

Distributed Cognition and the Task of Science*

P.D. Magnus

<http://www.fecundity.com/job>

Published in *Social Studies of Science*, 37(2): 297–310. April 2007.
This is my final draft and so does not reflect copyediting done by the staff of the journal.

Abstract

This paper gives a characterization of distributed cognition (d-cog) and explores ways that the framework might be applied in studies of science. I argue that a system can only be given a d-cog description if it is thought of as performing a task. Turning our attention to science, we can try to give a global d-cog account of science or local d-cog accounts of particular scientific projects. Several accounts of science can be seen as global d-cog accounts: Robert Merton’s sociology of scientific norms, Philip Kitcher’s 20th-century account of cognitive labor, and Kitcher’s 21st-century notion of well-ordered science. Problems that arise for them arise just because of the way that they attribute a function to science. The paper concludes by considering local d-cog accounts. Here, too, the task is the crux of the matter.

1 Introduction

Science, at least for limited creatures like us, is a social activity. It is not the sort of thing that someone could invent and perfect alone. The framework of *distributed cognition* (d-cog) seems like a natural way to characterize this social character of science. For one thing, it allows us resist the dichotomy between the social and the cognitive. We could recognize science as social without thereby reducing it to a matter of brute political force. Moreover, d-cog would give us resources to talk about how science allows more-or-less ordinary people to do collectively what they could not do alone. As Giere and Moffatt write, “[H]umans with normal human cognitive capacities manage to

*An earlier version of this paper was presented to the Experimental Philosophy Lab at the University of California, San Diego. I would like to thank Bill Bechtel, Jonathan Cohen, Ronald Giere, and John Jacobson for helpful discussions of these issues. Special thanks also to Bradley Armour-Garb, Matthew Brown, Ron McClamrock, Orin Zelenak, and a number of anonymous referees for feedback on previous drafts.

do modern science. One way ...[this is done]... is by constructing distributed cognitive systems that can be operated by humans possessing only the limited cognitive capacities they in fact possess" [GM03, p. 7].

In §2, I give a characterization of distributed cognition (d-cog). In §3, I argue that a system can only be given a d-cog description if it is thought of as performing a function. Turning our attention to science, we can try to give a global d-cog account of science or local d-cog accounts of particular scientific projects. In §4, I survey several accounts of science which can be seen as global d-cog accounts: Robert Merton's sociology of scientific norms, Philip Kitcher's 20th-century account of cognitive labor, and Kitcher's 21st-century notion of well-ordered science. Problems that arise for them arise just because of the way that they attribute a function to science. In §5, I consider local d-cog accounts. Here, too, the function is the crux of the matter.

2 The d-cog framework

Ed Hutchins' analysis of navigation aboard navy ships [Hut95] is a locus classicus for distributed cognition. Hutchins describes the process by which a ship's navigation team calculates the position of the ship over time by plotting a series of lines of position. The process involves several crewmen who use instruments to determine the bearings of certain landmarks, another who coordinates their work and records the bearings, and another who plots the bearings on a map and projects the ship's future position. The first several crewman are clearly involved in a computation that they do not fully understand. Since the plotter chooses the next set of landmarks, it may be tempting to see the rest of the system as a sort of instrument that he uses to determine the position of the ship. Yet it would be wrong to think that the things happening inside the mind of the plotter are sufficient to reconstruct the process; he employs heuristics and skillfully embodies specific computations in hoeyes and parallel rulers.

The *hoey* is a sort of straightedge-cum-protractor that is laid on the map to plot the line of position. As Hutchins describes, it does the work of converting the digital representation of bearing as such-and-so many degrees into the analog representation of a physical angle on the map [Hut95, pp. 124–125]. Its operation relies on the hoey representing angles in the same way as the gyro-compass that was used to produce the reported bearings. Because of these tools, the system performs cognitive operations not performed by any of the individual agents. Hutchins argues:

The whole cycle is something that emerges from the interactions of the individuals with one another and with the tools of the space. The structure of the activities of the group is determined by a set of local computations rather than by the implementation of a global plan. [Hut95, p. 200]

It would also be wrong to think that the whole process is understood by someone else, some master planner back at naval command. Yet, on Hutchins' analysis,

navigation is still an information processing task structured around a specific goal: to determine the position of the ship given available information. There is nothing essentially distributed about this goal. It does not *a priori* require either instruments or a community of actors. In order to capture the way in which the task is distributed, we need to characterize it at a different level of description.

Hutchins (following Marr [Mar82, pp. 24–27]) distinguishes the level of *computation*, the level of *algorithm*, and the level of *implementation*. A similar tripartite distinction is made by Zhang and Norman [ZN95, p. 290]. At the first level— computation— the system is characterized in terms of what it is expected to do. At the second level, the inputs and outputs are represented in a specific way and connected by a specific algorithm. At the third level— implementation— the algorithm is physically realized in some system.

Consider, for example, a system that does arithmetic multiplication. At the computational level, it has normative answers defined for all possible input pairs of natural numbers. At the algorithmic level, the inputs and outputs might be represented in decimal; the calculation might be executed using the usual, grade school algorithm. At the level of implementation, the system is typically someone with a pencil and scratch paper. Take me, for instance. I cannot, without pencil and paper, calculate 372×821 . Using the algorithm, I can multiply 372 by 8, record the result, multiply 372 by 2, record the result at an offset from the first, and so on until I can add up the result. This breaks the problem into three less daunting problems. Each subproblem is further broken down by the algorithm— 372×8 involves multiplying 2 by 8, carrying a 1, multiplying 7 by 8, and so on. The algorithm reduces the problem to a stultifying succession of one-digit multiplication problems, seasoned with addition. The computation— the cognitive activity of multiplying three-digit numbers— is not something that happens in my head. Instead, it is performed by the system that includes me, the pencil, and the paper.¹ So the *distribution* in distributed cognition occurs only at the level of implementation. Nothing *in principle* precludes the algorithm being performed without pencil and paper, entirely in a person’s head.

Notice also that the computation could be made even more distributed than it is. I perform one-digit multiplication by using a lookup table that I memorized in grade school. Imagine instead that I have the multiplication table in front of me, such that I consult it to determine each sub-product. Or imagine that I do not even have the procedure memorized, such that I must consult a list of instructions at each step.

For our purposes, we only need to distinguish two levels. To borrow terms from McClamrock [McC91][McC95, §1.3], we can distinguish between the *task* and the *process*. The computational level specifies the task as “an abstract and idealized specification of the behavior to be achieved” [McC95, p. 20]. The lower

¹Multiplication is a stock example of distributed cognition; see McClelland and Rumelhart [MRtPRG86, pp. 44–48], Zhang and Norman [ZN95], and Giere [Gie02b, p. 288]. Zhang and Norman [ZN95, pp. 284–289] show how the cognitive advantages of the standard multiplication algorithm interact with the advantages of using arabic numerals.

levels specify the process for accomplishing the task.² This binary distinction will allow us to give a succinct definition of distributed cognition.

An activity counts as *d-cog* only if the process is not enclosed by the epidermis of the people involved in carrying out the task. The implementation uses tools and social structures to do some of the cognitive work. In the example of multiplication, the cognitive work is distributed between me and my scratch pad. On the navy ship, it is distributed among a team of individuals and a collection of tools.

D-cog might be met with this objection: Cognition happens in minds. Although we might speak of a group acting with ‘one mind,’ it cannot literally be true.³ Therefore, distributed cognition is ridiculous claptrap.

This objection could be answered in any of several ways. First, it could be argued that minds can belong to systems larger than an individual person; for an account along these lines, see [Cla97]. Second, we could say that ‘cognition’ in the sense used here does not require a mind; cf. [Gie02b, p. 295]. Third, we could say that d-cog does not involve genuine cognition—rather, it involves something akin to ‘cognition’ in important, specifiable respects.⁴ Since this paper is not an excursion in metaphysics, I will not try to choose between these three approaches. It suffices that the objection can be answered.

Here is another objection: The task/process distinction might be applied to phenomena that are not cognitive even in an extended sense. For instance, it is possible to give an abstract characterization of what needs to be done by the excretory system and a further account of how various organs and chemicals carry out this task. Yet we do not want to say that someone attached to a dialysis machine has a d-cog excretory system. We need criteria to distinguish d-cog activities from other activities that are amenable to functional explanation.

Considering the activities of an isolated organism, we can distinguish between the cognitive and non-cognitive tasks that it performs. The metabolic activity of its kidneys does not perform a cognitive task. We do not need a precise definition of a cognitive task in order to say this. It is enough that we can apply the concept uncontentiously in some specific cases.⁵ The limits of a finite, organic brain are not at issue in such judgments. We recognize thousand-digit arithmetic as a cognitive task, even though a *human* mind cannot implement it except approximately.

We can say that, in order for an activity to count as d-cog, the task must be such that *it would be cognitive* if the process were contained entirely within the epidermis of one individual. That is: It is the sort of task that *could* be carried out in a mind. So an activity is d-cog if (1) the task is such that it would count as cognition if it were carried out entirely in a single mind or brain, and (2) the

²McClamrock offers the task/process distinction as a common thread that connects Marr’s levels with similar distinctions made by Chomsky, Simon, and Pylyshyn.

³See [WGH85] for a brief history of *group mind* theories and their fall into disrepute.

⁴Depending on how the meaning of ‘cognition’ is determined, the second and third options might just be verbal variants.

⁵At the same time, the following discussion is meant to be compatible with any plausible definition of ‘cognition.’

process by which the task is carried out is not enclosed within the boundary of a single organism.

These criteria are jointly sufficient for an activity to be d-cog. They do not provide a definition. Imagine there were a task which *could not* be performed by a single mind— even in principle— but which uncontestedly counted as cognitive. I would have no objection to calling the performance of that task d-cog. However, I know of no such tasks, and I suspect that tasks which seem to be like this really are not. One might look around a carpentry shop, for example, and see signs of distributed cognition: numbers on machine parts, color coding, and so on.⁶ Of course, the task of making a cabinet could not be carried out inside a single mind— but neither is cabinet making the cognitive task implemented by the numbers, colors, and so on. A workman must service the parts of the machine in a specific order. Remembering the appropriate order is cognitive, and it could be done in the mind of an attentive workman. The numbers mean that he does not need to remember the order himself. The proper order is produced by the numbers along with his ability to recognize numerals and his proceeding in numerical order. Similar analyses might be given for the color coding of tools, and so on.

The criteria I have suggested will apply to systems in which the process spans a single person using some instruments, in which the process spans several people, and in which the process spans both people and instruments. For example, some varieties of *transactive memory* will count as d-cog. Wegner et al. [Weg86] [WGH85] describe phenomena such that an intimate couple is capable of storing and recalling information that neither could recall separately. In some cases, this may best be treated as two feats of individual cognition— two individuals performing distinct but related cognitive tasks. Yet in other cases they describe, such an atomic description is inadequate; the best description treats the couple as implementing a single, unified cognitive task.

Informally, a d-cog account describes the distributed implementation of a cognitive task. So the d-cog category only applies relative to a cognitive task. This feature will be important when we try to describe science as d-cog.

3 Science as d-cog

Although the relationship between science and its social context is an issue of much debate, certainly *some* social context is required for science. Modern science could not practicably be developed by a solitary genius. So it may be helpful to consider science as d-cog and apply the methods developed by Hutchins and others. Giere has recently argued that philosophers of science should begin to think about science in terms of d-cog. This would, he suggests, “provide a useful framework for thinking about cognition in the sciences generally” [Gie02b, p. 286]. (See also [Gie02a].) As an example, he considers data gathered at the Indiana University Cyclotron Facility. Seeing the activity there as d-cog, he says, “[T]he data is being gathered by a complex cognitive system

⁶This example was suggested to me by Orin Zelenak.

consisting of the accelerators, detectors, computers and all the people actively working on the experiment” [Gie02b, p. 292]. This can clarify the way that instrumental and social structures are integral to the production of the data: A great deal of the computation happens without the explicit attention of any person.

Drug trials provide another promising example. Double-blind trials rely on a network of agents each with partial information, so as to mitigate unconscious bias and placebo effects. Considering a specific drug trial, we might say that following double-blind methods was the task and that protocols, doctors, nurses, patients, and numbered vials constituted the process by which the task was implemented. The organizers of the trial would be rightly chastised if the process failed to implement this task. Yet the trial described in this way would not count as d-cog, given the criteria I offered in the previous section. Double-blind studies *ipso facto* cannot be carried out within a mind. They *require* a social structure.

One might object that this only shows that studies so-described are not d-cog *by my criteria*, and it does not show that they are not d-cog *tout court*. Indeed, my criteria are meant to be sufficient but not necessary for d-cog. Nevertheless, I contend that double-blind studies described in this way should not count as d-cog. The task of performing a double-blind study, as the task of not revealing to the agents all of what is going on, is not especially interesting. It becomes a kind of parlor game: Each of the partially-informed agents is engaged in their own cognition, and the best account of what is happening will treat them as distinct individuals. Treating a double-blind study in this way, merely as a game of subterfuge, would miss the point of the study. We do not chastise studies for failing to follow double-blind methods because such methods are intrinsically good, but because the methods contribute to the reliability of a study.

We can give a different account of the very same drug trial by saying that the task is to discover whether the drug is efficacious or has side effects. The double-blind trial figures as the process for carrying out this task. The task of discovery could be performed without double-blind methods, and the explanation counts as d-cog. This brings out four important features of d-cog accounts.

First, as noted above, a d-cog account always requires the specification of a task. It is impossible to describe a process *qua* implementation without first specifying the task which the process was supposed to carry out. To describe a system as d-cog thus involves specifying the function that the system is supposed to perform.

Second, task specifications allow us to say when something has *gone wrong*.⁷ When you consider me as part of a system aimed at multiplying numbers—rather than drawing pictures or writing stories—then you can look over my

⁷Giere writes: “Thinking in terms of cognitive systems provides a natural way to incorporate *normative* considerations into a cognitive theory of science. Any individual cognitive system is designed to do a particular job. It can be evaluated as to how well it does that job” [Gie02b, p. 295]. He adds: “The main point I want to make here is that the evaluation of a cognitive system is largely an *empirical* matter. It is a matter for empirical investigation to determine how well a particular cognitive system does what it is supposed to do” [Gie02b, p. 295, italics in original]. It is unclear whether he means “empirical” in the sense of descriptive (in contrast with prescriptive) or in the sense of tangible (in contrast with speculative).

shoulder and see if I got the right answer. When we appreciate the dangers of experimenter bias, then we can appreciate how double-blind methods contribute to reliability.

Third, whether a system counts as d-cog is relative to the task specification. If the task is specified as conducting a double-blind trial, the drug trial is not d-cog; if the task is specified as determining how safe the drug is, then the trial is d-cog. In a strict sense, d-cog is a predicate that applies to a system under a description rather than to a description alone.

Fourth, picking out the right task description can sometimes be a tricky business. With the drug trial, it would not be *false* to say that a double-blind trial was the task—it would merely be infelicitous. This complicates the normative value of the d-cog account; if there are rival conceptions of task, then there may be rival assessments of whether the system is doing well.

What Giere has in mind is that we could repeat the discussion of the cyclotron and drug trials for other scientific projects *mutatis mutandis*. There is a wealth of literature on the rôle of instrumentation and social networks in science, much of which could be framed in the d-cog idiom. This might, as Giere suggests, be a fruitful enterprise. It would banish a seductive individualism that treats scientists as cognizing atoms, thinking their way through a world of intricate widgets. Or would it?

One might reply: Even if it could be shown that science *as practiced* is often or usually instantiated by d-cog, this might be an accidental feature of it. We define multiplication for all pairs of numbers even though actual calculators are limited to numbers of a certain size. Similarly—the reply continues—cognitive atomism is still the best way to think about science *as such*.

This reply cannot be dismissed out of hand; any d-cog account describes a task that might be carried out in an undistributed way. Yet it would be wrong to think that d-cog is merely an accidental feature of our science. Hippocratic angels might be able to perform unblinded drug trials, but *for us* double-blind methodology was an important development. It is not merely a curious custom of our pharmaceutical practices. D-cog accounts given in other domains might reveal important features about *our* science, about the best way we have of learning about the world.

4 SCIENCE! writ large

Case studies can only be about ‘science generally’ by being about a great number of particular scientific projects. We might use d-cog for thinking about ‘science generally’ in a different sense, by asking about *the scientific enterprise* as one enormous cognitive system. Consider this bold suggestion: Can we understand science altogether as one giant, distributed cognitive enterprise?

To my knowledge, no proponent of d-cog has explicitly posed this question. In the following sections, I’ll explore the suggestion by reflecting on an account of science given by Robert Merton and two given by Philip Kitcher.

4.1 Merton: The ethos of science as distributed cognition

In a widely-influential 1942 essay,⁸ Robert Merton portrays science aimed at truth and, although he would not have put it this way, he provides a d-cog account of it. Mertonian *norms* play the rôle of an algorithm for getting the truth, since (Merton thinks) the norms have both a moral force and a methodological rationale. Although Merton is typically criticized for misidentifying the norms of science, he identifies the norms that he does because he conceives of the aims of science in the way that he does. The inadequacy of his list of norms suggests reasons to think that he misidentifies the goal of science; science cannot aim at truth *simpliciter*.

Merton characterizes science as aimed at “the extension of certified knowledge” [Mer73a, p. 270]—that is, at the discovery of truth. I do not think Merton means to be merely stipulating that science aims at the extension of certified knowledge; nor does he intend it as an empirical claim. He bases the claim on his understanding of methodology, and he is not careful to defend it because his work is “an essay in the sociology of science, not an excursion in methodology” [Mer73a, p. 268]. This conception of the aim of science is sufficiently mundane that Merton need not have acquired it from particular sources. Indeed, later in his life, Merton suggested that it came from nowhere in particular [personal communication, 1998].

In popular myth, this extension of knowledge is performed by scientists selflessly pursuing truth. They will need to work together, of course, but only because there is so much to be learned—each step-wise extension of knowledge is secured by an individual who hopes to learn something. In this myth, scientists are distinguished as people who aim at truth rather than at personal gain, and science is the aggregate of those individual cognitive systems.

Merton rejects the myth. Accounting for scientific practice, he insists, does not require thinking that scientists are men of special character or that they are epistemic saints. Merton writes:

A passion for knowledge, idle curiosity, altruistic concern with the benefit of humanity, and a host of other special motives have been attributed to the scientist. The quest for distinctive motives appears to have been misdirected. It is rather a distinctive pattern of institutional control of a wide range of motives which characterizes the behavior of scientists. [Mer73a, p. 276]

On Merton’s account, scientists are more or less like anyone else, except that they are members of a scientific community— a community structured so as to fulfill the ends of science. An explanation of the scientific enterprise will thus not be in terms of lone scientists selflessly pursuing the truth, but rather in terms of communities and institutions.

⁸The essay appeared first as ‘Science and Technology in a Democratic Order’ [Mer42], subsequently as ‘Science and Democratic Social Structure’ [Mer68a], and later as ‘The Normative Structure of Science’ [Mer73a].

Merton offers what has become the canonical list of Mertonian *norms* or *mores*: universalism, common ownership of knowledge,⁹ disinterestedness, and organized scepticism. These norms promote, but are not strictly identical with, the extension of knowledge. Merton writes, “The mores of science possess a methodologic rationale but they are binding, not only because they are procedurally efficient, but because they are believed right and good” [Mer73a, p. 270]. He points to the goal of science, describes moral norms that are supposed to promote this goal, and concludes with the claim that scientific institutions impress these norms on the psychology of individual scientists.¹⁰ This moves between levels of description. At the highest level, science is characterized as an enterprise executing a specific function. At a lower level, Merton describes a procedure for executing that function: Conform to the norms, he suggests, and knowledge will be extended. Further down, Merton describes the procedure implemented in institutions and individuals: Scientists, socialized to value the norms, go about acting in accord with them.

We can parse this as a d-cog account. The task of science, Merton thinks, is extending knowledge. The process which does this involves many individual sciences motivated to pursue an array of norms.¹¹ The scientists do not self-consciously and directly do the task of science; they follow the norms. When scientists individually follow the norms, the community implements a process that performs the task of science.

These features of Merton’s account are usually overlooked. As a result, although many objections have been directed at Merton’s specific list of norms, few if any objections have been raised against his identification of the goal of science. In part, this is because Merton’s work is most often read as straightforward functionalist sociology. Even some otherwise savvy readers of Merton misunderstand him on this point, e.g. [Kin71]. Merton’s account is not functionalist in the usual sense. Consider a typical example of functional analysis, like the explanation of a rainmaking ceremony as a mechanism to promote social solidarity. A feature of a society is explained by the fact that it fulfills a need of the society, where a *need* is a prerequisite for social or biological survival.¹² For functionalist sociology, identifying a practice as functional for us *ipso facto* gives us a reason to adopt it. In the account of science, however, the function of science is specified independently of its promoting the fitness of any scientific community. Since the norms are identified only as executing the function of science, the norms of science might prove to be harmful; the function of science might be *dysfunctional* for a society. In fact, Merton specifically addresses potential conflicts between the functioning of science and the functioning of society at large.

⁹He bills this as *communism*, but carefully notes that this should be understood in a mundane, non-Bolshevik sense [Mer73a, p. 273].

¹⁰The same progression can be seen in other essays from the period; e.g. [Mer73b, esp. p. 258, fn. 15].

¹¹We might parse the account in a different way: If the task is to follow certain norms, then the process is the structure of institutions and psychology of scientists. This would also be a d-cog account. In general, three levels can be parsed in a number of ways; see [McC91].

¹²This follows Merton’s own characterization of functional analysis; see [Mer68b, p. 106].

Fortunately, since the goal of finding truth implicates certain norms, some criticisms of Merton’s list of norms reflect back on the goal itself. First, as Barnes and Dolby [BD70, pp. 10–11, 17] have argued, the concern for certified knowledge downplays the rôle of provisional acceptance and working hypotheses. In *certifying knowledge*, one must apply strict standards; but in *exploring an hypothesis*, one might follow hunches and wild guesses. Thus, Merton’s conception of science directs him to emphasize rigorous or sceptical aspects of scientific practice over risky or speculative aspects. Merton concludes that organized scepticism is a norm of science, but such a norm is already implicit in the end he sets for science. Second, there is no room in Merton’s conception to ask whether the new knowledge is *significant* or whether it is anything we care about, because science aims at the extension of knowledge *tout court*. There is no discrimination as to what the knowledge is about, and science as an institution disregards interests. Merton concludes that disinterestedness is a moral norm of science, but here again it seems to be smuggled in with his conception of science.

We might attempt to address these concerns by saying that science is not merely about finding truth, it is about finding *significant truths*.¹³ In the next two sections, I consider Philip Kitcher’s accounts of science that explicitly see it aimed at significant truth.

4.2 20th-century Kitcher: The distribution of cognitive labor

So there is a motivation to say that science aims at *significant truth*, where significance is a shared sense that points the scientific community to investigate some things rather than others. In *The Advancement of Science* [Kit93], Philip Kitcher identifies the function of science in this way.¹⁴ He argues that progress in science is possible, as judged by what he calls the *external standard*: Changes in scientific practice are rational if and only if they increase our prospects of acquiring significant truth [Kit93, p. 189].

As a contingent matter, we would be ill-advised to try and satisfy the external standard by self-consciously implementing it. If individual scientists independently did their own best to discover the significant truth about some matter,

¹³One might think that Merton’s “certified knowledge” is not truth *tout court*, but only truth about observables. Scientific knowledge, Merton writes, means “empirically confirmed and logically consistent statements of regularities”; the scientific demand for evidence promotes “sustained true prediction” [Mer73a, p. 270]. Merton might well sympathize with van Fraassen’s anti-realist insistence that “the defining criterion of success [in science] is empirical adequacy” [van03, p. 486]. I pass over the point, because reading Merton as a constructive empiricist changes neither of the features which interest me here. First, the explanatory structure is the same: Science is a d-cog enterprise in which the norms promote the function of science and in which socialization serves to implement the norms. Second, the problem that arose for truth also arises for empirical adequacy. There are just as many empirically adequate statements as there are truths and just as many *useless* ones. If the realist is forced to say that *significant truth* is the genuine aim of science, then the constructive empiricist is forced to say that *significant* empirically adequate claims are the genuine aim.

¹⁴A similar account of the aim of science is given by Haack [Haa98].

there would be egregious duplications of effort and disregard for less-promising research programmes. The problem requires a social solution.

One might imagine a solution in which labor is distributed by self-conscious coordination among scientists, but Kitcher suggests that there is a more practical solution that would do almost as well. Given the proper community structure, the self-interest of scientists would result in a nigh optimal distribution of labor [Kit92][Kit93, ch. 8]. If scientist Jones wants credit for making the discovery, he will be tempted to pursue the less promising line of research. The chance that the research will pay off is lower, but, since fewer other scientists will be competing with him, the chance that it will pay off *for him* is higher. Treating scientists as self-interested agents in the manner of classical economics, Kitcher provides schematic examples in which the equilibrium solution approximates the best solution.

The self-conscious solution is like the myth of scientist-saints; it supposes high-minded scientists in a utopian social arrangement. Kitcher's solution is rather like Merton's account; scientists are socialized into the right kind of community for producing knowledge.¹⁵ And it is d-cog. No individual would plan it, and no one even need have imagined it. The *community* would be rational in the sense of the external standard, even though no individual would be.

4.3 21st-century Kitcher: Well-ordered science

In more recent work, Kitcher has argued that there is no context-independent measure of significance [Kit01, ch. 6]. Which truths count as significant for us depends on our interests and projects. Our interests even serve to pick out which language we speak and, thus, which truths are available to us [Kit01, ch. 4].¹⁶ So Kitcher now thinks there is a more serious coordination problem: How should science be organized so that it asks the right questions, pursues answers in responsible ways, and fairly applies the resulting knowledge?

Kitcher imagines a process of ideal deliberation. Members of the community learn about the possibilities, express their priorities, and share their concerns. The course of science is then charted in light of the group's 'tutored preferences' [Kit01, ch. 10]. Importantly, Kitcher does not say that the aim of science is to carry out some version of this ideal process. Rather, Kitcher coins the phrase 'well-ordered science' to describe any science pursued in the way that *would* have been selected by the ideal process. The task of science is to be well-ordered, and it becomes an open question as what process would do best at carrying out the task. In principle, at least, it might be accomplished by a central planner giving orders to every scientist— that is, it might be a matter of explicit, individual cognition. It might instead be a matter of the community being organized in a

¹⁵Strevens [Str03] provides an analysis of the value of priority in science that follows the same pattern. He presumes that scientific knowledge is a social good and then shows how emphasis on first discovery can promote that end.

¹⁶Note that our interests do not determine what is true, according to Kitcher. Within a given language, truth is made by the way the world is.

certain way, such that citizens and scientists go about their own business but produce a total constellation of behavior that is well-ordered. This would be d-cog.

5 Scientific practices

The examples of Merton and Kitcher reveal that it is a tricky thing to specify the task of science-as-such. The question of whether we should give a d-cog account of it is a subsequent thorny issue. Yet the conception of science as d-cog is not necessarily esoteric or inert. Although we might disagree with Merton about the task of science, his insistence that scientists need not be epistemic saints is unquestionably an important insight. Just as double-blind trials are an important feature of our medicine, institutions are an important feature of our science.

Once we do try to specify the task of science, however, we are embrangled in difficulties. If some of these seem insurmountable, perhaps this is because there is no account to be given of science altogether. Just as there is no single task performed by *cognition*, there is no single task performed by *science*. It would be nice to conclude this paper with the correct account of the task of science, but I have no such account to offer. Although science is a collection of interrelated practices, I doubt that there is any substantive way to characterize the task of science-as-such.

We can always retreat to the position sketched out at the end of §3. It was easy enough to specify the task of a navy ship or drug trial. Specific scientific activities are characterized by specific tasks, too.

In recent work [Gie03], Giere gives the example of the Hubble Space Telescope. The telescope itself is part of a larger system that includes multiple communications satellites, facilities in New Mexico and Maryland, numerous computers, and— to simplify somewhat— we can take the “the cognitive output of this whole process to be just the claim that the image indicates the existence of galaxies 13 billion years ago.”

The space telescope looks off into the distance, making it functionally like our eyes or like the telescopic sights aboard navy ships. The sights on the ship are part of a system that determines the position of the ship, and Hubble is part of a system that determines facts about galaxies 13 billion years ago. It should be amenable to cognitive analysis, after the pattern of Hutchins’ analysis of shipboard navigation.

Yet consider another of Giere’s examples: What is the task performed by a cyclotron? We can describe the physical operation of a cyclotron and— insofar as something is not a working ‘cyclotron’ if it does not operate in this way— we have normative leverage for saying when and how well the cyclotron works. But this functional account is not d-cog; the mere physical operation of a cyclotron is not its cognitive task. The cognitive business of a cyclotron facility is various and changing.

The difficulty that arises is the same difficulty that arose in considering the

carpentry shop as d-cog. In order to understand a d-cog account, we need to attend to the right task. Consider work by Nancy Nersessian and her colleagues [NKMND03][Ner05]. Nersessian et al. studied researchers who worked with devices which were meant to simulate blood vessels. One such device is called the “construct.” In a deliberately provocative turn of phrase, Nersessian et al. describe the construct as a “mental model.” It is important that ‘mental’ not be taken as meaning ‘within a mind.’ Although we may interpret d-cog as involving extrapersonal minds (recall §2, above), that is not their point; cf. [NKMND03, p. 4, fn. 3]. Rather, they mean to underscore the similarity between the rôle that the construct plays in the cognitive business of the lab and the rôle that familiar mental models play in individual cognition. Is this d-cog? If the task of the lab is *studying properties of the construct*, then the device is an object in the environment that the researchers explore.¹⁷ Nersessian et al. argue against understanding the lab in this way. The construct is studied, but it is also put into interaction with other devices and revised over time so as to be more like a blood vessel. The task is better described as *understanding and simulating blood vessels*. The task of simulating blood vessels is cognitive, and it could in principle be done all in the mind of a single agent. In the lab that Nersessian et al. describe, it is done by a team of researchers using the construct and other devices. With the task specified in this way, the work in the lab is d-cog.

In this paper, I have tried to clarify some features of the d-cog approach. Whether we think of science altogether or piecemeal, d-cog requires us to specify a task that the system is supposed to perform. Where a task can be specified, d-cog is a promising approach. Where a task cannot be specified, d-cog is a non-starter.

References

- [BD70] S.B. Barnes and R.G.A. Dolby. The scientific ethos: A deviant viewpoint. *Archiv. europ. social.*, XI:3–25, 1970.
- [Cla97] Andy Clark. *Being There: Putting Brain, Body, and World Together Again*. The MIT Press, Cambridge, Massachusetts, 1997.
- [Gie02a] Ronald N. Giere. Discussion note: Distributed cognition and epistemic cultures. *Philosophy of Science*, 69(4):637–644, December 2002.
- [Gie02b] Ronald [N.] Giere. Scientific cognition as distributed cognition. In Peter Carruthers, Stephen Stich, and Michael Siegal, editors, *The Cognitive Basis of Science*, pages 285–299. Cambridge University Press, 2002.
- [Gie03] Ronald N. Giere. Computation and agency in scientific cognition. In *Proceedings of the 25th Annual Conference of the Cognitive*

¹⁷Since the researchers work as a team, the implementation might still be d-cog— but it will not be d-cog on account of interaction with the construct.

- Science Society*, 2003. Boston, Massachusetts. July 31– August 2.
- [GM03] Ronald N. Giere and Barton Moffatt. Distributed cognition: Where the cognitive and the social merge. *Social Studies of Science*, 33(2):1–10, April 2003.
- [Haa98] Susan Haack. *Manifesto of a Passionate Moderate*. University of Chicago Press, 1998.
- [Hut95] Edwin Hutchins. *Cognition in the Wild*. The MIT Press, Cambridge, Massachusetts, 1995.
- [Kin71] M.D. King. Reason, tradition, and the progressiveness of science. *History and Theory*, X(1):3–32, 1971.
- [Kit92] Philip Kitcher. Authority, deference, and the role of individual reason. In Ernan McMullin, editor, *The Social Dimensions of Science*, pages 244–271. University of Notre Dame Press, 1992.
- [Kit93] Philip Kitcher. *The Advancement of Science*. Oxford University Press, 1993.
- [Kit01] Philip Kitcher. *Science, Truth, and Democracy*. Oxford University Press, 2001.
- [Mar82] David Marr. *Vision*. W.H. Freeman and Company, San Francisco, 1982.
- [McC91] Ron McClamrock. Marr’s three levels: A re-evaluation. *Minds and Machines*, May 1991.
- [McC95] Ron McClamrock. *Existential Cognition: Computational Minds in the World*. University of Chicago Press, 1995.
- [Mer42] Robert K. Merton. Science and technology in a democratic order. *The Journal of Legal and Political Sociology*, 1:115–126, 1942.
- [Mer68a] Robert K. Merton. Science and democratic social structure. In *Social Theory and Social Structure* [Mer68b].
- [Mer68b] Robert K. Merton. *Social Theory and Social Structure*. The Free Press of Glencoe, revised and enlarged edition, 1968.
- [Mer73a] Robert K. Merton. The normative structure of science. In Storer [Sto73], pages 267–73.
- [Mer73b] Robert K. Merton. Science and the social order. In Storer [Sto73], pages 254–266.

- [MRtPRG86] James L. McClelland, David E. Rumelhart, and the PDP Research Group. *Parallel Distributed Processing*, volume 2. The MIT Press, Cambridge, Massachusetts, 1986.
- [Ner05] Nancy J. Nersessian. Interpreting scientific and engineering practices: Integrating the cognitive, social, and cultural dimensions. In M. E. Gorman, R. D. Tweney, D. C. Gooding, and A. Kincannon, editors, *Scientific and technological thinking*, pages 17–56. Lawrence Erlbaum Associates, Mahwah, NJ, 2005.
- [NKMND03] Nancy J. Nersessian, Elke Kurz-Milcke, Wendy C. Newstetter, and Jim Davies. Research laboratories as evolving distributed cognitive systems. In R. Alterman and D. Kirsh, editors, *Proceedings of the Twenty-Fifth Annual Conference of the Cognitive Science Society*, 2003.
- [Sto73] Norman W. Storer, editor. *The Sociology of Science*. University of Chicago Press, 1973.
- [Str03] Michael Strevens. The role of the priority rule in science. *The Journal of Philosophy*, 100(2):55–79, 2003.
- [van03] Bas C. van Fraassen. On McMullin’s appreciation of realism concerning the sciences. *Philosophy of Science*, 70(3):479–492, July 2003.
- [Weg86] Daniel M. Wegner. Transactive memory: A contemporary analysis of the group mind. In B. Mullen and G. R. Goethals, editors, *Theories of group behavior*, pages 185–208. Springer-Verlag, New York, 1986.
- [WGH85] Daniel M. Wegner, Toni Giuliano, and Paula T. Hertel. Cognitive interdependence in close relationships. In W.J. Ickes, editor, *Compatible and incompatible relationships*, pages 253–276. Springer-Verlag, New York, 1985.
- [ZN95] Jiajie Zhang and Donald A. Norman. A representational analysis of numeration systems. *Cognition*, 57:271–295, 1995.