

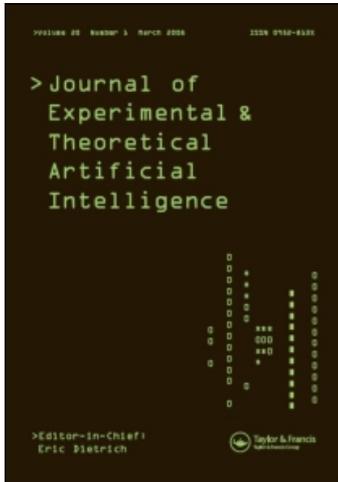
This article was downloaded by: [University of Memphis]

On: 14 August 2008

Access details: Access Details: [subscription number 788683960]

Publisher Taylor & Francis

Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK



Journal of Experimental & Theoretical Artificial Intelligence

Publication details, including instructions for authors and subscription information:

<http://www.informaworld.com/smpp/title-content=t713723652>

Pluralism, radical pluralism and the politics of the Big Bang

Eric Dietrich ^a

^a Department of Philosophy, Binghamton University, Binghamton, NY, USA

Online Publication Date: 01 September 2008

To cite this Article Dietrich, Eric(2008)'Pluralism, radical pluralism and the politics of the Big Bang',Journal of Experimental & Theoretical Artificial Intelligence,20:3,231 — 237

To link to this Article: DOI: 10.1080/09528130802319169

URL: <http://dx.doi.org/10.1080/09528130802319169>

PLEASE SCROLL DOWN FOR ARTICLE

Full terms and conditions of use: <http://www.informaworld.com/terms-and-conditions-of-access.pdf>

This article may be used for research, teaching and private study purposes. Any substantial or systematic reproduction, re-distribution, re-selling, loan or sub-licensing, systematic supply or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The accuracy of any instructions, formulae and drug doses should be independently verified with primary sources. The publisher shall not be liable for any loss, actions, claims, proceedings, demand or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

COMMENTARY

Pluralism, radical pluralism and the politics of the Big Bang

Eric Dietrich*

Department of Philosophy, Binghamton University, Binghamton, NY, USA

Imagine that, in your travels, you meet an earth systems scientist. She says that she is a *total* earth systems scientist. She is looking for one, single, concise, clean, elegant theory to explain Earth's magnetosphere and Van Allen radiation belts, all the different kinds of cloud formations, the flow and turbulence of Earth's atmosphere, the world-wide distribution of rainfall and snow, all of Earth's varied geology, including the dynamics of its molten iron core, the whys and wherefores of all of Earth's flora and fauna and the dynamics of the oceans at all different depths. Further, she states, her theory will incorporate atomic, molecular, chemical, etc. reductions of all the above phenomena. She goes to say that this sought-after theory has to mesh smoothly with the correct scientific account of Earth's history, as well as comport with the correct theory of the formation of the solar system, Milky Way galaxy and the universe. You, though not an earth systems scientist, seriously doubt that such a theory can ever be produced. Indeed your doubts are so strong that you wonder if you have happened on to a quack of some sort. Being forthright to a fault, you ask: 'Are you mad? Surely no agency is funding your research, for no one theory could ever provide such a comprehensive list of explanations'. Your interlocutor assures you she's tenured at a well-known university and the recent recipient of a large National Science Foundation grant. 'So what?' you say, dismissively, '... The state of Utah poured millions into cold fusion before the so-called experiment could even be replicated'. Unruffled, the earth systems scientist says that her project is not as crazy as it sounds, for all of the phenomena she reeled off are really manifestations of just one phenomenon – everything else can be explained in terms of it. 'Do you want to know what *it* is?' she asks. 'I certainly don't,' you say, and walk away muttering something about wasted tax dollars and the fact that the phrase 'theory of everything' was originally intentional physicist hyperbole, physicists being the kind of people that they are.

There are no such total earth systems scientists, thankfully, but there are cognitive scientists – actual scientists, tenured and funded, who really do think one theory is going to explain everything about the mind – or at least everything interesting. And typically they think this because they think they can reduce the vast realm of all thought to one thing, one fundamental thing, which itself admits of a single, concise, clean, elegant theory. The denial of this view is called *pluralism*. Phrased positively: pluralism is the view that 'explaining the mind' will really be explaining the many different levels and kinds of activity that go into producing a thinking, reflectively self-conscious human mind¹, and further, that such explanations will require many different kinds of theory quantifying

*Email: dietrich@binghamton.edu

over many different kinds of theoretical entities (e.g. neurons, synapses, ions, thresholds, differentials, connections, vectors, information, nodes, arcs, rules, percepts, concepts, self-models, etc.). The authors of the three target articles are to be commended for presenting the case for pluralism so well and so elegantly.

One can discern four major epochs in the short, unhappy life of cognitive science. First there were rule-based symbols, then neuralesque models, then situated embodiment, then dynamical systems. Each of these has failed for the exact same reason: they couldn't explain all of cognition (actually, they couldn't even explain much of it), so they each failed in their stated goal. The failures had a deeper similarity: looked at together over the last 60 years, it is clear that the dominant theories in each epoch were playing a zero sum game: theory T_i could explain small aspects of brain–mind process X , but only by giving up any chance for explaining brain–mind process Y . Y was then demoted, or a check was written for it. If one of these theories would have won, the others would have lost. And *all cognitive scientists wanted it that way* (this is why cognitive science's short life is unhappy: we're competing, working against each other, and we're doing this because we assume there's just one thing to explain). Indeed the large majority of cognitive scientists still want it that way, hence this special issue. But now, with four major failures and the accompanying data in our quiver, perhaps cognitive science is ready to get serious about explaining the brain–mind complex and to kiss off the antiquated and harmful approach of 'winner take all'.

Pluralism is an absolute necessity in cognitive science because of one brute fact: the working human brain is among the most complicated things in the universe – a universe, note, *filled* with complicated things. It is madness – collective madness, it turns out – to suppose that one theory can explain all that needs to be explained about the brain and its emergent mind.

Dale, at the end of his excellent paper, asks you to consider your reaction were you to bump into a person who has a 'theory of the Mississippi'. The fact that I have asked you to consider bumping into such a person *on steroids* exposes the fact that Dale and I appear to disagree over how complex the brain–mind is. I know he considers the mind very complex. He is careful about making this point in his paper. My view is that the brain–mind is *so* complex we have to use a mangling phrase lying athwart several existing disciplines just to refer to the subject matter: 'brain–mind'. When it comes to complexity, it is useful to consider a couple of very hard facts. We currently cannot even make a machine that can walk as well as a cockroach (these little beasts walk extremely well, so simulating them is a worthy goal, but, we've got a long way to go). Though there are robots which mimic cockroach locomotion somewhat well, they are quite slow...and the size of toasters, plus they cannot dodge or hide from predators, eat or mate. A cockroach can live for a month on the glue from the back of postage stamp. As far as I know, no one is even researching that one, though it would obviously be very nice to know how the little dears do that. We cannot make a machine that can even come somewhat close to passing the Turing test. We have no idea how the neurochemistry of the brain produces thoughts, let alone the whole symphony required for a paper like this one, let alone the kind of thought required to produce a book, like say, *On the Origin of Species*. We have no idea how emotions fold into thought, and only a dim idea of how they are rational (many still consider rational thought the antithesis of emotion). We are clueless about creativity. How did Einstein dream up his theories? And they were 'in the air', so the historians of science tell us. How in the blazes did Gödel figure out that arithmetic was incomplete when what was in the air was the direct opposite of what he wound up proving? How did Lennon and

McCartney dream up Ticket to Ride or Day Tripper? We are so in the dark about the self that it is currently quite the fad to deny it exists (Metzinger 2003). And consciousness... our abject cluelessness on this problem is positively frightening (Dietrich and Harcastle 2004). And, all of this is just for starters. The complexity I am calling your attention to is easily revealed in the vastness of what we don't know. But it is more readily revealed when what we don't know is contrasted with the vastness of what we do know. Cognitive science is huge. We in fact know tons about the brain–mind. But what we don't know dwarfs what we do know. Often, at conferences, I hear (only over beers, in the evening) that cognitive science is still awaiting its Darwin or Einstein. That is completely wrong. We are still awaiting our Da Vinci, Descartes, Galileo and Newton. We thankfully already have our Thales and Aristotle, maybe.

I sometimes wonder if cognitive scientists more or less deliberately blind themselves to the brain–mind's complexity because it is hard to do cognitive science if one really considers what one is up against. Had Mallory really known what Mount Everest was all about, he probably wouldn't have attempted to climb it, in wool knickers using hemp ropes. And he then wouldn't have fallen a few thousand feet to his death.

But Jilk and co-workers have not blinded themselves; they have faced the brain–mind's complexity head on. The SAL model of Jilk and co-workers is a beautiful example of a working pluralist computational model. I commend it to all cognitive scientists. Among its numerous good properties is that it highlights a central brain–mind process that I've been jumping up and down about for years (Dietrich 2000, 2001; Dietrich, Markman and Winkley 2003) – a process that doesn't get anywhere near the attention it deserves. I refer of course to *abstraction*.² Abstraction takes us from light bouncing onto one's retina to recognising the opportunity for a double play in baseball. Abstraction is at the heart of categorisation, no matter how primitive (some plants can even do it, apparently). Abstraction is crucial for cognition of any sort, and it is the key to human-level intelligence. One couldn't ask for a more central phenomenon. SAL doesn't tell us what abstraction is and how it works, but Jilk and co-workers point out that SAL makes the problem of abstraction *explicit* so that they can now work on it directly (see their section 6, Continuing Research). SAL is still a hybrid system, but future SALs may not be, and are certainly not intended to be – see, for example, their figure 7. If Jilk and co-workers build on their success, they may well solve the problem of abstraction.

Finally, I agree with everything Edelman says. But that's because I am an old computationalist from way back (see, e.g. Dietrich 1990). Edelman shows that computation can be the backbone of a robust pluralism. In this regard, I would like to point out what may have escaped another reader's attention (and which allows me to stress again the importance of abstraction): Edelman, in his paper (and in his book, in press), gives what is in fact a non-standard, new way of understanding computation. He says: 'A computation is a process that establishes a mapping among some symbolic domains', and 'a system instantiates a computation if its dynamics can be interpreted (by another process) as establishing the right kind of mapping'. He says that his view of computation is quite broad; broad equals abstract. But his definition is just the kind of abstracting of the fundamental notion of computation that cognitive science needs.

Still, I worry about the fact that computation is *one paradigm*. Edelman is committed to the view that when all is said and done, the multiple theories comprising a future cognitive science will form a nice hierarchy, each level grading into the other. Edelman's idea is that computation will be the unifying theme of *all* of the brain–mind's diverse theories. Indeed, in a deep sense, the theories' relation to one another will embody the very

abstraction processes they will explain. Philosophers call such a relation as envisioned by Edelman *supervenience*. (Definition: *A* properties supervene on *B* properties if no two possible situations have identical *A* properties while differing in their *B* properties; or put more epistemically, *A* properties supervene on *B* properties when specifying completely the *B* properties completely specifies the *A* properties.) Edelman's claim then is that, when the dust settles, cognitive science will reveal a beautiful hierarchy of computational theories, each supervening on the next one down because, in fact, the myriad components of the brain–mind all supervene on the next one down, beginning with the highest most. And, almost certainly the other contributors of the main target articles of this issue agree with Edelman, at least to this extent: there will be one unified hierarchy of brain–mind theories (Dale, perhaps, might protest the computational part).

This is indeed a grand view. . . . But what if it is false? What if the brain–mind is so complex that not even one paradigm can contain it? What if computation is only useful within one (presumably large) *research silo* (to use Jilk and co-workers' suggestive metaphor)? What if, in short, pluralism is not enough? What if we need *radical pluralism*: multiple paradigms of wildly different kinds explaining the myriad aspects of the brain–mind? Could the brain–mind be that complex?

The answer is Yes. Let's return to consciousness. Consciousness is still so mysterious that dualism is alive and well. (I always get screaming protests when I make this claim, but it is a brute fact, screaming about it won't help. Also, I recommend reading Chalmers (1996).) Understanding consciousness in any deep but still standard scientific way (i.e. reducing it to neural processes) is not in the cards, basically because dualism will always *seem* true (Dietrich and Hardcastle 2004). Yet for all we will ever know, consciousness could be a physical process. That's a pretty unhappy result (or happy, depending on your point of view). It suggests that a 'science' of consciousness will not look like other sciences and may in fact be limited to cataloging correlations between conscious states and various inputs and various neural states. If this is correct, then computation is worthless for understanding consciousness. But consciousness is pretty darn central. So, a radical pluralism is needed already.

But perhaps consciousness can be put to one side, ignored, winked at. After all, virtually all of cognitive science proceeds by doing just this. Dream on.

It is clear that being relative to, or only accessible from *points of view* is the fundamental property of all epistemic information. Indeed, objectivity itself turns out to be just one point of view, contrasting with subjectivity (see also Nagel 1979, 1986). So, if, as seems likely, points of view are essential to understanding anything, and if points of view differ necessarily, then we are stuck with different understandings. So far, ok . . . but couldn't these understandings align; couldn't they form a nice, well-behaved hierarchy – of the sort envisioned by Edelman and other hopeful cognitive scientists? I think not.

Many of the points of view influencing our knowledge are simply incommensurable. And they are obviously so. Just look at all the different kinds of logic there are. Many of these are incommensurable (e.g. paraconsistent and consistent logics); one chooses a logic based on what one is trying to accomplish. The rest of mathematics is crawling with examples of this phenomenon (e.g. geometry). This phenomenon shows up in physics – in the radical difference between quantum physics and the physics of the ordinary-sized. And, the same phenomenon shows up in psychology where it has been clear for decades that one cannot replace the science of the mind with the science of the brain, even though the mind supervenes on the brain. (So, in this sense, we've known for decades that at least a minimal

pluralism was going to be required to fully understand the brain–mind, even if we ignore consciousness.)

But this is just beating around the bush. Philosophy shows conclusively that there exist incommensurable points of view on matters of great significance. Examples include: are our actions free or determined? What is a moral act? What, if anything, is metaphysically basic? Is consciousness a material property of the world or not? Philosophy is the study and mapping of these incommensurabilities (again, see Nagel 1979; 1986). Philosophy has never conclusively solved a single important problem. And worse, psychology is more or less powerless to say what, really, philosophers are up to when they engage in philosophical thinking. So, psychology looks to be unable to explain why the points of view are incommensurable. Certainly, philosophy cannot offer such an explanation (though some have made valiant attempts, see McGinn 1993).

So, we are stuck with incommensurable points of view, and hence, we are stuck with incommensurable understandings of the world. Hence, radical pluralism is required for understanding the world (for more on this see, Cartwright 1999).

But so what? This doesn't seem to take us all the way to incommensurable theories of the understander – the brain–mind complex – itself. Cognitive science's explanations of cognitive phenomena work independently of the phenomena's contents. So, for one example, given a theory of reasoning it doesn't matter that Jack is cogitating about Jill, π , or world peace; his reasoning mechanisms postulated to be the same in each case. For another example: the current dominant theory of analogy-making proposes an algorithm for constructing them that is independent of the contents of the memory items over which the analogy is made (see, e.g. Gentner 1983). In short, and this is usually touted as one of the big advantages of computationalism, there is a big difference between syntax and semantics, and the point here is that it seems like a lot of cognitive science can be done by just worrying about the syntax, the semantics are irrelevant. If this is correct, then though radical pluralism is required for understanding the world, such robust pluralism may not be required for understanding the brain–mind.

But this is wrong. It forgets that cognitive science is a *science*, the fundamental task of which is to *explain cognition*. It is impossible to figure out how something works without knowing what it is working on. To see this, imagine that you get up one morning and there's a big, weird thing in your backyard (or living room or kitchen). You've never seen one before. Neither has anyone else. The thing has various behaviours: certain precipitating conditions lead to certain outcomes. In order for you to figure out what the big weird thing is doing, you have to know how it is interpreting the precipitating conditions. Perhaps it interprets everything you put in it as bread and tries to produce toast, or perhaps it interprets everything you put in it as an ordinal number and tries to produce that n th number in the decimal expansion of π , or perhaps it interprets inputs as English sentences and responds accordingly. Let's call the big weird thing's interpretations *contents* (though, clearly in certain cases calling them this is an abuse of the term (bread is not a content), but it is not in other cases, e.g. the case of using language). The point is that you cannot figure out what the big weird thing is doing unless you can figure out how it interprets its inputs. And, you cannot figure out that unless you know what it is doing. This vicious circle is broken, as in all sciences, by the well-known process of hypothesise and test. To explain the big weird thing, you will have to bootstrap yourself into an explanation, and this explanation *will have to refer to the contents of the big weird thing's hypothesised internal processes*. So, contents are required to do cognitive science.

So you do need to worry about the contents of the brain–mind’s thoughts. You need to have them to figure out what the brain–mind is doing. But those contents come from and refer to the blooming, buzzing, dappled world. Since radical pluralism is required to understand it, radical pluralism is going to be required to understand the contents. Put it this way: The radically pluralistic sciences required to explain the world are citizens of the brain and mind, or better: The radically pluralistic *universe* in fact *resides* in the brain and its mind. To bring points of view back in: The very points of view, radically differing, required to explain the universe, reside in the contents of the brain–mind. Understanding that complex head-shaped thing on the top of your body, therefore, requires deploying the same pluralism needed to understand the universe. So, radical pluralism is needed to understand the brain–mind, and that’s independent of consciousness.

It is a brute fact of the world that there is more to it than physics. It might seem to the non-philosopher that this fact is either not worth explaining or is easily explained. But it is neither. An understanding of the world requires radical pluralism. And why this should be so is a mystery. But minds are at the centre of the world (in spite of what the physicists and mathematicians say). So, there is no reason to hope for mere pluralism when it comes to understanding that thing atop your shoulders.

I understand that the politics of doing good science essentially precludes all of us from actually warmly embracing radical pluralism. But it seems to me that minds living in a universe which appeared *ex nihilo* and which is expanding for reasons passing understanding should neither be so trepidatious, nor so greedy. In a universe whose very existence is essentially beyond understanding, we can still do good cognitive science even though we have to be radical pluralists. Indeed, it is the only way we will do good cognitive science.

Notes

1. I focus only on humans and human minds in this article, but only for that reason – to keep things focussed.
2. There is an excellent collection of papers on abstraction in the *Philosophical Transactions of the Royal Society*, London, B, 2003. And from there, one can find other good research on abstraction.

References

- Cartwright, N. (1999), *The Dappled World: A Study of the Boundaries of Science*, Cambridge: Cambridge University Press.
- Chalmers, D. (1996), *The Conscious Mind*, Oxford: Oxford University Press.
- Dietrich, E. (1990), ‘Computationalism’, *Social Epistemology*, 4, 135–154, (with commentary).
- . (2000), ‘Analogy and Conceptual Change, or You Can’t Step in to the Same Mind Twice’, in *Cognitive Dynamics: Conceptual Change in Humans and Machines*, eds. E. Dietrich and A. Markman, Mahwah, NJ: Lawrence Erlbaum, pp. 265–294.
- . (2001), ‘AI, Concepts, and the Paradox of Mental Representation, with a Brief Discussion of Psychological Essentialism’, *Journal of Experimental and Theoretical Artificial Intelligence*, 2001, 1–7.
- Dietrich, E., and Hardcastle, V. (2004), *Sisyphus’s Boulder: Consciousness and the Limits of the Knowable*, Amsterdam: John Benjamins.

- Dietrich, E., Markman, A.B., and Winkley, M. (2003), 'The Prepared Mind: the Role of Representational Change in Chance Discovery', in *Chance Discovery by Machines*, eds. Yukio Ohsawa and Peter McBurney, Berlin: Springer-Verlag, pp. 208–230.
- Edelman, S. (in press), *Computing the Mind: How the Mind Really Works*, New York: Oxford University Press.
- Gentner, D. (1983), 'Structure-Mapping: a Theoretical Framework for Analogy', *Cognitive Science*, 7, 155–170.
- McGinn, C. (1993), *The Problems of Philosophy: The Limits of Inquiry*, Oxford: Blackwell.
- Metzinger, T. (2003), *Being No One: The Self-model Theory of Subjectivity*, Cambridge, MA: MIT Press.
- Nagel, T. (1979), 'Subjective and Objective', *Mortal Questions*, Cambridge, UK: Cambridge University Press, pp. 196–213.
- . (1986), *The View from Nowhere*, New York: Oxford University Press.