The Problem of Agency in Scientific Distributed Cognitive Systems

RONALD N. GIERE*

From the perspective of cognitive science, it is illuminating to think of much contemporary scientific research as taking place in distributed cognitive systems. This is particularly true of large-scale experimental and observational systems such as the Hubble Telescope. Clark, Hutchins, Knorr-Cetina, and Latour insist or imply such a move requires expanding our notions of knowledge, mind, and even consciousness. Whether this is correct seems to me not a straightforward factual question. Rather, the issue seems to be how best to develop a theoretical understanding of such systems as distributed cognitive systems. I argue that there is no need to attribute to such systems as a whole any form of cognitive agency. We can well understand the importance of such systems while restricting agency to the human components. The implication is that we think of these large-scale distributed cognitive systems not so much as unified wholes, but as hybrid systems including both physical artifacts and ordinary humans.

1. Scientific Cognition as Distributed Cognition

In previous publications I have argued that it is illuminating to think of much contemporary scientific research as typically taking place within distributed cognitive systems.[1] This is particularly true of research involving large-scale experimental and observational systems such as particle accelerators at CERN or The Hubble Telescope.

Many of those who have been developing notions of distributed cognition, or related notions, have been tempted to suppose that this way of thinking requires extending our conceptions of knowledge, mind, and even conscious agency, thus attributing to such systems as a whole attributes heretofore restricted to human agents. In this paper I argue against such extensions and conclude that it is preferable to restrict applications of such cognitive attributes as knowing and conscious intentional agency to the human components of distributed cognitive systems.[2] The result is that we should think of distributed cognitive systems not as completely unified wholes, but as hybrids including humans as the only cognitively active agents.

So as to have a common point of reference, I will begin by briefly describing The Hubble Telescope System. I will then survey the views of four prominent advocates of something like distributed cognition in science, all of whom suggest attributing cognitive agency to things other than humans. After raising some general problems for such extensions of ordinary cognitive concepts, I will present my alternative understanding of scientific distributed cognitive systems as hybrid systems. Which characterization of scientific distributed cognitive systems is correct seems to me not a straightforward empirical question. Rather, the issue seems to be how best to develop a theoretical understanding of such systems as distributed cognitive systems

2. The Hubble Space Telescope

The Hubble Space Telescope was launched on April 24, 1990 aboard the Space Shuttle Discovery. Since a dramatic mission to correct an embarrassing flaw in its mirror at the end of 1993, Hubble has produced genuinely revolutionary observations. In January of

2003, for example, the Space Telescope Science Institute released a remarkable image produced by the Advanced Camera for Surveys (ACS) aboard the Hubble.[3]

Among several remarkable features of this particular image is that it involved gravitational lensing. During the exposure, the Hubble Telescope was pointed directly at a massive cluster of galaxies, known as Abell 1689, estimated to be 2.2 billion light-years away. In accordance with the General Theory of Relativity, this mass acts like a lens by warping space around it and thus effectively bending light passing by. Scientists who have studied the data claim that the image captures galaxies from which light was emitted roughly 13 billion years ago, when the universe was only one billion years old, which is to say, about one fourteenth its present age.

The process that produced this image is far too complex for me to describe in any detail here. In addition to the infamous mirror, it involves electronic detectors sensitive to light in the infrared part of the electro-magnetic spectrum. The output of the detectors is fed into an onboard computer and put into a form in which it can be transmitted to a Tracking and Data Relay Satellite from which it is retransmitted to the White Sands Complex near Las Cruces, New Mexico, from which it is again retransmitted by domestic satellite to the Data Operations Control Center at the Goddard Space Flight Center in Greenbelt, Maryland. From there it is routed to the Data Capture Facility and finally on to the Space Telescope Science Institute where it is studied by astronomers and other space scientists. I will refer to this whole complex as the Hubble Space Telescope System (or just Hubble System), which includes the Hubble Space Telescope itself, the specific instrument launched in 1990. To keep things manageable, let us take the cognitive output of this whole process to be just the claim that the image

indicates the existence of galaxies 13 billion years ago.

3. Extending Agency

Let us now examine the views of four prominent advocates of something like distributed cognition: Andy Clark, Ed Hutchins, Karin-Knorr Cetina, and Bruno Latour. The first two explicitly advocate thinking of cognition as distributed, but do not consider scientific distributed cognitive systems in particular. Both consider themselves members of the cognitive science community, one a philosopher and the other a cognitive anthropologist. The latter two focus on scientific systems, but do not explicitly invoke a conception of distributed cognition as understood in the cognitive science community. Their primary identification is with science studies from a broadly sociological or anthropological point of view. In this section, I will merely expound the views and give some indication of the arguments behind them, saving criticism for the following section.

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. The 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, 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 (1997, Ch. IX). Clark claims that the person's mind should be thought of as including the notebook. For this person, the process of remembering something 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 to assault the person. The notebook is as crucial to this man's normal cognitive functioning as is the left part of his brain.

Ed Hutchins' study of navigation in his 1995 book, *Cognition in the Wild*, remains one of the major sources for the concept of distributed cognition within the cognitive sciences. Hutchins' book is an ethnographic study of traditional 'pilotage', that is, navigation near land as when coming into port.[4] 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 fulfill 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 pilothouse 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. Those recording the locations of landmarks have no reason to know the result of the process.

Hutchins, like Clark, includes not only persons but also instruments and other artifacts 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 artifacts. So, for Hutchins, cognition may be somewhat more broadly distributed than in the examples from Clark noted above.

Again like Clark, Hutchins argues that we should regard mind as extended beyond the human body. There is mind at work, he claims, 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 he 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.[5]

In *Epistemic Cultures* (1999), Karin Knorr Cetina argues that *different* scientific fields exhibit *different* epistemic cultures. In particular, she examines research in both High Energy Physics (HEP) and Molecular Biology. Her first and most extensive case is HEP; in particular, experiments done between 1987 and 1996 at the European Center 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 kilometers around. This collider is now being replaced by a Large Hadron Collider (LHC) coupled with a very large detector called ATLAS. The ATLAS detector itself is 44m wide, 22m high, and weighs 7000 tons. The ATLAS project involves hundreds of scientists, technicians, and other support personal.

My view, of course, is that, like the Hubble Telescope, the laboratory at CERN

should be thought of as a cognitive system in which cognition is even more broadly distributed than it is aboard Hutchins' ship. 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. Here are two examples:

> … the subjectivity of participants is ... quite successfully replaced by something like distributed cognition. (1999, 25)
>
> 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, 173)

Her uses of the expression "distributed cognition" are almost always qualified with expressions such as "something like" or "a sort of." There is never any further characterization of what distributed cognition might be, and these uses of the term are neither referenced nor footnoted.

Perhaps Knorr Cetina's most provocative idea is "the erasure of the individual as an epistemic subject" in HEP. She claims 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 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, 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, 179)

Here Knorr-Centia 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. But, being unable to find a traditional epistemic subject within the organization of experiments in HEP, she feels herself forced to find another epistemic subject, settling eventually on the experiment itself as the epistemic subject.

Knorr Cetina's herself provides an enlightening comparison of HEP with Molecular Biology, arguing that molecular biology laboratories exhibit a two-level structure. The lower level consists of individual researchers each working on their own project. The upper level consists of the whole laboratory usually managed by a single director. For Knorr Cetina, the individual nature of the lower level has important theoretical implications. She writes:

> This is perhaps molecular biology's first most important difference from experimental high energy physics: in the molecular biology laboratory,

> the person remains the epistemic subject. ... The laboratory, experimentation, procedures, and objects obtain their identity through individuals. The individual scientist is their intermediary—their organizing principle in the flesh, to whom all things revert. (1999, 217)

Accordingly, the chapters on molecular biology contain no mention of distributed cognition. Here I think Knorr Cetina mistakenly assumes that *distributed* cognition is the same as *collective* cognition, notions she seems to use interchangeably. This identification eliminates the possibility of the sort of local distributed cognition, emphasized by Clark and by Hutchins, where a *single* person operating with an instrument can already be a distributed cognitive system.

Turning finally to Bruno Latour, his paper, 'Visualization and Cognition: Thinking with Eyes and Hands' (Latour, 1986a), expresses views very close to those of Hutchins and other champions of distributed cognition. He cites many cases of what I would call locally (and some not so locally) distributed cognition. For example, the use of chemical formulas in organic chemistry was introduced by Berzelius in 1813 and developed to its mature form by Dumas in the 1830s. Assuming that the basic constituents in reactions are conserved, one can represent chemical reactions by equations in which the numbers of all constituents are the same on both sides of the equation. That is, the equation must *balance*. One can literally do theoretical chemistry by manipulating these symbols as in the following example:

C8 H8 + H4 O2 + Ch4 = C8 H8 O2 + Ch4 H4

Understood in terms of distributed cognition, these formulas are external representations that form part of a distributed cognitive system for exploring possible

reactions in organic chemistry. That is, the cognitive process of balancing an equation does not take place solely in the head of some person, but consists of interactions between a person and physical, external representations.[6]

Latour (1986a, 22) claims that “the history of science is the history of ... [such] innovations.” This is, of course, an exaggeration, but 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 types of distributed cognitive system. Thus 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.

Latour’s 1996 paper must have been written without knowledge of Hutchins’ *Cognition in the Wild*.[7] Here he still holds to the proposal in the Postscript to the second edition of *Laboratory Life* (Latour & Woolgar 1986, 280) that there be “a 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!” In a less quoted passage, the proposal continues: “If our French epistemologist colleagues are sufficiently confident in the paramount importance of cognitive phenomena for understanding science, they will accept the challenge.” This suggests that the kind of cognitive explanations being rejected are those to be found in books like Claude Levi-Strauss’ *La pensée sauvage* (1962) or in more general appeals to a scientific ‘mentalité’. Latour seeks simpler, more verifiable, explanations. As he says (1986a, 1),

“No “new man” suddenly emerged sometime in the sixteenth century ... . The idea that a more rational mind ... emerged from darkness and chaos is too complicated a hypothesis.” Clark and Hutchins would agree on this point. Appeals to cognitive capacities now studied in the cognitive sciences are meant to explain how humans with normal human cognitive capacities manage to do modern science. One way is by constructing distributed cognitive systems that can be operated by humans possessing only the limited cognitive capacities they in fact possess. Moreover, Latour himself now aggress with this assessment. In a review of Hutchins’ *Cognition in the Wild* (1986b, 62) he explicitly lifts his earlier moratorium claiming that “cognitive explanations … have been … made thoroughly compatible with the social explanations of science, technology and formalism devised by my colleagues and myself.”[8]

In spite of his endorsement of Hutchins’ work, there remain areas of strong disagreement, including disagreement over the nature of cognitive agency. The sources of this disagreement are too deep and complex easily to be summarized. They stem from Latour’s desire to overcome standard distinctions such as that between society and the world as well as that between subject and object. It may not be an exaggeration to say that for Latour there is no such thing as a cognitive agent. There are only “actants” connected in more or less tightly bound networks, transforming material representations, and engaged in agonistic competition with other networks. Actants include both humans and nonhumans in relationships that Latour insists are “symmetric.” Thus, in his recent book, *Pandora’s Hope* (1999, 90), Latour writes of Joliot trying to produce the first nuclear chain reaction: “Such was his scientific work: holding together all the threads and getting favors from everybody, neutrons, Norwegians,

deuterium, colleagues, anti-Nazis, Americans, paraffin … .”

Neutrons and Norwegians. There are at least two ways of understanding this symmetry from a more traditional standpoint. One way is as ascribing to nonhumans properties normally ascribed to human agents, such as granting favors or authorizing others to speak in their name. Another way is as reducing human agents to the status of nonhumans, so Norwegians are no more cognitive agents than are neutrons. Latour, of course, would reject both of these understandings since he rejects the categories in which they are stated. Nevertheless, in what follows I will consider both alternatives.

4. Problems with Extended Agency

The first thing to realize is that applying concepts associated with human agency (mind, consciousness, intentionality) to extended entities involving both humans and artifacts, 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 steeling the man’s notebook is as serious crime as assaulting him bodily, that 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. My view is that not only do these extensions not provide theoretical advantages, they introduce a host of theoretical problems that confuse more than enlighten. We are theoretically better off rejecting these supposed innovations.

The culture in scientifically advanced societies includes a concept of a human

agent. According to this concept, human agents are conscious of things in their environment and self-conscious of themselves as actors in their environment. Agents have beliefs about themselves and their environments. They may justifiably claim to know some things and not other things. Agents are capable of making plans and some times intentionally carrying them out. Agents are also responsible for their actions according to the standards of the culture and local communities. Agents are said to have minds as well as bodies. Now I do not claim that this concept of agents is necessarily instantiated or even, if examined closely, logically coherent or consistent with our best science. But, like idealized models in the natural sciences, it has proven very useful in organizing our individual and collective lives.

The main problem with extending a concept like mind or consciousness beyond human agents is that we don't know what to do with these concepts outside of the conceptual nexus that characterizes a full-fledged human agent. It may seem initially plausible to do so in localized examples such as Clark's man with his notebook. It is far less plausible in a less localized situation such as an experiment at CERN. Knorr-Cetina wants to take an experiment as an epistemic agent, indeed, a self-conscious agent. So are we to say that the experiment *knows* whether or not a Higgs Boson has been detected? Did the experiment *intend* to find a Higgs Boson? Did it believe it would find one? Was it disappointed? Would it be *responsible* for the success or failure of the experiment to detect a Higgs Boson?

Consider again the Hubble Telescope. As a distributed cognitive system it extends at least from the telescope in orbit through a series of intermediaries to the Space Telescope Science Institute in Maryland. If one adds the Abell 1689 cluster of galaxies

used as a gravitational lens, the system extends 2.2 billion light-years out into space. Are we to say that its mind extends from the telescope in orbit to Maryland, or 2.2 billion light years out into space? Do minds operate at the speed of light? Just how fast do intentions propagate? How long does it take for one part of its consciousness know what another part is thinking? These are clearly ridiculous questions. That they are ridiculous is a good reason to reject the proposed theoretical innovations that generate them.

Similar considerations militate against the first of my two standard understandings of Latour's actants. Norwegians can form *intentions* to create a representation of neutrons. Neutrons are not equipped to form any intentions whatsoever, let alone ones regarding Norwegians. Contrary to Latour, I think this asymmetry is fundamental for understanding science as a cultural activity.[9] This is not to say that neutrons are not *causally* active. One might even say that they are causal agents, but this does not make them comparable to human agents. On the other hand, being a thoroughgoing naturalist, I would not claim that human agents operate outside the causal order of nature in some sort of Kantian "kingdom of ends." It is enough that humans are a highly evolved form of life capable of sustaining a culture which has developed the conception of human agency outlined above. There is no need for a transcendent notion of free will.[10] In this sense, my second understanding of Latour's symmetry between humans and nonhumans is, after all, closer to the mark.

5. Distributed Cognition and Human Agency

One might now 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 always somewhat of a specialists’ term. We are thus more 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 is produces cognitive outputs. The operation of such a system is a cognitive process. There is no difficulty in thinking that the whole system, no matter how large, is involved in the process. But there is also no need to endow the whole system with other attributes of human cognitive agents. So thinking of cognition as distributed throughout the system raises no untoward questions.

But what makes the output of the system a *cognitive* output? Here I think the only answer can be that it is the kind of output we recognize as the result of human (or animal) cognition such as a belief, knowledge, memory, or a representation of something. This means that ultimately there must be a human agent in the system. Otherwise there are just a lot of physical goings on and no genuine cognition. Without the human agents, the images produced by The Hubble System are just patches on a computer screen produced by a bizarrely complex physical process. It requires the human agents, scientists who know how the system works, to turn those patches into knowledge about distant galaxies.[11]

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 means there will be a short period of time when

members of the research group who reached consensus on the result can each claim personally to know the result even though the result does not yet count as scientific knowledge. This implies that the cognitive system that produces scientific knowledge should really be taken to be a whole scientific community, including things like the institutions that make publishing possible. So scientific cognitive systems turn out to be quite heterogeneous systems with very fuzzy boundaries.

6. Distributed Cognitive Systems as Hybrid Systems

Latour has popularized the notion of hybrid systems consisting of combinations of humans and nonhumans, all called “actants.” I would also like to say that distributed cognitive systems are hybrid systems, but retain the genuine differences between humans and nonhumans. In more detail, I think distributed cognitive systems include at least three distinct kinds of systems, physical, computational, and human.[12] It is the humans, and only the humans, that provide cognitive and intentional agency to distributed cognitive systems. We need not extend our notions of cognitive agency to other components of distributed cognitive systems. Restricted to humans, our ordinary notions of human agency will do.

REFERENCES

CLARK, A.

1997 *Being There: Putting Brain, Body, and World Together Again*. Cambridge, MA: MIT Press.

Giere, R. N.

2002 'Scientific Cognition as Distributed Cognition.' In *The Cognitive Basis of Science*, ed. P. Carruthers, S. Stitch and M. Siegal. Cambridge: Cambridge University Press, 285-299.

2003a 'The Role of Computation in Scientific Cognition.' *Journal of Experimental & Theoretical Artificial Intelligence*, 15:195-202.

2003b 'Computation and Agency in Scientific Cognition.' *Proceedings of the Cognitive Science Society* 25, ed. R. Alterman & D. Kirsch. The Cognitive Science Society.

GIERE, R. N. & MOFFATT, B.

2003 'Distributed Cognition: Where the Cognitive and the Social Merge.' *Social Studies of Science*, 33: 301-310.

HUTCHINS, E.

1995 *Cognition in the* Wild. Cambridge, MA: MIT Press.

KLEIN, U.

1999 'Techniques of modeling and paper-tools in classical chemistry.' In *Models as Mediators: Perspectives on Natural and Social Science*, ed, M. S. Morgan and M. Morrison. Cambridge: Cambridge University Press, 146-167.

KNORR-CETINA. K.

1999 *Epistemic Cultures: How the Sciences Make Knowledge*. Cambridge, MA: Harvard University Press.

LATOUR, B.

1986a 'Visualization and Cognition: Thinking with Eyes and Hands', *Knowledge and Society* 6: 1-40.

1986b ‘Review of Ed Hutchins’ *Cognition in the Wild.*’ *Mind, Culture and Activity* 3(1): 54-63.

1999 *Pandora’s Hope: Essays on the Reality of Science Studies*. Cambridge, MA: Harvard University Press.

LATOUR, B. & WOOLGAR, S.

1986 *Laboratory Life: The Construction of Scientific Facts*, 2nd ed. Princeton: Princeton University Press.

LÉVI-STRAUSS, C.

1962 *La pensée sauvage*. Paris: Plon**.**

NERSESSIAN, N. J., KURZ-MILCKE, E., NEWSTETTER, W. C., & DAVIES, J.

2003 ‘Research Laboratories as Evolving Distributed Cognitive Systems.’ *Proceedings of the Cognitive Science Society* 25, ed. R. Alterman & D. Kirsch. The Cognitive Science Society.

PICKERING, A.

1995 *The Mangle of Practice: Time, Agency, and Science*. Chicago: University of Chicago Press.

WEGENER, D.

2002 *The Illusion of Conscious Will*. Cambridge, MA: MIT Press.

ZIMAN, J. M.

1968 *Public Knowledge: An Essay Concerning the Social Dimension of Science***.** London: Cambridge University Press.

* Department of Philosophy, Center for Philosophy of Science, University of Minnesota. E-Mail: giere@umn.edu

[1] See Giere (2002) and Nersessian (2003).

[2] I hasten to add that I do not advocate denying cognitive attributes to animals or infants. An evolutionary perspective dictates that the cognitive capacities of other animate creatures always remain a subject for empirical investigation. So far, the only scientists we know about are humans.

[3] Which color image, I, unfortunately, cannot reproduce here. Also, unfortunately, NASA has just (January, 2004) cancelled all further missions to keep the Hubble operating beyond the life of its current batteries and other support systems.

[4] Here the qualifier 'traditional' means before the advent of global positioning systems. This latter technology renders superfluous the cognitive skills Hutchins investigates in his study.

[5] I have not been able to find these claims about mind in *Cognition in the Wild.* I did personally hear him make these claims at a richly illustrated plenary lecture at the Annual Conference of the Cognitive Science Society in Boston, MA, July 31-August 2, 2003.

[6] Latour explicitly refers to this sort of case (1986a, 14), but this particular detailed example is taken from Klein (1999).

[7] He cites Hutchins earlier book but not *Cognition in the Wild*.

[8] I regret that I did not discover this review until my paper with Barton Moffatt (2003) was in press. We ended up urging Latour to adopt a position he had already accepted.

[9] This opinion is shared by Andy Pickering (1995, 9-20) who is otherwise quite

sympathetic to Latour's enterprise.

[10] This is no place for a discussion of free will. The latest and best word I know on this topic is due not to a philosopher but a psychologist, Daniel Wegener (2002).

[11] Here I am not claiming that the scientists who interpret the final images 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 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 the system.

[12] Here I am concerned primarily with the human-nonhuman distinction. In (Giere, 2003a, b) I am concerned also to distinguish physical from the computational systems.