

Introduction

Philosophy in and Philosophy of Cognitive Science

Andrew Brook

Abstract: Despite being there from the beginning, philosophical approaches have never had a settled place in cognitive research and few cognitive researchers not trained in philosophy have a clear sense of what its role has been or should be. We distinguish philosophy *in* cognitive research and philosophy *of* cognitive research. Concerning philosophy in cognitive research, after exploring some standard reactions to this work by nonphilosophers, we will pay particular attention to the methods that philosophers use. Being neither experimental nor computational, they can leave others bewildered. Thought experiments are the most striking example but not the only one. Concerning philosophy of cognitive research, we will pay particular attention to its power to generate and test normative claims, claims about what should and should not be done.

One of the features of our new journal, *TopiCS*, will be groups of papers on a single theme running across a number of issues. One of the themes is the place of philosophy in cognitive science. A number of cognitive researchers trained in philosophy who have made signal contributions to cognitive science will launch the theme, Dan Dennett, Bill Bechtel, Paul Thagard, Pierre Jacob, Tom Metzinger, and Zenon Pylyshyn.¹ There will be opportunities for others to enter the conversation. In this introduction, I will try to specify the theme in detail, say something about why it is interesting, introduce some of the ways in which philosophically-knowledgeable researchers contribute to cognitive research, and sketch some of the methods they use. (Note: Not all cognitive researchers trained in philosophy work in Departments of Philosophy and not all philosophically-knowledgeable researchers have formal training in philosophy.)

Philosophy has never settled into a stable position in cognitive science and few cognitive researchers without training in philosophy have a clear sense of what it has contributed and should contribute. Overall, it has never achieved the position of behavioural experiments, computational modelling or system-building and, more recently, cognitive neuroscience. There have of course been a wide range of collaborations between individuals trained in philosophy and individuals with other training. All the contributors launching our theme have done extensive work with researchers from other backgrounds. Others include Patricia Churchland, Jesse Prinz, Jerry Fodor, Ned Block, Martin Davies, Brian McLaughlin, Evan Thompson, Aaron Sloman, Nancy Nersessian, Georges Rey, Peter Carruthers, Rob Stainton, Kathleen Akins, Jeff Pelletier, Steven Stich, Michael Anderson, Rick Grush, ... – a complete list would be quite long. Concepts, colour perception, explanation/creation of theory, reference and assertion in language, and consciousness are some of the topics. Two bodies of collaboration have even come to have

¹. Pylyshyn trained as a psychologist but has made prodigious philosophical contributions.

names, the philosophy and neuroscience movement (Brook and Akins 2005) and experimental philosophy (the activity of testing philosophical theses and philosophers' intuitions empirically) (Knobe and Nichols 2008, Appiah 2008).

These achievements notwithstanding, the nature of philosophically-inspired research and its potential for interacting with other work to shed light on cognition is not well understood and relatively few researchers with other backgrounds make use of philosophical work in their research. Our aim in this theme is to explore what philosophy could and should contribute to cognitive research.

The unsettled place of philosophy in cognitive science is not because philosophy is a cognitive Johnny-come-lately. Philosophers have been a part of cognitive science since the activity was merely a twinkle in the eyes of a small but hardy group of pioneers in the 1960s. Hilary Putnam and Jerry Fodor come immediately to mind. In the 1960s those two did much to articulate the view that came to be known as *functionalism*. Functionalism (as the term is used in cognitive research) is the idea that cognitive processes are to be understood by what they do, i.e., how they function, rather than in terms of their structure or constituting mechanisms. Even when mechanisms are invoked, they are usually expressed as organized structures of sub-functions which, when put together, implement a more complex cognitive function. Functionalism is something like the official philosophy of mind of cognitive science and has been from the beginning. This important contribution notwithstanding, work done using philosophical tools did not go on to build a stable, well-understood role for itself in the new science.²

There is a tendency among those trained in philosophy to blame themselves for this state of affairs. Well, not *themselves* – other philosophers. The kinds of philosophy most relevant to cognitive science are, as we said, philosophy of mind, language, and science (including parts of epistemology). According to these critics, too many philosophers in these sub-fields have been at best indifferent, at worst hostile, to science. Indeed, they often reach conclusions without even knowing what science could tell them about the topic on which they are pronouncing. There is doubtless some merit to these charges. To take a personal example, I did a doctorate in philosophy at Oxford some decades ago. I have no recollection of ever setting foot in the experimental psychology building – even though I was working on consciousness! However, our goal in this series of papers is, in the words of the old Johnnie Mercer song, to e-lim-in-ate the negative and ac-cent-u-ate the positive.³ Even philosophers who *are* well-versed in the relevant science have never had a settled role in cognitive science. When philosophers' work on cognition has had shortcomings, it is far easier to beat up on them for their shortcomings than to identify and clearly articulate what work done using philosophical techniques could and should contribute. The latter will be the primary focus of the papers in the theme to which this paper is the introduction.

² Much the same could be said of linguistic theory. In the form of the towering contribution made by Chomsky's theory of syntax, it too was central in the early days and it too has not settled into a stable, fully interactive position in cognitive science.

³ And not messing with Mr. Inbetween? I'm not sure what that would be. The reference is shamelessly stolen from Fred Dretske (1995).

To make progress we need to distinguish philosophy *in* cognitive science and philosophy *of* cognitive science. The former embraces work done on topics such as mind and language that are also studied using other approaches such as behavioural experiments and theoretical linguistics, so philosophy of mind and language. The latter is a branch of the philosophy of science and is a meta-study. It studies what others do – rather than doing cognitive science, it studies cognitive science. Both play an important role. The nature and role of philosophy *of* cognitive science is clearer and better understood than the nature and role of philosophy *in* cognitive science. The first two papers in the current theme (Dennett, this number; Thagard, this number) deal mainly with philosophy *in* cognitive science so we will start with it, then turn to the role of philosophy of cognitive science. The hope is that these discussions will provide a foundation for the papers to come in this theme.

Philosophy *in* Cognitive Science

When philosophy plays a role in cognitive science, it is usually philosophy of mind, philosophy of language, or logic. We will focus on philosophy of mind and language. Logic has certainly made a contribution via its role in formal linguistic theory and some kinds of AI. And it gives rise to some interesting philosophical questions, for example whether patterns of valid inference reflect something deep about human cognition or merely the requirements of an invented system of norms. However, logic has a stable place in the research to which it contributes.

This is not true of philosophy of mind and language. The range of topics that philosophers of mind and language tackle is as broad as cognitive science itself, going all the way from formal semantics and pragmatics to memory, perception, reasoning, emotion and consciousness to how to make neuroscience more useful to cognitive research but most nonphilosophers have no clear sense of what or even whether philosophy can contribute to our understanding of them. To avoid cumbersome repetition, until further notice I will refer to the philosophy that deals with such issues, philosophy of mind and language, simply as philosophy or sometimes cognitive philosophy (on the model of cognitive psychology).

Attitudes to work on cognition done using philosophical tools often take one of two general forms, one dismissive, the other baffled. There are two important variants of each. All four are actual reactions that I have encountered but I will not name names. The first dismissive form goes like this:

‘Philosophers mounted some interesting speculations about the mind in times past but we are now in a position where we can get out of the armchair and do real science on these things. Philosophy, imaginative and entertaining though it can be, has been relegated to the dustbin of history. There is still something to logic and maybe ethics but the rest of philosophy has been superseded by science.’

This view of philosophy is very widely held among cognitive researchers. Here is a variant:

You philosophers with your relentless pursuit of the big picture exhaust me. What are the research payoffs? What we need at this point in time is disciplined work on specific issues, not big pictures painted in broad strokes. Even those of you who try to be

interdisciplinary and responsive to what is now known bite off more than any mortal can currently chew.

The key idea embedded in both variants of the response is that there is nothing distinctive to what philosophers do, logicians and ethicists maybe excepted. Philosophy is just speculative armchair science.

There is, however, an important distinction between them. The first variant takes cognitive philosophy, most of it anyway, to be doing exactly what science does now, only not very well. The second takes it, some of it anyway, to be doing something that science does relatively rarely, namely, integrate results in a bigger picture. The difference can be brought out by reference to the generate-and-test picture of science. On the generate-and-test picture, the work of science is to generate hypotheses and then test them to discover which ones remain standing when faced with data. (Reichenbach and then Popper captured this picture in the distinction between the context of [hypothesis] discovery and the context of [hypothesis] justification.) To these two add a third activity, namely, interpretation. The job of interpretation, broadly, is to make sense of the results of generate-and-test. Put in the language of these distinctions, the first variant views much of philosophy as ‘generate without test’, that is to say, discovery without justification.⁴ And the second views at least part of philosophy as premature attempts at broad, integrative interpretation.

Nor is it clear that the two things of which philosophy stands ‘accused’ here are entirely bad things. Even if philosophy is just speculative hypothesis generation in the generate-and-test process, your view of the merits of this activity will hang on whether you believe that the hypotheses afoot in cognitive science open up all the important alternatives. In the 2500 years that philosophers have been speculating about the mind, they have canvassed an extremely wide range of possibilities. Some of them may suggest important hypotheses about otherwise-puzzling data that current researchers have not even thought of, let alone tested. Among others, Prinz (2002) says this about empiricism, Chomsky (as recently as 2007) and Fodor (1975) say this about rationalism, and Kitcher (1990), Brook (1994), and others say this about the Kantian tradition (Kant 1781/7). Philosophy’s speculative hypotheses may expand the range of interesting hypotheses about particular bodies of data considerably.

Your view of the second variant will hang on whether you agree with Newell (1973) when he urged nearly forty years ago that we need to do more than generate-and-test, more than find new effects. We need to build big, generative theories that can tie together the thousands of effects that we have now identified.

The baffled form recognizes that often philosophical work in cognitive science at least appears to be quite different from other kinds of work on cognition, indeed from science in general – and is puzzled. Not infrequently the reaction will focus on the techniques used in cognitive philosophy that are the most unlike those used elsewhere in cognitive science. Thought-experiments are a leading example. The reaction then goes something like this:

‘Why are philosophers interested in thought experiments, *intuition pumps* as Dan Dennett (1991) calls them? Thought experiments at best explore possibilities. Why worry about what is possible? How could such a worry contribute to the generation of knowledge?’

⁴ This failing is precisely what experimental philosophy aims to fix.

What we should be doing is discovering what is actually going on in cognition. For this work, thought experiments are useless.’

One finds the same reaction to *reductio ad absurdum* (reduction to absurdity) arguments, hang-ups about clarity of concepts, and so on. The people reacting this way do find something distinctive in philosophical work in cognitive science – but don’t see why anyone would want to do these things or think that doing them might contribute to progress in cognitive research. Here is a variant of this reaction, to do with concepts:

‘Philosophers spend their time worrying about concepts. Why? The concepts of cognitive science are mostly just fine. What we need is to get on with discovering the facts.’

I have treated these two as variants of a single reaction because thought experiments are often used as a tool for investigating concepts.

For the second variant to hold water, it would have to be possible to do good investigation of the facts without paying attention to the concepts with which one is describing and categorizing the facts. Often, such a separation is not possible. But if it is not, thought experiments are more closely linked to investigation of the facts than is generally thought.

Again, it is not clear that exploration of thought experiments and analysis of concepts are bad things. We will consider thought experiments separately later. The suggestion that philosophers’ interest in concepts is baffling is itself a bit baffling. It is widely recognized that the conceptual apparatus of cognitive science is not in good shape. We have dozens of terms for memory, for attention, for consciousness, and for many other things of central interest to us. Each approach to cognition has its own proprietary terminology that is often mysterious to other approaches. We use the word ‘representation’ to name two very different kinds of thing, ‘attention’ to name three, ‘consciousness’ three, ‘information’ four, and so on. Our concepts are just fine? Not likely.

Methods used to do philosophy in cognitive science

Some of the methods that cognitive philosophers use are baffling for good reason. Philosophy in cognitive science does not do experiments, it is not computational, it seldom makes use of detailed findings about the brain ... – it seldom does any of the things that other cognitive researchers do most of the time. So how does it proceed?

Concerning methods in philosophy, there is something to all four of the reactions that we just explored. Yes, some philosophical work, especially in the philosophy of mind and language, generates and argues for hypotheses. In connection with the mind, this work goes back to Descartes (the mind is not the brain), indeed even to Aristotle (the mind is the living form of the body). And yes, some philosophers have not had as much interest in testing as one would like, preferring to look to arguments and intuitions to support their claims. In philosophy, sometimes imaginative appeal actually counts for more than justification. Yes, some philosophical work consists of broad integrative interpretations. Yes, thought experiments play a central role in some philosophical work, especially in the philosophy of mind, and it is true that the nature and merits

of that role are none too clear. And yes, people with philosophical training tend to pay more attention to the conceptual toolkit of cognitive science than is common in those with other kinds of training.

In an introductory way, here is what each of these methods is like:

1. *Speculative hypothesis generation.* There is nothing wrong with speculative hypothesis-generation as such. Indeed, since all hypothesis generation consists in the application of the imagination to some group of facts, all hypothesis generation has a speculative element to it. If so, there is also nothing distinctive to those with training in philosophy in this work. There is little speculation in the context of justification but the context of discovery is full of speculation.
2. *Integrative interpretation.* The activity of interpreting tested hypotheses is a large part of what the philosophical contribution to cognitive science consists in. One of the roles of philosophy has always been to show how 'things in the broadest possible sense of the term hang together in the broadest possible sense of the term', as the great philosopher of science Wilfrid Sellars (1963) once put it.

The results of these activities are theories and models offered by philosophers similar to theories and models offered by others in cognitive science. Indeed, Quine (1953), Sellars (1963), Castañeda (1980) and others take most of philosophy to be like this, including the philosophy of mind and language. The main difference is that philosophers tend to go after bigger and sometimes more abstract objects than researchers with other backgrounds: in cognitive research, representation as a whole rather than particular kinds of representation, rationality as a whole rather than particular activities of reasoning well, the nature of explanation in cognitive science as a whole rather than how to explain a given phenomenon before us. So far, except for level of generality and abstraction, there is nothing distinctive to this methodology. However, philosophers' interpretations are generally aimed at results achieved by others. This sometimes confers advantages of breadth and objectivity.

3. *Exploration of thought experiments.* Thought experiments consist of imagined manipulations of imagined scenarios. Philosophers (and, as we will see, others) sometimes claim that one can reach substantive conclusions by using them. We will return to this topic.

4. *Analysis of concepts.* Philosophically-grounded work in cognitive science has spent a lot of time clarifying concepts and making recommendations for how concepts should be used. Here is an example that shows why.

Pylyshyn (1988) and collaborators have shown that humans can track a number of objects at the same time. He suggested that we use what he calls visual indexicals to do so. How does such tracking relate to attention? There are two options: (1) It is itself a kind of attention, albeit less focussed and less conscious than a lot of attention. (2) It is a precursor to attention, a way of fixing on objects so that we can come to pay attention to them. Which is best?

Notice two things. First, this is not a trivial issue. If we say that tracking is a form of attention, then there is a form of attention that is dispersed over a number of objects and not fully conscious. And theory of attention will have to account for it. If on the other hand we say that only something subsequent to the kind of object tracking that Pylyshyn has in mind is attention, then we are free of these burdens. Second, the issue does not concern the facts and could not be

settled by reference to the facts. It is about the best way to group observed phenomena under a word. Is the more or the less expansive use of the word 'attention' simpler and cleaner? Which better captures what us interests here, our reasons for using the word? Which alternative is more neutral in the light of possible theories in the neighbourhood? And so on.⁵

Philosophers do a lot of this kind of work. Though it is not distinctive to philosophy – having a clean, trouble-free toolkit of concepts is a requisite of good science and so of interest to anyone doing science –, training in philosophy tends to makes one better at it. Also, it could be suggested, the rest of cognitive science suffers from not doing enough of it. In the context of this suggestion, it is striking that even when collaborative work between philosophers and others has occurred, as it has, for example, on concepts and how we process them, this work often does not affect cognitive science elsewhere, in this case the stifling profusion and confusion of concepts in cognitive research.

Moreover, this kind of investigation is not limited to words. The investigation of the semantic properties of different types of explanation, the investigation of how various scientific (and perhaps other) activities use a given word, and perhaps some other things deal with the same kinds of issues.

And thought experiments are one of the main tools used in this work, though they are used elsewhere, too, as we will see.

Thought experiments in cognitive philosophy

Thought experiments deserve some special attention. The activity of exploring them looks very peculiar to those doing normal science and is widely misunderstood.

Thought experiments, we said, are exercises of the imagination. Thus they contrast with hands-on experiments.⁶ In a thought experiment, we imagine, i.e., represent, a scenario, rather than manipulate a scenario that actually exists. And we are interested in figuring out what is possible rather than what is actually the case. Some of the most famous thought experiments in cognitive science are Searle's Chinese room, Putnam's twin earth, Jackson's Mary the colour-blind colour scientist, Dennett's qualia impasses, and Chalmers' zombie thought experiments. A major source of people's bewilderment about philosophers and thought experiments is that philosophers seem to hold that merely imagining scenarios could reveal things of importance.

Searle's (1980) Chinese room is the most famous thought experiment in cognitive research. Someone who knows no Chinese is put in a room. Sheets of paper with shapes on them come in through a slot. The person has an inventory of other shapes and a huge rulebook linking

⁵ Jackson 1998 is one of the very few philosophers who discuss work of this kind on concepts. Kripke 1972, Nozick 1981, Cohen 1986, and Dummett 1993 offer related discussions.

⁶ There is no short yet completely adequate way to mark the distinction. Question-begging options that don't work include: 'real experiment', 'physical experiment', ...

shapes to shapes. For each shape that comes in, she identifies and finds the linked shape(s) and shoves it(them) out a second slot. Unbeknownst to her, the shapes coming in encode questions in Chinese and the shapes going out are answers to these questions. What this is supposed to show is that what the person who knows no Chinese does in the room is all that a computer processing symbols by their physically-detectable properties could do.

Twin earth thought experiments go like this (Putnam 1975). Imagine a person here on Earth, Adam, and his completely identical twin, Twadam, on experientially-indistinguishable Twin Earth. Adam and Twadam both use the word 'water' and they use it in situations that are experientially indistinguishable. Yet on earth what is called 'water' is H₂O, on twin earth it is XYZ. Does the word 'water' as used by Adam and by Twadam have the same meaning? Evidently not. Yet everything in their heads is the same. Hence, in Putnam's memorable phrase, "meanings just ain't in the head" (Putnam, 1975, p. 227).

The thought experiment about Mary the colour scientist goes like this. Mary is a wonderful colour scientist. Indeed, she knows absolutely everything there is to know about colour experience. Yet she has never seen colour. One day the barrier is removed and she sees colour for the first time. It would seem that she would gain a new item of knowledge: what it is like to *experience* colour. Hence experience is not ... (draw your favourite conclusion). (Frank Jackson created this thought experiment but no longer finds the radical implications in it that he once did.)

Here is a leading example of a qualia impasse, Dennett-style (1988, 1991). Chase and Sanborn both notice that they don't like their favourite coffee as much any more. Chase says that the coffee tastes the same but he doesn't like that taste as much as he used to. Sanborn says, no, he would still like *that* taste as much but the coffee no longer tastes the way it used to. Since there would seem to be no way in which this putative difference could make a difference, we are invited to ask ourselves whether there is a real difference here.

Chalmers (1996) mounts zombie thought experiments to convince us that we can imagine a creature behaviourally, cognitively, and even physically exactly like us except that it is not conscious. If so, consciousness does not consist of anything behavioural, or cognitive, or (nonconscious and) physical.

Thought experiments are used more widely in science than is generally recognized. Physics, for example, is famous for its thought experiments (on thought experiments in science, see Brown 1991, Horowitz and Massey 1991, and Sorensen 1992). Schrodinger's cat and Galileo's tying a smaller and a bigger piece of matter together are two of the best known. Schrodinger argued, against the indeterminacy interpretation of quantum mechanics, that a cat in a box had to be in some determinate state even if we did not know what it was. Galileo argued, against the Aristotelians, that if Aristotle was right and a smaller mass A falls slower than larger mass B, then if we make up a new object C by joining A and B, C will have to fall both faster and slower than B. This was supposed to eliminate Aristotle's hypothesis, no experiments needed. Thought experiments played a key role in Einstein's discovery and formulation of relativity theory.

Sometimes psychological researchers use thought experiments, too. For example, some research into reasoning starts with getting the subject to do a thought experiment. The subject is

asked to determine whether more words begin with 'r' than end with it, for example, or to determine which of two alternatives is more probable. Presumably the way in which they do this is by imagining something and then using what they have imagined to answer the question.

Equally, linguistic theory is full of thought experiments. In linguistic theory, thought experiments usually concern much smaller matters than in philosophy – strings of words – and the objective is to test intuitions of grammaticality or meaningfulness. 'Does this string sound grammatical?' or 'Could this string mean such-and-such in the imagined context?' are the kinds of question asked. Chomsky's invitation to imagine the string, 'Colourless green ideas sleep furiously', to see if it seems grammatical is the most famous example. Smaller to be sure, but these too are thought experiments.

There are differences between thought experiments in physics, psychology, and linguistics and the kind of thought experiment that philosophers do. In physics, not only is the thought-experimental situation clear, its implications are clear, too, many of them anyway. The latter is notably *not* true of philosophers' thought experiments. And in linguistics and psychology, the thought experiments are done by the subjects, not by the researchers, and are meant to generate data, not sort out what is actual or possible with respect to data. (Brown 1991 says that they are real experiments done in the mind, not thought experiments at all.) Also, in the hands of philosophers, thought experiments sometimes play a normative role, which is not true elsewhere. Think again of the thought experiments in reasoning research that we mentioned. Psychologists get subjects to do these exercises to find out how they actually reason: what mistakes they make, what produces these mistakes, and so on. When a philosopher runs a thought experiment about reasoning, her interest is different. She wants to find out what *good* reasoning consists in. Nonetheless, thought experiments are not unique to philosophers, not even inside cognitive science. However peculiar they may seem to some people, they are used in science.

It could still be true, of course, that thought experiments play a bigger role in philosophy than in any other part of cognitive science. That they do. Indeed, they are central to some philosopher's work. Equally, no other discipline uses them to study normative issues. Nonetheless, thought experiments are not unique to philosophers.

What can thought experiments do for us? This is a topic to which I will return if there is an opportunity later but here are some of the claims made on their behalf:

1. *Thought experiments can isolate crisp examples of a phenomenon under investigation.*
2. *Thought experiments can tell us what we initially take something to be like and can sometimes play a role in us deciding what we should take something to be like.*
3. *Thought experiments are often central to hypothesis generation.*
4. *Thought experiments can sometimes play a role in hypothesis testing and elimination.*

The last claim is more contentious than the other three. In summary, thought experiments are quite different from hands-on experiments. However, they play a key role in the clarification and reconstruction of the concepts used in doing hands-on experiments and in the generation of

hypotheses for hands-on experiments to test. And they may sometimes play a role in hypothesis elimination.

Philosophy of Cognitive Science

So far, we have mainly discussed the kind of philosophical work that is *part* of cognitive science and the methods that researchers use when they do it. We turn now to philosophy of cognitive science. Since some of the papers to come in the theme deal with this role for philosophy splendidly, I will keep my remarks brief.

Unlike the reaction to philosophy in cognitive science that we examined earlier, few have serious doubts that philosophy of cognitive science plays an important role. The range of philosophical questions that can be asked about cognitive science is almost limitless. (From here on, the term ‘philosophy’ will no longer be short for ‘philosophy of mind and language’, as it was above.) A sample of some of the important ones:

- § Do cognitive models such as ACT-R merely fit the behavioural data or can they capture the actual structure of human cognition (Fodor’s 1968 and Pylyshyn’s 1984 weak and strong equivalence)? Closely related, could AI systems or robots operate as we do or will they always be just engineering?
- § What kinds of explanations work best on cognition, particularly in light of the massive complexity and interconnectedness of the human brain and the inherently statistical, rather than rule-like, nature of complex behaviour (Bechtel 2008 and his contribution to the theme, next number of this journal)?
- § When are correlations mere correlations and when can they be taken to reflect a causal or structural connection (Hardcastle 1996, Ch. 2)?
- § What does it take to confirm or refute a theory or model of cognition (see Thagard 2000 on the role of coherence in cognitive theory-building, for example)?
- § Can studies done in the highly artificial setting of a laboratory reveal how cognition works in more natural settings or do we need to study cognition ‘in the wild’ to achieve such understanding (Clark 1997)?
- § Does the structure of human cognition tell us something about the structure of all possible cognitive systems or is our cognitive structure just a tiny fragment of the space of possible cognitive systems? (Chomsky [2000, p. 16, and many other places] holds that the latter is true of human languages at any rate. If this is true of language, how could it not also be true of the rest of cognition?)
- § While the notion of computation provided a useful metaphor and operational framework for cognitive science in the early decades, is human cognition computational in anything like the way any computer, serial or parallel, is? (Searle 1980, Fodor 1983, 2000 and many others say no.) What implications if any does Gödel’s incompleteness theorem have for this question? (Note: The answer to this question probably will have implications for

the whole formal systems approach to human cognition, particularly prominent at the moment in theoretical linguistics.)

§ How can we integrate the multitude of different approaches to cognition into a single, unified picture (Hardcastle 1996, Chs. 6 and 7; Dawson 1998, Ch. 8 is a penetrating case study)?

Questions such as these not only raise important issues about guidelines and frameworks for doing good cognitive research, they also raise important meta-questions about the goals of a science like cognitive science in the first place.

Equally, philosophical reflections on cognitive science can take a number of forms. What I have tried to do in this paper, namely, understand how an activity in cognitive science works (philosophy in this case), is one form. Another is to apply general philosophy of science to cognitive science to determine what kind of science it is and what works in it (Bechtel, next number of this journal, Pylyshyn 1984, Dawson 1998, and Bechtel 2008 are some important examples). A third is what we earlier called the context of *interpretation*. Often it is unclear what an experimental finding means. What are its implications for the question before us? Does it help with other questions? ... and so on. A fourth consists of integrating results attained by different ways of doing cognitive research (again, Hardcastle 1996, Chs. 6 and 7 and Dawson 1998, Ch. 8). This is once again Sellars' 'how things in the broadest possible sense of the term hang together in the broadest possible sense of the term'. Yet a fifth is to make and defend claims of various kinds about how it *should* work. Indeed, in philosophical reflections on cognitive science, how it does and how it should work are often not far apart. Investigations of shoulds and should nots, goods and bads, oughts and ought nots are normative investigations.

What methods do philosophers use to do such work? This is a big and largely unstudied question, one much in need of some unhurried exploration. Here I will restrict myself to investigation of normative questions, as this is the part of philosophy of cognitive science most distinctive to it.

Normative investigations are sometimes thought to be restricted to ethics and aesthetics. Certainly ethical issues arise in cognitive science. In particular, applications of cognitive science, which can vary from human computer interfaces and virtual realities to pedagogy to prosthetic brain implants and cognitive robotics, raise all kinds of questions about rights and responsibilities. If bio-technology is controversial, 'cogno-technology' will be much more so. Indeed, cognitive science probes such ancient and central questions as what it is to be human, whether we have free will, and what the boundaries of the self are. There is lots of work for philosophers to do on ethical issues in cognitive science.

That said, normative questions are not limited to ethics and adjoining philosophical fields. Philosophy of science also concerns itself with norms, especially the part of philosophy of science that studies knowledge – such things as norms for good evidence, for warrant, for adequacy of theory, for good explanations, and so on. Such investigations are clearly important to cognitive science.

Because norms, what ought to be, cannot be settled by studying what is, scientific method, which consists of description and theory of what is, cannot by itself resolve normative

questions, not even the normative questions that arise about science itself. If normative work on cognitive science is essential, there is also a case for holding that philosophers' training in how to think about norms makes a contribution that is both distinctive and irreplaceable. Other disciplines may explore what explanations we in fact use, how we in fact access evidence for purposes of belief formation, how we in fact reason about probabilities, and so on. However, the activity of investigating what explanations we *should* use, how evidence *should* constrain belief, how we *should* reason about probabilities cannot be read off any facts about how we do in fact do these things.

As an example of how the descriptive/theoretical and the normative mix in philosophy of cognitive science, consider work on explanation in cognitive science. Often this work is, to begin with, broadly descriptive, aimed at questions such as: What styles of explanation do cognitive scientists use? How do they work? But before long a normative dimension often slides in: What styles of explanation *should* cognitive scientists use? What styles of explanation might get us somewhere and which probably would not?

How do philosophers build such normative assessments? By what methods, by what kinds of reasoning, can we validly/justifiably/adequately assess what we ought and ought not do? Even this question about the methods of one kind of philosophy of cognitive science is a big topic; I will restrict myself to just two suggestions. At least in epistemic contexts, two main methods are,

§ Inferring norms from general principles of some kind, principles that are themselves rationally compelling.

§ Shaping norms to serve our deepest epistemic interests.

In connection with epistemic norms, principles that preserve truth or that proportion degree of belief to degree of probability might be examples. Occam's Razor (postulate no more entities than are needed to explain something) and simplicity (postulate the smallest number of principles needed to explain something) might be others. Relevant interests might include our interest in prediction, in control, in understanding structure, in knowing origins, and so on.

Assessing normative claims against general principles of rationality or deep-seated epistemic interests does seem to be a method distinctive to philosophical training. Indeed, when cognitive scientists from other backgrounds engaged in this sort of epistemology and philosophy of science, they are no longer doing science, strictly speaking.

Thought experiments often have the purpose of helping us to help settle normative claims. When we mount thought experiments to investigate concepts, styles of reasoning, and so on, we want the result to be *good* concepts, *good* styles of reasoning, and so on. That is to say, we are interested in *normatively reconstructing* our conceptual tools so that they serve our epistemic interests better, not just in *finding out* what is built into them already.⁷

⁷ Our discussion does not exhaust the range of methods used by philosophers to study cognition. One method not mentioned so far is the view that philosophy is or should be the untying of conceptual knots, an idea associated particularly with Ludwig Wittgenstein (1953). Bennett and Hacker (2003) attempted to show its efficacy in cognitive neuroscience but the

Concluding remarks

I have tried to construct an overview of the kinds of contributions to cognitive science made by those with philosophical training – both philosophy *in* cognitive science and philosophy *of* cognitive science – and of some of the main methods used to make them.

Final note. Even if philosophy lacks distinctive *methods*, it could still play a distinctive *role* in cognitive science. In particular, philosophy could have distinctive preoccupations. And indeed it often does. Where experimentalists are interested in whether a claim is true, philosophers tend to be more interested in how concepts hang together, what the possibilities are, and so on. I am reminded of the 1950s television show *Dragnet*. Like Jack Webb, the lead detective in the show, an experimentalist would say, “Just the facts, ma’am, nothing but the facts”. A philosopher would be more likely to ask, “Wha’d’ya mean, ma’am, wha’d’ya mean?” If an experimentalist does turn to look at her conceptual toolkit or stands back to consider the general nature of what she is doing, this would be a sideline for her, maybe because something in the toolkit is not working well. For philosophers, doing such things is their occupation.

So, what could and should the role of philosophy in and of cognitive science be? Let the show begin!⁸

References

This paper is a distant descendant of Brook, A. 1999. Does Philosophy Have Distinctive Methods to Offer Cognitive Science? *Proceedings of the 21st Annual Conference of the Cognitive Science Society* (New York: LEA), pp. 102-8. The discussion of thought experiments draws on Brook, A. 2006. Thought Experiments. In Barber, A. and R. Stainton, eds., *Encyclopaedia of Language and Linguistics*. Amsterdam, Oxford: Elsevier.

Appiah, K. 2008. *Experiments in ethics*. Cambridge, MA: Harvard University Press.

attempt largely failed. I will not discuss it further. Another method not mentioned so far is the attempt to study the mind by identifying standards that a system must meet to be a mind. Standards of minimal rationality and certain intersubjective standards are popular. This approach was introduced by Davidson (1973) and Dennett (1973). In more recent times, McDowell (1996) and Davidson (2001) have championed it. Sometimes called the normative model of mind, this approach played a role in cognitive ethology for a while (Seyfarth and Cheney 2002) but has also not had much influence on general cognitive research, so I will not discuss it further, either.

⁸. My thanks (in alphabetical order) to Bill Bechtel, Wayne Gray, Jo-Anne LeFevre, Heidi Maibom, Rob Stainton, and Chris Viger for very helpful comments.

- Bechtel, W. 2008. *Mental mechanisms: Philosophical perspectives on cognitive neuroscience*. London: Routledge.
- Bennett, M. and P. Hacker. 2003. *Philosophical Foundations of Neuroscience*. Oxford: Blackwells.
- Brook, A. 1994. *Kant and the Mind*. New York: Cambridge University Press.
- Brook, A. and K. Akins. 2005. *Philosophy and Cognition: The Philosophy and Neuroscience Movement*. New York and Cambridge: Cambridge University Press.
- Brown, J. 1991. *The Laboratory of the Mind*. London: Routledge.
- Castañeda, H.-N. 1980. *On Philosophical Method*. South Bend, IN: University of Indiana Press.
- Chalmers, D. 1996. *The Conscious Mind*. Oxford: Oxford University Press.
- Chomsky, N. 2000. *New Horizons in the Study of Language and Mind*. Cambridge, UK: Cambridge University Press.
- Chomsky, N. 2007. Language and thought: Descartes and some reflections on venerable themes. In A. Brook, ed. *The Prehistory of Cognitive Science*. Houndsmill, Basingstoke, UK: Palgrave Macmillan, pp. 38-66.
- Clark, A. 1997. *Being There: Putting Brain, Body, and World Together Again*. Cambridge, MA: MIT Press.
- Cohen, L. J. 1986. *The Dialogue of Reason*. Oxford: Clarendon Press.
- Dawson, M. 1998. *Understanding Cognitive Science*. Oxford: Blackwells.
- Davidson, D. 1973. The material mind. In his *Essays on Action and Events*. Oxford: Oxford University Press, pp. 245-60.
- Davidson, D. 2001. *Subjective, Intersubjective, Objective*. Oxford: Clarendon Press.
- Dennett, D. C. 1973. *Brainstorms*. Montgomery, VT: Bradford Books.
- Dennett, D. C. 1987. Evolution, error and intentionality. In his *The Intentional Stance*. Cambridge, MA: MIT Press, pp. 287-322.
- Dennett, D. C. 1988. Quining qualia. In A. Marcel and E. Bisiach, eds., *Consciousness in Contemporary Society*. Oxford: Oxford University Press, pp. 43-77.
- Dennett, D. C. 1991. *Consciousness Explained*. New York: Little, Brown.
- Drekske, F. 1995. *Naturalizing the Mind*. Cambridge, MA: MIT Press.
- Dummett, M. 1993. *The Origins of Analytical Philosophy*. London: Duckworth.
- Fodor, J. 1968. *Psychological Explanation: An Introduction to the Philosophy of Psychology*. New York: Random House.
- Fodor, J. 1975. *The Language of Thought*. New York: Crowell.
- Fodor, J. 1983. *Modularity*. Cambridge, MA: MIT Press.

- Fodor, J. 2000. *The Mind Doesn't Work that Way*. Cambridge, MA: MIT Press.
- Hardcastle, V. 1996. *How to Build a Theory in Cognitive Science*. Albany NY: SUNY Press.
- Horowitz, T. and G. Massey. 1991. *Thought Experiments in Science and Philosophy*. Lanham, MD: Rowman and Littlefield.
- Jackson, F. 1998. *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Oxford University Press.
- Kant, I. 1781/7. *Critique of Pure Reason*, Norman Kemp Smith, trans. London, UK: Macmillan, 1927.
- Kitcher, Patricia. 1990. *Kant's Transcendental Psychology*. Oxford: Oxford University Press.
- Knobe, J. and S. Nichols. 2008. *Experimental Philosophy*. Oxford: Oxford University Press.
- Kripke, S. 1972. *Naming And Necessity*. In D. Davidson and G. Harman, eds., *The Semantics of Natural Languages*. Dordrecht, Holland: Reidel, pp. 253-355.
- McDowell, J. 1994. *Mind and World*. Cambridge, MA: Harvard University Press.
- Millikan, R. 1984. *Language, Thought, and Other Biological Categories*. Cambridge, MA: MIT Press.
- Newell, A. 1973. You can't play twenty questions with nature and win. In W. G. Chase, ed. *Visual Information Processing*. New York: Academic Press.
- Nozick, R. 1981. *Philosophical Explanations*. Cambridge, MA: Harvard University Press.
- Prinz, J. 2002. *Furnishing the Mind: Concepts and Their Perceptual Basis*. Cambridge, MA: MIT Press.
- Putnam, H. 1975. The meaning of 'meaning'. In his *Mind, Language and Reality: Philosophical Essays*. Cambridge, UK: Cambridge University Press.
- Pylyshyn, Z. 1984. *Computation and Cognition*. Cambridge, MA: MIT Press.
- Pylyshyn, Z. and R. Storm. 1988. Tracking of multiple independent targets: Evidence for a parallel tracking mechanism. *Spatial Vision* 3:1-19.
- Quine, W. v. O. 1953. Two dogmas of empiricism. In his *From a Logical Point of View*. Cambridge, MA: Harvard University Press, 1961, pp. 20-46.
- Searle, J. 1980. Minds, brains, and programs. *Behavioral and Brain Sciences* 3: 417-24.
- Sellars, W. 1963. Philosophy and the scientific image of man. In his *Science, Perception and Reality*. London: Routledge and Kegan Paul, pp. 1-40.
- Seyfarth, R. and D. Cheney. 2002. Dennett's contribution to research on the animal mind. In Brook, A. and D. Ross, eds. *Daniel Dennett*. New York and Cambridge: Cambridge University Press, Ch. 5.
- Sorensen, R. 1992. *Thought Experiments*. Oxford: Oxford University Press.

Thagard, P. 2000. *Coherence in Thought and Action*. Cambridge, MA: MIT Press.

Wittgenstein, L. 1953. *Philosophical Investigations*. G. Von Wright, E. Anscombe, and R. Rhees, eds., E. Anscombe, trans. Oxford: Basil Blackwell.