Social Epistemology Review and Reply Collective, 2015
Vol. 4, No. 6, 45-51. http://wp.me/p1Bfg0-28R

A Critical Appreciation of Ronald N. Giere's "Distributed Cognition without Distributed Knowing" [1]

Matthew J. Brown, University of Texas at Dallas

Ron Giere's "Distributed Cognition without Distributed Knowing" (Giere 2007) is a relatively short but nonetheless significant paper. Giere has, since 2002, been defending the use of distributed cognition (dcog) theory as the best theoretical framework for the cognitive science of science (Giere 2002a; 2002b; 2002c; 2004; 2006a; 2006b; 2009; 2012; Giere and Moffatt 2003). In his work, he has mainly defended the dcog approach theoretically and applied it as a framework for reinterpreting existing case studies (e.g., by Knorr-Cetina and Latour). Giere's work is a complement to the empirical work by Nancy Nersessian and her collaborators, who apply dcog in their mixed-methods empirical laboratory studies (N. J. Nersessian, Kurz-Milcke, et al. 2003; N. J. Nersessian, Newstetter, et al. 2003; Nersessian 2005; Osbeck et al. 2011).

Science-as-DCOG

Distributed cognition is a radical theory in contemporary cognitive science that holds that cognitive processes or systems exist not only, or perhaps not primarily, inside the head, but can be distributed across social groups and artifacts, and involve the body and parts of the environment. The theory was developed primarily by Edwin E. Hutchins (1995a), who first applied it to the study of navigation teams on large naval vessels. He argued that the cognitive process of navigation near land on these ships could not be identified with any one member of the crew; rather, it was distributed across the entire navigation team and their tools (though on other types of ship, it might be accomplished by a single individual).

It may seem obvious that such a theory could be fruitfully applied to cognitive studies of science, as much of contemporary science clearly involves complex technical artifacts and large groups of people, from the ordinary research lab to huge projects like the Large Hadron Collider or the Hubble space telescope. And indeed, the science-as-dcog approach has generated significant discussion and controversy (Alač and Hutchins 2004; Becvar, Hollan, and Hutchins 2005; 2007; Magnus 2007; List 2008; Vaesen 2011; Brown 2011; Toon 2014; Magnus and McClamrock 2014), in addition to the empirical work already mentioned.

In this paper, Giere attempts to soften the impression that dcog is *too* radical an approach; he argues that while *distributed cognition* is a useful theoretical posit, it need not commit us to the idea that *knowing* (or the related concepts of mind or agency) are so distributed. While "cognition" is something of a term of art, Giere argues, and thus malleable insofar as it provides useful to cognitive science, "knowing requires a subject [cognitive agent] with a mind" (316). Distributed knowing would thus require a collective mind or a "super agent." But the idea of a mind or agent distributed at the level of a dcog system like the

[1] This material is partially based upon work supported by the National Science Foundation under Grant No. 1338735. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

Hubble space telescope strains credulity. In any case, such a problematic supposition is unnecessary: Giere argues, correctly I think, that it is coherent to hold that distributed cognition can produce scientific *knowledge* (or can reliably produce accurate representations), while that knowledge is *known* only by individual human agents. Thus, "science-as-dcog" can proceed as a useful form of analysis without adopting the radical impulses of cognitive scientists like Hutchins or Andy Clark. (Krist Vaesen (2011) raises some important challenges for this argument, however.)

While I think Giere is correct that we *can* insist that scientific cognition is (often) dcog, but that knowing the results of that cognition is something only some human parts of the dcog system do, I don't think he gives us enough reason to think we *should* limit our dcog analyses in that way. I will thus focus the rest of my remarks on Giere's argument and motivations for restricting our dcog analyses in this way.

On the Restricted View

What is Giere's argument for the restricted view? His main concern is that the extension of "knowing" and related concepts ("mind," "agency") creates theoretical problems, but provides no apparent theoretical benefits. Unfortunately, Giere doesn't provide a careful argument that there are insurmountable theoretical problems; what he gives us amounts to a kind of incredulity argument. After setting out how the Hubble space telescope constitutes a dcog system, including using the very distance galaxy cluster Abell 1689 as a gravitational lens, Giere writes,

> 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 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 (317-8).

Is this an argument? The questions certainly create some discomfort and incredulity. But many have reacted with similar incredulity to the idea that minds operate in the wet, gooey substance of the nervous system; that thought amounts just to electrochemical signals amongst neurons; that such a meaty organ should have beliefs, expectations, intentions, hopes, dreams. I take it that Giere would not find this sort of incredulity to be particularly telling about the neural substrate of the human mind.

One could give a kind of Sellarsian interpretation of what is causing the incredulity in Giere's case. It is not so much the large spatiotemporal dimensions of the Hubble system that causes the problem; it is the importation of spatiotemporal talk at all, rendered all the

more stunning by the great scale. Giere switches back and forth willy-nilly between the "space of reasons" and the "space of causes" (or the "space of physical objects in Space and Time"). But as Sellars insists: "The essential point is that in characterizing an episode or a state as that of *knowing*, we are not giving an empirical description of that episode or state; we are placing it in the logical space of reasons, of justifying and being able to justify what one says" (Sellars [1956] 1997).

As Giere himself says, the application of concepts like "human agency" and "knowing" are "a matter of fairly high-level interpretation"(316). From the Sellarsian point of view, Giere has mixed up two very different lenses of interpretation: Giere creates the confusion by mixing up questions about the spatiotemporal and causal structure of the cognitive system and questions about epistemic relations like knowledge, justification, belief, agency, etc., which belong to another type of description altogether.

But surely, Giere may say, it is odd to extend concepts like "intention, belief, knowledge, responsibility"(317) and so on to systems distributed widely across time and space. These concepts are at home in describing individual human beings. "Belief" and "intention" are concepts from human psychology, not concepts applied to these sociocultural systems. So let's look at this not from the Sellarsian perspective, but from a neurocomputational one. Since the early days of cognitive neuroscience, we have been faced with the problem that the brain doesn't seem to do the things that we have assumed are fundamental to human cognition. As Rumelhart et al. (1987) argued, "the human information-processing system" constituted by the brain doesn't seem to work via logical operations and inferences, but rather by a form of pattern-matching. This provided perhaps the original motivation for the theories of distributed and extended cognition:

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

Likewise, Paul Churchland (1981) has argued that there is nothing in the brain that seems to answer to propositional attitude concepts like belief, desire, intention. Churchland argues that we must thus eliminate these concepts. But dcog gives us an alternative strategy. It seems near-impossible to deny that I have beliefs; furthermore, in certain contexts, both folk psychology and cognitive psychological theories of belief work quite well.

If we cannot find "beliefs" in the brain, maybe we can find them in my socioculturally embedded and distributed performances. I perform believing and belief-formation successfully in dialogue with others, in inquiry and problem-solving which may involve tools and collaborators, in the academic writing and revision process that helps me figure out what I believe, etc. A dcog analysis of belief is in this way natural; we find belief

where we see it in action, in one's actions and assertions in the sociocultural world, not hidden inside the brain of an individual. Dcog actually helps us recover the successes of folk psychology and cognitive psychology in the face of the neurocomputational and eliminativist challenges.

When we turn back to science, I think we can see the utility of thinking more radically in the dcog way about belief and knowledge. There are many cases today where we want to ask what scientists believe or know, e.g., about the human causes of climate change or the safe levels of some potentially toxic chemical. When we treat belief and knowing as a property of individual scientists, we focus on aggregating those beliefs to get some degree of consensus. That seems to me to have been at least rhetorically ineffective, as well as philosophically questionable.

Some distrust consensus because it looks like groupthink and dogmatism. We certainly wouldn't want to license the inference from the consensus of *any* epistemic community to truth or reliability. We might do better instead to talk about the distributed cognitive systems that produce the knowledge and the role of peer review, publication, and the discussion of the literature in certifying that knowledge, making it reliable. In other words, we might be better off talking about what the scientific community believes, is justified in believing, or knows in systematic (i.e., dcog) rather than aggregative ways.

Why does Giere want an abstemious version of dcog? Why does he not see the potential benefits of a broader dcog analysis? I think the fundamental problem here is that Giere fails to give due justice to Hutchins' original exposition of the idea of distributed cognition. Giere reinforces the common misconception that dcog is a different, special kind of cognition, an extension of the concept of individual cognition. But Hutchins argues that the home concept of cognition is actually socioculturally distributed, not locked inside the brain. Hutchins argues that the main concepts and theories of cognitive science actually fail to model individual cognition, because they are really aimed at processes of a sociocultural system rather than an individual mind (E. Hutchins 1995b, 363). And going further,

> Sometimes my colleagues ask me whether I feel safe metaphorically extending the language of what's happening inside people's heads to these worlds. My response is: "It's not a metaphorical extension at all." The computer was made in the image of the sociocultural system, and the human was remade in the image of the computer, so the language we use for mental events is the language that we should have used for these kinds of sociocultural systems to begin with (E. Hutchins 1995b, 363–4).

Hutchins here (writing in the mid-1990's) is engaging with the classic computationalism of twentieth century cognitive science. He points out that the computer itself is a model of a sociocultural system---a mathematician breaking a problem down into steps and rules, manipulating conventional mathematical symbols with pencil and paper. This is the process that Turing modeled, with the mathematician in a way removed or automated. Cognitive science then made a slide, from a model of the abstract symbol-manipulation with the symbol-manipulator removed, to treating it as a model of individual cognition

"in the head." And so it is with many of the metaphors with which we describe the human mind: representations, maps, propositions, languages, scripts, models, etc. The home meaning of all of these are as cultural artifacts we use to accomplish a variety of tasks in our social and natural environments. Far from dcog being an extension of concepts whose natural home is "in the head," dcog as Hutchins indented it reminds us that what we have come to regard as "in the head" mainly has its conceptual home in sociocultural processes and systems, and that this is no accident.

Giere consistently presents dcog in a way that is incompatible with Hutchins' theory. For instance, "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 ... " (314).

On the Primacy of the Individual

In defending the "cognitive" nature of distributed cognition, Giere here assumes that individual cognition is primary, and dcog a derivative extension. Likewise, Giere discusses Clark's theory of extended cognition in the context of arguing that we should not "extend" the concepts of mind, belief, agency, etc. to the "extended entities" constituted by dcog systems. The problem is that Giere adopts Clark's language of "extension," as if these concepts are at home in describing individual mental events, and we must decide which we can "extend" to dcog systems. Hutchins argues, to the contrary, that cognition, mind, belief, etc. were properly descriptions of sociocultural systems all along, and the philosophers and psychologists who read them into a hypothetical internal, private mental realm were mistaken. Rather than the connotation of being "extended" from the head out into the world, the mind has always been socioculturally distributed.

While Giere's view is coherent and interesting, he has not yet shown why it is preferable to this original dcog approach. Which is not to say that an argument for Giere's view over dcog could not be made. We have lots of options on the table in expiring the "science as dcog" approach. Clearly that not all collective decision-making is best analyzed as a case of distributed cognition; for instance, elections or jury-trials may best be analyzed using frameworks from political science. Likewise, we may indeed find ways of distinguishing species of dcog systems, perhaps even based on spatiotemporal scale that require different conceptual resources.

Time will tell whether the diversity of epistemic processes in science, from the inquiry of an individual lab, to large collaborative projects like the Hubble or the LHC, to the system of peer-review, publishing, and post-publication vetting of the literature, are all best analyzed in terms of dcog. We also have to compare dcog to more traditional approaches, as well as to more radical approaches such as Latour's "actor-network theory," which would seem to not only accept that cognitive and mental processes and properties might be distributed across the individual, social, artifacts, and environment, but might actually be attributable to the (apparently) inanimate "actants" that are part of the system/network. At this early stage, however, I can see no reason to take the option of distributed knowing off the table.

Contact details: mattbrown@utdallas.edu

References

Alač, Morana, and Edwin Hutchins. "I See What You are Saying: Action as Cognition in fMRI Brain Mapping Practice." *Journal of Cognition and Culture* 4, no. 3 (2004): 629-661.

Becvar, Laura A., James Hollan, and Edwin Hutchins. "Representational Gestures as Cognitive Artifacts for Developing Theory in a Scientific Laboratory." In *Artifacts in Workplace Practice*, edited Christine A. Halverson and Mark Ackerson, 117–143. Kluwer Academic Publishers, 2007.

Becvar, Laura A., James Hollan, and Edwin Hutchins. 2005. "Hands as Molecules: Representational Gestures as Cognitive Artifacts for Developing Theory in a Scientific Laboratory." *Semiotica* 156, no. ¼ (2005): 89-112.

Brown, Matthew J. 2011. "Science as Socially Distributed Cognition: Bridging Philosophy and Sociology of Science." In *Foundations of the Formal Sciences VII: Bringing Together Philosophy and Sociology of Science*, *Volume 32,* edited by Karen François, Benedikt Löwe, Thomas Müller, and Bart van Kerkhove. London: College Publications, 2011.

Churchland, Paul M. "Eliminative Materialism and the Propositional Attitudes." *The Journal of Philosophy* 78 (2 (1981): 67–90.

Giere, Ronald N. "Discussion Note: Distributed Cognition in Epistemic Cultures." *Philosophy of Science* 69, no. 4 (2002a): 637–644.

———. "Models as Parts of Distributed Cognitive Systems." In *Model-Based Reasoning: Science, Technology, Values*, edited by Lorenzo Magnani and Nancy Nersessian, 227-241. Kluwer Academic Publishers, 2002b.

———. "Scientific Cognition as Distributed Cognition." In *The Cognitive Basis of Science*, edited by Peter Carruthers, Stephen Stich, and Michael Siegal, 285-299. Cambridge University Press, 2002c.

———. "The Problem of Agency in Scientific Distributed Cognitive Systems." *Journal of Cognition and Culture* 4, no. 3 (2004): 759-774.

———. *Scientific Perspectivism*. University of Chicago Press: Chicago, 2006a.

———. "The Role of Agency in Distributed Cognitive Systems." *Philosophy of Science* 73, no. 5 (2002b): 710-719.

———. "Distributed Cognition without Distributed Knowing." *Social Epistemology* 21, no. 3 (2007): 313-320.

———. "Scientific Perspectivism: Behind the Stage Door." *Studies In History and Philosophy of Science Part A* 40, no. 2 (2009): 221-223.

———. "Scientific Cognition: Human Centered but not Human Bound." *Philosophical Explorations* 15, no. 2 (2012): 199-206.

Giere, Ronald N., and Barton Moffatt. "Distributed Cognition: Where the Cognitive and the Social Merge." *Social Studies of Science* 33, no. 2 (2003): 301–310.

Hutchins, Edwin. 1995a. *Cognition in the wild*. Cambridge, Mass.: MIT Press.

———. *Cognition in the Wild (Bradford Books)*. The MIT Press, 1995b.

List, Christian. "Distributed Cognition: A Perspective from Social Choice Theory." In *Scientific Competition: Theory and Policy, Conferences on New Political*

Economy, *Volume 24,* edited by Max Albert, Dieter Schmidtchen, and Stefan Voigt. Tuebingen: Mohr Siebeck, 2008.

Magnus, P. D. "Distributed Cognition and the Task of Science." *Social Studies of Science* 37 no. 2 (2007): 297-310.

Magnus, P. D., and Ron McClamrock. "Friends with Benefits! Distributed Cognition Hooks up Cognitive and Social Conceptions of Science." *Philosophical Psychology* (2014): 1-14. DOI:10.1080/09515089.2014.964857

Nersessian, Nancy J. "Interpreting Scientific and Engineering Practices: Integrating the Cognitive, Social, and Cultural Dimensions." In *Scientific and Technological Thinking*, edited by Michael E. Gorman, Ryan D. Tweney, David C. Gooding, Alexandra P. Kincannon, 17-56. Erlbaum, 2005.

Nersessian, Nancy J., Elke Kurz-Milcke, Wendy C. Newstetter and Jim Davies. "Research Laboratories as Evolving Distributed Cognitive Systems." In *Proceedings of The 25th Annual Conference of the Cognitive Science Society*, 857-862, 2003.

Nersessian, Nancy J., Wendy C. Newstetter, Elke Kurz-Milcke, and Jim Davies. "A Mixed-Method Approach to Studying Distributed Cognition in Evolving Environments." In *Proceeedings of the International Conference on Learning Sciences*, 307-314, 2003.

Osbeck, Lisa M., Nancy J. Nersessian, Kareen R. Malone, and Wendy C. Newstetter. 2011. *Science as Psychology: Sense-Making and Identity in Science Practice*. New York: Cambridge University Press, 2011.

Rumelhart, David E., Paul Smolensky, James L. McClelland and Geoffrey E. Hinton. 1987. "Parallel Distributed Processing." In *Parallel Distributed Processing, Volume 2,* edited James L. McClelland and David E. Rumelhart. MIT Press, 1987.

Sellars, Wilfrid. *Empiricism and the Philosophy of Mind*. Edited by Robert B. Brandom and Richard Rorty. Harvard University Press, [1956] 1997.

Toon, Adam. "Friends at Last? Distributed Cognition and the Cognitive/Social Divide." *Philosophical Psychology* 27, no. 1 (2014): 112-125.

Vaesen, Krist. "Gieres (In)Appropriation of Distributed Cognition." *Social Epistemology* 25, no. 4 (2011): 379-391.