

Philosophical Foundations of Neuroscience

Bennett & Hacker

[ISBN: 978-1-405-10838-6 - April 2003 - Wiley-Blackwell](#)

Foreword

Denis Noble CBE FRS hon FRCP

This book was simply waiting to be written. The reductionist agenda in biological science has generated so many conceptual difficulties that someone, sometime, had to analyse these problems in depth from outside the reductionist viewpoint. That a neurophysiologist and a philosopher should combine to do so is also a sign of the times. As biology moves on to address the complexity and extraordinary subtlety of life, now that it has broken it down into its smallest pieces, we will find this kind of combination of skills and ways of thinking even more necessary. As the authors make clear, philosophy (at least in the analytical form practised here) and empirical science are not in opposition. Rather they deal with different kinds of question. Yet, since a conceptual scheme is necessary to any fruitful experimentation, we cannot avoid asking both kinds. Keeping a clear head while we do so is not as easy as it may seem!

I must issue a warning: this book is highly controversial. Some of my scientific colleagues will strongly challenge, and will surely be deeply provoked by, the claim that neuroscience has frequently and systematically confused conceptual and empirical questions. To them I would say, first, that the authors clearly recognize the brilliance and phenomenal achievements of the scientists whose conceptual work they analyse. This is emphatically not a book debunking experimental science, any more than the fact that most physiologists now dismiss the dualist philosophy of Sherrington or Eccles detracts in any way from recognizing the immense significance of their scientific achievements. We find it perfectly possible to admire the experimental and associated analytical skills while wincing when we see how completely trapped they were in their outdated and indefensible philosophical position.

Second, I would appeal for some patience and humility. Patience, because as a physiologist who has interacted with (and published with) professional philosophers of various persuasions for over 40 years,¹ I have to say that I find scientists unthinkingly debunking philosophy more often than the other way round. Humility, because the issues are of the utmost social importance. Some of the claims of reductionist science are not only conceptually incorrect or even unintelligible, they have major social implications. The words we use, the concepts by which we analyse and present biological discovery, deeply affect the way in which we see ourselves as human beings. For that reason, if for no other, a critical debate is necessary. The authors of this book have thrown down a major challenge in that debate.

The controversial nature of this book arises because the particular reductionist philosophical position it criticizes is very widely held today within the scientific community (and also by some well-known philosophers). Moreover, for most of them, this position is a *methodological* necessity, perceived to be the only paradigm for science to successfully explain things. The first reaction to the counter-argument, as presented here, will be to protest that somehow science is being (unnecessarily?) circumscribed; that some problems are, as it were, being taken from its grasp. I would argue the other way. The first step to scientific progress is to ask the right questions. If we are conceptually confused, we will ask the wrong questions. The authors illustrate this in detail with many examples. It is hard to

escape the confines and confusions of the culture in which one finds oneself. The history of philosophy shows that, just as much as the history of science. The central appeal of this book is to throw off the remaining legacy of the Cartesian confusions, first expressed as a duality of mind and body, but latterly expressed as a duality of brain and body. The authors show that, although the first required belief in a non-material substance, while the latter is wholly materialistic, many of the conceptual problems (essentially those of the 'ghost in the machine') are the same. For our dualist predecessors the ghost was an actual immaterial substance, for us it is 'the "I"' (or 'inner eye' or whatever) that 'sees' the qualia that 'form our experience'. This is what may lead us to ask which group of cells, or even which neurone(!), is doing the 'seeing'. The point here is that simply replacing 'I' or 'inner eye' by the brain or a part of the brain doesn't avoid the problem.

The key to understanding the confusions here lies in an analysis of the logical conditions for ascribing mental and psychological properties. This is not easy. It involves one of the most difficult of twentieth-century philosophical ideas, that of the 'private language argument': what it is to say things like 'I feel pain' or 'I see red'. I struggled through the ramifications of this argument many years ago before writing my own contributions to the philosophy of biology. I wish I had had the benefit of the relatively easy path that Bennett and Hacker have provided. Even those who fundamentally disagree with their arguments (and I look forward to seeing them engage in debate) must surely acknowledge that this is a sustained and valuable exposition of an important and influential philosophical position.

Although I would describe that position as philosophically radical (in the correct sense of that word: going back to basic roots and eradicating those that shouldn't be there), it is often dismissed by scientists as conservative because it may appear to restrict using language in new ways. Yet, they would argue, science cannot advance without doing that. And what better way to achieve it than to start with metaphor or *façons de parler*, consolidate with dead metaphor (metaphors that become part of everyday language - constructivists argue that that is the way language evolved) and finally end up with a change in our conceptual scheme? Indeed, why not, if that is what will enlighten us, lead us into new conceptual territory, formulate new theories. But there is a simple test for whether that could work in any particular case. For each such metaphorical (or similar) change in use or meaning, or novel piece of terminology (such as 'qualia' or 'memes'), imagine stating its opposite, and then ask whether any conceivable experiment could test empirically between the two. The deep problem for many 'novel' concepts and language uses in reductionist approaches is that this test totally fails. The novel use of language is then not so much a scientific as a political or social tool. If you doubt this, try imagining an experiment to test between the existence or non-existence of qualia. Or for whether or not the brain makes representative maps (which are not homunculi incidentally). Or for brain states that 'explain' rational thought (rather than being a necessary physical basis for its existence). Surely we should only introduce new terminology where, as with quarks and black holes, we provide the empirical criteria for determining their existence?

Perhaps the problem for many scientists is to imagine what would happen if we abandoned the universality of the reductionist approach. For sure, the nature of science would change. But so it should! We would have to recognize that causation and explanation do not always run upwards from lower to higher levels. And, surely, at a time when we have already come to understand the extent to which causation runs in the opposite direction (higher-level states in biological systems even influence something as fundamentally lower-level as gene expression), how can we possibly imagine that we will progress without recognizing the validity of explanations at all levels? One of the criteria for determining the level at which explanation succeeds is to ask what can sensibly be ascribed at different levels. It does not make sense to look for explanations at levels lower than that for the applicability of the

relevant predicates.

This is particularly true of rational behaviour, including the use of language. The argument is basically very simple. We cannot, coherently, deny our own rationality. Otherwise we would have difficulty meaning what we say or being convincing in saying it, which is precisely what happens in the sad cases of those mentally ill people who nevertheless are aware of, but can't help, their irrationality. If we really could succeed in 'reducing' rational behaviour simply to molecular or cellular causation then we would no longer be able meaningfully to express the truth of what we had succeeded in doing. But, thankfully, no such reduction is conceivable. We know what it is to be rational, and what it is to lose that capacity. That knowledge has nothing to do with the question whether there exist specific and causally sufficient neural states and interactions while I am writing this review, for example. Of course they do. And, if we can discover them, they may well provide a complete explanation for the mechanisms of my brain while thinking and writing. The main claim of anti-reductionism in science is that such a complete explanation of mechanisms at one level does not necessarily explain what exists and happens at higher levels. Indeed we may need to know about the higher levels in order to explain the lower-level data that form an input to the mechanisms involved (which is what must be the case in writing this review! - one of the inputs was my reading this book, but the book is not thereby 'inside' my brain).

The most spectacular case of this need for higher-level understanding in modern biology is, of course, the genome, whose sequences will only be understood eventually in terms of higher-level function (genes don't come with functional names attached - nor do neurones!).

I started my life in physiological research as a fully paid-up member of the reductionist club. In the 1960s you couldn't get much more reductionist than to discover ion channels in excitable cells and then to simulate their activity in a bottom-up approach. I did for heart cells what Hodgkin & Huxley did for nerve cells. It is through trying to extend this approach to higher physiological levels that I have come to see the conceptual and computational problems that arise in practice. I have written elsewhere on the impossibility of a completely bottom-up reconstruction of living systems.

Concurrently I also interacted extensively with professional philosophers (of different schools of thought - including those who would support the reductionist agenda). Coming to adopt an integrationist agenda was not an easy road, either scientifically or philosophically. But it is a far richer position. The integrationist does not deny the validity or immense achievements of successful reduction. For some reason (political, social, philosophical?) reductionists seem to need to claim universality for their approach. This book will give them some cause to re-think that position - or so I would hope.

Denis Noble CBE FRS hon FRCP
Professor of Cardiovascular Physiology at Oxford University
Secretary-General, International Union of Physiological Sciences (1993 - 2001)

Philosophical Foundations of Neuroscience
Bennett & Hacker
Blackwell, 2003

Introduction

Philosophical Foundations of Neuroscience presents the fruits of a cooperative project between a neuroscientist and a philosopher. It is concerned with the conceptual foundations of *cognitive* neuroscience - foundations constituted by the structural relationships among the psychological concepts involved in investigations into the neural underpinnings of human cognitive, affective and volitional capacities. Investigating logical relations among concepts is a philosophical task. Guiding that investigation down pathways that will illuminate brain research is a neuroscientific one. Hence our joint venture.

If we are to understand the neural structures and dynamics that make perception, thought, memory, emotion and intentional behaviour possible, clarity about these concepts and categories is essential. Both authors, coming to this investigation from very different directions, found themselves puzzled by, and sometimes uneasy with, the use of psychological concepts in contemporary neuroscience. The puzzlement was often over what might be meant by a given neuroscientist's claims concerning the brain and the mind, or over why a neuroscientist thought that the experiments he had undertaken illuminated the psychological capacity being studied, or over the conceptual presuppositions of the questions asked. The unease was produced by a suspicion that in some cases concepts were misconstrued, or misapplied, or stretched beyond their defining conditions of application. And the more we probed, the more convinced we became that, despite the impressive advances in cognitive neuroscience, not all was well with the general theorizing.

Empirical questions about the nervous system are the province of neuroscience. It is its business to establish matters of fact concerning neural structures and operations. It is the task of *cognitive* neuroscience to explain the neural conditions that make perceptual, cognitive, cogitative, affective and volitional functions possible. Such explanatory theories are confirmed or infirmed by experimental investigations. By contrast, conceptual questions (concerning, for example, the concepts of mind or memory, thought or imagination), the description of the logical relations between concepts (such as between the concepts of perception and sensation, or the concepts of consciousness and self-consciousness), and the examination of the structural relationships between distinct conceptual fields (such as between the psychological and the neural, or the mental and the behavioural) are the proper province of philosophy.

Conceptual questions antecede matters of truth and falsehood. They are questions concerning our *forms of representation*, not questions concerning the truth or falsehood of empirical statements. These forms are presupposed by true (*and* false) scientific statements and by correct (*and* incorrect) scientific theories. They determine not what is empirically true or false, but rather what does and what does not make sense. Hence conceptual questions are not amenable to scientific investigation and experimentation or to scientific theorizing. For the concepts and conceptual relationships in question are *presupposed* by any such investigations and theorizings. Our concern here is not with trade union demarcation lines, but with distinctions between logically different kinds of intellectual inquiry. (Methodological objections to these distinctions are examined in chapter 14.)

Distinguishing conceptual questions from empirical ones is of the first importance. When a conceptual question is confused with a scientific one, it is bound to appear singularly refractory. It seems in such cases as if science should be able to discover the truth of the matter under investigation by theory and experiment - yet it persistently fails to do so. That is not surprising, since conceptual questions are no more amenable to empirical methods of investigation than problems in pure mathematics are solvable by the methods of physics. Furthermore, when empirical problems are addressed without adequate

conceptual clarity, misconceived questions are bound to be raised, and misdirected research is likely to ensue. For any unclarity regarding the relevant concepts will be reflected in corresponding unclarity in the questions, and hence in the design of experiments intended to answer them. And any incoherence in the grasp of the relevant conceptual structure is likely to be manifest in incoherences in the interpretation of the results of experiments.

Cognitive neuroscience operates across the boundary between two fields, neurophysiology and psychology, the respective concepts of which are categorially dissimilar. The logical or conceptual relations between the physiological and the psychological are problematic. Numerous psychological concepts and categories of concepts are difficult to bring into sharp focus. The relations between the mind and the brain, and between the psychological and the behavioural, are bewildering. Puzzlement concerning these concepts and their articulations, and concerning these apparent 'domains' and their relations, has characterized neurophysiology since its inception (we shall begin our investigations in chapter 1 with a historical survey of the early development of neuroscience). In spite of the great advances in neuroscience at the beginning of the twentieth century at the hands of Charles Sherrington, the battery of conceptual questions popularly known as the mind-body or mind-brain problem remained as intractable as ever - as is evident in the flawed Cartesian views embraced by Sherrington and by such of his colleagues and protégés as Edgar Adrian, John Eccles and Wilder Penfield. Brilliant though their work unquestionably was, deep conceptual confusions remained - as we show in chapter 2. Whether the current generation of neuroscientists has successfully overcome the conceptual confusions of earlier generations, or whether it has merely replaced one conceptual entanglement by others, is the subject of our investigation in this book.

One such tangle is evident in the persistent ascription of psychological attributes to the brain. For, while Sherrington and his protégés ascribed psychological attributes to the mind (conceived as a peculiar, perhaps immaterial, substance distinct from the brain), contemporary neuroscientists tend to ascribe the same range of psychological attributes to the brain (commonly, although not uniformly, conceived to be identical with the mind). But the mind, we argue (c 3.10), is neither a substance distinct from the brain nor a substance identical with the brain. And we demonstrate that ascription of psychological attributes to the brain is incoherent (chapter 3). Human beings possess a wide range of psychological powers, which are exercised in the circumstances of life, when we perceive, think and reason, feel emotions, want things, form plans and make decisions. The possession and exercise of such powers define us as the kinds of animals we are. We may enquire into the neural conditions and concomitants for their possession and exercise. This is the task of neuroscience, which is discovering more and more about them. But its discoveries in no way affect the conceptual truth that these powers and their exercise in perception, thought and feeling *are attributes of human beings*, not of their parts - in particular, *not of their brains*. A human being is a psychophysical unity, an animal that can perceive, act intentionally, reason and feel emotions, a language-using animal that is not merely conscious, but also self-conscious - not a brain embedded in the skull of a body. Sherrington, Eccles and Penfield conceived of human beings as animals in whom the mind, which they thought of as the bearer of psychological attributes, is in liaison with the brain. It is no advance over that misconception to suppose that the brain is a bearer of psychological attributes.

Talk of the brain's perceiving, thinking, guessing or believing, or of one hemisphere of the brain's knowing things of which the other hemisphere is ignorant, is widespread among contemporary neuroscientists. This is sometimes defended as being no more than a trivial *façon de parler*. But that is quite mistaken. For the characteristic form of explanation in

contemporary cognitive neuroscience consists in ascribing psychological attributes to the brain and its parts *in order to explain* the possession of psychological attributes and the exercise (and deficiencies in the exercise) of cognitive powers by human beings.

The ascription of psychological - in particular, cognitive and cogitative - attributes to the brain is, we show, also a source of much further confusion. Neuroscience can investigate the neural conditions and concomitants of the acquisition, possession and exercise of sentient powers by animals. It can discover the neural preconditions for the possibility of the exercise of distinctively human powers of thought and reasoning, of articulate memory and imagination, of emotion and volition. This it can do by patient inductive correlation between neural phenomena and the possession and exercise of psychological powers, and between neural damage and deficiencies in normal mental functions. What it *cannot* do is *replace* the wide range of ordinary psychological explanations of human activities in terms of reasons, intentions, purposes, goals, values, rules and conventions by neurological explanations (reductionism is discussed in chapter 13). And it *cannot* explain how an animal perceives or thinks by reference to the brain's, or some part of the brain's, perceiving or thinking. For it makes no sense to ascribe such psychological attributes to anything less than the animal as a whole. It is the animal that perceives, not parts of its brain, and it is human beings who think and reason, not their brains. The brain and its activities *make it possible for us* - not for *it* - to perceive and think, to feel emotions, and to form and pursue projects.

While the initial response of many neuroscientists to the accusation of conceptual confusion is to claim that the ascription of psychological predicates to the brain is a mere *façon de parler*, their reaction to the demonstrable fact that their explanatory theories *non-trivially* ascribe psychological powers to the brain is sometimes to suggest that this error is unavoidable due to the deficiencies of language. We confront this misconception in chapter 14, where we show that the great discoveries of neuroscience *do not require* this misconceived form of explanation - that what has been discovered can readily be described and explained in our existing language. We demonstrate this by reference to the much discussed phenomena resultant upon commissurotomy, described (or, we suggest, misdescribed) by Sperry, Gazzaniga and others (¶ 14.3).

In Part II we investigate the use of concepts of perception, memory, mental imagery, emotion and volition in current neuroscientific theorizing. From case to case we show that conceptual unclarity - failure to give adequate attention to the relevant conceptual structures - has often been the source of theoretical error and the grounds for misguided inferences. It is an error, a *conceptual* error, to suppose that perception is a matter of apprehending an *image* in the mind (Crick, Damasio, Edelman), or the production of a *hypothesis* (Helmholtz, Gregory), or the generation of a *3-D model description* (Marr). It is confused - a *conceptual* confusion - to formulate the binding problem as the problem of combining data of shape, colour and motion to form *the image* of the object perceived (Crick, Kandel, Wurtz). It is wrong, *conceptually* wrong, to suppose that memory is always of the past, or to think that memories can be *stored* in the brain in the form of the strength of synaptic connections (Kandel, Squire, Bennett). And it is mistaken, *conceptually* mistaken, to suppose that the investigation of thirst, hunger and lust is an investigation into the emotions (Rolls) or to think that the function of the emotions is to inform us of our visceral and musculoskeletal state (Damasio).

The initial reaction to such critical remarks may well be indignation and incredulity. How can a flourishing science be fundamentally in error? How could there be unavoidable conceptual confusion in a well-established science? Surely, if there are problematic concepts, they can easily be replaced by others that are unproblematic and that serve the

same explanatory purposes. Such responses betoken a poor understanding of the relation between form of representation and facts represented, and a misunderstanding of the nature of conceptual error. They also betray ignorance of the history of science in general, and of neuroscience in particular.

Science is no more immune to conceptual error and confusion than any other form of intellectual endeavour. The history of science is littered with the debris of theories that were not simply factually mistaken, but conceptually awry. Stahl's theory of combustion, for example, was *conceptually* flawed in ascribing, in certain circumstances, negative weight to phlogiston - an idea that made no sense within its framework of Newtonian physics. Einstein's famous criticisms of the theory of electromagnetic aether (the alleged medium by which light was thought to be propagated) were directed not only at the results of the Michelson-Morley experiment, which had failed to detect any effect of absolute motion, but also at a conceptual confusion concerning relative motion involved in the role ascribed to aether in the explanation of electromagnetic induction. Neuroscience has been no exception - as we show in our historical survey. It is true enough that the subject is now a *flourishing* science. But that does not render it immune to conceptual confusions and entanglements. Newtonian kinematics was a flourishing science, but that did not stop Newton from becoming entangled in conceptual confusions over the intelligibility of action at a distance, or from bafflement (not remedied until Hertz) over the nature of force. So too, Sherrington's towering achievement in explaining the integrative action of synapses in the spinal cord, and thereby eliminating, once and for all, the confused idea of a 'spinal soul', was perfectly compatible with conceptual confusions concerning the 'cerebral soul' or mind and its relation to the brain. Similarly, Penfield's extraordinary achievements in identifying functional localization in the cortex, as well as in developing brilliant neurosurgical techniques, were perfectly compatible with extensive confusions about the relation between the mind and the brain and about the 'highest brain function' (an idea borrowed from Hughlings Jackson).

In short, conceptual entanglement *can* coexist with flourishing science. This may appear puzzling. If the science can flourish despite such conceptual confusions, why should scientists care about them? Hidden reefs do not imply that the seas are not navigable, only that they are dangerous. The moot question is how running on these reefs is manifest. Conceptual confusions may be exhibited in different ways and at different points in the investigation. In some cases, the conceptual unclarity may affect neither the cogency of the questions nor the fruitfulness of the experiments, but only the understanding of the results of the experiments and their theoretical implications. So, for example, Newton embarked on the *Optics* in quest of insight into the character of colour. The research was a permanent contribution to science. But his conclusion that 'colours are sensations in the sensorium' demonstrates failure to achieve the kind of understanding he craved. For, whatever colours are, they are not 'sensations in the sensorium'. So in so far as Newton cared about *understanding* the results of his research, then he had good reason for caring about the conceptual confusions under which he laboured - for they stood in the way of an adequate understanding.

In other cases, however, the conceptual confusion does not so happily bracket the empirical research. Misguided questions may well render research futile (examples will be examined in relation to mental imagery (¶ 6.3.1) and voluntary movement (¶ 8.2)). Rather differently, misconstrual of concepts and conceptual structures will sometimes produce research that is by no means futile, but that fails to show what it was designed to show (examples will be discussed in relation to memory (¶¶ 5.2.1-5.2.2) and to emotions and appetites (¶ 7.1)). In such cases, the science may not be flourishing quite as much as it appears to be. It requires conceptual investigation to locate the problems and to eliminate

them.

Are these conceptual confusions *unavoidable*? Not at all. The whole point of writing this book is to show how to avoid them. But, of course, they cannot be avoided while leaving everything else intact. They *can be* avoided - but if they are, then certain kinds of questions will no longer be asked, since they will be recognized as resting on a misunderstanding. As Hertz put it in the wonderful introduction to his *Principles of Mechanics*: 'When these painful contradictions are removed, . . . our minds, no longer vexed, will cease to ask illegitimate questions.' Equally, certain kinds of inferences will no longer be drawn from a given body of empirical research, since it will be realized to have little or no bearing on the matter which it was meant to illuminate, even though it may bear on *something else*.

If there are problematic concepts, can they not be replaced by others that serve the same explanatory function? A scientist is always free to introduce new concepts if he finds existing ones inadequate or insufficiently refined. But our concern in this book is not with the use of new technical concepts. We are concerned with the misuse of old, nontechnical concepts - concepts of mind and body, thought and imagination, sensation and perception, knowledge and memory, voluntary movement, and consciousness and self-consciousness. There is nothing inadequate about these concepts relative to the purposes they serve. There is no reason for thinking that they need to be replaced in the contexts that are of concern to us. What are problematic are neuroscientists' misconstruals of them and the misunderstandings consequently engendered. These are remediable by a correct account of the logico-grammatical character of the concepts in question. And this is what we have tried to supply.

Granted that neuroscientists may not be using these common or garden concepts the way the man in the street does, with what right can philosophy claim to correct them? How can philosophy so confidently judge the clarity and coherence of concepts as deployed by competent scientists? How can philosophy be in a position to claim that certain assertions made by sophisticated neuroscientists *make no sense*? We shall resolve such methodological qualms in the following pages. But some initial clarification here may remove some doubts. What truth and falsity is to science, sense and nonsense is to philosophy. Observational and theoretical error result in falsehood; conceptual error results in lack of sense. How can one investigate the bounds of sense? Only by examining the use of words. Nonsense is often generated when an expression is used contrary to the rules for its use. The expression in question may be an ordinary, non-technical expression, in which case the rules for its use can be elicited from its standard employment and received explanations of its meaning. Or it may be a technical term of art, in which case the rules for its use must be elicited from the theorist's introduction of the term and the explanations he offers of its stipulated use. Both kinds of term can be misused, and when they are, nonsense ensues - a form of words that is excluded from the language. For either nothing has been stipulated as to what the term means in the aberrant context in question, or this form of words is actually excluded by a rule specifying that there is no such thing as . . . (e.g. that there is no such thing as 'east of the North Pole'), that this is a form of words that has no use. Nonsense is also commonly generated when an existing expression is given a new, perhaps technical or quasi-technical use, and the new use is inadvertently crossed with the old - for example, inferences are drawn from propositions containing the new term which could only licitly be drawn from the use of the old one. It is the task of the conceptual critic to identify such transgressions of the bounds of sense. It is, of course, not enough to show that a certain scientist has used a term contrary to its ordinary use - for he may well be using the term in a new sense. The critic must show that the scientist intends using the term in its customary sense and has not done so, or that he intends using it in a new sense but has inadvertently crossed the new sense with the old. The wayward scientist

should, whenever possible, be condemned out of his own mouth. We address methodological qualms in detail both in chapter 3, section 3, and in chapter 14.

The final misconception against which we wish to warn is the idea that our reflections are unremittingly negative. All we are concerned with, it might be thought, is criticizing. Our work may appear at a superficial glance to be no more than a destructive undertaking that promises neither assistance nor a new way forward. Worse, it may even appear to be engineering a confrontation between philosophy and cognitive neuroscience. *Nothing could be further from the truth.*

We have written this book in admiration for the achievements of twentieth-century neuroscience, and out of a desire to assist the subject. But the *only* ways in which a conceptual investigation can assist an empirical subject are by identifying conceptual error (if it obtains) and by providing a map that will help prevent empirical researchers from wandering off the high roads of sense. Each of our investigations has two aspects to it. On the one hand, we have tried to identify conceptual problems and entanglements in important current theories of perception, memory, imagination, emotion and volition. Moreover, we argue that much contemporary writing on the nature of consciousness and self-consciousness is bedevilled by conceptual difficulties. This aspect of our investigations is indeed negative and critical. On the other hand, we have endeavoured, from case to case, to provide a perspicuous representation of the conceptual field of each of the problematic concepts. This is a constructive endeavour. We hope that these conceptual overviews will assist neuroscientists in their reflections antecedent to the design of their experiments. However, it cannot be the task of a conceptual investigation to propose empirical hypotheses that might solve the empirical problems faced by scientists. To complain that a philosophical investigation into cognitive neuroscience has not contributed a new neuroscientific theory is like complaining to a mathematician that a new theorem he has proved is not a new physical theory.

It is improbable that many neuroscientists will wish to read a 450-page conceptual investigation from cover to cover. Consequently, we have tried to make our chapters on select psychological concepts as self-contained as possible. We have done this in the hope that the book will serve as a conceptual reference work for cognitive neuroscientists who wish to check the contour lines of the psychological concept relevant to their investigation. This has, of course, meant that there is a degree of repetition between certain chapters. This is, we hope, warranted by the objective.

The chapters of the book are accompanied by italicized marginalia indicating the subject under discussion in the correlated paragraph or paragraphs. The purpose of this is to facilitate surveyability, to make it easier to follow the steps in the argument, and to assist in locating arguments. The section headings in the table of contents are accompanied by the italicized names of neuroscientists (and occasionally philosophers who concern themselves with neuroscientific and cognitive scientific matters) whose theories are either discussed in some detail or mentioned *en passant* in the course of the chapter. This will, we hope, help the reader to locate the themes and discussions that are of specific interest with ease.