Giere's (in)appropriation of distributed cognition

Krist Vaesen
Eindhoven University of Technology, The Netherlands
k.vaesen@tue.nl

forthcoming in Social Epistemology

Abstract

Ronald Giere embraces the perspective of distributed cognition to think about cognition in the sciences (Giere, 2007, this journal). I argue that his conception of distributed cognition is flawed in that it bears all the marks of its predecessor, namely individual cognition. I show what a proper (i.e. non-individual) distributed framework looks like, and highlight what it can and cannot do for the philosophy of science. (73 words)

1 Introduction

In a series of writings, Giere argues that many of our scientific practices involve distributed cognition (Giere, 2002a,b,c; Giere and Moffatt, 2003; Giere, 2006a,b, 2007). The basic idea behind distributed cognition (d-cog) is that cognition often distributes across individuals and/or epistemic aids such as instruments, diagrams, calculators, computers, and so forth. Distributed cognitive systems may consist of (i) a single person using some instrument(s); (ii) a group of people; (iii) a group of people using some instrument(s) (Magnus, 2007). With its insistence on the distribution of cognitive labor, d-cog is supposed to be an antidote to more traditional accounts, according to which cognition only takes place in the head of (isolated) individuals (hence, individual cognition, or i-cog). D-cog is closely related to other recent movements away from i-cog, such as embodied cognition, embedded cognition, and situated cognition.[1]

Given the sciences' strong reliance on teamwork and instrumentation, d-cog seems at least intuitively appealing. That it has won wide support in cognitive science (its place of origin), makes it all the more attractive. Nevertheless, I think that we should evaluate Giere's proposal critically. It might indeed alter our conceptions of science and scientific practice, but it as well might impoverish rather than improve our understanding. These matters should be sorted out, I believe, before we follow Giere's example. This paper means to do just that, providing a critical assessment of Giere's ideas, and more generally, of the potential of d-cog for the philosophy of science.

My argument unfolds as follows. In the next section I consider the scope and content of d-cog—as Giere construes it. In Section 3 I give some historical background and derive a set of features marking the difference between d-cog and its predecessor i-cog. In light of these, I evaluate Giere's proposal (Section 4), and conclude that it bears the marks of i-cog, not d-cog. In Section 5, I anticipate and counter a set of objections. Finally, Section 6 sketches the implications of my argument for the philosophy of science.

2 D-cog according to Giere, take 1

Giere uses a set of examples to introduce his readers to the notion of distributed cognition. For similar purposes, I will briefly rehearse two of them. The first is Ed Hutchins' landmark study of navigation, presented in his *Cognition in the Wild* (1995).

Hutchins shows that ship navigation near land, as when entering a port, requires a complex system of people interacting with each other and with a set of technological devices. For instance, sailors record telescopically the locations of landmarks relative to the ship's gyrocompass, and these

[1] For an overview of the overlaps and differences between these approaches, see Wilson, 2002.

readings are communicated to the pilot-house; there, they are combined by the navigator to plot the ship's location. And it is not just that more than one person is needed to perform all these tasks, the social relationships between crew members matter greatly. For smooth navigation, navigators must be higher-ranked than the sailors producing the landmark readings; navigators, in turn, must follow the orders of pilots and captains, delivering locations and bearings when requested. Hierarchy (rather than horizontal organization) is crucial to success. Finally, the system can operate only by the grace of sophisticated instruments and charts, onto which much of the cognitive work is offloaded. Among them there are, for instance, the "alidade" (used to observe the bearings of a landmark) and the "hoey" (a ruler-like device to draw the bearings on a navigational chart). All this leads Giere to the conclusion that 'no one human could physically do all the things that must be done to fulfill the cognitive task [of determining a ship's position]' (Giere, 2002c, p. 286). Instead, humans are merely components in a complex cognitive system, where parts of the cognitive process take 'place not in anyone's head but in an instrument or on a chart' (*ibid.*, p. 287).

The second example concerns scientific practice proper. It is about the Hubble Space Telescope (HST), launched on April 24, 1990, and since 1993, producing spectacular images of far-off galaxies and other cosmic ornaments. These images come about by a dazzlingly complex process, involving electronic detectors, on- and off-board computers, data transmission devices, relay satellites, and so forth. Each step in the process slightly modifies the original signal and contributes to the final images. And of course, without the joint cognitive efforts of an entire team running the system, no such images would ensue. So again, Giere invites us to think of the HST as a distributed cognitive system, a system 'involving both humans and artifacts that produces cognitive outputs [...] that could not be produced by any one person or group of people without the complex instrumentation employed' (Giere, 2006a, p. 712-713).

There is still one element missing. Giere not only maintains that we *can* frame Hutchins' navigational systems and the HST in terms of distributed cognition, he also argues that in a sense we *should* do so. He advises us to abandon the traditional approach of individual cognition (i-cog), according to which the instrumentation of the HST simply produces 'evidence (inputs) for the humans who produce knowledge claims (outputs)' (*ibid.*, p. 713)—meaning that the cognition is solely 'localized in human brains (or minds)' (*ibid.*). Instead, we should embrace distributed cognition, for two reasons. First, d-cog provides the hope for a reunion of several disparate approaches to the study of the sciences: that is, d-cog provides 'a conception of cognition in science that unifies much thinking about science in the history, philosophy, and sociology of science within a general cognitive science framework' (*ibid.*). And second, we should embrace d-cog, since d-cog—at least according to Giere—fits better with contemporary scientific approaches to cognition than i-cog (see Giere, 2002c, 2006b). As a naturalist, Giere is of course eager to import these approaches into the philosophy of science. Question is, however, if he is importing the right kind of thing.

3 D-cog according to history

The widespread adoption of the term distributed cognition (well beyond scientific circles) comes with the risk that it is used so loosely that it tends to mean little at all. This in turn may distract us from the core features of the phenomenon which the original term was meant to describe. So to appreciate what is genuinely distinct about d-cog (properly conceived), it is worthwhile to have a look at its agenda as expressed in its early days:

> [distributed cognition] is a new branch of cognitive science devoted to the study of: *the representation of knowledge* both inside the heads of individuals and in the world ...; the propagation of *knowledge* between different individuals and artifacts ...; and the transformations which external structures undergo when operated on by individuals and artifacts ... By studying cognitive phenomena in this fashion it is hoped that an understanding of how intelligence is manifested *at the systems level*, as opposed to the individual cognitive level, will be obtained. (Flor and

Hutchins, 1991, all italics mine)

Two important claims are made here. *First*, d-cog espouses a methodological principle of parity between human and artifact. Human and artifact are equivalent in the sense that they both transform and propagate representations. Both are considered as "nodes" in an extended computational network; they play a similar functional role in the realization of the system's overall capacity. In that sense, *Cognition in the Wild* is an exercise in reverse engineering: the overall function of the navigational system and its physical (component-level) description are given, Hutchins' task is to derive and model the system's information-transforming subfunctions and the particular ways in which these are organized so as to realize the system's functional goal. These subfunctions are spread across humans and artifacts—whereas for i-cog, they would only be operative inside the human head.

This logic may well be reversed for the purpose of cognitive systems design. Given an overall function, the d-cog scientist decomposes it in a set of subfunctions and determines a way of organizing them, working her way toward a material implementation (i.e. a design) of the system—hence, d-cog's popularity among scholars having a foot in design, such as, among others, Ed Hutchins, Donald Norman and Yvonne Rogers.

An implicit assumption made in such functional decompositions is the following: subfunctions differ from the overall function they subserve (Vermaas and Garbacz, 2009). As put succinctly by William Bechtel (2009, p. 162): 'The parts of a watch do not themselves keep time but perform actions that enable the watch to keep time.' So in matters of cognition, a system's overall cognitive task is not assigned to any one subfunction, but realized through different subfunctions working together. Put differently, let ϕ_1, ϕ_2, ..., ϕ_n be subfunctions of overall function Φ_1 of a system S_1. If any single subfunction ϕ_x would perform overall function Φ_1, all other supposed subfunctions of Φ_1 would stop being subfunctions of Φ_1; rather, they would be (at best, in pain of being redundant or obsolete) subfunctions of a separate overall function Φ_2 of a separate system S_2.[2] This assumption is

[2]At a fairly general level, the description of a subfunction ϕ_x may be said to coincide with the description of the overall function Φ_1. For instance, in a sense it is true that seeing

indeed quite straightforward, even trivial perhaps, but I highlight it in light of what is to come.

The *second* important commitment of d-cog (apparent in the quote, and already touched upon in passing) is an extension of the traditional cognitive unit of analysis (that is, the individual cognizer of i-cog). Systems have properties (such as, "intelligence at the systems level") that are in need of explanation, and it is believed that such cannot be done by simply aggregating the cognitive properties of individual agents (the "individual cognitive level"). So in ship navigation, according to Hutchins, there is no central unit orchestrating the cognitive endeavor, there is no central blueprint for successful navigation, there is no crew member who has internalized all relevant knowledge and skills. The captain may set the goals, but it is the system as a whole (crew plus artifacts) that accomplishes the navigation task. Therefore, Hutchins writes (and he is at pains to stress the point throughout *Cognition in the Wild*): 'Depending on their organization, groups must have cognitive properties that are not predictable from a knowledge of the properties of the individuals in the groups' (Hutchins, 1995, p. xiii). A group of people with the same individual knowledge base but organized differently (say, hierarchically versus horizontally), will perform differently at a given cognitive task. Simply adding up properties of individuals, thus, will not account for success or failure; instead, we need to consider the wider socio-technical system. It is the system *as a whole* (crew plus artifacts) that *knows* how to enter the port.

I cannot move to the next section before I have addressed one final, quite thorny, issue. The quote of Flor and Hutchins contains one particularly tricky word: knowledge. D-cog studies 'the propagation of knowledge between different individuals and artifacts', the authors claim. So apparently, they apply the parity principle across the board: artifacts are cognizing entities, and may, just like humans, be attributed *knowledge*. And in fact,

is a function of both eyes and the overall visual system. However, if we would functionally decompose the visual system, we would say that the visual system is for seeing (or visual perception), whereas the eye is for capturing light and focusing it onto an array of visual receptors. By and of itself, an eye is blind; it takes the entire visual system to see.

the same goes for socio-technical systems as a whole: the crew-plus-artifact system is said to *know* how to navigate the ship.

Should we take this literally? Does Hutchins really believe that artifacts and systems may have knowledge, or is he proposing just a metaphorical extension of our epistemic concepts? I agree with Giere (2007, p. 315): we should read Hutchins indeed literally here.[3] Here is some evidence:

> 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. (Hutchins, 1995, p. 364)

Even on this literal reading, however, one could deflate the point Hutchins is making. Perhaps Hutchins uses the term knowledge in an information-theoretic sense: to know that p is to have the information that p, where "information" entails truth (for a discussion see (Goldman, 2002, p. 185)), and where "to have" does not require any active endorsement by a (human) agent—the computer knows that p in the sense that the machine bears information that p. Even the staunchest JTB-theorist would arguably admit that artifacts and systems indeed may *know* things in this weak sense.

Giere doesn't accept such a deflationary account of knowledge—science produces *knowledge*, not just information—but on the other hand finds inappropriate attributions of knowledge to artifacts. His alternative is to evacuate knowledge from the d-cog framework—deleting the word "knowledge" in Flor and Hutchins' earlier statement, as it were. In doing so, though, Giere makes d-cog indistinguishable from i-cog, its natural rival—or so I will argue.

[3] Hutchins even goes as far as granting the system a *mind*, Giere (2007, p. 319) suggests, based on Hutchins' plenary talk at the Annual Conference of the Cognitive Science Society in Boston, MA, 31 July2 August 2003.

4 Giere on (distributed) knowledge

So the simple question is this: can one embrace d-cog and yet deny artifacts and socio-technical systems attributions of knowledge? Giere thinks we can. Attributions of belief and knowledge should be 'restricted to the human components[4] of a distributed cognitive system' (Giere, 2006a, p. 715; see (Giere, 2006b) for a similar point).

To convince us, Giere invites us to consider the example of a system consisting of one single person and her desktop computer (*ibid.*, p. 716). This system can perform cognitive tasks (such as multiplying 10-digit numbers) that the person alone would not be able to perform—at least not as swiftly as the joint human-computer system. On the other hand, Giere claims (correctly perhaps) that it would defy common sense to say that the system as a whole "remembers" or "forgets" (or "regrets forgetting") the answer to any assigned calculation. Hence, 'the term "cognitive" can be applied to the whole system without also applying such terms as "know", "believes", etc.' (*ibid.*, p. 718); these can only concern the human operator.

This indeed seems plausible, in case we, like Giere, adopt a strong notion of knowledge.[5] But let us carefully consider the implications.

If a system is a distributed cognitive system, it must have some kind of system goal—a cognitive one at that. Giere acknowledges this: d-cog systems produce outputs that we would 'recognize as the result of human cognition, such as a belief, knowledge, or a representation of something' (Giere, 2006b, p. 113). Now, I think that the production of representations may indeed be a systemic achievement, but the production of *belief* and *knowledge* is *not*—at least not, if we assume, like Giere, that belief and knowledge is restricted to humans. In that case, it is more useful to consider two systems: one that produces data for the second [the human] producing

[4] Here, I take the term "components" to refer to individual human beings, not to collectives of humans, since the author writes: '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' (Giere, 2007).

[5] Nevertheless, I think that systemic human-computer knowledge (in a strong sense) has been made quite palatable elsewhere, the case of Otto and his notebook, both starring in (Clark and Chalmers, 1998), being a particularly powerful example.

knowledge—indeed, the image Giere set out to replace. To see why, let us return to functional decomposition and to the Hubble Space Telescope.

Suppose we define the HST's overall function (Φ_{HST}) in terms of knowledge (as opposed to mere representation) production. A very (very!) rudimentary functional decomposition could look like this:

$$(\Phi_{HST}, \quad \phi_1 \rightarrow \phi_2 \rightarrow \phi_3 \rightarrow \phi_4), \tag{1}$$

where ϕ_1 is "data capture", ϕ_2 is "data transmission to planet earth", ϕ_3 is "delivering data in a format accessible to humans", ϕ_4 is "data to knowledge conversion",[6] and "$\rightarrow$" indicates the ordering of the different subfunctions. It is subfunction ϕ_4 that we are particularly interested in. Such a function, however implemented, must be operative *somewhere* in the system, even if distributed over several parts. Somewhere representations have to be promoted to knowledge, for if not, the system's overall function Φ_{HST} isn't realized. Given his insistence on the "non-parity" between humans and artifacts/systems in matters of knowledge, it is clear where Giere thinks (or must think) this happens: ϕ_4 is carried out within the traditional human skinbag... which in fact brings us back to i-cog. For in the end, it appears that the instrumentation of the HST simply produces, repeating Giere's own characterization of i-cog, 'evidence (inputs) for the humans who produce knowledge claims (outputs)' (Giere, 2006a, p. 713).

At this point one might object that my argument glosses over the distinction between *producing* knowledge and *knowing*. One might argue that the non-parity between humans and artifacts concerns only the latter, not the former. Humans and artifacts and distributed cognitive systems as a whole are all (jointly) able to *produce* knowledge, it is only humans that have epistemic agency, in the sense of being able to *know* the things produced. If so, Giere isn't forced to localize the production of knowledge in the human head.

The problem with this sort of reasoning is this. What is knowledge? If

[6] For our purposes, it doesn't matter how we spell this out. We may invoke whatever we think makes the difference between data and knowledge, processes of justification and endorsement presumably being the most obvious difference-makers.

knowledge is more than correct information (or, say, true belief), one usually introduces some requirement of justification. For instance, for Giere knowledge in science requires knowing that '[the] results are reliably produced' (Giere, 2007, p. 318) or of knowing 'the capabilities of the whole [HST] system' (Giere, 2006a, p. 718). What, then, is the most salient feature explaining the difference between correct data and knowledge? Indeed, human agency. While data might be stored without being endorsed by any one, to yield knowledge, they must pass through the human cognitive engine (for purposes of interpretation and justification). So what is producing knowledge? The human agent; for the human agent is the relevant difference-maker. Artifacts and systems as a whole might produce data, but knowledge, no. In other words, the non-parity between humans and artifacts/systems applies just as much to *knowing* as to knowledge.

Given this, it is hard to see in what sense subfunction ϕ_4 differs from the overall function we assigned to the HST (i.e. Φ_{HST}). And if ϕ_4 and Φ_{HST} coincide, subfunctions ϕ_1, ϕ_2 and ϕ_3 become redundant to Φ_{HST}, implying that they can be evacuated to a different system—or vice versa, ϕ_4 is assigned to its own dedicated system, which in this case appears to be the human cognizer (the scientist). So instead of (1) we would rather get:

$$(\Phi_{HST}, \quad \phi_1 \rightarrow \phi_2 \rightarrow \phi_3) \rightarrow (\Phi_{Scientist}, \quad \phi_4), \tag{2}$$

where $\Phi_{Scientist}$ ($= \phi_4$) amounts to scientific knowledge production,[7] and Φ_{HST} is the HST's overall function of producing data. Hence, instead of one system, we would have two, the first (the instrumentation) providing inputs to the second (the human). Enter i-cog, thus, once again.[8]

Consider how natural a picture this is. Given (2), both systems have a certain autonomy. They can be decoupled, without losing their capacity to perform their respective overall functions. Of course, the content of the

[7] Almost certainly, ϕ_4 is decomposable in a set of subfunctions, but that should not detain us here.

[8] Note that to argue my case, I only have to show that model (2) is just as plausible and explanatory salient as model (1), since if they are equivalent, the methodological principle of conservatism tells us to stick to model (2), and hence, to i-cog.

knowledge claims produced in the human cognitive system is highly dependent on what the HST presents as data. But dependency is insufficient to treat two systems as one (more on this in Section 5).

Here is a final way of framing the problem. According to d-cog, systemic outputs cannot be explained in terms of the cognitive properties of the individuals in the system. The crew-plus-artifact-system as a whole knows how to safely navigate a ship, and this capacity isn't reducible to the capacity of one. Compare this with Giere. If we are interested in how the HST-plus-scientist system was able to produce knowledge (rather than merely generate data), we know where to look: to the scientist interpreting the cosmic imagery. Even if the artifactual system would not only make pictures and transmit data, but would also do all calculations, and perhaps even write a brilliant scientific report about the age of the universe, knowledge is produced only when the human scientist thinks it is, when the report would get her stamp of approval (to be clear: I am still working under Giere's assumptions concerning knowledge). And this knowledge-conferring act is fully explainable in terms of the *individual* scientist's cognitive endowments; we would look at her reasons for believing and endorsing what is reported on, and/or at the reliability of her judgmental skills.[9] Knowledge, thus, is not an emergent system property, but reducible to the properties of an individual cognizer. Just like i-cog would have it.

In sum, Giere says he espouses d-cog, and that he can do so while rejecting d-cog's particular perspective on knowledge. I have argued that such cannot be done without deflating the notion of d-cog up to the point that it gets identical to its supposed rival, viz. i-cog.

[9] I grant that this is a particularly individualistic take on scientific practice. But again, remind that I am just elaborating on a Gierean theme: knowledge primarily is an individual affair (see footnote 4). If one finds the idea of collective knowledge appealing, things might be different. But I guess that even then we can discern two systems: one producing data for a second (now a *collective* of scientists) producing knowledge.

5 Objections and replies

To add some more precision to my argument and to convince the skeptic, let me anticipate some counterarguments.

Objection. A first objection might run as follows: perhaps Giere doesn't advocate the most radical version of d-cog, but it isn't i-cog either. Fact is—and this is the main point—that in case of the HST a substantial part of the cognition takes place not inside (as with i-cog) but outside the scientist's head.

Reply. It is undeniable that in case of the HST much of the processing takes place outside the heads of scientists. To turn this into a case for d-cog, however, one needs to argue in addition that these external processes are genuinely cognitive. But why would one want to do that, if i-cog is perfectly suited to characterize the HST (see e.g. model (2), Section 4), and, importantly so, whithout the need for such an extra argument? In other words, the most straightforward route is to rely on i-cog, thereby sidestepping the conceptual morass surrounding the notion "cognitive".

In fact, by calling the HST's external processes cognitive, one takes quite a radical position—even more radical than some of the people Giere claims to be too radical, such as extended mind (EM) theorists. Andy Clark, for instance, proposes a *disciplined* extension of the cognitive: only those processes that we would call cognitive if they were located inside the head, deserve the label "cognitive" (intergalactic data transmission definitely not being one of them). Mark Rowlands, EM-proponent since the late nineties, argues in a similar vein that some token cognitive processes may indeed extend into the cognizing subject's environment, but adds the following condition : the process in question must *belong* to a particular subject (to whom belongs HST's intergalactic data transmission?). Why impose this condition? To exclude 'processes occurring inside my calculator, or my computer, from counting as cognitive' (Rowlands, 2009, p. 14).

Objection. A second objection is to point out that these external processes not need to be cognitive; the fundamental notion is not so much that of distributed cognition, but rather that of a distributed cognitive system (as put in (Giere, 2006b, p. 112)). What matters is simply the fact that the system manages to produce cognitive outputs that an un-aided individual could never produce.

Reply. I agree (and most of us would, I guess) that the scientist couldn't achieve what she can achieve when scaffolded by external aids. So indeed, her scientific mind is critically dependent on features of her (sometimes fairly remote) environment. But as Margaret Wilson puts it: 'Distributed causality [...] is not sufficient to drive an argument for distributed cognition' (Wilson, 2002, p. 630). Indeed, the claim that the forces driving cognition do not reside exclusively within skin and skull is trivially true. The sun, by illuminating the objects we perceive, certainly causally contributes to our cognitive outcomes, just like food contributes to the proper functioning of our mind; but these entities do typically not belong to the systems studied in cognitive science. The point generalizes: the fact that the human cognitive system is an open system, doesn't mean it encompasses the things it is open to.

So there is always the question how far into the environment we extend the human cognitive system. Model (2) is conservative on this score and respects the traditional bounds of cognition; but apparently it is capable of explaining everything that model (1) was supposed to explain. By dint of a methodological principle of conservatism, then, we should give preference to model (2).

Objection. In another line of counterargument, one grants that knowledge production itself isn't explainable in terms of d-cog. That is, one agrees to remove $\Phi_{Scientist}$ out of model (2). All the rest, however—that is: *every* process promoting Φ_{HST}—is still up for a d-cog analysis. And that is still quite a lot.

Reply. Three points in response. First, it is a bit awkward to remove from a cognitive analysis the thing that is taken to be cognitive *par excellence*, namely the production of knowledge. Second, what is remarkable about science is not so much its capacity to produce *data*, but its capacity to generate *knowledge*; hence such an analysis would be (very) partial at best. Third, even if one thinks that instrumentation has been underplayed in the philosophy of science, one could well analyze it without invoking d-cog (illuminating examples include Giere's own non-d-cog treatment of scientific observing (Giere, 2006b, Ch. 3), and the excellent volume on scientific experimentation edited by Radder (2003)).

Objection. Finally, one could try to defend Giere's position by claiming that d-cog is scientifically more respectable than i-cog. Being in line with current findings in cognitive science is a benefit one should not dismiss lightly.

Reply. I think we should be careful here. Even if d-cog is scientifically respectable,[10] or useful in some domain, extending it to phenomena outside its intended range of application may lead to "unscientific" results. Now, d-cog was originally devised as a framework to model cognition guiding *action* (such as playing tetris, ship and airplane navigation, taking customer orders at a bar, sketching); the assumption that these models can be simply scaled up to higher-level forms of cognition (such as cognition producing quite abstract knowledge claims) is, until further notice, unsupported. The title of Hutchins' monograph is instructive: the book addresses cognition taking place *not* in an armchair, but *in the wild.* This implies that action-guiding types of cognition must be *studied* in the wild, which in turn explains d-cog's strong reliance on ethnographic field studies. So if one wishes to argue for the fruitfulness of d-cog in matters of scientific cognition, some detailed field work (e.g. of the HST) is needed.

Now, work of this sort has been conducted by Osbeck and Nersessian (2006). These authors indeed apply the methods of d-cog to track the distribution of representations in biomedical engineering. Their research, which they call 'very much in progress' (p.155), definitely shows that *it is possible* to apply d-cog to scientific settings. This, however, is insufficient to establish the much stronger claim, viz. that d-cog *outperforms* earlier methods. To prove the usefulness of d-cog, one needs to specify where the extra explanatory leverage is to be had. Specifically on this score, d-cog still has to

[10]Even this will be questioned by quite a few cognitive scientists, in particular those favoring non-representational accounts of cognition (e.g. Thelen and Smith, 1994, Port and van Gelder, 1995).

earn its credentials.

6 (Distributed prospects and) conclusion

‘Science provides arguably the best example of a higher cognitive activity’, Giere (2002c, p. 285) writes. What is typical about such a higher cognitive activity is the fact that it is able to produce knowledge—not just knowledge, but *scientific* knowledge, which is KNOWLEDGE WRIT LARGE (as Giere puts it). Now let that be the part of the equation d-cog remains silent on.

At least, that conclusion is forced on us, I have argued, if we accept Giere’s “knowledge is a human affair only”-clause, for under that assumption, whatever makes the difference between data and knowledge must be located within the traditional human skinbag. And for those difference-makers, we already had a toolkit: i-cog.

There are three alternatives. First, we accept d-cog’s parity principle, implying that scientific knowledge (not just data) may be produced by artifacts (in principle: no humans needed). Second, we re-engineer the concept of scientific knowledge: we deflate the notion (scientific knowledge *is* just data), or redefine it any other way that brings it in alignment with d-cog.[11]

When choosing either alternative, the payoff should be (made) clear. *Perhaps* d-cog explanations of scientific practice outperform i-cog explanations, and *perhaps* d-cog is capable of unifying the concerns of historians, sociologists and philosophers of science. To date, however, we lack the necessary evidence to validate these claims. Therefore, I suggest the following, third alternative: let us try to demonstrate d-cog’s payoffs empirically, but stick to i-cog for the time being.

[11] In some places, Giere seems to try this idea. He refers for instance to Popper’s idea of knowledge without a knowing subject (see Giere, 2007, p. 318), but immediately returns to more standard accounts of knowledge, such as Goldman’s epistemic reliabilism.

Acknowledgments

Research by Krist Vaesen was supported by the Netherlands Organisation for Scientific Research (NWO). He thanks Ronald Giere for his detailed comments on an earlier draft of this paper. He also benefited greatly from discussions with Andy Clark, Hans Radder, Wybo Houkes, Joel Katzav, Anthonie Meijers, Joel Anderson, Lieven Decock, Olle Blomberg, Marieke van Holland and Auke Pols.

References

Bechtel, W. (2009). Explanation: mechanism, modularity and situated cognition. In Robbins, P. and Aydede, M., editors, *Cambridge Handbook of Situated Cognition*. Cambridge University Press.

Clark, A. and Chalmers, D. (1998). The extended mind. *Analysis*, 1:7–19.

Flor, N. V. and Hutchins, E. L. (1991). Analyzing distributed cognition in software teams: a case study of team programming during perfective software maintenanc. In Koenemann-Belliveau, J., Moher, T. C., and Robertson, S. P., editors, *Proceedings of the Fourth Annual Workshop on Empirical Studies of Programmers*, pages 36–59. Ablex Publishing, Norwood, N.J.

Giere, R. N. (2002a). Discussion note: Distributed cognition in epistemic cultures. *Philosophy of Science*, 69(4):637.

Giere, R. N. (2002b). Models as parts of distributed cognitive systems. In Magnani, L. and Nersessian, N., editors, *Model based reasoning: science, technology, values*. Kluwer Academic Publishers.

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

Giere, R. N. (2006a). The role of agency in distributed cognitive systems. *Philosophy of Science*, 73(5):710–719.

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

Giere, R. N. (2007). Distributed cognition without distributed knowing. *Social Epistemology*, 21(3):313–320.

Giere, R. N. and Moffatt, B. (2003). Distributed cognition: Where the cognitive and the social merge. *Social Studies of Science*, 33(2):301–310.

Goldman, A. I. (2002). *Pathways to Knowledge: Private and Public.* Oxford University Press.

Hutchins, E. L. (1995). *Cognition in the wild.* MIT Press.

Magnus, P. (2007). Distributed cognition and the task of science. *Social Studies of Science*, 37:297–310.

Osbeck, L. M. and Nersessian, N. J. (2006). Distribution of representation. *Journal for the Theory of Social Behaviour*, 36(2):141–160.

Port, R. F. and van Gelder, T. (1995). *Mind as motion.* MIT Press.

Radder, H., editor (2003). *The Philosophy Of Scientific Experimentation.* University of Pittsburgh Press.

Rowlands, M. (2009). Extended cognition and the mark of the cognitive. *Philosophical Psychology*, 22(1):1–19.

Thelen, E. and Smith, L. B. (1994). *A Dynamic Systems Approach to the Development of Cognition and Action.* MIT Press.

Vermaas, P. E. and Garbacz, P. (2009). Functional decomposition and mereology in engineering. In Meijers, A., editor, *Handbook of the Philosophy of Technology and Engineering Sciences.* Elsevier Science.

Wilson, M. (2002). Six views of embodied cognition. *Psychonomic Bulletin & Review*, 9(4):625–636.