The Philosophy of Cognitive Science

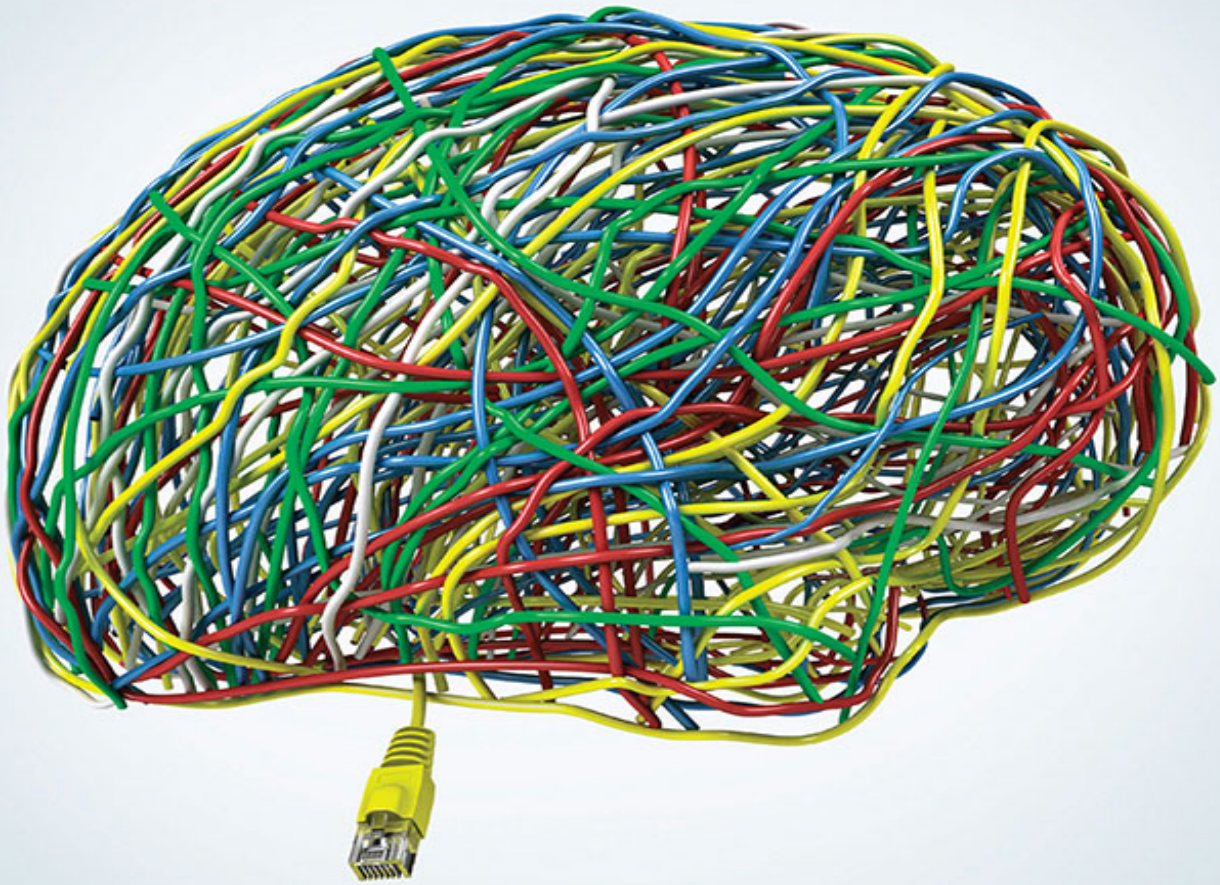

M. J. CAIN

Contents

To Julie, Maya and Theo

The Philosophy of Cognitive Science

M. J. CAIN

polity

First published in 2016 by Polity Press

Polity Press
65 Bridge Street
Cambridge CB2 1UR, UK

Polity Press
350 Main Street
Malden, MA 02148, USA

ISBN-13: 978-1-5095-0205-9

A catalogue record for this book is available from the British Library.

Library of Congress Cataloging-in-Publication Data

Cain, M. J.
The philosophy of cognitive science / Mark Cain.
pages cm
Includes bibliographical references and index.
ISBN 978-0-7456-4656-5 (hardcover : alk. paper) -- ISBN 0-7456-4656-5 (hardcover : alk. paper) -- ISBN 978-0-7456-4657-2 (pbk. : alk. paper) -- ISBN 0-7456-4657-3 (pbk. : alk. paper) 1. Philosophy and cognitive science. 2. Philosophy of mind. I. Title.
BD418.3.C34 2015
128′.2--dc23

2015008643

The publisher has used its best endeavours to ensure that the URLs for external websites referred to in this book are correct and active at the time of going to press. However, the publisher has no responsibility for the websites and can make no guarantee that a site will remain live or that the content is or will remain appropriate.

Every effort has been made to trace all copyright holders, but if any have been inadvertently overlooked the publisher will be pleased to include any necessary credits in any subsequent reprint or edition.

For further information on Polity, visit our website: politybooks.com

Acknowledgements

Writing this book has taken much longer than I anticipated when I first sketched its outline several more years ago than I care to dwell upon. I therefore need to thank my editors Emma Hutchinson and Pascal Porcheron at Polity for their continued patience, support and encouragement. Thanks also to my philosophy colleagues at Oxford Brookes University, Steve Boulter, Constantine Sandis, Dan O'Brien and Bev Clack. Rumour has it that they find many of my philosophical views slightly bonkers, but that hasn't stopped them providing a warm and stimulating environment to work in. I've also had the pleasure of supervising the PhD research of Dana Vilistere, Annie Webster, Marcus Westberg, Andrew McIntyre, Colin Beckley, Linda Fisher-Hoyrem, Esther Shallan and Istvan Zardai, all of whom have provided very helpful feedback on my developing ideas. Massive thanks to my daughter Maya Reid-Cain for helping me wrestle with the various diagrams used in this book and for keeping her disdain at my IT ineptitude in check. Most of all thanks to my wife Julie and children Maya and Theo for all their love and support. This book is dedicated to them.

1 Cognitive Science and the Philosophy of Cognitive Science

1 Introduction

This is a book about the philosophy of cognitive science. That topic immediately raises two questions: what is cognitive science and what is the philosophy of cognitive science? In one respect the answers to these questions are obvious: cognitive science is the scientific study of cognition, and the philosophy of cognitive science is that branch of philosophy that addresses philosophical questions generated by the scientific study of cognition. But these answers are hardly illuminating as they raise a number of subsidiary questions. What is cognition? What are the key assumptions and methods adopted by those who attempt to study cognition scientifically? How does cognitive science relate both to other sciences and to our commonsense understanding of ourselves as minded agents? When and how did cognitive science emerge as a distinct discipline? How does the philosophy of cognitive science relate to other branches of philosophy such as the philosophy of mind and the philosophy of science? In this chapter I will address these general questions and so provide the foundations for the more specific discussions of the later chapters.

2 Cognition

What exactly is cognition? In order to answer this question it is helpful to begin from a commonsense perspective. According to commonsense – or at least the commonsense perspective of most twenty-first-century Westerners –

human beings can be distinguished from inanimate physical objects in having a mind. What, according to commonsense, is involved in our having minds? Here is a by no means exhaustive list: we think; we perceive the external world by means of our senses; we experience pain and other sensations when our bodies are appropriately stimulated; we experience moods such as depression and light-heartedness; we experience emotions such as anger, joy and jealousy; we are conscious both that we think and feel and how we specifically think and feel; we act on the basis of our decisions and intentions, which in turn often reflect how we perceive the world to be, what we think and what we want; when we act we are conscious of what we do without having to rely on our external sense organs; we recollect our past thoughts, actions and experiences; we imagine particular scenarios; and so on.

In so describing the activities that are central to human mentality from the commonsense perspective I did not use the term ‘cognition’. But if one were to ask which of the above activities were best described as involving cognition, most people would answer ‘thinking’. So one might say that, to a first approximation, cognition is thinking. But what exactly is thinking? Thinking is a mental process or activity that results in having a thought. Such thought processes range from the intellectually demanding and abstruse (as when one thinks when doing philosophy) to the banal and everyday (as when I address the question of whether I have enough time before my next lecture to buy a coffee). Some concern matters of great importance, others not. Some issue in thoughts that have an immediate impact on action and others don’t. Some issue in thoughts that become deeply entrenched and influence many subsequent thought processes, while others issue in thoughts that are fleeting and inconsequential. Some concern counterfactual or hypothetical matters (would I

have been late for my lecture had I queued for a coffee?) whereas others are straightforwardly factual.

If thinking is a mental process that issues in a thought, then what exactly is a thought? In everyday talk the term ‘thought’ is most commonly used to refer to the mental state of considering a particular hypothesis or answer to a question or that of committing oneself to a particular hypothesis or answer. Such mental states are, or are closely related to, beliefs. For when one considers a particular answer to a given question one is paving the way to holding a particular belief, and when one commits oneself to a particular answer one has thereby acquired a belief. Beliefs are examples of what philosophers call propositional attitudes, as believing that p involves adopting the belief attitude to the proposition p. But there are propositional attitudes that are not beliefs, for, just as one can believe that p, one can desire that p, intend (to make it the case) that p, hope that p, fear that p, expect that p, and so on. Consequently, although it slightly strains everyday usage, one might say that thoughts are propositional attitudes, so that thinking is the mental process that results in the acquisition of a propositional attitude, be it a belief, a desire, an intention, or whatever.

In virtue of being relations to propositions, propositional attitudes have meaning or semantic properties. They are therefore akin to declarative sentences of a natural language. Just as such sentences are about particular things (or types of thing) or states of affairs and represent them as being a particular way, so do propositional attitudes. For example, my belief that aardvarks eat termites is about aardvarks and represents them as being termite eaters. What a propositional attitude is about and how it represents its object are elements of its meaning, but philosophers tend to use the term ‘content’ when talking about the meaning of propositional attitudes. So, for

example, the sentence ‘aardvarks eat termites’ has a particular meaning, and the belief that aardvarks eat termites has a particular (corresponding) content. There is a further commonality between declarative sentences and propositional attitudes. Sentences are made up of simpler components, namely words, and the meaning of a sentence is a product of the meaning of the words that belong to it and the way they are put together (the syntactic structure of the sentence).[1] It is natural to think that something similar is true of propositional attitudes. The simpler components of propositional attitudes are concepts, and the content of a propositional attitude is a product of its constituent concepts and the way they are put together. Hence, one cannot form the belief that aardvarks eat termites without employing each of its constituent concepts.

I have suggested that, to a first approximation, cognition is thinking. Allied with the idea that thinking is a mental process that results in the formation of propositional attitudes, the implication might seem to be that much of what goes on in the mind falls outside of the domain of cognition. This is because philosophers standardly distinguish between propositional attitudes and other mental states such as sensations, emotions and perceptual experiences. However, even if we want to uphold such a distinction (say, on the grounds that such mental states differ from propositional attitudes in that they have an essential phenomenal, qualitative or ‘what it is like’ aspect), it would perhaps be a little rash, for several reasons, to conclude that only propositional attitudes and the processes that generate them count as cognitive. First, consider perception. Perception is a process where the external world stimulates an individual’s sensory organs, resulting in a perceptual experience. In the case of vision, light reflected off external objects is focused onto the

retina, a light-sensitive surface at the back of the eye that sets off a mental process which results in a visual experience. Even if, as many philosophers have thought, perceptual experiences are very different from propositional attitudes in not involving the deployment of concepts and in having an intrinsic qualitative character, they are like them in a key respect. For perceptual experiences are representational in the respect that they are typically about objects located in the external world and represent them as being a particular way. For example, my current visual experience was not only caused by the computer directly in front of me but is also about that computer and represents the outer world immediately before me as containing a grey, rectangular object. In virtue of their being representational, it seems reasonable to think of perceptual states as belonging to the domain of cognition if propositional attitudes are paradigmatic cognitive states.

A second point relates to the nature of the mental process of thinking. Typically, when one acquires a propositional attitude by means of thinking, a process takes place in one's mind that is extended in time. This process has stages each of which involves moving from one or more propositional attitudes to another, where the earlier propositional attitude(s) in each stage justify the later. Here is an example. Feeling a little de-energized, I wonder if I have enough time to buy a coffee before my next lecture. I start from the beliefs that it will take me ten minutes to walk from the café to the lecture hall and that, given the current length of the queue, it will take me five minutes to purchase a coffee. Given that my lecture starts at 1.00 p.m., I infer that I must leave the café by 12.45 p.m. at the latest to be on time. I look at my watch and come to believe that it is 12.50 p.m. and draw the implication that I cannot

buy a coffee and arrive at my lecture on time. Given my strong desire not to be late, I decide to give coffee a miss.

When we think, we are often aware of the stages we go through to reach the propositional attitude that is the end product of the thinking process. Or, if the thought process is so quick and routine that we don't have such awareness, we can retrace our steps by deliberately seeking to justify our conclusion and so gain awareness of our thought process retrospectively. But, with respect to perception, things don't seem to be like that. I open my eyes, orient them to the world and have a perceptual experience without having any awareness of executing an extended process the earlier elements of which justify the later elements. So there would appear to be a substantial contrast between thinking and perception once again. But on second thoughts the contrast might not be so great. Just because we are not aware or conscious of any stages of inference in the process of perception, it does not follow that there are no such stages once we countenance the possibility that much might go on in the mind that is unconscious. Indeed, as we shall see in later chapters, in the 1950s and 1960s psychologists began to hypothesize that perception does involve unconscious inference.[2] If one takes this idea seriously, then one has grounds for including perception in the domain of cognition. Moreover, similar points can be made concerning other mental states and processes. Suppose I am in a café having a cappuccino. I reach out to grasp the cup, raise it to my lips and take a sip. This is a routine everyday event, and when it takes place I am not aware of my having done much in the way of thinking. But if we try to make sense of how we carry out such actions it is clear that an awful lot of complex internal processing must be involved. Once again, in the 1950s and 1960s the idea developed that our actions are driven by thought-like inferential processing that draws upon and

coordinates constantly changing information about the state of the external world and one's own body coming from a range of distinct sources. So we might have grounds for regarding action and motor control as belonging to the domain of cognition.

A third point is that, even if there is a significant difference between thinking and other aspects of the mind such as perception and action, the latter must interface with the former. Consider perception first. Perceptual experiences might not be propositional attitudes or the upshot of thought processes, but they do play a key role in determining our propositional attitudes. For we routinely form beliefs and update our stock of beliefs on the basis of our perceptual experiences. Were this not the case, then perception would be of little use to us, as its value resides in its ability to provide us with knowledge about the world that we can utilize in deciding how to act so as to satisfy our needs and desires. For example, suppose you are in the supermarket looking for a tomato and while in the fruit and veg section looking at a tomato you have a visual experience as of a spherical, red object. If this experience is going to help you to satisfy your desire for a tomato, you must take it at face value and come to believe that there is a spherical, red thing before you, and you must further infer the belief that that spherical red thing is a tomato. All this clearly requires the output of perception to be taken up by thought processes. A parallel point can be made about motor control and action. Even if the movements that we make when we act are not directly driven by thought processes, those movements must be routinely and systematically motivated by our propositional attitudes. Were this not the case then our propositional attitudes would be robbed of their central function of enabling us to act so as to satisfy our desires. For example, there wouldn't be much benefit in desiring an elderflower pressé and

(correctly) believing that there was a bottle of elderflower pressé in the fridge if that belief and desire pair were incapable of making an impact upon how one acted. In short, then, it is a fundamental property of thinking that it interfaces with both perception and action. This suggests that one should resist drawing a fundamental divide between perception and action, on the one hand, and thinking, on the other, and view only the latter as belonging to the domain of cognition.

I began this section by asking what cognition is and put forward the suggestion that cognition is thinking understood in commonsense terms as a process involving inferring propositional attitudes from other propositional attitudes. However, I argued that, although such thought processes are clear-cut cases of cognition, there are other processes that take place in the mind, namely those involved in perception and action and motor control, whose similarities and relationship to thinking are such as to imply that they count as cognitive processes. Thus, if cognitive science is the scientific study of cognition, it should be concerned as much with perception and action as it is with central cases of thinking.

3 Science and cognition

In the previous section I gave a preliminary account of cognition. That account implies that cognitive science is the scientific study of such phenomena as thinking, perception and action. Thus, the concern of the cognitive scientist is to explain how we humans are able to think, perceive and act, to uncover precisely what goes on within our minds when we exercise such cognitive capacities. If successful, cognitive science will reveal the core properties of humans that enable us to cognize and thus differentiate us from all the inhabitants of the universe that are not capable of

cognition. Put this way, it is clear that cognitive science rests on the assumptions both that cognition is the kind of phenomenon that is amenable to scientific investigation and that we humans are cut out to execute such a scientific investigation. However, such assumptions are not universally held within the philosophical community. In this section I will describe some of the most prominent reasons for scepticism about the viability of cognitive science.

A first reason for scepticism relates to a metaphysical view about the mind most associated with Descartes (1985) and is generally known as substance dualism. According to Descartes, the human individual is a two-component system consisting of a body and a mind. The body is an inhabitant of the physical world and so, reflecting the nature of the physical world in general, is a mechanical system whose behaviour is governed by laws of nature that can be stated in mathematical terms. Moreover, the body is essentially spatially extended; that is, it has a spatial location and takes up physical space. In virtue of all this, the human body is precisely the kind of system whose nature and workings can be investigated from the scientific perspective. The mind, on the other hand, is essentially a thinking thing. It does not inhabit the physical world and, in virtue of having free will, it is not a mechanical system governed by laws of nature. Nevertheless, the mind is harnessed to the physical body and engages in a systematic two-way process of causal interaction with it. For example, the body can affect the mind, as when the physical stimulation of the sense organs causes a perceptual experience, and the mind can affect the body, as when a decision to execute a particular action causes the body to move. For Descartes, in virtue of being a fundamentally different kind of thing than the body, the mind, and therefore cognition, is not amenable to scientific study. If

Descartes is correct on this point, cognitive science is not a viable enterprise.

How should we respond to this Cartesian line of thought? Descartes produced a number of arguments for dualism, but the general view of the philosophical community is that none of these are successful. Moreover, it is widely held that dualism cannot make sense of the existence of causal interaction between the mind and the body: if mental phenomena reside outside of the physical domain, how could they cause or be caused by physical phenomena taking place within the body? Indeed, as we shall see, most mainstream cognitive science is underwritten by a commitment to the anti-dualist view that the mind is in some substantial respect the brain, so the study of cognition is the study of the workings of a physical system.

A second source of scepticism regarding the viability of cognitive science accepts that cognition is the kind of phenomenon that could in principle be investigated scientifically but questions whether we humans are up to the task of conducting such an investigation. One way of characterizing this line of thought is in terms of Noam Chomsky's (2000) distinction between problems and mysteries. A problem is a question that is difficult to answer but which we have some hope of answering. A mystery is a question which we have no hope of answering because it is beyond our cognitive powers to do so. Which questions are problems and which mysteries can be species relative. For example, questions we humans find easy to answer might be utterly mysterious to a rat. We humans are bound to have cognitive limitations which make some questions that we can frame mysterious to us. One popular suggestion as to what might be a mystery for us is phenomenal consciousness. Colin McGinn (1989b) argues that phenomenal consciousness, the 'what it is like' aspect of our experiences (Nagel, 1974), is a physical phenomenon

but that we are incapable of explaining how physical phenomena can generate consciousness.[3] Now I haven't characterized phenomenal consciousness as belonging to the cognitive domain, but it wouldn't be outlandish to argue that questions about human cognition are mysteries for us. For, in attempting to answer such questions, we are attempting to use our cognitive powers to understand our cognitive powers.

I don't think we should expect to be able to produce a knock-down response to this kind of scepticism. Perhaps the most sensible response would be to argue that we should evaluate the power of the objection in light of the success of our ongoing attempts to explain cognition. Thus, the power of the scepticism would be undermined to the extent that cognitive scientists produced theories and explanations that were productive and successful when judged by the criteria operative in other scientific domains.

A third source of scepticism concerning the viability of cognitive science relates to my characterization of the discipline in terms of our commonsense understanding of ourselves. For, one might argue, the phenomena that seem significant from a commonsense perspective might not be tractable from a scientific perspective, in that they are too complex and messy in being the products of many distinct but interacting factors, each of which belong to different scientific domains. Oddly enough, this idea can be found in the work of Chomsky, who, to my mind, has made one of the most valuable contributions to cognitive science.

Suppose that I utter a particular sentence. To understand me you will have to cognize what individual words I produced and how they were put together to make the sentence in question (that is, the grammatical or syntactic structure of the sentence). This is no mean feat. For one thing, there is no neat correspondence between words and

sounds. This is because distinct utterances of one and the same word can be very different at the sonic level; think of how different the word ‘dog’ sounds when uttered by a small child, a young woman and an elderly man. And, with respect to syntactic structure, that is not an immediately perceivable property of our utterances, as is indicated by the phenomenon of structural ambiguity. For example, the sentence ‘he chased the dog with the stick’ could have a structure such that it means that the dog had the stick or one such that it means that the man had the stick. But suppose that you overcome such problems and work out what individual words I uttered and how I structured them in the sentence. If you know the meaning of the individual words of the sentence, you will be able to further cognize the literal meaning of my utterance, assuming that the literal meaning of a sentence is exhaustively determined by the meaning of its component words and its syntactic structure.

However, we normally understand more than the literal meaning of the sentences we hear, and successful communication depends upon this. For we appreciate the communicative intentions of the speaker (Grice, 1975). For example, suppose that you come into my office and I say ‘There’s a nice fresh breeze coming through the door.’ I might be using that sentence with the literal meaning that it has to say any one of several things. I might be aiming to state a fact that does not require you to do anything. Or I might be aiming to point out that you have left the door open and be requesting that you close the door. If communication between us is to be successful, you must appreciate what I am aiming to do in producing the sentence – what my communicative intentions are. But if you are to do this you need to draw upon an appreciation of a potentially wide and disparate range of information. In this case such information might relate to my facial

expression, my tone of voice, the temperature conditions in the room, your history of interactions with me, the likelihood of my desiring privacy for our meeting, the conventions governing meetings in my institution, and so on. Now much of this potentially relevant information lies outside of the body of linguistic knowledge that one might think we need to be capable of producing syntactically well-formed sentences and appreciating the literal meaning of the sentences we encounter. Chomsky's point is that, because understanding communicative intentions involves appreciating such a potentially wide and disparate range of information from outside the linguistic domain, it is going to be impossible for us to explain how we understand communicative intentions.

A related point is that sciences carve the world into domains of enquiry within which the scientist has some hope of making progress. But in the real world such domains often interact with one another, so that phenomena that are salient from a commonsense perspective are often massive interaction effects from the scientific perspective. Consequently, explaining such phenomena would require us to engage in 'the science of everything' – something that is not possible for us.

Chomsky's own engagement in cognitive science suggests a response to this sceptical line of thought that involves conceding the point that lies at its core. Our commonsense conception of ourselves as cognizers motivates our engagement in cognitive science and provides the discipline with a starting point of questions to address. But it is an open question as to which of those questions remain on the agenda and what questions are to replace those that are abandoned. How these issues are to be settled is something that cannot be determined in advance of our actual engagement in cognitive science. This situation with respect to cognitive science is typical of science in general,

so it motivates scepticism concerning the viability of cognitive science only if it motivates scepticism about the viability of science in general.

4 Science

The upshot of the discussion of the previous section is that it would be somewhat hasty to conclude that cognitive science is not viable. But this does raise a further question: what characteristics should we expect cognitive science to have in virtue of being a science? In this section I will address this question.

Science is one of the great triumphs of Western civilization and has provided us with a systematic body of knowledge of the workings of the natural world. I am very wary of attempting to provide a general account of the nature of science. Rather, I will describe some prominent features of science that will prove to be very relevant when understanding the core commitments of cognitive science and when addressing philosophical questions about cognitive science.

First, most scientists assume that happenings in that portion of the natural world with which they are concerned are not entirely random and irregular; rather, they are governed by laws. Thus, one of the main goals of science is to discover the laws that govern the workings of the natural world. Here are some examples of such laws of nature: freely falling bodies accelerate at a uniform rate (Galileo); planets have an elliptical orbit (Kepler); the volume occupied by a gas is inversely proportional to the pressure on it (Boyle); the strength of the gravitational attraction between two bodies depends on the product of their masses and is inversely proportional to the square of the distance between them (Newton). Note that some laws (particularly those operating at the quantum level) are

probabilistic rather than deterministic. That is, they are of the form that, if x is the case, then, with probability P (where P is less than 1), y will be the case. And some laws are *ceteris paribus* (all else equal) laws. Such laws have the form that, if x occurs, then y will occur, all else equal. Hence, the claim that such a law holds is not necessarily undermined by the observation that an x has occurred without being followed by a y. For, all else might not have been equal in such a case. Jerry Fodor (1987) has provided a well-known example of a *ceteris paribus* law from geology: a meandering river erodes its outer bank, *ceteris paribus*. A case where a meandering river might fail to erode its outer bank would be one where the bank had been concreted over; in such a case all else wouldn't be equal. Given the general role of laws in science, we should seriously entertain the possibility that the cognition is law governed and that one of the main tasks of cognitive science is to uncover such laws.

Second, another of the major goals of science is to explain features of the natural world. This involves not merely describing the natural world but saying why it is the way it is and how things work in the natural world. Laws are relevant here as they have an important role in explaining natural events. For example, Galileo's law plays a key role in explaining why a cannonball and a marble when dropped from a particular height at a given time will hit the ground at the same point in time, in violation of our commonsense expectations that the cannonball, in virtue of being heavier, would hit the ground first. Indeed, explaining why a particular law holds typically involves appealing to more basic laws. For example, Newton appealed to his laws of motion and the inverse square law in order to explain Kepler's law of planetary motion. My initial characterization of cognitive science highlighted the role of

explanation, as I said that cognitive science is concerned with explaining our cognitive abilities.

Third, science is an empirical discipline in that observation and experiment play a central role in the scientific project. This contrasts science with any purely *a priori* or armchair activity. One prominent account of the precise way in which science is empirical was developed by Karl Popper (1959). Popper was concerned with the so-called demarcation problem, with specifying precisely what distinguished science from non-science. He argued that it was a mark of a scientific theory that it was falsifiable – that is, a scientific theory has implications as to what observable phenomena will take place in particular circumstances such that there are, in principle, observations we could make that would definitively show the theory to be false. Thus, the scientist proceeds by a two-stage process of conjecture and refutation. This involves producing a conjecture or hypothesis in order to explain some target phenomenon. The observational implications of the conjecture are then worked out and an experiment conducted in which the relevant observations are made. If the observations are inconsistent with the observational implications of the conjecture, then the conjecture is rejected as false and the scientist goes back to the drawing board. If the observations do not falsify the conjecture, then the scientist conducts further experiments in order to test and falsify the conjecture.

Fourth, in the course of explanation, scientists often postulate theoretical entities. These are entities that are not observed but are invoked to explain phenomena that are observable. For example, suppose a biologist wanted to explain why organisms generally bear many similarities to their parents but have some differences from them. The explanation proffered by the biologist will appeal to genes. The key point is that the genetic explanation[4] is not

produced on the basis of observing genes and their activities. Rather, the genes are postulated in order to make sense of the otherwise mysterious observed phenomena. Other prominent examples of theoretical entities are atoms, quarks and photons. Given the role of theoretical entities in science, we should expect cognitive scientists to postulate theoretical entities in theorizing about human cognition.

Fifth, although science is empirical in nature, the role of observation in science is not as straightforward as implied by Popper's work as described above. W. V. O. Quine's 'The Two Dogmas of Empiricism' – first published in 1951 – is probably the most important philosophical article written in English since the Second World War. In this paper Quine seeks to undermine two central theses of the empiricism of his day. The second of these theses is that the meaning of a sentence is its method of verification – that is, the empirical means of determining its truth-value. Thus, any meaningful sentence can be translated without loss into a sentence about experience or observations (this doctrine is known as reductionism). In rejecting verificationism, Quine champions a position that has become known as the Duhem–Quine thesis.[5] It has this name because Quine's basic idea was anticipated by the French physicist and philosopher of science Pierre Duhem (1954), in his book *The Aim and Structure of Physical Theory*.

Suppose that one wants to determine by empirical means whether a particular (declarative) sentence is true or false. One will then have to make relevant observations using one's senses. It is natural to think that, with respect to any true sentence, it is in principle possible to establish that it is true by making relevant observations; and that, with respect to any false sentence, it is in principle possible to establish that it is false by making relevant observations. Applied specifically to science, the idea would be that, with

respect to any sentence expressing a scientific theory, it is in principle possible to establish whether it is true or false by observational means. Quine's basic point is that this is not the case because it is not possible to verify a sentence or theory in isolation. This is because individual sentences or theories do not by themselves have the kind of observational implications that enable them to be verified. Rather, it is only collections of sentences or theories that have such implications. To see this, consider a topical example. Attention Deficit Hyperactivity Disorder (ADHD) is a condition attributed to many children in the Western world today. The condition manifests itself in a child's inability to concentrate on schoolroom tasks and a tendency to be disruptive. What causes ADHD and why is it so prevalent today? One theory is that it is the result of a deficiency in Omega 3. How are we to verify this theory or the sentence that expresses it? Here is an answer. Give children with the condition an Omega 3 supplement and observe their subsequent behaviour. If the children are observed to undergo an improvement in their powers of attention and concentration, then that tells in favour of the theory, suggesting that the theory may well be true or is a serious candidate for truth. If, however, the children are observed to undergo no such improvement, then that tells against the theory, suggesting that it is (probably) false. The problem is that attempting to verify the theory/sentence in this way involves making a whole load of assumptions that are independent of the theory/sentence. For example, the assumption that the supplement contains Omega 3 in a form that can be readily absorbed by the body; the assumption that if the condition is due to a deficiency in Omega 3 then it can be remedied by taking a supplement; assumptions about how concentration/attention powers manifest themselves in observable behaviour; and so on. Making such assumptions will involve committing oneself to the truth of substantial

scientific theories. If one doesn't make these assumptions, then one will regard the above way of verifying the theory as illegitimate. If one makes alternative assumptions, then one will be committed to adopting a different means of verifying the theory. And if one makes no supplementary assumptions at all, then one will have no idea how to verify the theory, as the theory in itself will not tell one how to go about verifying it.

Now suppose that one does make the above described assumptions and that no improvement in the children is observed. Does this tell against the theory? According to Quine, the answer to this question is negative. For, in order to deal with the recalcitrant observational data, one can hold onto the theory but reject some of the associated assumptions. Thus Quine says, 'any statement can be held true come what may, if we make enough adjustments elsewhere in the system' (1951: 43). All this is captured in the Duhem–Quine thesis, according to which:

> any theory can be held onto in the face of recalcitrant observational data by making suitable adjustments elsewhere in one's system of commitments.

The Duhem–Quine thesis is widely held by philosophers of science. With respect to cognitive science, an implication of the thesis is that it may well be difficult to adjudicate between different theories about our cognitive lives, as their advocates disagree on the significance of the empirical data they gather.

A sixth feature of science concerns its relationship to commonsense. If a human individual is going to prosper in a challenging and potentially dangerous world, they will need to have a battery of concepts in order to categorize the phenomena with which they interact. And they will need to employ those concepts to form knowledge or beliefs concerning how such phenomena behave. Armed

with such knowledge, they will then be able to predict and explain worldly events and so enhance their survival prospects. When philosophers and cognitive scientists talk about commonsense, they are talking about the relatively untutored and unsophisticated conceptual schemes and associated belief/knowledge systems that ordinary people routinely employ in their everyday lives. Commonsense is therefore distinguished from mature and sophisticated scientific theories. However, that is not to say that commonsense bears none of the characteristics of science. Questions about commonsense are prominent within cognitive science, and a widely held view represents it as bearing some of the characteristics of mature science. First, it has several components each of which relate to a distinct subject matter, so that commonsense is underwritten by the assumption that there are different kinds of phenomena in the world that behave in their own distinctive way. In other words, commonsense, like science, has its component disciplines.[6] Moreover, some of them correspond to prominent scientific disciplines. For example, there is commonsense physics that is concerned with inanimate physical objects, commonsense psychology that is concerned with minded agents, particularly humans, and commonsense biology that is concerned with biological entities such as animals. Second, commonsense concepts can be quite abstract in the respect that they group together phenomena that differ widely in terms of their perceivable properties, such as their shape, size and colour, and in the respect that they refer to the unobservable. For example, commonsense physics employs a quite general concept of a physical object that is utilized in such beliefs that unsupported physical objects fall to the ground. And commonsense biology employs the concept of a hidden essence, in that we think of creatures as having characteristics we cannot directly perceive that determine

the kind of thing they are and which are causally responsible for the non-essential surface characteristics that we can perceive (Keil, 1989).

One quite natural view of the relationship between science and commonsense is that the former is born of the latter. Science is a cultural phenomenon which in an organized and rigorous form has a relatively short history of approximately 500 years and so is very much pre-dated by commonsense. Commonsense constituted the starting point for science, in that it provided it with its core questions and conceptual scheme. For example, if commonsense physics assumes that physical objects causally interact with one another and are governed by generalizations that advert to their physical properties, then those assumptions generate questions for the physicist – questions as to the underlying nature of these physical properties (for example, what is heat?), the underpinnings of the generalizations (for example, why do unsupported bodies fall?) and the identity of any further physical properties and generalizations not currently recognized by commonsense. However, even if science is born of commonsense, it doesn't follow that in its mature form a science shares many of the characteristics of its parent. Science is a self-conscious research endeavour that is driven by a relentless search for progress and the truth in a way that commonsense is not. As a result, science often exposes the limitations of commonsense: the parochial and inadequate nature of its conceptual scheme, the falsity of its assumptions, the limits of its explanatory powers and ambitions (Churchland, 1979). To see this, just think of how relativity theory and quantum mechanics have left commonsense physics behind. In short, then, individual sciences have a basis in commonsense but, in their mature form, have often moved a considerable distance from their parent.

In this section I have described a number of key characteristics of science in general. In virtue of its status as a science, we should expect cognitive science to share these characteristics. Thus we should expect cognitive science to (i) seek to uncover laws governing the workings of the cognitive mind; (ii) utilize such laws to explain cognitive phenomena; (iii) appeal to unobservable entities in its laws and explanations; (iv) utilize observation and experimentation; (v) confirm theories in a Duhemian–Quinean manner; and (vi) be born of yet to have moved beyond commonsense psychology.

5 The birth of cognitive science

So far I have been talking quite generally about cognitive science, but it is now time to be a little more specific about its origins and core commitments. Attempts to study cognition from a scientific perspective are hardly new. For example, David Hume ([1738] 1978) described the project executed in his work *A Treatise of Human Nature* as being an attempt to do for the mind what Newton had done for the external physical world. Indeed, some of the concepts and ideas employed by Hume have their analogues in contemporary cognitive science. Nevertheless, it is generally assumed that cognitive science came into existence as a discipline with a distinct identity only in the late 1950s and early 1960s with an intellectual turn known as the cognitive revolution.

The cognitive revolution was a revolt against a movement in psychology that dominated that discipline, at least in the English-speaking world, throughout the first half of the twentieth century. The movement in question was behaviourism, which had J. B. Watson (1913) as one of its pioneers and B. F. Skinner (1953) as its most prominent advocate in its maturity. Psychology as the study of mind