



Distributed Cognition without Distributed Knowing

Ronald N. Giere

5

10

AQ1 *In earlier works (2002, 2006), I have argued that it is useful to think of much scientific activity, particularly in experimental sciences, as involving the operation of distributed cognitive systems, since these are understood in the contemporary cognitive sciences. Introducing a notion of distributed cognition, however, invites consideration of whether, or in what way, related cognitive activities, such as knowing, might also be distributed. In this paper I will argue that one can usefully introduce a notion of distributed cognition without attributing other cognitive attributes, such as knowing, let alone having a mind or being conscious, to distributed cognitive systems. I will first briefly introduce the cognitive science understanding of distributed cognition, partly so as to distinguish full-blown distributed cognition from mere collective cognition.*¹

15

20

AQ2 *Keywords:*

25

1. Distributed Cognitive Systems


A now canonical source for the concept of distributed cognition within the cognitive sciences is Ed Hutchins' (1995) study of navigation in his book *Cognition in the Wild*. This is an ethnographic study of traditional "pilotage"; that is, navigation near land, as when coming into port. Hutchins argues that individual humans may be merely components in a complex distributed 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 designated landmarks relative to the ship's gyrocompass. These readings are then passed on by the ship's telephone to the pilothouse, where they are combined

30

35

Ronald N. Giere is Professor of Philosophy Emeritus as well as a member and former Director of the Centre for Philosophy of Science at the University of Minnesota, USA. He is the author of *Understanding Scientific Reasoning* (5th edition, 2006), *Explaining Science: A Cognitive Approach* (1988), *Science Without Laws* (1999), and *Scientific Perspectivism* (2006). Correspondence to: Ronald N. Giere, Department of Philosophy, University of Minnesota, 749 Heller, Minneapolis, MN 55455, USA. Email: giere@umn.edu

40





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. Those recording the locations of landmarks have no reason to know the result of the process.

5

One might wish to treat Hutchins' case merely as an example of *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 workplaces 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 coordinated action, including sharing information, by several different people. So Hutchins is inviting us to think differently about many 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.

10

15

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 the 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.

20

25

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 has remained vague. By expanding the concept of cognition to include these artefacts, Hutchins provides a clearer account of what things as different as instruments and charts have in common. They are parts of a distributed cognitive process.

30

A crucial feature of distributed cognitive systems is that they contain "external representations"; that is, representations of aspects of the world that are not localized in a person's brain or in a computer, but somewhere external to these locations. Thus, for example, the telescopes used by Hutchins' sailors contain representations of the angular location of landmarks relative to the forward motion of the ship. Similarly, as soon as the readings regarding the location of the landmarks are recorded on the navigator's chart, it contains a representation of the location of the ship as of roughly one minute earlier. These external representations are created and manipulated by the human actors in the course of the operation of the whole distributed cognitive system.

35

40

2. The Problem of Epistemic Agency

What makes a distributed cognitive system "cognitive" is that it produces an output that, attributed to an individual person, would clearly be a cognitive achievement, such





as acquiring knowledge of where the ship was one minute ago. The problem of epistemic agency is whether we should say that possession of the resulting knowledge resides in the system as a whole or merely in the human components of the system. Further problems concern the ascription to the system as a whole of such attributes as having beliefs, being responsible for knowledge claims, having a mind, or being conscious, even self-conscious. I will now briefly review the claims of some advocates of distributed cognition who take these further steps and then argue that this is a mistake. We should, I conclude, restrict these further cognitive attributions to the human components of distributed cognitive systems.

5
10

3. Extending Epistemic Agency

Foremost among those who would extend cognitive agency to cognitive systems as a whole is Hutchins himself, who argues that we should regard *mind* as extended beyond the human body. There is mind at work, he has claimed, on the visible surface of the chart as the navigator and his assistant point to representations of various landmarks and decide which landmarks to use for the next sightings. Apparently Hutchins thinks that these decisions are to some extent literally made on the chart rather than just in the heads of the navigator and his assistant.²

15





Andy Clark's (1997) spirited advocacy of cognition as being distributed focuses on what I would call *locally distributed cognition*. A person working with a computer would be a paradigm example. A person with a computer can perform cognitive tasks, such as complex numerical calculations, that a person alone could not possibly accomplish as accurately or as fast, if at all. When it comes to arguing for extending the concept of *mind* beyond the confines of a human body, however, he invokes a more primitive example—that of a man with a defective memory who always keeps on hand a notebook in which he records all sorts of things that most of us would just remember (Clark 1997, 213–18). Clark claims that the person's mind should be thought of as including the notebook. For this person, the process of remembering something important typically involves consulting his notebook. The notebook is part of his memory. A major component of Clark's argument for this position is that for someone else deliberately to damage the notebook would be equivalent to assaulting the person. The notebook is as crucial to this man's normal cognitive functioning as is the left part of his brain.

20
25
30

In *Epistemic Cultures*, Karin Knorr Cetina (1999) argues that *different* scientific fields exhibit *different* epistemic cultures. Her most extensive case is High Energy Physics (HEP); in particular, experiments done between 1987 and 1996 at the European Centre for Nuclear Research (CERN). The scale of this laboratory is suggested by the fact that CERN's Large Electron Positron Collider, located on the border between France and Switzerland, was 27 kilometres around and the facility employs hundreds of scientists, technicians, and other support personal.

35
40

My view is that the laboratory at CERN should be thought of as a large distributed cognitive system. Knorr Cetina, in fact, indirectly suggests this approach. In at least a half dozen passages she uses the term "distributed cognition" to describe what is going on in a HEP experiment. Thus, she writes, "... the subjectivity of participants is ... quite



successfully replaced by something like distributed cognition” (Knorr Cetina 1999, 25). In fact, she goes on to argue for “the erasure of the individual as an epistemic subject” in HEP (Knorr Cetina 1999, 166–71), claiming that one cannot identify any individual person, or even a small group of individuals, producing the resulting knowledge. The only available epistemic agent, she suggests, is the extended experiment itself. Indeed, she attributes to the experiment itself a kind of “self-knowledge” generated by the continual testing of components and procedures, and by the continual informal sharing of information by participants. In the end, she invokes the Durkheimian notion of “collective consciousness”. Speaking of stories scientists tell among themselves, she says that “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” (Knorr Cetina 1999, 178).

Here Knorr Cetina seems to be assuming that, if knowledge is being produced, there must be an epistemic subject, the thing that knows what comes to be known. Moreover, knowing requires a subject with a mind, where minds are typically conscious. Being unable to find a traditional individual or even collective epistemic subject within the organization of experiments in HEP, she feels herself forced to find another epistemic subject, settling eventually on the experiment itself.

4. Problems with Extending Epistemic Agency

The first thing to realize is that applying concepts associated with human agency (mind, consciousness, intentionality) to extended entities involving both humans and artefacts, or to inanimate entities themselves, is a matter of fairly high-level *interpretation*. One cannot even imagine an *empirical* test of Clark’s claim that his man’s notebook is part of the man’s mind. Even if a court of law determined that stealing the man’s notebook is as serious crime as assaulting him bodily, this would not show the notebook should be included as part of his mind. Similarly for the navigator’s chart on Hutchins’ ship or for an experiment at CERN. If we are to adopt these interpretations, it can only be because they provide theoretical benefits of some sort, benefits that cannot be obtained without these innovations.³ My view is that these extensions do not provide theoretical advantages for the study of science. On the contrary, they introduce a host of theoretical problems that confuse more than enlighten. We are theoretically better off rejecting these supposed innovations.⁴

The culture in scientifically advanced societies includes a concept of a human agent. According to this concept, agents are said to have minds as well as bodies. Agents are conscious of things in their environment and are self-conscious of themselves as actors in their environment. Agents have beliefs about themselves and their environments. Agents have memories of things past. Agents are capable of making plans and sometimes intentionally carrying them out. Agents are also responsible for their actions according to the standards of the culture and local communities. And they may justifiably claim to know some things and not other things. There is much more that could be said about the details of this concept of an epistemic agent, but this is enough for my purposes here.



A traditional difficulty with the ordinary conception of human agency is that it seems to presuppose a notion of freedom of choice and action that is incompatible with a scientific (naturalistic, materialistic) understanding of humans as biological organisms. To stay within a naturalistic framework, I am willing to grant that the underlying causes of our actions are largely hidden from us, and that the subjective impression that one can in a strong (metaphysical?) sense freely cause things to happen (particularly one's own bodily motions) is an illusion.⁵ So I am willing to regard the ordinary concept of an agent as an idealized model, like that of a point mass in classical mechanics. Such things do not physically exist, but it is a very useful model nonetheless. The same could be true of our ordinary notion of human agency. Like idealized models in the sciences, it has proven very useful in organizing our individual and collective lives. Indeed, our systems of morality and justice are built upon it.

5
10

There is no doubt that *some* extensions of concepts originating with humans beyond the bounds of biological agents are natural, even helpful. Memory is a prime example. Modern civilization, as well as modern science, would be impossible without various forms of record-keeping that are usefully characterized as external memory devices. The modern computer hard drive, for example, is a recent, very powerful, memory device. The organization of such devices—the software, so to speak—is equally important. Thus a person with a computer, like Clark's man with his notebook, is usefully thought of as a cognitive system with an effective memory much more powerful than a system consisting of a person alone, without any such devices. But even here it is pushing the applicability of the concept of memory to say that the extended system as a whole "remembers" something as opposed simply to having the capacity to store large amounts of what to appropriately educated humans is recognized as "information".

15
20
25

More severe problems arise, however, when one attempts to extend concepts such as intention, belief, knowledge, responsibility and consciousness to extended entities. These concepts are all bound up in the more general concept of "mind". Here it is helpful to move beyond examples of locally distributed cognition to consider examples of cognitive systems even more widely distributed than CERN. The Hubble Telescope provides a particularly striking example of a distributed cognitive system. It extends at least from the telescope in Earth orbit through a series of intermediaries to the Space Telescope Science Institute in Greenbelt, Maryland. In 2003 the Hubble system produced a remarkable series of images purporting to show the universe as it was 13 billion years ago. A special feature of these images is that they were produced by utilizing a cluster of galaxies, Abell 1689, as a gravitational lens. Abell 1689 is itself 2.2 billion light-years out into space, yet it was cleverly incorporated into the distributed cognitive system that produced the final images. If we treat the Hubble system as itself an epistemic agent with a mind of its own, it seems we would have to say that its mind extends from the Earth 2.2 billion light years out into space? Do minds operate at the speed of light? Just how fast do intentions propagate? Is the Hubble System as a whole epistemically (as opposed to just causally) responsible for the final conclusions? Did the system as a whole expect to find galaxies as far as 13 billion

30
35
40



light years away? Did it then believe it had found them? These questions do not make much sense. We should not have even to consider them. So we should resist the temptation to ascribe full epistemic agency to distributed cognitive systems as a whole.⁶

5

5. Distributed Cognition and Human Agency

One might now wish raise similar questions about the initial conception of distributed cognition. Does not this notion invite similar unanswerable questions? I think not. The word “cognition” was part of the English language before the field of cognitive science was invented, but I think was always somewhat of a specialists’ term.⁷ We are thus free to develop it as a technical term of cognitive science. The fundamental notion, I think, is not so much that of distributed cognition as that of a *distributed cognitive system*. A distributed cognitive system is a system that produces cognitive outputs, just as an agricultural system yields agricultural products. The operation of a cognitive system is, of course, a cognitive process, but there is no difficulty in thinking that the whole system, no matter how large, is involved in the process. There is, however, no need to endow the system as a whole with other attributes of human cognitive agents. So thinking of cognition as distributed throughout the system need not raise untoward questions.

10

15

20

Actually, as has long been taken for granted in the general science studies community, a claim does not count as *scientific* knowledge until it is publicized and accepted by the relevant scientific community. In John Ziman’s (1968) terms, scientific knowledge is “public knowledge”. This implies that the cognitive system that produces scientific knowledge should really be taken to be a whole scientific community, including things such as the institutions that make publishing possible. So scientific distributed cognitive systems turn out, finally, to be quite heterogeneous systems with very fuzzy boundaries.

25

30

35

40

Karl Popper (1968) long ago introduced the idea of “epistemology without a knowing subject”, although embedded in a metaphysics including a realm of pure problems and ideas. Shorn of its metaphysical excesses, the idea has considerable merit. The institution of science, including its many distributed cognitive systems, produces scientific knowledge. There is, however, no need to postulate a super agent, or a collective agent, *Science*, and endow it with other attributes of human agency. It is enough that the processes of science are generally reliable, and known to be so. Here strains in science studies, as found, for example, in John Ziman’s (1978) *Reliable Knowledge*, and reliabilist strains in philosophical epistemology going back to Goldman (1986), coincide. Individuals can claim to know the conclusions of scientific inquiries because they know these results are reliably produced. This implies there will typically be a short period of time after members of a research group have reached consensus on a conclusion, and can thus each personally claim to know the conclusion, even though this result does not yet count as public scientific knowledge. Later, anyone knowing that the conclusion has been certified by accepted scientific procedures can legitimately claim to know the conclusion.



6. Implications for Collective Cognition

Most of the literature on collective agency and collective knowing, going back to Margaret Gilbert's (1989) innovative *On Social Facts*, has focused on collective cognition, which involves only people, rather than fully distributed cognition, which includes such things as external representations and instruments as well. Much effort has gone into articulating acceptable collective notions such as that of joint knowledge or joint intention that do not reduce to individual knowledge or intentions plus communicative interactions. To some extent, the plausibility of such autonomous notions derives from consideration of small numbers of people, often only two, in close proximity engaging in direct communication. Much of this plausibility would vanish if, like the components of some highly distributed cognitive systems, the participants were separated by light years.

The Hubble System, like the system at CERN, or even the system aboard Hutchins' ship, includes numbers of individuals organized so as to accomplish tasks necessary for the proper functioning of the whole system. Collectively they make the system work. To understand how the members of the groups collectively make the system work, it is not necessary, and I think definitely unhelpful, to introduce the concept of a super, or collective, agent. Our ordinary concept of individual human agents, plus concepts from such fields as psychology, sociology, organization theory, and anthropology, is sufficient to provide as good an account as we need of how all the individuals interact to produce the final cognitive output. The basic notion is simply that the individuals, acting together, make the system work.⁸

7. Distributed Cognitive Systems as Hybrid Systems

Bruno Latour (1993) has popularized the notion of hybrid systems consisting of combinations of humans and non-humans, all called "actants". I would also like to say that distributed cognitive systems are hybrid systems, but retain the genuine differences between humans and non-humans. In more detail, I think distributed cognitive systems include at least three distinct kinds of systems: physical, computational, and human. It is the humans, and only the humans, that provide intentional, cognitive agency to scientific distributed cognitive systems. We need not extend our notions of cognitive agency to include other components of these distributed cognitive systems. Restricted to humans, our ordinary notions of human agency will do. The ultimate justification for this theoretical choice is that it provides a productive theoretical framework that all contributors to science studies, including historians, philosophers, psychologists, sociologists and anthropologists, can share.⁹

Notes

- [1] The present article is adapted from chapter 5 of my *Scientific Perspectivism* (Giere 2006).
- [2] I have not been able to find these claims about mind in *Cognition in the Wild*. I did, however, personally hear Hutchins make these claims in a richly illustrated plenary lecture at the Annual Conference of the Cognitive Science Society in Boston, MA, 31 July–2 August 2003.

320 R. N. Giere

- [3] Here I discount the idea that some sort of linguistic or conceptual analysis could demonstrate that irreducibly collective knowledge either does or does not exist.
- [4] For a detailed critique of the idea of extended minds from the standpoint of analytical philosophy of mind, see Rupert (2004).
- 5 [5] The latest and best I know on this topic is due not to a philosopher but to a psychologist, Daniel Wegener (2002).
- [6] My argument would be similar, although less dramatic, if we leave out the use of gravitational lensing in the Hubble system and regard the system as extending only from Earth orbit to the Space Telescope Science Institute in Maryland. There is no denying that the system is at least that large.
- 10 [7] Among Anglo-America philosophers, for example, “cognitive” has typically been associated with “rational” or “reasons”. Thus, a philosopher would distinguish between a person’s cognitive grounds (reasons) for a particular belief and mere causes of that belief.
- [8] Here I am not claiming that the scientists who interpret the final images produced by the Hubble Telescope System need to know every detail of the system or could successfully perform most of the tasks required to keep the system operating. Knorr Cetina is surely correct about the distribution of expertise needed to operate large-scale experimental systems. I question only her attribution of agency and self-consciousness to the operations of such systems as a whole.
- 15 [9] I suspect that most historians of science and technology, for example, would have little difficulty with notions of distributed or collective cognition. But most, I think, would balk at the notion of extended minds and be very suspicious of talk about collective consciousness.
- 20

References

- Clark, Andy. 1997. *Being there: Putting brain, body, and world together again*. Cambridge, MA: MIT Press.
- 25 Giere, R. N. 2002. Scientific cognition as distributed cognition. In *The cognitive basis of science*, edited by Peter Carruthers, Stephen Stich, and Michael Siegal. Cambridge: Cambridge University Press.
- . 2006. *Scientific perspectivism*. Chicago, IL: University of Chicago Press.
- Gilbert, Margaret. 1989. *On social facts*. London: Routledge.
- Goldman, Alvin. 1986. *Epistemology and cognition*. Cambridge, MA: Harvard University Press.
- 30 Hutchins, Edwin. 1995. *Cognition in the wild*. Cambridge, MA: MIT Press.
- Knorr Cetina, Karin. 1999. *Epistemic cultures: How the sciences make knowledge*. Cambridge, MA: Harvard University Press.
- Latour, Bruno. 1993. *We have never been modern*. Cambridge, MA: Harvard University Press.
- Popper, K. R. 1968. Epistemology without a knowing subject. In *Proceedings of the third international congress for logic, methodology and philosophy of science*, edited by B. van Rootselaar and J. F. Staal. Amsterdam: Reprinted in K. R. Popper. 1972. *Objective knowledge: An evolutionary approach*. Oxford: Oxford University Press.
- 35 Rupert, Robert D. 2004. Challenges to the hypothesis of extended cognition. *Journal of Philosophy* 101 (8): 389–428.
- Wegener, D. 2002. *The illusion of conscious will*. Cambridge, MA: MIT Press.
- Ziman, John. 1968. *Public knowledge: An essay concerning the social dimension of science*. Cambridge: Cambridge University Press.
- 40 ———. 1978. *Reliable knowledge*. Cambridge: Cambridge University Press.

AQ3