15 Scientific cognition as distributed cognition

Ronald Giere

After introducing several different approaches to distributed cognition, I consider the application of these ideas to modern science, especially the role of instrumentation and visual representations in science. I then examine several apparent difficulties with taking distributed cognition seriously. After arguing that these difficulties are only apparent, I note the ease with which distributed cognition accommodates normative concerns. I also present an example showing that understanding cognition as distributed bridges the often perceived gap between cognitive and social theories of science. The paper concludes by suggesting some implications for the history of science and for the cognitive study of science in general.

1 Introduction

Pursuing the philosophy of science within the context of philosophy has long been justified on the grounds that scientific knowledge is the best example of knowledge there is. It is KNOWLEDGE WRIT LARGE, as the saying goes. And epistemology is one of the main areas of philosophical inquiry. Philosophers of science have tended to regard the relationship as asymmetrical. Philosophy of science illuminates the problems of epistemology, but not much the other way around.

One can now say something similar about cognition in general. Science provides arguably the best example of a higher cognitive activity. Since the emergence of cognitive science as a recognizable discipline, however, the relationship has definitely been symmetrical, if not somewhat asymmetrical in favour of the cognitive sciences. A number of philosophers of science have explicitly drawn on developments in the cognitive sciences in attempting to illuminate the practice of science (Nersessian, 1984b, 1992a, 1999; Giere, 1988, 1999a; Thagard, 1988, 1992, 1999; Churchland, 1989; Darden, 1991; Bechtel, 1996a). In this paper I wish to continue this tradition by suggesting that some new developments within the cognitive sciences provide a useful framework for thinking about cognition in the sciences generally. These developments may be classified under the title of *distributed cognition*.

2 Distributed cognition

The idea of distributed *processing* has long been a staple in computer science. A dramatic contemporary example is the project *SETI at Home*, in which currently a million and a half participants contribute spare time on their personal computers to analyse data from the Arecibo Radio Telescope in Puerto Rico as part of a Search for Extra-Terrestrial Intelligence. Distributed *cognition* is an extension of the basic idea of distributed processing. I will focus on two of several contemporary sources of the notion of distributed cognition within the cognitive sciences.

2.1 Hutchins' cognition in the wild

One source for the concept of distributed cognition within the cognitive sciences is Ed Hutchins' study of navigation in his 1995 book, *Cognition in the Wild*. This is an ethnographic study of traditional 'pilotage', that is, navigation near land as when coming into port. Hutchins demonstrates that individual humans may be merely components in a complex cognitive system. No one human could physically do all the things that must be done to fulfil the cognitive task, in this case repeatedly determining the relative location of a traditional navy ship as it nears port. For example, there are sailors on each side of the ship who telescopically record angular locations of landmarks relative to the ship's gyrocompass. These readings are then passed on, e.g., by the ship's telephone, to the pilot-house where they are combined by the navigator on a specially designed chart to plot the location of the ship. In this system, no one person could possibly perform all these tasks in the required time interval. And only the navigator, and perhaps his assistant, knows the outcome of the task until it is communicated to others in the pilothouse.

In Hutchins' detailed analysis, the social structure aboard the ship, and even the culture of the U.S. Navy, play a central role in the operation of this cognitive system. For example, it is important for the smooth operation of the system that the navigator holds a higher rank than those making the sightings. The navigator must be in a position to give orders to the others. The navigator, in turn, is responsible to the ship's pilot and captain for producing locations and bearings when they are needed. So the social system relating the human components is as much a part of the whole cognitive system as the physical arrangement of the required instruments. One might say that the social system is part of the overall cognitive system.

Now one might treat Hutchins' case as an example of 'socially shared cognition' (Resnick, 1991) or, more simply, *collective cognition*. The cognitive task – determining the location of the ship – is performed by a collective, an organized group, and, moreover, in the circumstances, could not physically be carried out by a single individual. In this sense, collective cognition is ubiquitous in modern societies. In many work-places there are some tasks that are clearly cognitive and, in the circumstances, could not be carried out by a single individual acting alone. Completing the task requires co-ordinated action by several different people. So Hutchins is inviting us to think differently about common situations. Rather than simply assuming that all cognition is restricted to individuals, we are invited to think of some actual cognition as being distributed among several individuals.

How does this differ from the simple 'pooling' of knowledge possessed by several individuals, an already well-known phenomenon? Why introduce the new notion of 'distributed cognition'? Part of the answer is that categorizing some activities as a type of cognition is a more fruitful way of thinking about them. Among other things, it brings them within the scope of cognitive science. Another consequence of thinking in terms of distributed cognition is a focus on the *process* of acquiring knowledge rather than the static possession of knowledge. The knowledge acquired by the members of the crew is not 'pooled' in the sense of their simply bringing their individual bits of knowledge together. Rather, the individual bits of knowledge are acquired and co-ordinated in a carefully organized system operating through real time.

Hutchins' conception of distributed cognition, however, goes beyond collective cognition. He includes not only persons but also instruments and other artefacts as parts of the cognitive system. Thus, among the components of the cognitive system determining the ship's position are the alidade used to observe the bearings of landmarks and the navigational chart on which bearings are drawn with a ruler-like device called a 'hoey'. The ship's position is determined by the intersection of two lines drawn using bearings from two sightings on opposite sides of the ship. So parts of the cognitive process take place not in anyone's head but in an instrument or on a chart. The cognitive process is distributed among humans and material artefacts.

(See also Latour, 1986.)

The standard view, of course, has been that things such as instruments and charts are ‘aids’ to human cognition which takes place only in someone’s head. But the concept of an ‘aid to cognition’ had remained vague. By expanding the concept of cognition to include these artefacts, Hutchins provides a clearer account of what things so different as instruments and charts have in common. They are parts of a distributed cognitive process.

2.2 The PDP research group

Hutchins comes to distributed cognition through anthropology and ethnography. Another source comes through the disciplines usually regarded within the core of cognitive science: computer science, neuroscience, and psychology. This is the massive, two-volume ‘Exploration in the Microstructure of Cognition,’ titled simply *Parallel Distributed Processing* produced by James McClelland, David Rumelhart and the PDP Research Group based mainly in San Diego during the early 1980s (McClelland and Rumelhart, 1986). Among many other things, this group explored the capabilities of networks of simple processors thought to be at least somewhat similar to neural structures in the human brain. It was discovered that what such networks do best is recognize and complete *patterns* in input provided by the environment. The generalization to human brains is that humans recognize patterns through the activation of prototypes embodied as changes in the activity of groups of neurons induced by sensory experience. But if something like this is correct, how do humans do the kind of *linear* symbol processing apparently required for such fundamental cognitive activities as using language and doing mathematics?

Their suggestion was that humans do the kind of cognitive processing required for these linear activities by creating and manipulating *external representations*. These latter tasks *can* be done well by a complex pattern matcher. Consider the following simple example (1986, Vol. 2, pp.44-8). Try to multiply two three-digit numbers, say 456 × 789, in your head. Few people can perform even this very simple arithmetical task. Figure 1 shows how many of us learned to do it.

Insert Figure 1 about here

This process involves an *external representation* consisting of written symbols. These symbols are manipulated, literally, by hand. The process involves eye-hand motor co-ordination and is not simply going on in the head of the person doing the multiplying. The person's contribution is (1) constructing the external representation, (2) doing the correct manipulations in the right order, and (3) supplying the products for any two integers, which can be done easily from memory.

Notice again that this example focuses on the *process* of multiplication; the task, not the product, and not knowledge of the answer. Of course, if the task is done correctly, one does come to know the right answer, but the focus is on the *process* rather than the *product*. The emphasis is on the *cognitive system* instantiating the process rather than cognition simpliciter.

Now, what is the cognitive system that performs this task? Their answer was that it is not merely the mind-brain of the person doing the multiplication, nor even the whole *person* doing the multiplication, but the *system* consisting of the person *plus* the external physical representation. It is this whole system that performs the cognitive task, that is, the multiplication. The cognitive process is distributed between a person and an external representation.

Here is a more complex example making a similar point. The diagram in Figure 2 embodies a famous proof of the Pythagorean Theorem. The area of the large square is $(a + b)^2$. That of the small square is c^2. Remembering that $(a + b)^2 = a^2 + b^2 + 2ab$, and that the area of a triangle is ½ab, we see that $(a + b)^2$ is also equal to $c^2 + 2ab$, from which it follows that $a^2 + b^2 = c^2$.

Insert Figure 2 about here

The claim is that this reasoning is perfectly valid and that the diagram is essential to the reasoning. One literally *sees* that $(a + b)^2 = c^2 + 2ab$. So, here again, we seem to have a case where the cognitive task involves an external representation as an essential part of completing the task successfully. This is an example of what is now often called 'diagrammatic reasoning' (Chandrasekaran *et al*., 1995). It is also a case of what the logicians Jon Barwise and John Etchemendy (1996) call 'heterogeneous inference' because it involves *both* linguistic and visual representations.

2.3 Other sources

Although I have focused on just two sources of recent interest in distributed cognition, the idea is much more widespread, as one can learn from Andy Clark's recent *Being There: Putting Brain, Body, and World Together Again* (1997). Lucy Suchman's *Plans and Situated Action* (1987) and *The Embodied Mind* by Varela *et al*., (1993) are among earlier influential works. A common theme here is that the human brain evolved primarily to co-ordinate movements of the *body*, thereby increasing effectiveness in activities such as hunting, mating, and rearing the young. Evolution favoured cognition for effective action, not for contemplation. *Cognition is embodied*. This point of view argues against there being a single central processor controlling all activities and for there being many more specialized processors. Central processing is just too cumbersome. Thus, an emphasis on distributed cognition goes along with a rejection of strongly computational approaches to cognition. In fact, some recent advocates of Dynamic Systems Theory have gone so far as to argue against there being any need at all for computation as the manipulation of internal representations (Thelen and Smith, 1994).

Clark, in particular, invokes the notion of 'scaffolding' to describe the role of such things as diagrams and arithmetical schemas. They provide support for human capabilities. So the above examples involving multiplication and the Pythagorean Theorem are instances of scaffolded cognition. Such external structures make it possible for a person with a pattern matching and pattern completing brain to perform cognitive tasks it could not otherwise accomplish.

The most ambitious claim made in behalf of distributed cognition is that *language* itself is an elaborate external scaffold supporting not only communication, but thinking as well (Clark, Ch. 10). During childhood, the scaffolding is maintained by adults who provide instruction and oral examples. This is seen as analogous to parents supporting an infant as it learns to walk. Inner speech develops as the child learns to repeat instructions and examples to itself. Later, thinking and talking to oneself (often silently) make it seem as though language is fundamentally the external expression of inner thought, whereas, in origin, just the reverse is true. The capacity for *inner* thought expressed in language results from an internalisation of the originally *external* forms of representation. There is no 'language of thought'. Rather, thinking in language is a manifestation of a pattern matching brain trained on external linguistic structures (see also Bechtel, 1996b).

This distributed view of language implies that cognition is not only embodied, but also *embedded* in a society and in a historically developed culture. The scaffolding that supports language is a cultural product (Clark, 1997). Interestingly, an externalised and socialized view of language was advocated in the 1920s by the Soviet psychologist, Lev Vygotsky (Vygotsky, 1962, 1978; Wertsch 1985). A product of the intellectual ferment inspired by the Russian Revolution, Vygotsky explicitly appealed to Marxist thought for the idea that the key to understanding how language developed lay not so much in the mind as in society. Nevertheless, his major 1934 book (translated in 1962 as *Thought and Language*) was suppressed within the Soviet Union from 1936 to 1956, and has only since the 1960s received notice in the English-speaking world.

A view similar to Vygotsky's, with no apparent connection, has recently been developed under the title of 'cognitive linguistics' or 'functional linguistics' (Velichkovsky and Rumbaugh, 1996). Like Vygotsky, some advocates of cognitive linguistics emphasize comparative studies of apes and humans (Tomasello, 1996). The most striking claim is that the major difference between apes and human children is *socialization*. Of course there are some differences in genotype and anatomy, but these are surprisingly small. The bonobo, Kanzi, raised like a human child by Duane Rumbaugh and Sue Savage-Rumbaugh (1993), is said to have reached the linguistic level of a two-year-old human child. Tomasello argues that what makes the difference between a natural and an enculturated chimpanzee is developing a sense of oneself and others as *intentional* agents. Natural chimpanzees do not achieve this, but enculturated ones can.

From the standpoint of computational linguistics, there has always been a question of how the neural machinery to support the necessary computations could possibly have evolved. From the standpoint of cognitive linguistics, this problem simply disappears. Language is not fundamentally computational at all, but the product of a pattern matching neural structure, which biologically could evolve, supported by an elaborate scaffolding of social interaction within an established culture.

3 Distributed cognition in scientific research

I will now consider several examples for which the concept of distributed cognition provides resources for understanding aspects of scientific research (see also Bechtel, 1996a). I take these examples from my own earlier research partly because I am familiar with them and partly because I wish to avoid the risk of re-interpreting someone else's work in ways they might not appreciate.

3.1 Instrumentation

In the mid-1980s I spent considerable time observing the operations of the Indiana University Cyclotron Facility. This facility was designed for the investigation of nuclear structures using light ions such as hydrogen and helium to bombard heavy nuclei. At the time I was very impressed by the role of technological artefacts in scientific research, but, never having heard of distributed cognition, I did not know what to say about it. The best I could do, in just three pages of *Explaining Science* (1988, pp.137-40), was to suggest we think of technology as 'embodied knowledge'. Thus, what philosophers typically regard as 'background knowledge' or 'auxiliary assumptions' may not best be understood in terms of symbolic representations but as being physically encapsulated in experimental apparatus.

I took many photographs around the laboratory, many of them showing the instrumentation being employed. In fact, I sometimes took great pains to picture specimens of a type of instrument in isolation both from other instruments and from the human researchers. I now think that this orientation was completely mistaken. On the contrary, it is particularly enlightening to think of the whole facility as one big cognitive system comprised, in part, of lots of smaller cognitive systems. To understand the workings of the big cognitive system one has to consider the human-machine interactions as well as the human-human interactions.

In thinking about this facility, one might be tempted to ask, *Who* is gathering the data? From the standpoint of distributed cognition, that is a poorly framed question. A better description of the situation is to say that the data is being gathered by a complex cognitive system consisting of the accelerator, detectors, computers, and all the people actively working on the experiment.

Understanding such a complex cognitive system requires more than just enumerating the components. It requires also understanding the organization of the components. And, as with Hutchins' ship, this includes the *social* organization. It is not irrelevant to the operation of this cognitive system that the people monitoring the data acquisition are most likely the PhD faculty members who participated in designing the experiment. Those tending the detectors may have PhD's, but may well be full-time laboratory employees rather than faculty. Those keeping the accelerator in tune are probably technicians without PhD's. One cannot adequately understand the operation of the whole cognitive system without understanding these differences in roles and status among the human participants.

3.2 Visual representations

Another result of my time in the laboratory was an appreciation for how big a role *visual* representations play in the actual doing of science. Here again I did not at the time know what to say about this. I devoted just one page of *Explaining Science* (1988, p.190) to describing how physicists judge goodness of fit *qualitatively* by visually comparing two-dimensional curves representing, respectively, data and theoretical predictions. In fact, the laboratory had a light-table in a corner of the control room so one could visually compare curves and data points by superimposing one sheet of graph paper over another.

In a later publication on visual models and scientific judgment (Giere, 1996a), I speculated that one of the roles of visual representations is to help scientists organize their knowledge of the situation so as to be able to judge the fit between abstract models and the world. I now think that this suggestion erred in putting all of the cognitive activity in the heads of scientists. It is more enlightening to think of the scientist plus the external visual representation as an extended cognitive system that produces a judgment about the fit between an abstract model and the world. The visual representation is not merely an aid to human cognition; it is part of the system engaged in cognition.

3.3 Models and theories

For those philosophers of science who understand scientific theories to be axiom systems ideally reconstructed in a formalized language, there seems little new to learn from at least the computational parts of cognitive science. For those philosophers of science who, like myself, think of theories more in terms of *families of models*, this connection is not nearly so strong. But there is other work in cognitive science that is relevant, namely, that on concepts and categorization. This is because a model is a lot like a concept. Like concepts, models themselves make no claims about the world. But they can be *used* to make such claims. That is, both concepts and models can be applied to things in the world.

There is, however, a conflict between how most philosophers of science think of models and how most cognitive psychologists think of concepts. Philosophers typically hold a classical view of models in which a model can be defined in terms of a set of necessary and sufficient conditions. Cognitive psychologists typically hold a graded view of concepts in which there are no necessary and sufficient conditions for application, only more and less central cases (Smith and Medin, 1981).

My own resolution of this conflict is to use the classical account for individual models, but to realize that a family of classical models radiating out from a central model can produce a graded structure among the things to which various members of the family of models can be applied (Giere, 1994, 1999b).

There is another difference between the notion of model employed in the philosophy of science and the account of concepts in cognitive psychology. The idea of distributed cognition provides a way of resolving this difference. Almost all studies of concepts in psychology deal only with concepts like ‘bird’ or ‘chair’ which are simple enough to be deployed by normal humans without relying on external representations. Most models in science, even in classical mechanics, are too complex to be fully realized as mental models. Not even authors of science textbooks can have in their heads all the details of the models presented in their texts. Rather, the details of these models are *reconstructed* as external representations when needed. These reconstructions typically take the form of equations or diagrams. What scientists have inside their skins are representations of a few general principles together with bits and pieces of prototypical models. They also possess the *skills* necessary to use these internal representations to construct the required *external* representations that are then part of an extended cognitive system.

4 Objections

I will now consider two objections to taking seriously the notion of distributed cognition for understanding scientific research.

4.1 Individuation

Once the boundaries of a cognitive system are allowed to extend beyond the skin of a human being, where do they stop? I earlier claimed that the cyclotron facility should be regarded as a large cognitive system. Now this facility requires much electricity to operate. The electricity comes over wires from a generating plant, which is fired by coal, which is mined in the western United States. Are we to suppose that the boundary of this cognitive system thus extends all the way from Indiana to Montana? Hardly.

Here I think we can distinguish those features of the system that *differentially* influence the output of the system in scientifically relevant ways from those features that merely make it possible for the system to generate any output at all. The person at the control panel maintaining the coherence of the beam contributes to the quality of the output. The electricity makes it possible for the machine to run, and thus for there to be output, but it does not differentially influence the particular output obtained. Coal mines in Montana are thus not part of this cognitive system. On the other hand, if the data were to be transmitted directly to a computer at the University of Montana and returned after appropriate processing to the facility in Indiana, then the boundary of the system in Indiana would indeed extend to Montana.

This response can be generalized. Anything, anywhere that is *designed* to contribute to the specific quality of the output of a cognitive system, and functions as designed, is surely part of that system. I am inclined to go further and affirm that anything that *actually* contributes to the specific quality of the output of a cognitive system, whether or not it was designed to do so, and even if it is completely unknown to anyone, is part of that system. However, I do not think it is worth arguing over such niceties. The main point is that, for scientific purposes, one can *identify* particular cognitive systems in space and time. They are not unmanageably amorphous.

4.2 Mind

In countenancing extended cognition, are we not thereby committed to the existence of something like 'group minds'? And has not that notion long been discredited?

Now it is true that some advocates of distributed cognition are tempted by the notion of *distributed minds* (Clark, 1997, chapter 9). The presupposition behind this move seems to be that cognition is necessarily a property of minds. So, where there is cognition, there is also a mind. But I do not think that one is forced to take this route. I do not see why one cannot maintain an *ordinary* conception of mind, which restricts it to creatures with brains, that is, humans and perhaps some animals. So, much cognition does involve minds. Perhaps we can also agree that all of our currently existing cognitive systems have a human somewhere in the system. So every cognitive system has a mind somewhere. But as a concept in cognitive science, cognition need not be restricted exclusively to places where there are also minds. There is no reason why we cannot extend the *scientific* concept of cognition to places where there are no minds.

I do not think that the question of whether there are extended minds is one for which there is now a determinate answer. Rather, the question will be settled, if it is ever settled at all, by many individual decisions about the best way to develop a science of cognition.

5 Normative considerations

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. And evaluation may require consideration of tradeoffs among several desirable characteristics. A low energy cyclotron, for example, cannot tell us much about elementary particles, but it can provide high-resolution determinations of various characteristics of heavy nuclei.

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. Determining the desirable characteristics of any system is a matter of experienced professional judgment. There are no perfect cognitive systems and thus no absolute standards.

6 Between cognitivism and constructivism

The main competitor for a cognitive philosophy of science is some form of social constructivism. Constructivism is currently the dominant meta-theory of science in the sociology of science and constitutes a widespread orientation in the history of science. Theories of science emphasizing cognitive operations have been marginalized in the broader science studies community because so much of what seems important to understanding the workings of modern science takes place in a public arena, and not only in the heads of scientists. This includes particularly the technology employed in theoretical as well as in experimental contexts.

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

6.1 Epistemic cultures

That the concept of distributed cognition potentially provides a bridge across the gap between the cognitive and the social is illustrated in a recent book by one of the founders of social constructivist sociology of science, Karin Knorr-Cetina. In *Epistemic Cultures: How the Sciences Make Knowledge* (1999), Knorr moves beyond promoting social constructivism to examine once again just how scientific work produces knowledge. She answers that it is done partly through the creation of *epistemic cultures*, cultures that may be different in different sciences. Her primary example is the culture of High Energy Physics (HEP) at CERN. She argues that the scale of these experiments in time (several years) and numbers of participants (several hundreds) has produced a communitarian structure in which the individual is effectively replaced by 'the experiment' as the *epistemic subject*.

She even evokes the concept of distributed cognition, although tentatively and without any elaboration. For example, in the introduction, she writes (1999, p.25): In HEP, 'the subjectivity of participants is ... successfully replaced by something like distributed cognition.' Much later, when discussing the management of large experiments, she remarks:

> Discourse channels individual knowledge into the experiment, providing it with a sort of distributed cognition or a stream of

> (collective) *self-knowledge*, which flows from the astonishingly intricate webs of communication pathways. (1999, p.173)

Again, commenting on how experiments are named, she makes the point emphasized by Hutchins:

> Naming, then, has shifted to the experiment, and so has epistemic agency – the capacity to produce knowledge. The point is that no single individual or small group of individuals can, by themselves, produce the kind of results these experiments are after.... It is this impossibility which the authorship conventions of experimental HEP exhibit. They signify that the individual has been turned into an element of a much larger unit that functions as a collective epistemic subject. (1999, pp.167-8.)

In spite of the tentativeness with which she refers to distributed cognition, I think she has got it exactly right. In HEP, new knowledge is produced by distributed cognitive systems consisting of both humans and machines.

Knorr clearly assumes that if knowledge is being produced, there must be an epistemic subject, the thing that knows what is known. The standard assumption that permeates modern philosophy is that the subject is an individual human being. But Knorr's deep understanding of the organization of experiments in HEP makes this assumption problematic in that setting. Feeling herself forced to find another epistemic subject, she settles on the experiment itself. But could we not think of scientific knowledge more impersonally, say in the form: 'It has been scientifically established that p' or: 'These scientific experiments indicate that p.' These forms of expression free us from the need to find a special sort of epistemic subject that knows *p*. Individuals cannot *produce* the knowledge in question, but they can in some ordinary sense come to *know* the final result.

A similar problem arises for Knorr because the traditional epistemic subject is a *conscious* subject. But if we take the extended experiment to be the epistemic subject, do we have to introduce an extended form of consciousness as well? Knorr is tempted by this implication. Speaking of stories scientists tell among themselves, she writes:

> The stories articulated in formal and informal reports provide the experiments with a sort of consciousness: an uninterrupted hum of

> self-knowledge in which all efforts are anchored and from which new lines of work will follow. (1999, p.178)

And on the following page, she continues:

> Collective consciousness distinguishes itself from individual consciousness in that it is public: the discourse which runs through an experiment provides for the extended 'publicity' of technical objects and activities and, as a consequence, for everyone having the possibility to know and assess for themselves what needs to be done. (1999, p.179)

Again, Knorr is not alone in making such connections. Such speculations could be avoided, however, if we adopt a more impersonal attitude toward scientific knowledge.

7 Implications

7.1 Distributed cognition in the history of science

It is often claimed that the scientific revolution introduced a new way of thinking about the world, but there is less agreement as to what constituted the 'new way'. The historiography of the scientific revolution has long included both theoretical and experimental bents. Those on the theoretical side emphasize the role of mathematics, Platonic idealization, and thought experiments. The experimentalists emphasize the role of experimental methods and new instruments such as the telescope and microscope. Everyone acknowledges, of course, that both theory and experiment were crucial, but these remain a happy conjunction.

The concept of distributed cognition provides a unified way of understanding what was new in the way of thinking. It was the creation of new distributed cognitive systems. Cartesian co-ordinates and the calculus, for example, provided a wealth of new external representations that could be manipulated to good advantage. And the new instruments such as the telescope and microscope made possible the creation of extended cognitive systems for acquiring new empirical knowledge of the material world. From this perspective, what powered the scientific revolution was an explosion of new forms of distributed cognitive systems. There remains, of course, the historical question of how all these new forms of cognitive systems happened to come together when they did, but understanding the source of their power should now be much easier.

7.2 The cognitive study of science

Taking seriously the role of distributed cognition in science obviously has implications for the cognitive study of science as a whole. Here I will only note several of the most prominent. In general, insisting on a role for distributed cognition does not automatically invalidate other approaches. It does, however, call into question their often implicit claims to be the whole, or most of, the story about scientific cognition.

In a burst of enthusiasm for parallel distributed processing, Paul Churchland (1989, chapter 9) suggests that we 'construe a global theory as a global configuration of synaptic weights' (p.188) in a very high dimensional vector space. He seems to be claiming, for example, that the general theory of relativity can be represented in such a vector space. Now that might be a theoretical possibility, but it is almost certainly not how it is represented by real scientists. Rather, given our current understanding of distributed cognition, it is much more likely that the structure of scientific theories is a feature of *external* representations, such as systems of equations, diagrams and prototypical models. These are manipulated using the kind pattern matching brain Churchland describes. All that need be explicitly represented in the brain is bits and pieces of the external representations necessary to facilitate external manipulations. The idea of parallel distributed processing, however, retains its relevance for understanding how science works.

A decade ago, Paul Thagard (1992) attempted to explain scientific revolutions as due to the greater explanatory coherence of a new theory, where explanatory coherence depends on logical relations among statements of theories. He suggested that these relationships might even be instantiated in the minds of scientists, leading them to choose the new theory over the old. Missing from this picture, of course, is any place for the influence of a new form of external representation or new instrumentation. The whole dynamics of theory choice is reduced to logical relationships among statements.

One of the most active areas of the cognitive study of science is cognitive development. The underlying assumption, often explicit as in Gopnik (1996a) and Gopnik and Meltzoff (1997), is that what we learn about conceptual development in infants and children can be applied to conceptual development in science. If it is true, however, that a major factor in the development of modern science has been the invention of elaborate means both for creating external representations and for observing the world, then the cognitive development of children in less scaffolded conditions loses much of its relevance. There may be much to learn about human cognition in general by studying cognitive development in children, but not much to learn specifically about the practice of science (Giere, 1996b).

Acknowledgments

I thank Linnda Caporael and Nancy Nersessian for introducing me to the work of Ed Hutchins and others developing the concept of distributed cognition in the cognitive science community. Nersessian, in particular, shared with me her grant proposal on cognitive and social understandings of science, and even dragged me off to hear Ed Hutchins at a Cognitive Science Meeting.