

# Minds & Machines

How does the brain work? One of the world's leading figures in philosophy of mind is **Daniel Dennett**, who describes himself as a cognitive scientist, and who is deeply interested in the development of artificial intelligence. **Tim Madigan** recently interviewed him for *Philosophy Now*.

*In your latest book Brainchildren there is a chapter entitled 'Self-Portrait', in which you state that one must have a theory of intentionality before attempting to develop a theory of consciousness. This seems central to understanding a lot of what you've been doing over the years. Can you;*

*(a) explain in layman's terms what 'intentionality' means to you? And*

*(b) explain in layman's terms why we have to understand this before attempting to understand consciousness?*

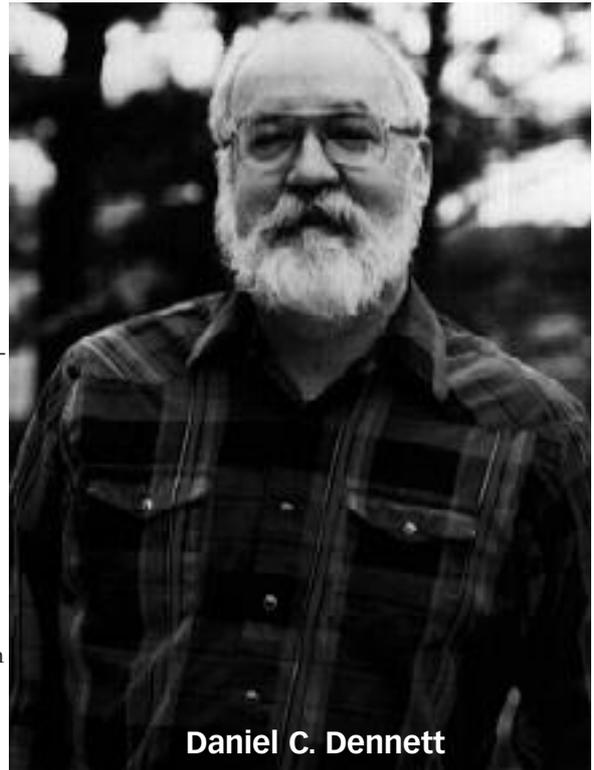
Does consciousness depend on intentionality, or is it the other way around? Intentionality is just *aboutness*, the way one 'thing' can be 'about' another 'thing' – a map of Dublin is about how the streets are laid out in Dublin, and *War and Peace* is about (among other things) Napoleon in Russia – and right now my thoughts are about this answer, and so are yours. The last case seems to many people to be 'original' or 'fundamental' or 'intrinsic' in a way the others aren't. They are merely cases of 'derived' intentionality, derived aboutness. According to this view, you should first explain consciousness, and the intentionality of conscious thoughts and experiences and *then* explain the lesser varieties of derived aboutness (or even 'worse': mere AS IF intentionality, like the aboutness of the patterns of natural selection, whimsically viewed as "what Mother Nature had in mind." I have long maintained that these simple, modest, imperfect forms of aboutness are both the ultimate historical source of, and the crucial ingredients in, the fancier forms of intentionality discernable in our streams of consciousness or phenomenology. So first you should have a theory of intentionality (of intentional systems, conscious or not), and then on that foundation you can build a theory of consciousness (of one sort or another).

*How accurate do you think the analogy is between the mind and a computer?*

*Would it be possible for instance, to someday download the mind of Tim Madigan onto some sort of storage medium and then load it up on another brain at a later date? Could one manufacture multiple copies?*

The analogy is excellent, but imperfect. The idea of a mind as a perfectly substrate-neutral, medium neutral system of information processing – which is what a computer is – is an idealization of great power for many purposes, but it can also be deeply misleading. To take your example, the mind of Tim Madigan is so inextricably tangled up with the body (and the endocrine and paracrine systems of hormones and neuromodulators) that although we can lift it out of its substrate *in our imagination*, we certainly couldn't do it in practice. Only a molecular level duplicate of you would have much real chance of preserving the psychological you, your mind. 'In principle' that is possible, but not in practice. Actual computers, on the other hand, are already perfectly storable and movable. Every quirk, every memory, every habit and glitch of your laptop can be not only saved, but sent effortlessly and at low cost over the Internet. But in the end this is a difference of degree only. Imagine a 3d fax machine that could use CT scanning to copy very thin 'slices' of a sculpture and send them electronically to a manufacturing plant that would

thereupon construct duplicates that even the sculptor couldn't tell from the original. Possible. Yes indeed. Now imagine one that can do the same thing with a potted geranium, and do it at such a fine-grained level (the atomic level) that the duplicate went right on growing! Possible? Well, probably not, but only because it's a matter of degree. It's the same with minds.



**Daniel C. Dennett**

*If this analogy is a reasonably accurate one, do you think that in principle we could survive the deaths of our physical bodies in some way?*

Well, I've more or less answered this question above. "In principle" we could survive the deaths of our physical bodies, and probably information storage is more practical, in the end,

than cryogenics (cold storage), but both are technically beyond us, probably forever.

*How do you distinguish 'self-consciousness' from general consciousness. Do you feel other living things besides humans possess the former?*

I don't like the term 'self-consciousness', since it means many different things that are typically run together. The lobster has a simple variety of self-consciousness: when it's hungry, it doesn't eat itself. (It could; it can rip off its own legs, and it will happily eat the parts of other lobsters in the vicinity). A teenager is self-conscious in an entirely different sense. Then there are the capacities for indefinitely many layers of reflection: I can notice my pain, notice my noticing, wonder about whether I would have noticed the pain if, ... etc. That is a very important phenomenon, and it is, so far as we know, a human-only talent. I wouldn't call it self-consciousness, but others might.

*What is your view about the changing fashions of philosophy of mind over the past few years? Is materialism more accepted today than when you first entered the field?*

More important than fashions in materialism (which I think has been the ambient or default view for the last quarter century, with some brave exceptions) is the fashion in how closely philosophy works with the empirical sciences of the mind. In my student days one was deemed eccentric at best if one took a serious interest in the brain or in empirical research in psychology. Now it is just the opposite: although some hardy iconoclasts still do what they call philosophy of mind in the old-fashioned armchair way with only a sidelong glance at science, they have a hard time being taken seriously. I think this is an excellent development. I think philosophers of mind who remain ignorant of the sciences of the mind are like aestheticians who have never studied art or music. Of course it is possible in principle to do good work in

philosophy of mind entirely untutored and untrained in science, just by astute and acute reflection. It is also possible in principle to do good work in the aesthetics of literature without ever having learned to read, and without ever having any novels or poems read to you. Anybody want to take any bets, though, on how likely this stunt would be? Possibilities in principle are too numerous to take seriously.

*What philosophers have most influenced your own work? How important was Gilbert Ryle as an influence, and why do you feel his work is rather neglected?*

Quine and Ryle were my mentors, my main teachers, during my undergrad and graduate career, of course, but the writings of Wittgenstein, Sellars, Putnam, and others were major guides. More than most philosophers, I think, I've been inspired by non-philosophers; by Turing, Darwin, and others, I did a seminar on Ryle and Wittgenstein a couple of years ago, and was astounded to discover that my students found Ryle tough going. They didn't think it was as lucid as I did. I have lots of thoughts about why this might be, too many to write down here. A project for another occasion.

*You are often described as a 'cognitive scientist'. What is the connection, if any, between cognitive science and philosophy, and why should the latter field be more concerned about the former? How does cognitive science differ from philosophy proper?*

I've already addressed this above, in the main. I view philosophy as just the most abstract, most conceptual way of doing cognitive science. Philosophy of mind done well is cognitive science. Philosophy of mind done poorly is a variety of fantasy of unknown value – typically not very amusing or even interesting.

*How close are we to achieving artificial intelligence? Can you fill us in on the latest news regarding Cog?*

Cog is running behind schedule, but not because of any major theoretical or

## COG

Since 1993 a team at MIT has been trying to design and build a humanoid robot capable of speech, eye-coordinated movement of objects, and various self-protective, self exploring activities. "The aim of the project is not to make a conscious robot, but to make a robot that can interact with human being in a robust and versatile manner in real time, take care of itself and otherwise tell its designers things about itself that would otherwise be extremely difficult if not impossible to determine by examination." The project was initiated by Rodney Brooks and Lynn Andrea Stein, and Daniel Dennett is an active participant. Cog is roughly the same size as a human adult, and has two human-length arms, but no legs. The robot's processing power comes from what is essentially a collection of Macintosh computers.

(for more info see 'The Practical Requirements for making a conscious Robot' in *Brainchildren* by Daniel C. Dennett)

scientific setbacks. The principal participants have just been unable to devote more than part time to the project. Still, the progress is heartening, and impressive. Cog's infancy is still only partially in place – for instance, the magnificent hand that Yuki Matsuo made for Cog turned out to be too heavy and not quite strong enough, and she went off to do other things, leaving Cog without hands – a major shortcoming in a robot who needs to manipulate things the way babies do. Until new and better hands can be designed and built, a lot of the most fascinating developments have to be postponed.

We are still a long way off from the sort of AI that could pass the Turing Test, but it's exciting, again, in AI, thanks to 'situated robotics' in its various flavours, and Artificial Life (and also connectionism, of course, in its various flavours). Lots of good things are happening, after some years in the doldrums.

“ Philosophy of mind done well is cognitive science. Philosophy of mind done poorly is a variety of fantasy of unknown value... ”

*Why did you feel the need to write a book defending Darwin, and why is it that the implications of evolution still are not fully understood?*

I've been taking Darwinian evolution for granted since my student days, and I was often astonished to discover pockets of ignorance and resistance that just shouldn't have been there. It finally became clear to me that the reasons for this anxiety and hostility were interesting in themselves, not *merely* obstacles to understanding good theories. That's when I decided to write the book. The reception it has had surely has confirmed my hunches, by flushing some remarkable and unexpected Darwin-dreaders out of the bushes. This is good. One can no longer be an anti-Darwinian without risking a challenge. I think that many people view Darwinism as the ultimate expression of a vision (often called 'reductionism') that they think is optional in science, and that they view as a threat to whatever they hold dear. It is not optional, and it is not the

threat they think it is (though it does subvert some of the traditional ways of shoring up the things we hold dear).

*Finally, you have been one of the few philosophers who constantly demonstrates a sense of humour. Why are most philosophers – at least in print – such dour folks?*

I've often wondered. Bertrand Russell once said something like this (not a direct quote, alas):

"People often make the mistake of thinking and 'humorous' and 'serious' are antonyms. They are wrong. 'Humorous' and 'solemn' are antonyms. I am never more serious than when I am being humorous."

That perfectly expresses my sentiments, though I know that I pay a price in some quarters: some colleagues refuse to take seriously anything that makes them (or – perhaps more important – their students) laugh. It is certainly true that a joke can distract our attention from the flaws in an argument, but I think that's a rare danger, well worth living with. If you

think, as I do, that the main obstacle to progress in philosophy is the tenaciousness of bad habits of thought, then you think that laughter, shock, surprise, are good levers for dislodging these habits.

*Thanks again!*  
You're welcome.

*Professor Daniel C. Dennett is Director of the Centre for Cognitive Studies at Tufts University, Massachusetts. He is the author of numerous books, including Consciousness Explained, Darwin's Dangerous Idea and Brainstorms.*

*Timothy J. Madigan is the Editor of FREE INQUIRY, the international humanist magazine based in Amherst, NY, and is also a US Editor of PHILOSOPHY NOW.*

