François, K., Löwe, B., Müller, T., Van Kerkhove, B., editors,
Foundations of the Formal Sciences VII
Bringing together Philosophy and Sociology of Science

Science as socially distributed cognition: Bridging philosophy and sociology of science

Matthew J. Brown*

School of Arts & Humanities, The University of Texas at Dallas, 800 W Campbell Road, JO31, Richardson, TX 75080, United States of America.

E-mail: mattbrown@utdallas.edu

In this paper, I want to make the following claim plausible:

> Analyzing scientific inquiry as a species of socially distributed cognition has a variety of advantages for science studies, among them the prospects of bringing together philosophy and sociology of science.

This is not a particularly novel claim; indeed, Paul Thagard has been suggesting something like this for well over a decade, while philosophers like Ronald Giere and Barton Moffat have been stumping for the distributed cognition approach in more recent years, and Nancy Nersessian's Cognition and Learning in Interdisciplinary Cultures research group at the Georgia Institute of Technology has been fruitfully applying this approach to the study of research laboratories and other scientific institutions.

I will retrace some of the major steps that have been made in the pursuit of a distributed cognition approach to science studies, paying special attention to the promise that such an approach holds out for bridging the rift between philosophy and the social studies of science. Such an approach is not without its pitfalls, and I will consider several problems, both for distributed cognition as a theory and for its applications to science. I will argue that there is a path out of the woods, and try to point the way. Ultimately, I argue that we shall have to widen the scope of the distributed-cognition approach.

*Many thanks to Michael Cole, Edwin Hutchins, Nancy Nersessian, and Paul Churchland for encouraging and discussing this work. Thanks also to P. D. Magnus, for helpful discussions about d-cog and science, to the members of the UCSD D-Cog/HCI Lab for allowing me to present an earlier version of this project, and to the organizers of the *Foundations of the Formal Sciences VII* conference, especially Benedikt Löwe, who encouraged and helped me to make the long journey from San Diego to Brussels.

Received by the editors: 30 January 2009.
Accepted for publication: 20 August 2009.

1 What is d-cog?

Distributed cognition (d-cog) is a radical theory in cognitive science, created primarily by researchers at the University of California, San Diego (UCSD), which maintains that one can fruitfully analyze activities taking place between one or more people along with technological artifacts as *cognitive* in the same way that traditional cognitive science has analyzed certain intrapersonal processes. The beginnings of the approach can be seen in the *Parallel Distributed Processing* research group on connectionist (neural-network) models of cognition. When it came to an explanation of how a neural-network architecture can do science, mathematics, and logic, they made an intriguing suggestion:

> If the human information-processing system carries out its computations by "settling" into a solution rather than applying logical operations, why are humans so intelligent? How can we do science, mathematics, logic, etc.? How can we do logic if our basic operations are not logical at all? We suspect the answer comes from our ability to create artifacts—that is, our ability to create physical representations that we can manipulate in simple ways to get answers to very difficult and abstract problems. (Rumelhart et al., 1987, p. 44)

This is quite the break from classical cognitive science research in two ways. First, cognitive science has traditionally treated "cognition" as a matter of computational operations on symbolic states, not unlike the operations of logic or the architecture of an ordinary computer. The move towards a connectionist architecture, where the basic computational processes are more like pattern-recognition than applying logical rules, is the radical step that the *Parallel Distributed Processing* group was most keen to argue for. Second, the cognitive sciences ordinarily focus on what goes on with an individual person, and "cognition" is what goes on in their head. It is this second break that spurs d-cog.

D-cog takes a wider perspective than classical cognitive science. It is concerned with use of "cognitive artifacts" such as pen and paper, longhand mathematical calculation, and digital computers. It is also interested in the cognitive role of social interactions, cultural institutions, and forms of social organization. D-cog attempts to move the boundaries of our concept of "cognitive activity" out from the head and into the world of social and technological interactions. The foundational text for research in distributed cognition is *Cognition in the Wild,* the work of another UCSD cognitive scientist, Edwin Hutchins, who was trained in cognitive anthropology (Hutchins, 1995a).

There are two ways we might understand the project of distributed cognition research. A more cautious definition—the one preferred by Giere, for example[1]—would be that some socially and technologically distributed activities can profitably be understood as "cognitive", while allowing that many elements of cognition—including agency—remain in the realm of the individual. The more radical definition of d-cog that one could adopt, and the one most supported by Hutchins's own statements, is more sweeping; as Hutchins says, "I hope to show that human cognition [. . .] is *in a very fundamental sense* a cultural and social process" (1995a, p. xiv, emphasis mine).

Two examples have become pervasive in papers about d-cog.[2] Nevertheless, it is worth briefly discussing each. The first originates in (Rumelhart et al., 1987). When we multiply large numbers, we rarely if ever do it in our heads. With a pencil and paper, multiplying even very large numbers is transformed into a simple task, requiring no more than the ability to do one-digit multiplication and addition. A nearly impossible task for an individual human cognitive system becomes perfectly easy when distributed across the human-pencil-paper system (cf. Figure 1).

```
   12044
x    432
--------
   24088
  36132
 48176
--------
 5203008
```

FIGURE 1. Longhand Multiplication

A second key example comes from Hutchins's work on ship navigation in the US Navy (1995a). Navigation on a large naval vessel is not the job of a single individual (as it is for Hutchins's contrast case of the Micronesian canoer), but rather the work of a team of people performing various jobs using various instruments. Here is a somewhat simplified account: A crewman (called a pelorus operator) is given a landmark to identify by a plotter in the pilothouse. The pelorus operator then uses a piece of equipment called an "alidade" to determine the bearing of the landmark. Generally, there

[1] Giere (2002a, p. 6); Giere (2002b, p. 238); Giere (2002c, p. 295); Giere (2006, pp. 112–114); Giere and Moffatt (2003, p. 4).
[2] Cf. Hutchins (1995a); Giere and Moffatt (2003); Magnus (2007).

is more than one pelorus operator, and they all relay their information to a plotter. The plotter uses that information to determine the location of the ship and its bearing. The plotter relies on a specially structured map, compasses and protractors, etc. in order to use the information about the bearing of landmarks to compute the ship's location and bearing.

In this example, it is physically impossible for a single human being (given the size of the ship, the location of various vantage points, and the time in which the task must be completed) to do the cognitive work of figuring out the location and bearing of the vessel. Of course, on a *different* kind of ship, it is possible for a single person, and indeed, in the case of the Micronesian navigator, it is possible for a lone individual to do so without instrumentation. Nevertheless, the navigation team on a large naval vessel completes the cognitive task as a team using artifacts. The essence of a d-cog analysis is in treating this network of individuals and artifacts as a single cognitive system.

2 Science as d-cog

Can science be analyzed using d-cog? Consider a case discussed by Giere and Moffatt (2003), originally due to Dumas (1834) as discussed by Klein (1999). Chemical formulae were originally introduced by Jacob Berzelius in 1813 (Klein, 2001, p. 7). A Berzelian formula like the one in Figure 2 allows one to do theoretical chemistry by manipulating such symbols on paper, replacing the need to directly manipulate chemicals. All one needs to do in order to determine what is going on in a reaction, knowing something about the products and the reactants, is assume conservation and balance the equation. Just as doing long multiplication by hand transforms a complex calculation into a set of simple pattern-matching problems, so the use of chemical formulae as a cognitive artifact transforms the complex theoretical or experimental analysis into a simple exercise in pattern matching (cf. Figure 2).

$$\mathrm{C_8H_8 + H_4O_2 + Ch_4 = C_8H_8O_2 + Ch_4H_4}\ [\textit{sic}]$$

FIGURE 2. Chemical Formula for reaction of alcohol and chlorine according to Dumas (1834)

Consider another example, the Hubble Space Telescope, an important piece of scientific equipment in contemporary research in astronomy, astrophysics, and cosmology (Giere, 2006, p. 99–100). The telescope is a large and complex instrument that must be operated remotely. It is used by an organized group of people, and that use is mediated by further instruments and computer equipment on earth. To draw out the sense in which d-cog analysis is appropriate, think of the telescope as the eyes of a large cognitive system that also includes the group of scientists and the earth-bound computer equipment. Just as cognitive science can study ordinary perception, distributed cognitive science can look at this distributed system of "perception."

A final example, due to Nancy Nersessian and her collaborators at Georgia Tech (Nersessian et al., 2004) comes from the cognitive-ethnographic study of research in a biomedical engineering laboratory. Nersessian discusses how certain lab equipment are used to *model* actual biological processes. For example, the lab she studies uses a piece of equipment called a "bioreactor," which, among other things, models blood flow in a way which can be used to study arterial cells and biomedical devices. Nersessian's work explicitly treats the bioreactor, along with the skills one needs in order to use it in certain ways as a "mental model" for the distributed system of the biomedical laboratory. Doing so reveals interesting facts about the system that aren't available if you treat it just like a device or an instrument.[3]

3 The cognitive and the social

Latour and Woolgar (1986) issued their famous "ten-year moratorium on cognitive explanations of science," promising "that if anything remains to be explained at the end of this period, we too will turn to the mind!" (Latour and Woolgar, 1986, p. 280; quoted by Giere and Moffatt, 2003, p. 301). Of course, the moratorium has run out, much remains to be explained, and they never turned solely to the mind to provide the missing explanations—but that's not the point. What is interesting are the motivations and implicit assumptions behind this rhetorical flourish.

Part of the reason they issued such a moratorium, as Giere and Moffatt (2003, p. 301), Nersessian (2005, p. 18), and others have argued, is that they held to a rigid dichotomy of cognitive and social factors. Because their primary goal was to get a serious sociology of science going, they regarded such a moratorium as necessary. In order to make room for social explanations of science, we must, they thought, bracket all cognitive issues and explanations.

D-cog provides an alternative to this way of thinking. It shows us how to treat the cognitive and the social as the same thing for certain purposes. Because cognitive structures need not exist only in the mind (and perhaps never do so, if the radical version of d-cog is correct), but instead can exist in the complex interactions of social groups and technological artifacts, one can study social groups cognitively, or cognitive systems sociologically. There need be no unbridgeable divide between social and cognitive explanations.[4]

[3] I will return to this case in further detail below, to indicate some of the major gains of such an analysis.

[4] If I read him correctly, Bruno Latour has come around to this more sophisticated

What's most interesting about the possibility of seeing the social in terms of the cognitive and vice versa is that it might just help heal the wounds of the Science Wars and bring the various parts of science studies which are often at loggerheads—especially philosophy and sociology of science—together towards a more common purpose. Because of the perceived incompatibility of the cognitive and the social, the terms of analysis of much recent sociology of science—negotiation, authority, power, mobilizing resources—seem to have a cynical cast, dismissive of the virtues of science. By contrast, the normative concerns of philosophers of science—justification, realism, objectivity—seem divorced from the obvious social reality of science.

There are plenty of philosophers nowadays—such as Helen Longino (2002a; 2002b; 2002c) and Philip Kitcher (2001, 2002a,b)—trying to reconcile the cognitive and the social, the normative issues of philosophy of science with descriptive sociological analyses. Their arguments are mainly about the very possibility of such a reconciliation, and focus more on the reformulation of traditional philosophical issues (e.g., objectivity) in ways that involve social relations and institutions rather than focusing on the properties of individual scientists or the *abstract structure* of science. D-cog presents more than the mere possibility of an in-principle or a post-hoc reconciliation. It allows one to fruitfully re-interpret the excellent and extensive body of sociological and historical studies in line with cognitive and epistemic concerns.

Consider again the case of chemical formulae. Bruno Latour (1986) has emphasized the importance of such innovations in the history of science. According to Giere and Moffatt, Latour thinks that something like chemical formulae are important because they concentrate information in a way that

> confers authority and power on those who control it. And it leads others to align themselves with such powers, thus increasing still further their authority and power. In a struggle for dominance, whether in science, politics, or war, those with the most and strongest allies win. (Giere and Moffatt, 2003, p. 305)[5]

What d-cog allows Giere and Moffatt to do is to look at the specifics of Latour's analysis of the social-technological aspects of science and point out the *cognitive* function of various parts of the process. What might be cast

view of the cognitive-social relation. Cf. Latour (1991, 1996).

[5]Having read Latour, this seems a slight exaggeration of his point. Already he seems to recognize the d-cog perspective when he says: "An average mind or an average man, with the same perceptual abilities, within normal social conditions, will generate totally different output depending on whether his or her average skills apply to the confusing world or to inscriptions" (Latour, 1986, p. 22). I take this to imply that inscriptions function as a cognitive artifact that change the functioning of the mind independent of the particular agent's perceptual abilities.

by a sociologist in terms of exerting power and gaining allies can be cast in terms of improving cognitive capacities of a distributed system over a naked cognitive agent. As Giere and Moffatt say about this particular case,

> The invention of new forms of external representation and of new instruments for producing various kinds of representations has played, and continues to play, a large role in the development of the sciences. From a cognitive science perspective, both sorts of invention amount to the creation of new systems of distributed cognitive system. So, for us, the notion of distributed cognition brings under one category such things as Cartesian coordinates and the telescope, both of which are widely cited as major contributions to the Scientific Revolution. (Giere and Moffatt, 2003, p. 305)

Even more promising than the idea of reinterpreting sociological and historical work in cognitive-epistemic terms is empirical work being done by cognitive scientists, philosophers of science, and researchers in science studies *using* the methods and theoretical frameworks of d-cog to analyze science. To return to another example from above, Nersessian and her collaborators have been studying work in a biomedical engineering laboratory, applying cognitive, historical, and ethnographic methods and understanding the organization of the lab and the function of artifacts within the lab as parts of a distributive cognitive system. Such an analysis allows researchers to understand how a bioreactor is both a "significant cultural artifact... [and] a locus for social interaction" (Nersessian, 2005, p. 50) with a history of different kinds of roles in the culture of the laboratory and also as a model that plays a role in distinctive types of representation and reasoning. Without such an analysis, the fact that a single object plays both of these roles (and the pervasiveness of such objects in science) is a colossal coincidence and a total mystery. One might even be lead to deny that the object has an important cognitive side (as sociologists of science are often led to do) or to claim that its cultural history and social roles are inessential to its role in representation and reasoning (as philosophers have often done). D-cog analysis makes better sense of what is going on in such cases, and makes better sense of how the social and the cognitive are integrating in science as a whole.

4 Challenges

Applying Hutchins's d-cog theory to the study of science is not without its problems. I will focus on two major challenges to the applicability of the theory.

The first problem is that d-cog looks like a theory applicable to fairly static systems.[6] The paradigm applications of d-cog in Hutchins's work

[6] I believe this problem was first pointed out to me by Yrjö Engeström in conversation.

(1995a; 1995b)—airplane cockpits determining their speed, crews of Navy ships navigating through a harbor, even pencil-and-paper multiplication—respond to dynamic situations where "problem-solving situations change in time" (Nersessian, 2005, p. 36), but the organization and nature of the technological artifacts in play are treated as static. The analysis is thus "dynamic but largely *synchronic*" (Nersessian, 2005, p. 36). But really, these systems are evolving, if slowly, both from external pressures (invention of new technology, new safety protocols, changes in policy) and internal developments (shifts in Navy culture, new pilots gaining skills, invention of new techniques). Further, many other kinds of cognitive activities are much more *diachronically* dynamic, involving creativity, innovation, and rapid changes in technology and social structure.[7] This is especially true of systems like scientific laboratories, where innovation, new discovery, and creative problem-solving are essential parts of the activity. Another aspect of this problem is that d-cog analyses tend to treat relatively well-bounded systems, with low-bandwidth information flow from outside the system and high-bandwidth information flow within the system. In order to straightforwardly apply Hutchins's d-cog framework, the nature of the task at hand and the system that carries it out must be rather well-bounded. On the other hand, many activities, including scientific activities, have quite vague and porous boundaries. What counts as part of the system might change rapidly as the activity goes on.

The second problem comes from a direct critique of Giere's appeals for treating science as d-cog by Magnus (2007). Magnus's critique turns on a particular move in Hutchins's (1995a) account of d-cog, where he relies on the tripartite distinction from Marr (1982) between *computation*, *algorithm*, and *implementation* (Hutchins, 1995a, pp. 50–52), which Magnus simplifies into the distinction between *task* and *process*,[8] where *task* is an abstract description of the computational goal or behavior that the cognitive system is to satisfy, and the *process* is just a specification of how the task is to be accomplished. This furnishes Magnus with a compellingly succinct definition of d-cog:

> 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. (Magnus, 2007, p. 299)

Where the task in question "would be cognitive if the process were contained entirely within the epidermis of one individual" (Magnus, 2007, p. 300) .

[7] "Although there are loci of stability, during problem-solving processes the components of the systems undergo development and change over time" (Nersessian, 2005, p. 36).

[8] He is following Ron McClamrock (1991); cf. also McClamrock (1995, § 1.3).

So, an activity is d-cog only if the process is not located inside the skin of an individual carrying out a cognitive task. Is science like that? It is easy enough to see that on Magnus's interpretation, if we are going to be able to analyze scientific activity as a species of d-cog, we must be able not only to analyze the scientific *process*, but we must also be able to specify the *task* of science. This poses two types of concern. First, at a local level, can we always abstractly specify a *task* for science? Does a biomedical engineering laboratory have a well-specified task? Does a physics journal? What about a conference on global warming? While it seems likely that there are some scientific activities which might be amenable to such an analysis, it seems dubious that one could specify the kind of *computational* task necessary for d-cog analysis for all or even most scientific activities.

The second worry that Magnus raises is whether one can specify a *global* task for science, and thus do a d-cog analysis of science "writ large". That is, "Can we understand science altogether as one giant, distributed cognitive enterprise?" Such an interpretation is already suggested by Hutchins (1991, p. 288). It would be a lucky thing if we can do so, for we could then give a clear explanation of the common view in science studies that it is the large-scale institution of science, rather than individual scientists, which produce or are responsible for scientific knowledge. To this end, Magnus analyzes three candidates for giving a *task* analysis for science-as-a-totality: Merton's *ethos* of science, Philip Kitcher's ideal of the distribution of cognitive labor, and his more recent image of well-ordered science.

As you might imagine, the prognosis is dire; Magnus is rightfully pessimistic about the possibility of specifying the task of science-as-such. After all, the range of activities of science, the differences in approaches in different research traditions, the variety of uses to which science is put, and so on make it highly unlikely that there is one simple task that all of science aims at. One need only compare high-energy physics to molecular biology to pharmaceutical trials to see how unlikely such a project is.[9] Given the poor prospects of rescuing d-cog analyses of science in this way, I will suggest we look elsewhere.

5 Prospects for a d-cog theory of science

Here's where we stand: d-cog holds great promise for analyzing science in a way that makes the relation of the social-technological nature of science to its cognitive-epistemic virtues most perspicuous, and thus joining together what the Science Wars hath put asunder, of healing the rift between philosophy and sociology of science. However, using d-cog to analyze science faces some severe difficulties: it treats systems whose basic structures and

[9]Cf. Knorr-Cetina (1999); Giere (2002a).

resources are fairly static, while science is not only synchronically but diachronically dynamic. Scientific systems evolve, and d-cog provides little in the way of resources for analyzing that evolution. D-cog applies to well-bounded systems, whereas the boundaries of science aren't so clear. D-cog analysis requires the specification of a computational task that can be implemented by a distributed process, while it seems doubtful that a global task can be specified for science, and even unlikely that a more local task can be specified for many important cases. Does this spell doom for the d-cog approach to science studies? Is there any way forward?

I think there is, and that way depends most importantly on going well beyond Hutchins's work from the mid-1990s. Of course, Hutchins himself is an active researcher and has gone beyond that work himself, into areas like conceptual change, learning, and the embodiment of cognition. Likewise, Nersessian, for example, relies on d-cog, but has taken it beyond Hutchins's original formulations. There are also traditions and research programs related to d-cog, such as neo-Vygotskian psychology, cultural-historical activity theory, and situated action theory, that have things to offer a *broadly* d-cog account of science. On the basis of the criticisms so far discussed, I will conclude by indicating the ways we must modify our understanding of d-cog in order for it to have positive prospects as an account of science.

The first important point to make, as against Magnus's interpretation and some of Hutchins's formulations of d-cog, is that cognition is *not* computation (in the classical sense). Certainly, computation is one kind of thing that cognitive agents and cognitive systems do, but it isn't the case that cognition is identical to computation.[10] Cognition is not a single algorithm or program, though it may use algorithms. Human cognitive capacities at their best are flexible and responsive to particular situations, creative and dynamic. Cognition is a multi-purpose capacity in humans, and likewise in any other sort of cognitive system.[11] While this may be a controversial point in some circles in cognitive science, those circles are shrinking precipitously. It is hard today not to agree with the point that was radical when proposed by the *Parallel Distributed Processing* group decades ago, that human cognition bears little if any resemblance to classical computation.[12]

While the approach may be controversial, there may be some valuable lessons to be learned from cultural-historical activity theory (CHAT) for

[10] Alternatively, as an anonymous referee pointed out to me, one might regard this as a radical extension of the idea of *computation*, rather than a denial of the identity of cognition and computation.

[11] Even those who regard cognition as having a modular architecture must admit that the human cognitive system at large is a complicated, multi-purpose, dynamic, and flexible system.

[12] Even if one does not accept *Parallel Distributed Processing*-type models.

providing a d-cog analysis of science.[13] CHAT provides a tripartite distinction between *operations, actions,* and *activities* that adds a useful layer to the talk about task vs. process. Operations are the basic components of actions; they are generally routinized human behaviors or mechanical operations, carried out under certain conditions, instrumental to engaging in some action. Actions are conscious, goal-directed processes, undertaken by individuals or small groups. For example, Leont'ev (1978, p. 66) describes learning to drive a car with a manual transmission. At first, all the processes of driving the car—breaking, using the clutch, shifting gears—require conscious attention. They are the focal, goal-directed activities. For the accomplished driver, these processes become unconscious, subordinated to actions like speeding up, driving up a steep incline, driving to work. In the end, the unconscious operations are actually off-loaded to a machine, the automatic transmission.

Beyond the level of action is the activity. Actions are goal-directed, relatively short-lived and well-bounded in time and space. Activities exist in and evolve over longer periods of time; they have a history. They are associated with a culture or a community, and they are often embedded in institutions or forms of social organization. While actions are simply goal-directed, activities are aimed at a more general, less-bounded, and changeable object or motive. While the particular actions of a welder in a factory have a quite well-defined goal (joining two metal pieces together), the activity of the whole factory has a more nebulous object of gaining profit, and the way that motive is conceived over time may change (for example, a change to a more socially-conscious, "green" corporate mission may alter both the ways that profit is got and the way that gaining profit is understood). The object of the activity system need not be at all available to the individual members of the system; indeed, the workers need not have any ideas about the economic purposes of the factory—they need only be in it to get a paycheck for themselves.

This set of distinctions may prove fruitful for thinking about science as d-cog. In particular, the *task* as Magnus (2007) seems to understand it seems identical to the *goal* to which actions are directed. The task-process distinction may thus make perfect sense at the level of action, but to get the whole sense in which science is a d-cog *activity*, we may need to think of it at the level of activity directed at an object which is partially constituted by the evolution of the activity itself.

Another potentially necessary turn is to supplement Hutchins's cognitive ethnography with Nersessian's cognitive-historical method. Nersessian is keenly aware of the problem of evolving systems for Hutchins's (1995a;

[13]Cf. Leont'ev (1978); Engeström (1987); Cole and Engeström (1993); Cole (1988); Engeström et al. (1999).

1995b) account, especially as applied to science. Indeed, the argument that Hutchins's account does not naturally accommodate the evolution of cognitive systems I gave above is her argument. In their own d-cog research on biomedical labs, Nersessian and her collaborators combine ethnographic investigation of the particular system with *cognitive-historical* analysis, which looks at different scales of history to understand the evolution of problems, concepts, cognitive artifacts, etc. (Nersessian et al., 2004).

So, *is* science a distributed cognitive system? This has been challenged on the basis of it being an evolving, messy, less-bounded system. Magnus (2007) has challenged it on the basis of whether there is a particular task that science carries out. But *what is a cognitive system* anyhow, even in the traditional sense of "cognitive system?" This shouldn't stand or fall on the details of a certain framework of cognitive analysis. After all, 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. Certainly, the limitations of a particular approach to d-cog shouldn't disqualify the more general notion. Rather, this points the way towards the need for better, more complex models of distributed cognition that might do a better job of applying to science. I have gestured towards some possibilities that seem particularly fruitful in the face of these difficulties. There is much more work to be done, and the possibilities are inspiring.

Bibliography

Cole, M. (1988). Cross-cultural research in the sociohistorical tradition. *Human Development*, 31:137–151.

Cole, M. and Engeström, Y. (1993). A cultural-historical interpretation of distributed cognition. In Salomon, G., editor, *Distributed cognitions. Psychological and educational perspectives*, Learning in doing. Social, cognitive, and computational perspectives, pages 1–46, Cambridge. Cambridge University Press.

Dumas, J. B. (1834). Recherche de chimie organique. *Annales de Chimie de Physique*, 56:113–150.

Engeström, Y. (1987). *Learning by expanding: An activity-theoretical approach to developmental research.* Orienta-Konsultit, Helsinki.

Engeström, Y., Miettinen, R., and Punamäki, R.-L. (1999). *Perspectives on activity theory.* Learning in doing. Social, cognitive, and computational perspectives. Cambridge University Press, Cambridge.

Giere, R. (2002a). Distributed cognition in epistemic cultures. *Philosophy of Science*, 69:637–644.

Giere, R. (2002b). Models as parts of distributed cognitive systems. In Magnani, L. and Nersessian, N., editors, *Model-based reasoning: Science, technology, values*, pages 227–241, Dordrecht. Kluwer.

Giere, R. (2002c). Scientific cognition as distributed cognition. In Carruthers, P., Stitch, S., and Siegal, M., editors, *Cognitive bases of science*, pages 285–299, Cambridge. Cambridge University Press.

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

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

Hutchins, E. (1991). The social organization of distributed cognition. Perspectives on socially shared cognition. In Resnick, L. B., Levine, J. M., and Teasley, S. D., editors, *Perspectives on socially shared cognition*, volume XIII of *APA Science Volume Series*, pages 283–307, Washington, DC. American Psychological Association.

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

Hutchins, E. (1995b). How a cockpit remembers its speeds. *Cognitive Science*, 19:265–288.

Kitcher, P. (2001). *Science, truth, and democracy*. Oxford University Press, Oxford.

Kitcher, P. (2002a). Reply to Helen Longino. *Philosophy of Science*, 69(4):569–572.

Kitcher, P. (2002b). The third way: Reflections in Helen Longino's *The fate of knowledge*. *Philosophy of Science*, 69(4):549–559.

Klein, U. (1999). Techniques of modelling and paper-tools in classical chemistry. In Morgan, M. S. and Morrison, M., editors, *Models as mediators: Perspectives on natural and social science*, pages 146–167, Cambridge. Cambridge University Press.

Klein, U. (2001). Berzelian formulas as paper tools in early-nineteenth century chemistry. *Foundations of Chemisty*, 3(1):7–32.

Knorr-Cetina, K. (1999). *Epistemic cultures: How the sciences make knowledge*. Harvard University Press, Cambridge, MA.

Latour, B. (1986). Visualization and cognition: Thinking with eyes and hands. *Knowledge and Society*, 6:1–40.

Latour, B. (1991). *Nous n'avons jamais été modernes. Essai d'anthropologie symétrique*. La Découverte, Paris. Page numbers refer to the English translation (Latour, 1993).

Latour, B. (1993). *We have never been modern*. Harvard University Press, Cambridge, MA.

Latour, B. (1996). On interobjectivity. *Mind, Culture, and Activity*, 3(4):228–245.

Latour, B. and Woolgar, S. (1986). *Laboratory life*. Princeton University Press, Princeton, NJ. Second Edition. Introduction by Jonas Salk. With a new postscript by the authors.

Leont'ev, A. N. (1978). *Activity, consciousness, and personality*. Prentice Hall, Englewood Cliffs, NJ.

Longino, H. E. (2002a). *The fate of knowledge*. Princeton University Press, Princeton.

Longino, H. E. (2002b). Reply to Philip Kitcher. *Philosophy of Science*, 69(4):573–577.

Longino, H. E. (2002c). Science and the common good: Thoughts on Philip Kitcher's *Science, truth, and democracy*. *Philosophy of Science*, 69(4):560–568.

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*. Plenum, New York.

McClamrock, R. (1991). Marr's three levels: A re-evaluation. *Minds and Machines*, 1:185–196.

McClamrock, R. (1995). *Existential Cognition. Computational Minds in the World*. University of Chicago Press, Chicago.

Nersessian, N., Kurz-Milcke, E., Newstetter, W. C., and Davies, J. (2004). Research laboratories as evolving distributed cognitive systems. In Alterman, R. and Kirsch, D., editors, *Proceedings of the 25th annual conference of the Cognitive Science Society. July 31st – August 2nd, 2003. Boston, Massachusetts, USA*, pages 857–862. Taylor & Francis.

Nersessian, N. J. (2005). Interpreting scientific and engineering practices: Integrating the cognitive, social, and cultural dimensions. In Gorman, M. E., Tweney, R. D., Gooding, D. C., and Kincannon, A. P., editors, *Scientific and technological thinking*, pages 17–56, Mahwah NJ. Lawrence Erlbaum.

Rumelhart, D. E., Smolensky, P., McClelland, J. L., and Hinton, G. E. (1987). Schemata and sequential thought processes in PDP models. In Rumelhart, D. E. and McClelland, J. L., editors, *Parallel distributed processing*, volume 2, pages 395–412, Cambridge, MA. MIT Press.