cognitive explanations they want and the explanations d-cog provides.

There are several possible answers to the criticism.

First, the gap could be closed by adopting a much extended conception of mind, so that all cognition would be permeated by external and social mechanisms. The revolutionary conception of mind would comfortably take in d-cog. As we explained above, however, we adopt an anodyne conception of d-cog and remain agnostic with respect to deeper issues in the metaphysics of cognition and mind.

Second, a parallel problem arises for any technical jargon. We can ask, for example, whether the epistemologist ends up saying anything which applies to ‘knowledge’ in an ordinary sense. The mere possibility of a mismatch is a general feature of philosophical explication. So what?

Third, for any d-cog which meets the criterion we defended in section 1.2, there is a more direct answer. The criterion for d-cog is silent as to whether or not an operation of a community is cognitive simpliciter, but the task must be one which would count as ‘cognitive’ in the ordinary sense if it were carrying out just in an individual thinker. So, although the d-cog system might not strictly-speaking *be* cognitive — depending on your view on extended as opposed to merely embedded mind — the characterization has a built-in connection to the ordinary conception of cognition.

It remains to be shown that d-cog which meets this criterion will do the work of reconciling the cognitive and the social. In the next section, we illustrate one way that it can by looking at recent research on the division of cognitive labor.¹⁰

¹⁰ Other examples given by Magnus (2007) include the machinery in a carpentry shop and double-blind clinical trials.

3 D-cog at work

Consider scientists working together to discover what the world is like.

If we focus only on individual cognition, scientists' tasks are specific ones about experimental design, data collection, analysis — and even more quotidian things like filling out paperwork. Putting all that into soft focus, writers in the epistemology literature tend to privilege a broader individual task: Evaluate the theoretical possibilities so as to believe the theory most likely to be true.

If we retain the broad perspective but focus on the community, then the collective task might be something like: Explore the space of possibilities so that all are sufficiently interrogated. This may not be something that is done well by each scientist pursuing their own personal, epistemic good. One individual only has limited resources; she sets up a lab with some particular equipment, she only has finite time and attention, etc. So different individuals might believe in different theories and approaches, pursuing them independently.

Exploring new possibilities needs to be balanced against further articulating established approaches. This would be difficult for a single scientist to do simply in her own work, but it can be solved at the community level by having a mixed community of more conservative and more adventurous scientists. Most scientists should work within the dominant approach, both because it has a proven track record of problem-solving and because its success relies on their cooperating to advance it. Yet some scientists should pursue novel approaches, even though most of those will be dead ends. This point was seen clearly by Thomas Kuhn. In the 1969 postscript to *The Structure of Scientific Revolutions*, he writes that

individual variability in the application of shared values may serve functions essential to science. The points at which values must be applied are invariably also those at which risks must be taken. Most anomalies are resolved by normal means; most proposals for new theories do prove to be wrong. If all members of a community responded to each anomaly as a source of crisis or embraced each new theory advanced by a colleague, science would cease. If, on the other hand, no one reacted to anomalies or to brand-new theories in high-risk ways, there would be few or no revolutions. (Kuhn, 2012, pp. 185–6)

So this explanation was available to all of the parties to the Science Wars, even if its implications were not recognized. Kuhn himself

simply employs the idea to motivate thinking of scientific desiderata as values rather than rules.

Since Kuhn, this has come to be called the *division of cognitive labor*. Michael Weisberg (2010) and Ryan Muldoon (2013) distinguish three approaches in the recent literature.

First: the marginal contribution/reward approach exemplified by the work of Philip Kitcher (1990; 1993, ch. 8). Scientists are modeled as self-interested agents after the fashion of models in classical economics. Without any central planning, a reward structure which pays scientists for making discoveries can cause them to divide between mining the established approach and exploring novel approaches. The established approach is more likely to pay off, but with so many scientists on the job it is less likely to payoff *for any particular individual* than a speculative approach.

Second: the epistemic network approach exemplified by the work of Kevin Zollman (2010). Scientists are modeled as only having partial information about what other scientists are doing, and they pursue various approaches because of what seems appropriate given their limited information. It is possible for this community arrangement to achieve better outcomes than if each scientist was always fully informed. Zollman concludes, “Looking at these scientists from the perspective of individualistic epistemology, one might be inclined to criticize the scientists’ behavior. However, when viewed as a community, their behavior becomes optimal” (2010, p. 33).

Third: the epistemic landscape approach exemplified by work of Weisberg and Muldoon (2009). It models the space of possible approaches as a landscape and scientists as either mavericks (who eagerly push out to new regions of the landscape) or followers (who go where others have had the most success). This is a formal version of Kuhn’s idea that some scientists will explore new possibilities aggressively while others will stick to the established way of doing things.

All of these approaches are recognizably socially-distributed varieties of d-cog. An individual might divide their time and attention between exploring different approaches, but it is impractical for this to actually be done within an individual. A scientist spends time developing the skills suited to one approach or another, and their enthusiasm and effectiveness may depend on not dividing their attention. There are also material as well as psychological constraints. A scientist equips their lab for particular approaches and cannot mount parallel labs in which they would pursue other approaches. So the problem of dividing attention becomes a social problem of dividing

labor. The cognitive task is implemented not in the work of a single universal scientist but across a community of scientists.

The example also underscores the task-relativity of d-cog. Above, we characterized the task as *exploring the space of possible theories*. That is something that an individual might try to do, although it is better done by a community. Once it is approached as a collective enterprise, one might consider the task to be *determine the optimal distribution of labor*. Too, this might be attempted by a single individual — a central planner assigning scientists to various tasks. The upshot of the literature on the distribution of cognitive labor is that it is better to configure the community so that individuals arrive at something near an optimal distribution without such central planning.

Finally, one might characterize the task as something like *distribute the cognitive labor*. That is obviously not something that a single individual could do. A partitioned mind which could function as a community of individuals would just be a community, rather than an individual. Yet the former two task specifications do describe important aspects of science, and the process which implements those tasks makes it d-cog.

4 Conclusion

By giving examples of substantive and rewarding analyses of the scientific community as solving what are recognizably cognitive problems with solutions which meet the criterion for d-cog, we have deflected the last of Toon's worries. Contra Toon, d-cog does not sneak by with some ersatz sense of 'cognition'.

At the outset, we noted that the promise of d-cog can be understood in a more modest or in a more bold way.

The modest promise was that d-cog would provide "a way of overcoming some of the opposition that many have felt between social and cognitive understandings of science" (Giere, 2006, p. 114). A worked example, like the literature on the division of scientific labor, is sufficient to fulfill that promise. It shows how social organization can be understood as doing epistemic work.

The bold promise was that d-cog would provide "a productive theoretical framework" (Giere, 2007, p. 319). This is a harder promise to discharge, because it is forward looking. Whether a present research program will be productive and rewarding is something to be judged in retrospect, by future scholars. It is a programmatic commitment, namely that d-cog looks promising. For this, there is little value in characterizing d-cog in a metaphysical way which stymies application

to cases. A readily applicable conception of d-cog — like the one we have defended — is most useful.

Acknowledgements

We would like to thank Adam Toon and several anonymous referees for their helpful comments on earlier drafts.

References

- Adams, F.R., & Aizawa, K. (2008). *The Bounds of Cognition*. Malden, Massachusetts: Blackwell Publishing.
- Bloor, D. (1991). *Knowledge and Social Imagery*. Chicago: University of Chicago Press. Second edition.
- Boag, Z. (2013) Searle: it upsets me when I read the nonsense written by my contemporaries. *New Philosopher*, 1(2), November 2013. Posted on-line 25jan2014. URL: <http://www.newphilosopher.com/articles/john-searle>
- Brown, M. (2011). Science as socially distributed cognition: Bridging philosophy and sociology of science. In K. François, B. Löwe, T. Müller, & B. van Kerkhove (Ed.), *Foundations of the Formal Sciences VII*, volume 32 of *Studies in Logic* (pp. 17–30). College Publications.
- Cheon, H. (2013). Distributed cognition in scientific contexts. *Journal for General Philosophy of Science*. Published on-line. doi:10.1007/s10838-013-9226-4
- Clark, A. (2008). *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford University Press.
- Danks, D. (2008). Rational analyses, instrumentalism, and implementations. In N. Chater & M. Oaksford (Eds.), *The Probabilistic Mind: Prospects for Bayesian Cognitive Science* (pp. 59–78). Oxford University Press. URL: <http://repository.cmu.edu/philosophy/89>
- Danks, D. (2013). Moving from levels & reduction to dimensions & constraints. In M. Knauff, M. Pauen, N. Sebanz, & I. Wachsmuth (Eds.), *Proceedings of the 35th Annual Meeting of the Cognitive Science Society* (pp. 2124–2129), Austin, Texas, August 2013.
- Giere, R.N. (2002a). Discussion note: Distributed cognition and epistemic cultures. *Philosophy of Science*, 69(4), 637–644.

- Giere, R.[N.] (2002b) Scientific cognition as distributed cognition. In P. Carruthers, S. Stich, & M. Siegal (Eds.), *The Cognitive Basis of Science* (pp. 285–299). Cambridge University Press.
- Giere, R.N. (2003a). Computation and agency in scientific cognition. In *Proceedings of the 25th Annual Conference of the Cognitive Science Society*, 2003, Boston, Massachusetts, July 31– August 2.
- Giere, R.N. (2003b). A new program for philosophy of science. *Philosophy of Science*, 90(1), 15–21.
- Giere, R.N. (2003c). The role of computation in scientific cognition. *Journal of Experimental and Theoretical Artificial Intelligence*, 15(2), 195–202.
- Giere, R.N. (2006). *Scientific perspectivism*. Chicago, Illinois. University of Chicago Press.
- Giere, R.N. (2007). Distributed cognition without distributed knowing. *Social Epistemology*, 21(3), 313–320.
- Giere, R.N. (2013). Scientific cognition: human centered but not human bound. *Philosophical Explorations*, 15(2), 199–206.
- Giere, R.N., & Moffatt, B. (2003). Distributed cognition: Where the cognitive and the social merge. *Social Studies of Science*, 33(2), 1–10, April 2003.
- Hutchins, E. (1995). *Cognition in the Wild*. The MIT Press, Cambridge, Massachusetts, 1995.
- Hutchins, E. (2014). The cultural ecosystem of human cognition. *Philosophical Psychology*, 27(1), 34–49, January 2014.
- Kitcher, P. (1990). The division of cognitive labor. *Philosophy of Science*, 87(1), 5–22.
- Kitcher, P. (1993). *The Advancement of Science*. Oxford University Press.
- Kuhn, T.S. (2012). *The Structure of Scientific Revolutions*. Chicago, Illinois: University of Chicago Press. Fourth edition.
- Longino, H. (1990). *Science as Social Knowledge*. Princeton, New Jersey: Princeton University Press.
- Longino, H. (2002) *The Fate of Knowledge*. Princeton, New Jersey: Princeton University Press.
- Latour, B, & Woolgar, S. (1986). *Laboratory Life: The Construction of Scientific Facts*. Princeton, New Jersey: 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*. San Francisco: W.H. Freeman and Company.

- McClamrock, R. (1991). Marr's three levels: A re-evaluation. *Minds and Machines*, 1(2), 185–196.
- McClamrock, R. (1995). *Existential Cognition: Computational Minds in the World*. Chicago, Illinois: University of Chicago Press.
- Mirowski, P. (2011). *Science-mart: Privating American science*. Cambridge, Massachusetts: Harvard University Press.
- Muldoon, R. (2013). Diversity and the division of cognitive labor. *Philosophy Compass*, 8(2), 117–125.
- Rupert, R.D. (2004). Challenges to the hypothesis of extended cognition. *Journal of Philosophy*, 101(8), 389–428.
- Shapin, S. (1996). *The Scientific Revolution*. Chicago, Illinois: University of Chicago Press.
- Toon, A. (2014). Friends at last? Distributed cognition and the cognitive/social divide. *Philosophical Psychology*, 27(1), 112–125.
- Vaesen, K. (2011). Giere's (in)appropriation of distributed cognition. *Social Epistemology*, 25(4), 379–391.
- Weisberg, M. (2010). New approaches to the division of cognitive labor. In P.D. Magnus & J. Busch, editors, *New Waves in Philosophy of Science* (pp. 250–269). Basingstoke, Hampshire: Palgrave MacMillan.
- Weisberg, M., & Muldoon, R. (2009). Epistemic landscapes and the division of cognitive labor. *Philosophy of Science*, 76(2), 225–252.
- Zollman, K.J.S. (2010). The epistemic benefit of transient diversity. *Erkenntnis*, 72(1), 17–35.