



Thoughts and Experience

Issues facing contemporary philosophy of mind

Sean Crawford

In April 2004 I talked to Jerry Foder, who's Professor of Philosophy at Rutgers University in New Jersey. He's the author of many important and influential articles and books in philosophy of mind and cognitive science and is a well-known defender of the computational view of the mind. I began by asking him about the materialist metaphysics that is largely taken for granted by contemporary philosophers of mind and cognitive scientists.

One of the few orthodoxies in contemporary analytic philosophy of mind is the view that the solution to the mind body problem lies in the doctrine of materialism or physicalism, the view that roughly speaking the mental is not something over and above the material, of course questions arise as to what exactly not over and above means but anyway roughly the mental is just in some sense the material and so materialism and physicalism predominates in analytical philosophy of mind. The famous linguist and philosopher Noam Chomsky however, has argued in opposition to this that the mind body problem really has no coherent formulation because we have no coherent concept of matter or body or the physical with which to contrast the mental, what do you think of that line of thought?

Jerry Foder

Well, I think it's a polemically elegant move but I actually think it's a bit of a red herring and for the following reason. As far as anybody knows, and this is probably not a point of universal consensus but I'll bet it's pretty close, as far as anybody knows if there's going to be a physicalistic account of how the mind works it's going to be in terms of the macrostructure of the brain, by macrostructure I mean something like at the level of cells or proteins or fibres or something of that sort, but not at the level of the microstructure of matter, if it turned out that the mental phenomena went all the way down, down to as it were protons or electrons or quarks or whatever they've got down there, I think physicalism would be in deep trouble. Part of the story about materialism or physicalism or whatever one ought to call that stuff is that the relation between psychological phenomena, mental phenomena and neurological phenomena is sort of like the relation between, I don't know, geology and chemistry or something of that kind, if it turned out that before you got something that had both psychological and physical properties you had to go down to quarks that would be very puzzling indeed, in fact part of the story is that psychological laws or explanations are not supposed to take place at the level of basic science and what's being contemplated I suppose is that it's the internal basic structure of matter that might matter for the psychological and that seems to me very unlikely.

Sean Crawford

Is it just a mind brain problem then?

Jerry Foder

I think to all intents and purposes it is, yes, that is if it's not then nobody has any idea what to do about it at all. The reasonable assumption I think is that if it's going to be a physicalistic story about the mind it's going to be at the level of middle-sized objects like the brain. I mean after all there are a lot of other things that are made of the same kind of stuff that the brain is made of like for example, rocks or finger nails of whatever but they don't think and they're not conscious so it's a reasonable guess I would suppose that it's going to be what's characteristic of the structure of brains rather than what's characteristic of the structure of physical objects as such that's going to tell us the story about psychology or the mental if anything does.

Sean Crawford

You're often described as a paradigmatic functionalist, how would you describe the basic idea behind the kind of functionalism that you favour?

Jerry Foder

I think I'm inclined to plead innocent of the charge, if I'm a functionalist I think I'm a pretty unparadigmatic one. A functionalist is somebody who says, I guess, that mental states and processes and so on are susceptible of functional analysis, where functional analysis is one in which you characterise whatever it is that you're analysing in terms of its causal properties, that is what kinds of things it causes to happen, what kinds of things cause it to happen, rather than say its mechanical or physical structure or something of that kind. So the idea would be look if you want to know what a fan is, its something, as it might be, its something that's used to move the air to a first approximation anyway, it doesn't matter what its made out of, you can have electrical fans and hand fans and feather fans and so forth and so on, what makes a thing a fan is not what its made out of but its function in moving the air around. So now the question is whether what makes something a belief or a desire or whatever is its causal structure too, its causal properties and I guess if you're really an honest to God true blue functionalist like say Dan Dennet, then you think that the answer to that question is yes, so that belief that its going to rain or something is to be in some sort of dispositional state to carry your umbrella or to look out the window and say I hope it doesn't rain or whatever, I actually think that view is probably pretty hopeless. I think whatever it is that gives a mental state the content that it does, whether it's the content of a belief or as one says qualitative content of a sensory state isn't something that has to do with function, I don't think that the theory of content can be functional. On the other hand I do think that probably the difference between believing that it's going to rain and wanting it to rain is, or hoping that its going to rain, probably is a difference in functional role. So, my own view is that probably semantics and the intentional and conscious contents of mental state you're not going to get a functional analysis, which leaves open of course the question what can't they do. So, the amount of functionalism that one ends up with on this kind of story is pretty minimal.

Sean Crawford

Your latest book is called Hume Variations and it's a reference to the 18th century Scottish philosopher David Hume. What's Hume's contribution to the study of the mind and how significant is it?

Jerry Foder

What I stressed in my book is that he's running what would these days be called a representational theory of mind, that is the basic idea is that mental states are relations to some sort of, what other folk, Hume's terminology is extremely misleading but what other philosophers have called Ideas (with a capital letter), so what you have in your head is a lot of ideas, ideas of cows and ideas of typewriters and ideas of people and so on, concepts these days. These have two interesting properties, one is the sort of referential property, that is in some way or other they apply to things in the world, and the other property is causal, so they interact causally with one another and Hume's project was to make a psychology, quite explicit project in fact, was to make a psychology using those basic notions of ideas, semantical properties, referential properties of ideas and causal properties of ideas as the basic materials. Now there are obviously lots of differences in the details and some differences even in the main picture but to a first approximation that's what we're still doing. In that sense I think cognitive science we're writing these days is sort of footnote to Hume, not in the sense that one accepts exactly his notion of how mental mechanisms work but rather in that the idea of having a naturalistic theory of mind built around the notion of mental representation, was I guess the one that he stated, that he was the first to state with any degree of clarity and its still the idea that we're running on.

Sean Crawford

Where did Hume go wrong?

Jerry Foder

Oh, in all sorts of ways that is unless we've gone wrong. One thing is that he thought that ideas were pictures and that their semantic properties are to be understood in terms of something like resemblance for all sorts of famous reasons now not plausible. He also thought, however, and I think this is where there has been some real honest to God progress, he also thought that the causal relations among these mental objects, ideas or mental representations, or whatever, are fundamentally associationistic. The basic idea is that

mental processes, at least in so far as they are shaped by learning, and most associationists took for granted, in fact claimed as a virtue of the theory, that basically everything that's in your head is learned, so the idea is that mental processes in so far as they are shaped by learning depend upon associated bonds and associated bonds are just mechanical couplings that come into being as a result of frequency of pairings and that kind of relation among ideas so if somebody encounters, I don't know, salt often enough in the company of pepper then he'll come to think pepper when he thinks salt. Various elaborations of that basic suggestion are possible and some of the elaborations are quite elaborate elaborations but that's the basic idea. So frequency, continuity, that sort of thing produces causal relations among ideas and these causal relations among ideas are the basic structures out of which mental processes are made. There are again all sorts of reasons why that won't work and in fact I think what happened to Hume's project was that the associationism seemed more and more clearly unsatisfactory, this is a point that was made against association, against empiricists like Hume by Kant and then again by Frager for all sorts of rather deep reasons and an associationistic account of the causal relations among mental states and objects is hopelessly inadequate and nobody knew what to replace it with. I think the main idea that anybody's had in cognitive science since Hume was in effect Turing's idea that you don't need an associated picture of mind you can have a computational picture. That's really a very different account than was available to Hume of the character of mental processes and a whole lot of questions that seem hopeless from the associationist point of view seem a lot better from the computational point of view that we've inherited from Turing. So there was plenty wrong but still I think Hume was maybe the first person to have a clear picture of what ought to be done if not how one ought to go about doing it and that's a great achievement.

Sean Crawford

Why do you think that the mind is like a computer, or that some mental processes are computational processes as it's sometimes put and why should we give up associationism and go for a computational view, what's the driving motivation there?

Jerry Foder

There's a complex of considerations but one of them is this. A way to put it that might be familiar to a philosophical audience might be something like this. Philosophers are sort of used to the idea that thoughts have logical form, there are conjunctive thoughts and disjunctive thoughts and quantified thoughts and so on. There isn't any way to reconstruct that sort of picture in an associationist model and that's very worrying for an associationist because it looks like it's in virtue of their logical forms that thoughts have their inferential roles, it's because John and Mary went swimming as a conjunctive thought that one infers from it that John went swimming and Mary went swimming. That whole picture of logical form and thought and inferential consequences of logical form simply can't be recovered on any known associationist picture. So the computational model allows us to say things about how the mind works that an associationist simply seems to be unable to say.

Sean Crawford

I would have preferred to have summarised that by saying that we know that some mental processes or a lot of mental processes are in a sense rational processes and that the computational view is one of the only ones that can, in your view, that can make sense of the rationality of mental processes.

Jerry Foder

Yeh, that's right. I mean to a first approximation anyway that seems exactly right. The point is associative relations aren't rational relations and to reconstruct the notion of rational relations you need to give an account of the rationality of mental processes and those kinds of properties and thoughts simply aren't available to an associationist.

Sean Crawford

You often cite the thought processes of Conan Doyle's famous fictional detective Sherlock Holmes as an example of the rationality of mental processes and of course the computational theory is supposed to be very good as explaining the rationality of mental processes. Holmes' thought processes are rational and of course he's often called a master of deduction and we know that logically deductive inferences can be implemented computationally. We

know this because of the work by Alan Turing and various others but of course Holmes' great powers of thought did not really lie in their logically deductive nature, after all Holmes isn't really a brilliant magician, he's a brilliant detective and what makes him a brilliant detective is his power of abduction as its called sometimes, or his power of inferring to the best explanation. Now, much of both scientific reasoning and everyday reasoning of us lesser mortals too is abductive in nature, are computational models of the human capacity to reason by inference to the best explanation likely to be forthcoming do you think?

Jerry Foder

Well, here as in the discussion on functionalism I'm inclined to be a good deal less optimistic than is currently fashionable. I would be very surprised if where computational means something like Turing's notion of computational that is some operations that apply to representations in terms of something like their logical form or their logical syntax, if that's the picture of computation one has and it's the only one that I know of then it seems to me very unlikely and for just the reasons you say that much of what goes on in the mind is going to in the very long run turn out to be computational. For a number of reasons of which the most striking is that if you start to look at non-deductive rationality for example, confirmation, belief formation, that sort of thing how evidence is martial to constrained beliefs, it doesn't look like it's fundamentally a syntactic relation and logical syntax as I say is the force that drives the Turing picture. It also doesn't look like it depends on local relations and if you think for a while about what Turing did it seems to me that its reasonable to say about his notion of computation that it's actually restricted to a very small set of, even of syntactic properties of mental representations, that is what Turing-type processes operate on is relations between complex symbols and their constituents. So Turing will tell you how to get from P and Q to P and part of the reason he can do that is that the inference from P and Q to P depends on nothing except the syntactic structure of the premise. P's a constituent of P and Q, it's a syntactical part of P and Q. I doubt that that kind of locality is a general property of mental processes, in fact a lot of the most characteristic and most puzzling properties of mental processes seem to be sensitive to non-local global properties of belief systems, this is why when one tries to build an artificial intelligence machine one starts running into frame problems and globality problems all of which have the same basic structure, namely I believe an enormous number of things, what I believe is relevant to what inference I'm going to make from whatever it is that I currently have in mind but I don't want to search all the enormous number of things I believe to determine what the relevance relation is. That's a highly non-local kind of worry and Turing's notion of computation at least as far as anybody knows and at least on the surface doesn't seem to be well-equipped to handle those kinds of issues and since I agree with you that its abduction or inference as the best explanation its pretty global in those kinds of ways that is a large part of what's characteristic of not only what Sherlock Holmes was good at but what everybody's good at, namely sort of getting around in the world on the basis of fragmentary information. I'd be very surprised if any computational model recognisably of the kind that we have now turned out to be most of the story about most of the mind. However, having said that the computational picture works in a lot of places where what was previously available some variety of associationism didn't, and there's some point to looking where the light is and we know roughly speaking how computational explanations go and there do seem to be cognitive processes of some interest, perceptions of a plausible candidate, what they call low level perception that is early stages in perceptual analysis. There are processes going on in the mind of some complexity and a great deal of interest which do look very much like they might be susceptible to computational treatment. There's been some real progress in the attempt to analyse such processes computationally. So there's something to work on and I think we've got hold of a part of the story but I suspect that we're a long way from the rest of it. What I suspect is that there's some reconstructed notion of computation or mental process other than the one that we got from Turing and that we're going to need to face the real problem about the mind which is of course how people think.

Sean Crawford

So the computational theory of mind isn't really a theory about how people think and it can possibly give an account of deductive thinking, and it can possibly give an account of various processes that we might want to call mental processes such as perceptual processes but these processes aren't conscious really are they?

Jerry Foder

That's certainly true. I don't think that the distinction between what's computationally tractable and what isn't is the same as the distinction, its co-extensive with the distinction between what's conscious and what isn't. I think some perfectly straight forward conscious thinking probably is computational and certainly a lot of the computational processes that go on in the mind, unless we've got things terribly, terribly wrong are equally certainly unconscious. So I suspect that these are more or less orthogonal distinctions. It's roughly true that the big successes with computational analysis have been concerned with processes which one can't introspect like the analysis of visual form and perception.

Sean Crawford

Does the computational theory of mind have anything to contribute to the study of consciousness and does that matter to you?

Jerry Foder

I think the answer to the first question is pretty clear, namely no. In fact I suspect that nobody has anything to contribute to the study of consciousness as of the latest tally. Chomski, somewhere, we started with him, Chomski somewhere makes a distinction between what he calls problems and mysteries. Problems are things you can work on, you may not know how to solve them but at least you know how to work on them, or at least you know how to formulate some of the questions you'd like to have answered about them. Mysteries meet none of those conditions and consciousness is a mystery. We not only don't know what it is, and not only don't have a theory of it, but we don't even know what it would be like to have a theory of it. It's very popular in the current philosophical literature and to some extent in the cog-psy literature but the popularity if you look at it closely it consists of dozens, scores, hundreds of people writing papers saying in effect consciousness is terribly important, somebody really ought to have a good idea about consciousness and indeed I agree somebody ought to but nobody has and I don't see myself how anything like a computational story is going to do any good. Does it bother me? Sure, I'd like to have something interesting to say about consciousness but I haven't. I think of that as sort of like I'd like to be rich and famous but I'm not. What you can't do you can't do. I have a friend that is in the consciousness line of work and he once said something like that the nice thing about being in the consciousness line of work is that you can drop out for 30 or 40 years and not miss a thing, and I think that's exactly the current situation.

Sean Crawford

Can the idea of computation help in any way to explain the nature of creativity and imagination, or are they really to be lumped with consciousness?

Jerry Foder

Pessimism is my normal state and here I think what makes one pessimistic isn't that we don't know that state hardly any clear issues, hardly anywhere in cognitive science and there's a sort of sense that if we did we'd have thereby solved the problems. No, I think the problem about creativity and imagination and emotion is that it's not clear that they're the kinds of things that you can have a science about, it's not clear that they're as some glossaries like to say natural kinds, that they're the sorts of states over which lower-like reliable counterfactual supporting generalisations can be stated and in terms of which theories can be elaborated. I like to refer to the sad career of a guy who thought he was going to develop a science of Tuesdays and it sort of worked for a while, he discovered some generalisations about Tuesdays, that is they come after Mondays, they come before Wednesdays, they last about 24 hours and so forth and so on, and then the subject seemed to dry up and the reason it dried up is that the property of being a Tuesday doesn't provide a scientific domain nor unfortunately do most of the properties of most of the things that we're humanly interested in. What I would guess is that emotion and creativity and imagination and all that stuff that we humanly and properly care about just aren't going to prove appropriate domains for the kind of theory construction that scientists do though they may be a perfectly appropriate domain for, as it were writing novels. I mean there were things to do other than science and the peculiar restrictions that hold on scientific theory construction really are restrictions, I mean there are just aspects of the world which scientific theory construction doesn't seem to be the appropriate way to approach. If you want to know about emotions read Henry James is my

view and I suspect that when we've gotten as far as we can in a science of the mind my guess is it'll still be true if you want to find out about emotions read Henry James.

Sean Crawford

Don't people have genuine imaginative abilities that science might be able to explain?

Jerry Foder

Imagination's a sort of symcategoric notion, I mean you can be an imaginative chef and an imaginative mathematician and it's very hard to believe that among the various things that are going on in their head there's something that counts as it were as what makes them imaginative in their various ways and which is identical over these differences. It's a bit like being good. You can get good knives and you can get good philosophers and it's not at all clear that there's anything that they have in common quai good though it's a form of evaluation that applies to both. My guess is that to look for a process of imagination as such I suspect is going to be like looking for a faculty of intelligence as such, very unlikely that there is such a thing or at least quite likely that there isn't. There's just intelligent ways of doing one kind of thing or another and when you look closely at what gets done the notion of intelligence drifts out of the explanation and what actually has an explanation is how you do the particular kind of thing that you're either good or bad at.

Sean Crawford

The imagination has historically been seen as a kind of faculty, I mean Aristotle and Aquinas and Descartes and Hume all bring in the imagination and they seem to treat it as a kind of faculty.

Jerry Foder

Right, well one of the things that tends to confuse the discussion is that there's a notion of imagination as the capacity for forming mental images. That's close to what it is I think in Kant and Hume, that is it's something like being able to form a mental representation of a particular which satisfies a certain kind of abstraction. So you've got the concept 'dog' and imagination comes in say in the Kantian story as what allows you to construct mental representation of the kinds of individuals that fall under that concept. So that sense of imagination you know, either there is some use to it or not, is clearly debatable but it's not what people have in mind when they say he's an imaginative chemist or an imaginative chef, or she's an imaginative mathematician or whatever. I doubt there's any science of the ability to think up new ideas which is really what the imagination in this informal non-Kantian, Humian sense is supposed to do. I don't think we're going to do any better at answering questions like how do you think up new ideas than by saying well, you know, concentrate, learn as much as you can, keep your fingers crossed and keep at it until something occurs to you. It's like you know, how do you get to Carnegie Hall, practise. There's not going to be a theory I think of creativity in that, or of imagination in that sense, it's not that all the creative processes form a domain which have something interesting in common from the point of view of explanation.

Sean Crawford

Why not say the same thing about consciousness?

Jerry Foder

That may be true. I mean the intuition about consciousness is that there is something that all the conscious states have in common, pains and sensations and so forth and so on, in virtue of being conscious. In a way that seems to just present itself as a fact of the mental life in a way that is not obvious that there is something that all the creative or intelligent or imaginative mental processes have in common. This could be an illusion, I mean it could be that there's really nothing that as it were a conscious visual sensation and a conscious thought have in common simply quai conscious, maybe that's true, it seems counter intuitive and I wouldn't bet on consciousness disappearing as a problem in the way that I'm sort of inclined to think that creativity and emotion and so on will.

Sean Crawford

Generally speaking what do you think is the role of philosophers in studying the mind? Are they just junior partners in the enterprise of cognitive science, or do they have something special to contribute?

Jerry Foder

Well, if you think of the problems about the mind as the traditional metaphysical and ontological problems then I suppose the philosophical problem is only because nobody else wants them. But if the problems about the mind are Hume's kind of problems, namely constructing an in some sense empirical model of how it works I would have thought that there's no particular field that that undertaking belongs to. There are all sorts of things that might help, there are all sorts of investigations that might help and as things now stand those investigations cut across a whole variety of traditionally distinct disciplines from philosophy to physiology and back. What I think philosophers are pretty good at is thinking about methodological and conceptual issues; issues that arise in theory construction and issues of what actually follows from what and what might be a reasonable thing to hypothesise given what else. That's the kind of thing that philosophers are professionally trained to do, it's what they cut their teeth on and in the cognitive sciences at least it's almost true, I think that what we have is a sort of pre-science, I mean were not really at the point where we have well entrenched theories and well entrenched research methodologies and we know what to do if only we could get the grant to do it with. There's a great deal of methodological confusion and a great deal of conceptual confusion and we simply don't know what we ought to take to follow from what, how to map out the consequences of the kinds of theories that we do hold and the relations between the theories, that it seems to me is something that philosophers might do if they can't think of anything else to do of an afternoon and at this stage at least in the development of theories of cognition its extremely important. If you read systematic books in psychology, I mean books where the author wants to set out as it were the big picture as he sees it, the way it works is you get a first chapter on methodology and it tells you about the nature of science and the relation between psychology and biology and the relation between the mind and the body and between the mental and behavioural and all that sort of stuff, very grand, which makes very little sense actually and which is hard to take very seriously as philosophy and which the author tends to ignore properly so through the rest of the book this is the law of the relevance of first chapters which holds pretty well in cognitive science publication, well that's a gap in which the way people, the resources people have at hand for doing the kind of things that psychologist and cognitive scientists are trying to do, those resources are pretty inarticulate and their sense of how the project should be described is often pretty inarticulate. To take one example the notion of representation plays an absolutely central role in standard theories of cognitive processes and nobody really has any very clear idea, either about just what kind of role it's playing or just what kind of a representational relation is, or what kind of demand psychological theory is likely to place on it. That seems to me the kind of project which its natural for a philosopher to look at and in my own case it's actually the kind of project I'm interested in, I really would like to know how the mind works and some of the tradecraft that you pick up in doing philosophy seems to be of some help in approaching that issue.

Sean Crawford

Jerry Foder, thanks very much for talking to us.