



Cognitive Science Online

A journal published by the
Department of Cognitive Science, UCSD

Vol 1, Issue 1. Winter, 2003

Departments

Articles
From the editors
The big chair
Interview
Ask Marta
Art and humor

About us

Information
Subscription
Submission
Contact us

Home
Download
Back issues

about this column

Column: Interview with David Chalmers



David Chalmers is a Professor of Philosophy and Director of the Center for Consciousness Studies at the University of Arizona. His work focuses on the philosophy of mind, and he is especially interested in consciousness, artificial intelligence and computation, philosophical issues about meaning and possibility, metaphysics, and the foundations of cognitive science.

His undergraduate degree was in mathematics and computer science at the University of Adelaide in Australia. After graduating, he was a graduate student in mathematics at the University of Oxford, but soon switched to Indiana University, where he worked in the Center for Research on Concepts and Cognition, obtaining his Ph.D. in 1993 in Philosophy and Cognitive Science. He then spent two years as a McDonnell Fellow in the Philosophy, Neuroscience, and Psychology program at Washington University in St. Louis. After nearly four years as a professor in the Department of Philosophy at UC Santa Cruz, Chalmers shifted to the Philosophy Department at the University of Arizona where he has been since early 1999. At the moment, he is working on a book concerning the connections between reason, meaning, and possibility, tentatively entitled *On What Might Be*.

David Chalmers is perhaps best known for his monograph, *The Conscious Mind: In Search of a Fundamental Theory* (1996, Oxford University Press), in which he examines the problem of consciousness and presents in rigorous detail his own non-reductive theory linking the physical and experiential, as well as providing support for a strong version of AI.

He is also the editor of a collection of readings in philosophy of mind, *Philosophy of Mind: Classical and Contemporary Readings* (2002, Oxford University Press), as well as the motivation for another book, *Explaining Consciousness: The Hard Problem*, Edited by Jonathan Shear (1997, The MIT Press). The latter includes Chalmers' seminal paper introducing the 'hard problem' of consciousness, "Facing Up to the Problem of Consciousness," and a collection of 26 papers responding to the 'hard problem' as well as Chalmers' ultimate response.

David Chalmers' web page can be found [here](#).

The following interview was conducted by **Christopher Lovett** with support and ideas from **David Gropp**.

Q: For those of us unfamiliar with your work, could you briefly say a bit about your major areas of research, what you have been doing since the publication of your book, *The Conscious Mind*, in 1996, and what you're working on now? Has the "hard problem" gotten any easier?

I work in a number of areas of philosophy and cognitive science, but I am especially interested in the philosophy and science of consciousness. I'm concerned both with philosophical questions about the place of consciousness in nature, and scientific

questions about how consciousness can be studied empirically. In my 1996 book I argued that consciousness is an irreducible part of nature, and in particular is not reducible to brain processes, though it is systematically correlated with brain processes.

Since then, I've been working on many different things, but I've been especially concerned with the question of how a rigorous science of consciousness can develop. I don't think the hardest problems of consciousness will be solved overnight, but one thing we've learned is that we can make a lot of progress in the science even without solving those problems. For now, the science of consciousness is a science of the systematic correlations between what we might call the "third-person data," about brain processes and behavior, and "first-person data" about consciousness. The centerpiece is the search for neural correlates of consciousness (or NCCs). Understanding this correlation is vital to understanding the relationship between physical processes and consciousness, and it raises a lot of foundational questions: what is an NCC? what are the criteria and methods for finding such a thing? I've written a couple of papers on these issues.

Q: Is the work you do now the same type of work you were doing as a graduate student? How have your work and interests evolved since the beginning of your career? What were the major factors that influenced this evolution?

I started out as an undergraduate student in mathematics in Australia, and I always thought of myself as oriented toward science. But over time, I came to feel that the most important problems in mathematics had already been solved. I gradually became obsessed with the problem of consciousness, which seemed to me to be the biggest unsolved problem in science. I ended up leaving my graduate work in mathematics at Oxford to do a Ph.D. in philosophy and cognitive science at Indiana. While I was there I did a lot of work on different topics, including some papers on connectionist models of language and evolutionary approaches to learning. But consciousness stayed my greatest interest, and I wrote a dissertation on the topic that eventually became my book. Over time, I'd say I've become more and more interested in an ever broader range of issues both in philosophy and in cognitive science. One nice thing about being a philosopher is that one is allowed to be professionally interested in a very diverse range of topics.

Q: Which people have most influenced your thinking about cognition and in what way(s)?

I'd say my interest in thinking about the mind was spurred especially by reading Hofstadter's *Godel, Escher, Bach: An Eternal Golden Braid* as a teenager, and also Hofstadter and Dennett's *The Mind's I*. These provided an enormous amount of food for thought for a philosophically inclined kid who didn't know what philosophy was. Later on, I worked closely with Doug Hofstadter as a graduate student, so I'd say he's certainly the greatest influence on my thinking, even though it turns out that my views about the mind are pretty different from his.

Q: What kind of background do you think best prepares someone to do interdisciplinary research in your area?

Where the study of consciousness is concerned, I think it's important to have a background in cognitive psychology (especially perception and memory), neuroscience (especially the cognitive neuroscience of vision), and the philosophy of mind. Linguistics and computer science have been somewhat more ancillary in this field, at least so far: the connection between consciousness and language is hard to pin down, and it's hard to know the criteria for modeling consciousness computationally. On the other hand, I have a strong background in computer science and have never regretted it.

Q: In a previous interview with Andrew Chrucky (Philosophy Now - 1998), you say that "I argue that neuroscience alone isn't enough to explain consciousness, but I think it will be a major part of an eventual theory." Can you explain why you don't

think neuroscience is enough and suggest what other research is needed to complement work in neuroscience?

Both of these are big questions. Basically, I think that neuroscience is well-suited in principle to explaining behavioral and functional capacities: learning, language, and so on. But explaining consciousness isn't just a matter of explaining behavior and explaining functioning. Even once one has explained things, there's always a further question: why is this accompanied by subjective experience? And a neural account alone leaves this question unanswered. So one might say that neuroscience provides correlation with consciousness, but correlation is not explanation.

I've argued that there must be certain primitive principles, or laws, bridging physical processes and consciousness, that play a role analogous to fundamental laws in physics. The question then is how we find these laws. To do this, I think one needs a combination of neuroscience and psychology, gathering objective "third-person data" about brain and behavior, with phenomenological and introspective methods gathering "first-person" data about consciousness. One then needs to systematize these data and abstract general principles. First there will be quite specific principles about specific neural correlates of certain sorts of specific subjective experiences (of the sort we're just starting to get at now), but eventually we can hope to find the very general and primitive principles that underlie these correlations. These principles will be at the core of our theory of consciousness.

Q: Suppose the "hard problem" were suddenly solved and we understood consciousness (defined as subjective experience) fully, or at least to the extent that most scholars agreed that it is a fundamental property of the universe, as is gravity. Then what further questions would you be interested in rigorously investigating? Are there any other unanswered questions that have piqued your interest as much as the question of consciousness?

Well, I think that the general "hard problem" is just one of the things that makes consciousness intriguing. Even if we somehow understood with clarity why consciousness arises, there would still be all sorts of interesting questions about, say: the relationship between conscious and unconscious cognition, the representational structure of consciousness, the unity of consciousness, the relationship between sensory and nonsensory consciousness, and so on. All of these are things I've been thinking about a lot.

But setting consciousness aside, I'm also interested in questions about the meaning of language and the representational contents of thought. I've spent a lot of time recently trying to develop a framework in which meanings and content can be formally represented in a way that helps us to understand the relationship between language, thought, and the world.

Q: What do you think are the most important "big questions" in Cognitive Science? What are some smaller versions or special cases of those that we can work on now that you think will shed some light on the big ones eventually?

Not surprisingly, I think the biggest question involves developing a scientific understanding of consciousness. My favorite "smaller" question is how we can understand the mechanisms and the experience of humor. I think we don't really have a clue of why humor exists, or what it really involves. Why did evolution bother to make things funny? What was the point? What did mechanisms in our brains make us appreciate funniness? Why is the experience of funniness so distinctive? I think a lot of questions about consciousness arise again in this apparently much smaller microcosm.

Q: How much does our understanding of cognition hinge upon our understanding of consciousness? Is it possible for us to achieve a deeper understanding of more specific aspects of cognition such as attention, memory and learning, pattern recognition, decision making and language without first having a deeper understanding of what consciousness is?

I think we've gotten a long way in understanding many of these things without having a deep understanding of consciousness, and I don't see why an even deeper understanding shouldn't be possible this way. If we want to explain the behavior and functioning associated with these phenomena at a mechanistic level, it's not clear that one needs to bring consciousness in. And in all these cases, one can get a long way just by focusing at a mechanistic level. Still, in many of these cases, that won't be the full picture. Attention, for example, has a rich phenomenological nature, and one might say that understanding the mechanisms of attention doesn't explain that nature. Similarly for the conscious aspects of decision. So eventually we will need to integrate our understanding of mechanisms here with an understanding of subjective experience.

Q: What do you think have been some of the greatest advances in your area of expertise over the past 10-20 years? What do you think will be some of the greatest breakthroughs in your area of expertise over the next 10-20 years?

In the science of consciousness, the greatest breakthroughs have probably been in the study of neural correlates of visual consciousness: we don't have anything like a full understanding here, but there is suddenly an active and thriving research program where before there was very little. There are limitations on what can be discovered using just brain imaging studies on humans (at a very coarse grain) and neuron-level studies on monkeys (who can't give verbal reports of conscious states), but people have developed ingenious methods to circumvent these limitations as well as they can. I expect the biggest future progress to also be in the neuroscience. If someday we have technology that allows us to get around the limitations just mentioned, e.g. by being able to noninvasively monitor neuron-level processing in humans, then I think the field will be revolutionized. Probably that would lead to the golden age for the science of consciousness, by allowing the simultaneous collection of really specific and informative third-person data (at the neural level) and first-person data (via verbal report).

In the philosophy of consciousness, progress is more incremental, and issues are hardly ever settled once and for all. But I think in the last few years people are coming to understand much better the landscape of options for understanding the place of consciousness in the natural order. An increasingly active topic of research just recently is the relationship between consciousness and representation: just how can we understand the representational content of conscious experience? There was a very exciting six-week institute on this topic in Santa Cruz earlier this year, and I expect this topic to be at the center of the philosophy of mind for at least the next decade.

Q: How do you define cognitive science? What are the advantages and disadvantages to having a single interdisciplinary field rather than a collection of individual and separate fields of scientific and philosophical inquiry? Are we losing precision while gaining breadth in answering questions about the mind?

I think cognitive science is the scientific study of the mind. More specific definitions in terms of specific research programs or methodologies (e.g. the computational research program) are overly limiting. I think the ultimate data that cognitive science seeks to explain are the data of both behavior and consciousness. And to do this, it can use any methods available. Cognitive science is still more of a collection of fields than a single field, and I think this is not such a bad thing: one needs careful attention to fine details to make progress, and that requires specialization. There's hardly anyone who's a real expert in all the relevant disciplines, though increasingly many who are expert in more than one. Still, communication between the fields has gotten much better, with the result that there's now a sense of a collective understanding in a way that would have been harder to make out a few decades ago. Certainly in the science of consciousness, it's now possible to see connections between all sorts of neuroscientific and psychological work in a way that was poorly understood just a decade or two ago. And this big picture helps to guide individual work in turn. I think this sort of two-way interplay between big picture and details is the most beneficial consequence of the move toward "cognitive science".

Q: How would you say your early training as a mathematician has influenced your

work? Knowing what you know now, would you have pursued your undergraduate and graduate studies differently? Was there any specific event or conversation or piece of writing that solidified your decision to shift from mathematics to philosophy?

I always thought I would be a mathematician until I wasn't any more. I loved mathematics as a student, and still have a soft spot for it today. I think my change of fields was due to two things. First, the more I advanced in mathematics, the more it seemed to me that current research was less fundamental and in a way less important than the research of centuries ago. Mathematics is just so well-understood these days. Even though there's a lot that isn't answered, I had the feeling that I'd have loved to have been working around the time of Newton when everything was up for grabs. By contrast, the study of the mind right now is wide open, and there's so much unexplored territory. At the same time, I'd had a sort of amateurish interest in consciousness for a long time, and would come up with various speculative theories of it from time to time. Before I went to Oxford for graduate study in mathematics, I spent six months hitch-hiking around Europe, and spent a lot of time thinking about consciousness. By the time I got to Oxford, I was obsessively trying to work things out. This occupied my attention at the expense of mathematics, and I ended up deciding to switch fields. A lot of my friends and family thought I was crazy at the time, probably with some justification, but I guess it worked out OK in the end.

My current work really isn't very mathematical, though occasionally issues about mathematics come up, which I always enjoy. I'd say that my mathematical training is more relevant as training in a certain sort of thinking: highly analytic and rigorous thinking in abstract areas, guided by a sort of intuitive ability to see patterns and connections. That's how a good mathematician works, and it's the sort of thinking that a philosopher should aspire to. Of course rigor comes much more naturally in mathematics than in philosophy: it's inescapably built into the core methods of mathematics, where things in philosophy are never so black and white. People often think that a formal subject like mathematics must be much harder than a less formal subject such as philosophy, but my experience is the reverse: formalism makes a subject much easier, because a formalism does so much of the work for you. So philosophy is harder in a way, and you have to work hard to keep yourself honest, but I think one can at least aspire to rigor.

Q: If you could have a discussion with any philosopher or scientist no longer living, who would that person be? What are some of the issues you would like to discuss or questions you would like to ask?

I think it would have to be Descartes. He was both such an interesting philosopher and so far ahead of his time as a scientist. I think what I'd enjoy most is telling him about all the developments in the last few centuries in both philosophy and in cognitive science, and hearing his reactions. Somehow I suspect that he'd get up to speed very quickly, and would have all sorts of quite unexpected insights.

Q: Regarding Consciousness Studies, do you foresee it as a possible interdisciplinary concentration or major on either the graduate or undergraduate level in the near future? Would this major or concentration, as you envision it, be most likely offered by a philosophy department? How would such a major differ from current programs in cognitive science?

I don't really see a need for degree programs in consciousness studies distinct from those in cognitive science. Consciousness studies is an area of cognitive science and should be treated as such. People occasionally ask why our Center for Consciousness Studies at Arizona doesn't have its own graduate program. Apart from the fact that we don't have the resources, it would be crazy to set up such a program, since students who graduate from it would have nowhere to go. Instead, it's better for an interested student to get a really good grounding in a traditional discipline (philosophy, psychology, neuroscience), and at the same time gain a lot of interdisciplinary exposure to work on consciousness through courses, lab work, and so on. But that's just the sort of thing that a good cognitive science program (typically, a cognitive science minor or double major with another discipline) will provide. It's just a matter of orienting the selection of courses and so on in the right way.

[back to top](#)

In this section of *Cognitive Science Online*, we will feature interviews with cognitive scientists. If you have suggestions for future issues please **contact the editors**. Please also feel free to interview your favorite scientists and help out with this section of the journal!