

Corballis's thoughts on the split-brain sharpened up chapter 6. In addition the Amnesia chapter benefits greatly from my continuing collaboration with Nicola Hunkin. However, as always, the final responsibility for content rests entirely with me. I would also like to thank the numerous students who have taken the course from which this book evolved – their comments on the material have been extremely valuable in shaping the book. I am deeply grateful to Ella Squires for all her editorial work and to Frances for staying up half the night figuring out how Endnote prints out 670 references. Finally I hope that readers of this book will spare a thought for the subjects of cognitive neuropsychology. I have always been impressed by how readily brain-injured people offer themselves as subjects in full knowledge that the experiments they take part in will not help them personally in any way. I think this type of contribution, often within the context of adverse or even tragic circumstances, is something both to be acknowledged and admired.

Alan Parkin
Brighton June 1995

1

Cognitive Neuropsychology as a Science

We have our sandwiches on a hill outside Weston with a vast view over Somerset. She wants to say, 'What a grand view' but her words are going too. 'Oh,' she exclaims. 'What a bit lot of About.' There are sheep in the field. 'I know what they are,' she says, 'but I don't know what they are called.' Thus Wittgenstein is routed by my mother.

Alan Bennett (*Writing Home*, 1994)

Cognitive neuropsychology is founded on the principle that one of the easiest ways to understand how a system works is to observe what happens when it goes wrong. By carefully recording and analysing the various errors that can occur in a system you can build up a picture of how its components are organized and how they operate. Before considering how this approach can be applied to the complex functions of the human brain consider how it might be applied to the relatively simple task of understanding how a television works. In an intractable television you are not aware of the various components that contribute to the output. However, if you observed enough faulty televisions you might begin to develop ideas about how they work. You would, for example, notice that televisions can lose sound but not vision and *vice versa*. This is an important observation because it tells us that these two components are independent because they can each function when the

From Parkin (1996)

other is not working. Another observation you might make is that the picture can lose its colour. However, the opposite observation, colour without picture, never occurs. This means that the systems governing picture generation and colour are not independent because of the production of colour, seems critically dependent on the presence of picture. Nonetheless the continual observation of colour loss with a normal black and white picture suggests that the production of colour involves a separate component of the television.

The above example is based on the assumption that the observer has no prior knowledge of televisions or any other concepts that might help understanding of how the television works. In practice many observers would have both prior knowledge and intuitions about what to look for. Thus their observations would not be naive but guided by various assumptions such as an existing knowledge about the separation of audio and visual channels and the fact that televisions have colour controls that enable the colours to be manipulated independent of picture quality. In short your observations of faulty televisions would be theory-driven in that you use pre-existing knowledge to both guide and interpret your observations.

The aim of cognitive neuropsychology is to provide a greater understanding of how the brain carries out mental operations based on the observation of people who have developed specific deficits as the result of brain damage. Cognitive neuropsychology is heavily dependent on careful observation of the behaviour exhibited by brain damaged people but it is also guided by the theoretical framework provided by cognitive psychology.

Cognitive Psychology and Modularity – A Very Brief Overview

Before 1960 the dominant approach to understanding human mental processes was provided by behaviourists such as Skinner. Behaviourism takes the view that behaviour cannot be explained by making any appeal to internal structures or processes within the brain. Instead behaviourists sought to account for all human behaviour in terms of the relations between inputs (stimuli) and outputs (responses). In 1959 the linguist Chomsky wrote a powerful critique of the behaviourist approach and, in essence, came to the conclusion that the nature of mental life could not be investigated effectively without some theory about how the structures and processes of the brain were organized and the principles by which they operated (Chomsky, 1959).

Experimental psychologists face a challenge encountered in no other

science except perhaps sub-atomic physics. There is little disagreement that mental processes take place in the brain and, from microscopic examination, we can see that the brain comprises nerve cells and fibres which interact with one another within a complex network. However, knowing that the mental processes that we seek to understand are a product of the interaction between nerve cells does not help us comprehend how mental processes take place. In short the explanation of psychological processes cannot, at present, be reduced to an understanding of how the brain works at a physiological level. Because of this psychologists are forced to use analogy or metaphor in explanation; that is they attempt to explain the workings of the mind in terms of something else that we do understand.

Cognitive psychology can be defined as the branch of psychology which attempts to provide a scientific explanation of how the brain carries out complex mental functions such as vision, memory, language and thinking. Cognitive psychology arose at a time when computers were beginning to make a major impact on science and it was perhaps natural that cognitive psychologists should draw an analogy between computers and the human brain. The computer analogy was used frequently to draw up a model of the brain in which mental activity was characterized in terms of the flow of information between different stores.

Perhaps the most well known of these models was the model of memory put forward by Atkinson and Shiffrin (1968). In this model (see figure 1.1) the human memory system comprises three stores: New information (external stimulation) first enters sensory store where, after initial processing it passes to short-term store. While in the short-term store information can be actively processed guided by the intentions of the individual. Thus new information can be stored in the permanent repository of memory known as long-term store or be discarded. The system also needs to be able to retrieve stored information at a later date and for this reason the model embodies a two-way flow of information between short- and long-term store. The interaction envisaged between short- and long-term store very much resembles the interactions that occur between the central processing unit of a computer and its permanent database. It is important to stress that the computer analogy attempted to emphasise the similarity between the organization of the human mind and that of a computer. It was evident from the outset that the precise manner in which computers operated, that is the binary coding system employed, was not an appropriate basis for understanding cognitive processes themselves.

The computer analogy, in modified form, continued to dominate cognitive psychology for the next 20 years. A vast range of models were




Figure 1.1 The multistore model of memory.

produced each attempting to explain some form of mental activity in terms of a series of processing stages intervening between an input and an output. Despite their widespread use, information processing accounts of mental abilities have not been without criticism. Indeed this approach has been characterized by some commentators as nothing more than 'boxology' – defined by Stuart Sutherland as 'The construction and ostentatious display of meaningless flow charts by psychologists as a substitute for thought' (Sutherland, 1989, p. 58). These criticisms arose because cognitive psychologists became increasingly happy to specify mental processes in terms of information processing systems without offering any explanation of how any of the supposed stages worked.

Modularity

More recently the cognitive approach has been strongly influenced by the concept of modularity. This has gone some way to addressing the accusation of boxology because properties are assigned to the processing elements, or modules, that comprise any system. The term modularity derives from computer programming and refers to the principle that it is important to make the different components of the program as independent from one another as possible. Modularity will then make any debugging operations much easier to carry out because the nature of the fault will tend to be an accurate indicator of which program module is at fault. An important feature of a modular system, therefore, is that the components are autonomous in that they remain functionally intact when other components of the same system become corrupted.

Marr (1976) proposed that the brain might also have a modular organization because modularity imparts considerable advantage to any complex system that is attempting to evolve. The reason for this is that

a modular system is easier to correct or improve because changes can be made to specific parts of a system without the need for parallel changes elsewhere. In vision, for example, improvements might occur in modules concerned with extracting information about depth without the need to improve the modules integrating the resulting depth information into the final percept.

The idea of modularity has been extensively developed by Fodor (1983) in his book *The Modularity of Mind*. Fodor's major contribution was to specify a number of properties that he assumed modules had.

- 1 **Informational encapsulation:** Modules carry out their operations in isolation from what is going on elsewhere. These operations are not amenable to 'cognitive penetration'. Fodor illustrates this by pointing out that we are unable to overcome visual illusions even though we know full well that they are illusions.
- 2 **Domain specificity:** Each module can only process one type of input. Modules therefore deal in only one source of information.
- 3 **Mandatory:** Each module operates in an all-or-none fashion. Once activated it will carry out the entire processing operation for which it is responsible.
- 4 **Innate:** The modules of the cognitive system are innate and not acquired through development. This is a highly controversial claim which we can highlight by considering this statement from the neurologist Geschwind:

the overwhelming majority of humans who have ever lived have been illiterate, and even today I believe it is the case that a very large percentage . . . of the world's population have never had the opportunity to learn to read. Most of us come from families that four generations ago did not possess the ability to read

We can interpret the above fact in two basic ways. First we could argue that the short history of reading in humans means that reading is carried out by modules that pre-existed to carry out other functions. Support for this view can be found in the study of developmental dyslexia where it can be found that dyslexia often occurs in association with deficits in more basic abilities such as sequencing and left-right discrimination. Alternatively we could argue that the innate assumption is wrong and that new modules can evolve within a developing brain. Arguments supporting this view would be the existence of highly specific acquired dyslexias in which the process of reading seems selectively disrupted (see below and chapter 7).

For our present purpose it is fortunate that the two aspects of modularity most important to cognitive neuropsychology, infortification encapsulation and domain specificity, are the least controversial. It is important to note, however, that Fodor's definition of modules is not strictly adhered to within cognitive neuropsychology in that modules are often defined as accepting different types of input whereas Fodor considers modules to be input-specific.

Identifying modules cannot, however, be the only goal of cognitive psychology. For a proper explanation we also need to understand how the processes operating within hypothesized modules are carried out. Fodor's view in this area is highly controversial. He accepts that some scientific explanation of how modules operate can be achieved. However, only more basic cognitive processes are held to be modularized and more complex processes such as thinking and deciding are thought to be highly interactive with one another. For Fodor this level of interactivity precludes effective scientific investigation because it is not possible to offer any precise theory about how more complex activities might be carried out.

In general cognitive psychology has not gone along with Fodor's viewpoint and theories abound to explain all manner of human cognitive processes. A good example of a cognitive model embodying both modular and non-modular components is the working memory model of Baddeley and Hitch (Baddeley and Hitch, 1977; Baddeley, 1990). This model is a development of Atkinson and Shiffrin's concept of short-term store (see figure 1.2). It comprises a central executive and two auxiliary components, an articulatory loop and a visuo-spatial scratch pad. The central executive represents the locus of control of the cognitive system. It is responsible for determining how inputs should be dealt with and for retrieving information from long-term storage. It is also multimodal and not domain-specific, able to combine information from both different sensory modalities and different types of stimulus input.

In contrast the two auxiliary components have very much the character of modules. The articulatory loop is considered a limited capacity system specifically dedicated to the processing of phonological information. Similarly the visuo-spatial scratch pad is solely dedicated to the representation of a limited amount of information in terms of its visuo-spatial characteristics. Working memory is just one example of a large number of models that currently exist to explain human cognition at various levels of complexity. The scale of this enterprise is too vast to even attempt a summary here. Instead the reader will have to be content in knowing that relevant theories will be explained at appropriate points in the text.




Figure 1.2 The multistore model of memory revised to incorporate the concept of working memory.

Parallel Distributed Processing

Most recently the study of cognitive processes has undergone an almost revolutionary change with the advent of parallel distributed processing (PDP) or, as it is more often known, connectionism (McClelland and Rumelhart, 1986). PDP models are implemented as computer programs known as networks. Unlike the models we have considered so far, in which each component of the model is explicitly stated (e.g. central executive, articulatory loop), connectionist models learn by themselves and set up their own representation of any given set of information.

A network comprises a series of nodes or units linked to each other by connections which can be either inhibitory or excitatory in nature. In addition the strength of any connection can be modified by assigning a particular weight. It is an important feature of these models that information about any concept is distributed across many nodes and not the property of a single node. Different pieces of information thus correspond to different patterns of activity within the same network (see figure 1.3).

PDP networks are able to learn by using rules or algorithms to change interconnections between the nodes of the network. One method of learning is the back-propagation algorithm which is a means of making an output conform to a desired state. For any input, the system is informed about what the resulting output should be. This information is compared with the actual state of the network when the stimulus is present and an adjustment is made. This process is incremental and

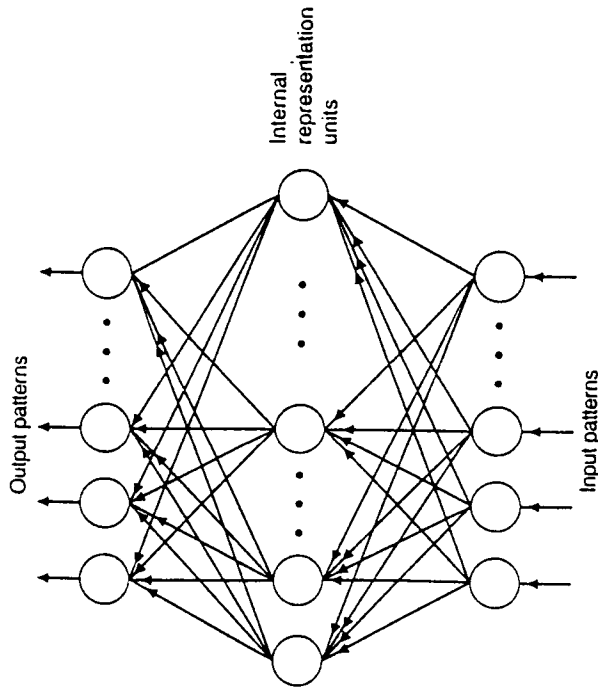


Figure 1.3 A prototype connectionist network comprising a layer of input units, a layer of internal 'hidden units', and a layer of output units. The network learns to associate a particular input pattern with an output pattern. This is achieved via the hidden units which recode input patterns so that they can be more easily differentiated from each other. Thus in a network learning to read the inputs CAT and CAR would be distinguished by different patterns of hidden unit activity. From McClelland and Rumelhart (1986), *Parallel Distributed Processing*. Reproduced by permission of MIT Press.

networks need many exposures to a stimulus before they learn the correct output. Methods like back propagation allow a network to represent information. But because human learning is not characterized by having pre-existing knowledge of the desired outcome, algorithms like back-propagation have obvious limitations. However, proponents of PDP models argue that it is the properties of the established network that are of interest to psychologists rather than the manner in which they become established.

Because PDP models are active learning networks cognitive neuropsychologists have explored the effects of damaging them in a way that might be considered analogous to the effects of a brain lesion on a real neural network. Thus a network that has learnt to read a set of words might be deprived of 30 per cent of its nodes and the resultant effects on reading the target word set re-observed. Later in the book we will

examine several instances where the results of degraded PDP networks produce impaired response patterns similar to a particular pattern of deficit caused by brain damage and consider the implications they have for developing theories about real brain function.

Neurological Specificity and Double Dissociation

Within modern cognitive neuropsychology the assumption that different mental functions occupy different brain regions is known as neurological specificity. For some the assumption of neurological specificity is purely abstract in that no attempt is made to describe the brain regions underlying mental functions that appear dissociated. However, the majority of cognitive neuropsychologists do provide neurological data on the brain injuries suffered by their patients and in a number of cases these data may be important for the arguments being made.

The assumption of neurological specificity, combined with acceptance of the modularity principle (see above) underlies a good deal of the work taking place in cognitive neuropsychology. Central to this is the method of double dissociation. In a neuropsychological investigation a psychologist wishes to specify the independent existence of the various modules comprising a system. To do this it is necessary to have demonstrations of each module working in the absence of the other. Thus, if we believe that a process has two modules, A and B, we need evidence of A operating in the absence of B and vice versa.

With the television example we considered earlier we are left in no doubt about dissociability because we can easily observe no picture with sound or vice versa. This constitutes what is known as a classic double dissociation (Shallice, 1988) in that performance on one function is deficient while the other is normal. In cognitive neuropsychology double dissociations are usually claimed on the basis of differing degrees of relative impairment. Thus Patient 1 might be better on a task tapping module A and Patient 2 better on a task reflecting module B. However, in neither case do the patients perform normally on any of the tasks. This is a difficult issue for neuropsychologists because dissociations are commonly claimed on this sort of evidence but they could be artefactual. It may be that uncontrolled differences between the two patients underlies their differential impairment rather than real qualitative differences in the nature of their cognitive deficit. Classic double dissociation is much harder to explain in this way but, unfortunately, instances of this are hard to come by.

By observing patterns of dissociation the cognitive neuropsychologist can develop an account of the modular organization of a particular

cognitive ability. However, it must be stressed that using the logic of double dissociation to make increasingly fine distinctions between processing elements in the brain is not a sufficient goal in cognitive neuropsychology. This point was addressed by Roediger (1990) when he noticed a report by Gazzaniga (1989):

One patient, for example, was unable to name the colour of fruits that were red. Thus, in response to the question, 'What is the colour of an apple?' there was a chance performance. Yet, the same patient was easily able to name the colour of a banana. Additionally, the same patient was easily able to name the colour of fire engines and school houses. It is conceivable that her incapacity to identify the colour of red fruits was due to crucial fibres being interrupted within her left hemisphere that connected together the appropriate information sources. (p. 949)

Roediger expresses concern at the rapidity with which modular dissociations can grow and the possibility that modular theorists might discover hundreds of similar fine-grained distinctions in the absence of any organizing principles. He draws a comparison with Gestalt psychology in which, at one time, there were 114 laws of perception. However, as greater understanding developed many of these laws became subsumed under more general headings. It is not enough, therefore, to argue for increasingly complex patterns of modular organization. There must, in addition, be a theoretical framework within which modular accounts of the cognitive system can be evaluated.

Phonological and Surface Dyslexia – An Example of Investigation in Cognitive Neuropsychology

Figure 1.4 presents a rather simplified idea of how the modules responsible for reading a word aloud might be organized in the brain. The word is first identified at module A but from here it can go along a pathway to module B or module C, before a pronunciation of the output is achieved at module D. This model has not been determined by cognitive neuropsychology but by pre-existing theory and observation based on normal readers. Fluent readers of English have the ability to pronounce both real words, for example, MINT, ROAD, and nonwords, for example, SINT, FOAD. Another ability is to deal with the large amount of irregularity in English pronunciation. Estimates vary but around 20 per cent of English words contain some irregularity in that if you tried to pronounce them by rule you would make a mistake, examples of these irregular words are PINT and BROAD.

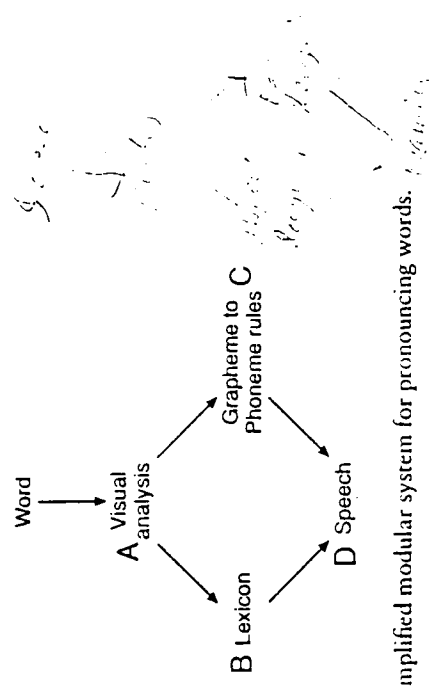


Figure 1.4 A much simplified modular system for pronouncing words.

The demands of being able to read both nonwords and irregular words mean that there must be at least two systems available for converting the printed word into sound. First we must have a set of rules, usually called letter-to-sound or grapheme-to-phoneme correspondence rules, which enable us to produce a pronunciation for nonwords such as SINT and FOAD. However, because they are based on the commonest correspondences between letters and sounds, these rules would also deal effectively with regularly pronounced words such as MINT and ROAD. Where these rules would come unstuck is with irregular words because here the incorrect 'regularized' pronunciation would be produced (e.g. PINT →; BROAD → BRODE). Because of this difficulty it is necessary to propose that we have an additional word-specific or lexical basis for specifying the pronunciation of individual words and it is only this system that allows the correct pronunciation of irregular words.

It is generally thought that these two paths to pronunciation operate in parallel. For regular words the correct response can be achieved via either path and nonwords can only be pronounced using grapheme to phoneme rules. In contrast irregular words will only be correctly pronounced if the lexical path is used. This model is based on both intuition and experimental evidence. We know, for example, that it takes longer to start pronouncing irregular words compared with regular words under speeded conditions and that when doing this subjects will tend to 'regularize' irregular words (Parkin, 1982a). However, perhaps our strongest motive for believing in two routes to pronunciation comes from our observations of acquired dyslexia – the sudden loss of reading ability due to a brain lesion. This evidence also provides us with a good introduction to the logic underlying much current neuropsychological investigation.

Norman's deficit does not, on its own, allow us to conclude the separable existence of lexical and grapheme-to-phoneme correspondence rules. One could, for example, argue that there is a single process that converts all types of print into sound and that, for some reason, it works less efficiently for unfamiliar patterns of letters – hence Norman's greater difficulty with nonwords. In order to argue for two routes, therefore, we need evidence for a double dissociation – a pattern of disturbed reading that can be explained in terms of a defective lexical route.

TOB developed language problems very slowly due to an evolving dementing illness centred on his left temporal lobe. He gradually lost vocabulary to the point where he could only name very common objects. Thus although he was able to drive his car perfectly well he could not name any of the parts. Along with the problem of naming things TOB also developed a reading problem. He found it hard to define what words meant even when his word-finding problems were by-passed by asking him to match words to pictures. In an experiment TOB was asked to read different types of words and nonwords aloud. TOB had little difficulty reading aloud regular words and nonwords, but had great difficulty reading aloud irregular words. This problem was emphasized by TOB's inability to read even the simplest irregular words (e.g. *chord*, *ache*) despite fluent pronunciation of very uncommon regular words such as *indiscoverable*, *chitterling*, and *buckaback*.

The reading problems shown by Norman and TOB are very different. Norman reads most words well but cannot read anything that is not a word. TOB can read words without difficulty as long as the words obey English grapheme-to-phoneme correspondence rules; he also has no problem reading nonwords. These two patterns of reading impairment have been observed many times and are described as *phonological dyslexia* and *surface dyslexia* respectively.

In phonological dyslexia there is a primary problem in reading nonwords. In terms of our simple modular account of reading this suggests that Norman no longer possesses any ability to convert graphemes to phonemes. However, Norman can read most words thus indicating that he does have a lexical basis for reading words. This evidence suggests that the simple modular account put forward in figure 1.4 may well be correct because this would predict the pattern we have observed. However, Norman's data alone do not clinch the argument because we need evidence for a double dissociation. This is, of course, provided by TOB because he reads regular words and nonwords equally effectively but has gross difficulty reading any irregular word accurately. TOB's reading errors thus suggest that he is reliant on grapheme-to-




Figure 1.5 Norman's ability to read common and uncommon words differing in imageability (data courtesy of Elaine Funnell).

There are many forms of acquired dyslexia and we will examine some of them in detail at a later point. For now we will consider only two types of reading problem exemplified in these case histories:

Norman suffered a stroke which, due to damage in his left cerebral hemisphere, has resulted in considerable language difficulties even though his ability to communicate verbally is generally good. Norman's problem is that he can no longer write, or read books. On tests of reading simple familiar words (e.g. *dog*, *man*, *tree*) he has no difficulty and, although he finds it more difficult, he can also read less meaningful words such as *it*, *when* and *how*. However, not all words can be read. We can divide words very roughly into those that enable us to form an image easily (e.g. *door*, *bird*) and those for which an image does not spring easily to mind (gist, desire). We can also divide words into common and uncommon on the basis of how frequently we see or hear them. Norman was given a simple experiment in which he had to read words that differed along these two dimensions and the results are shown in figure 1.5. He does reasonably well in three of the categories but has obvious difficulty with uncommon words low in imagery. It is also notable that when he fails to read a word he either says nothing or offers another more common word instead (syllable → label).

In a second experiment Norman was asked to read nonwords. This he could not do at all and when he did respond he produced a real word that resembled the nonword (e.g. *vich* → *wirch*; *fage* → *age*).

phoneme correspondence rules because he has no access to lexically-based pronunciation.

The above evidence thus provides us with a good example of double dissociation and supports the modular theory outlined in figure 1.4. However, it should be noted that the model does not fully account for the available facts. Why, for example, does Norman read high-imagery words better than low-imagery words? To explain this we must refine our ideas of how the lexical pathway operates. In addition our investigation has not told us anything about how the various modules work. How does the grapheme-to-phoneme module work? What rules does it use? This can only be achieved by more detailed exploration.

Another important point to notice is that the modular theory examined in this example pre-existed on the basis of evidence from normal subjects. In this instance, therefore, cognitive neuropsychology's role is confirmatory in that it serves to reinforce theories we have about the cognitive system. This is a common role for cognitive neuropsychology and we will see many instances of this throughout this book. However, it must be emphasized that cognitive neuropsychology can do more than this by providing insights into brain organization that are not apparent from studying normal people. Blindsight, the topic of the next chapter, is one example, and the work conducted on language (chapters 7 and 8) also provides instances of cognitive neuropsychology leading the development of theory.

Methods in Cognitive Neuropsychology – Single Cases or Groups?

In the preceding example we compared data from two people and came to a particular theoretical conclusion about the organisation of the reading system. However, we noted that both phonological and surface dyslexia had been observed in other people. To some scientists this latter fact would come as some relief because they would be uneasy about accepting a particular theoretical position on the basis of observing just two individuals with different deficits. However, within cognitive neuropsychology, this is not a universal position and many believe that only single cases should be studied.

To understand the arguments in this section let us first imagine an experiment in which there are two conditions A and B. Our hypothesis is that A will produce better performance than B. If we were to then test only one randomly selected subject in condition A and one in condition B would we be in a position to argue about the relative effectiveness of the two conditions? No. Because experimental subjects naturally vary in

their ability to do any experimental task and, with only one subject in each condition, there is an obvious danger that any apparent difference arises from sampling error. To control for this problem we test a substantial number of subjects in each condition and use inferential statistics to check whether any group difference we observe is significant given the overall variability in task performance.

At first sight it does not seem a difficult step to apply the same logic to experiments in cognitive neuropsychology. However, a number of authors have argued that extension of the group study logic to brain damaged populations is invalid (e.g. Badercker and Caramazza, 1985; Caramazza, 1984; Caramazza, 1986; Caramazza and Badercker, 1989; Caramazza and Badercker, 1991; McCloskey, 1993; McCloskey and Caramazza, 1988; Sokol, McCloskey, Cohen and Alimnosa, 1991). The basis of the argument is that group studies on normal subjects rest on the reasonable assumption that the cognitive processes under investigation are homogeneous – it being unlikely that a basic ability such as perception or memory would be carried out in a qualitatively different manner across individuals. In a normal group study, therefore, all one has to do is control for variability in a function whose neural representation can be considered similar across individuals.

Turning to brain damaged subjects one can still assume homogeneity of pre-morbid cognitive function but what can be assumed about the patients' current cognitive abilities? Opponents of group studies argue that very little can be gained from group studies because lesions cause such variation in damage that it is impossible to be sure that any two patients have the same form of cognitive deficit. This criticism can certainly be applied to a good number of neuropsychological studies where groups of subjects have been loosely defined. Many studies, for example, have simply compared the effects of left and right hemisphere damage – a crude classification which, without doubt, results in the grouping together of patients with very different patterns of impairment.

Supporters of group studies (e.g. Robertson, Knight, Rafal and Shimamura, 1993; Zurif, Swinney and Fodor, 1991; Zurif, Gardner and Brownell, 1989) have countered these arguments by proposing various ways in which the heterogeneity problem can be overcome. One approach is to group patients together on the basis of syndromes. A syndrome can be defined as a set of symptoms which co-occur with sufficient regularity to suggest that they reflect a single underlying deficit. A large number of syndromes have been described in neuropsychology and many experimental studies have been carried out using syndromes as a basis for specifying groups of subjects.

The success of the syndrome-based approach to group classification depends critically on how the syndrome is defined. If the criteria are too

loose and patients with qualitatively different deficits are grouped together, conclusions based on averaging their performance will be misleading. Opponents of the group-based approach have argued that syndrome-based groups often show dramatic individual variation which undermines treating them as a single group. Supporters of the group approach argue that syndromes have proved very useful in medicine even though not all patients show all the defining features of a given syndrome.

An alternative argument favouring group studies can be termed the significance implies homogeneity rule. Put simply this rule states that if a group of subjects exhibits significant heterogeneity then they will not be capable of generating statistically significant group differences. However, this rule has also been attacked because of mistaken assumptions it makes about the nature of statistical significance. Intuitively one considers that demonstration of a significant effect in a group of subjects means that the majority of subjects show the effect. However, analysis of statistical procedures shows that a significant effect can sometimes be obtained when less than 50 per cent of subjects show the effect (providing, of course, that the remaining subjects show no effect rather than an effect in the opposite direction).

A further attempt at strengthening the rationale of group studies is to use neurological data. As we have seen cognitive neuropsychology assumes a degree of neurological specificity. So, on the assumption that patients all had similar neurological specificity prior to their brain injury, it is valid to assume that all those with damage to the same part of the brain also have similar functional problems. An optimal arrangement, therefore, would be to construct experimental groups each defined as having a discrete lesion in a different part of the brain and then examine their patterns of cognitive impairment. Although in theory possible, this type of experiment is extremely uncommon for a number of reasons. First many of the dissociations that have been observed arise from damage to the same general brain area and neuroradiology does not enable us to map structure to function in fine enough detail. Moreover, our assumption that neurological specificity is consistent across subjects may not hold in all circumstances. For example, people show great individual variation in how the blood supply to their brain is organized. Thus, two patients might be considered similar in terms of suffering damage to the same brain artery but the consequent damage might be very different.

The alternative approach, sometimes called 'ultra' or 'radical' cognitive neuropsychology, is to consider only data from single-case studies. This approach overcomes problems concerning heterogeneity because any study only involves one subject but it opens up a different range of

problems. The most obvious of these is the problem of generalization. How can we be sure that the conclusions derived from the study of one individual reflect the population as a whole? Supporters of the radical approach employ the following logic:

The radical single-patient approach avoids the population heterogeneity problem in a simple way: The concept of a patient population is not invoked in linking data to theory... In a radical single-patient study, the results are first used to draw conclusions about the patient's previously normal cognitive mechanisms and functional damage to these mechanisms. Then, on the assumption that the cognitive mechanisms are shared by some normal population, the conclusions about the patients' premorbid cognitive mechanisms are generalised directly to that normal population. (McCloskey, 1993, p. 728)

There are a number of points to notice about this methodology. Most importantly there is the assumption that the patient's pre-morbid cognitive functions were similar to those of the normal population. A problem here is that the patient may have been atypical from the start (something Caramazza has dubbed 'The Martian among us problem'). While this might not be a frequent problem there is, in my opinion, at least one clear instance where data from an atypical patient has been central to a neuropsychological theory:

Speaking and spelling disorders tend to go together and this supported a traditional view in which spelling evolved on a pre-existing speech production system. A key piece of evidence against a theory of this kind would be someone who could write but not speak. In considering the issue of a separate spelling system (Ellis and Young, 1988) laid considerable emphasis on data from a patient known as EB (Levine, Calvanio and Popovics, 1982). EB continued to be able to read and write despite total mutism and an inability to appreciate even inner speech (the natural tendency to hear words in your head as you are reading). This seems compelling support for a modular separation of speaking and spelling systems but it is notable that Levine et al. considered their patient to be somewhat special:

Another possible factor, perhaps related to his talent for mental imagery, was his unusual premorbid reading strategy. While a young man (age 20) ... he took a 16 week course in speed reading, which he practised hard and used ever since ... always read fast. In this course he was taught to avoid saying the printed words to himself and to focus instead on their meanings. He wrote 'Learning to and practising speed reading required super concentration on sense of words within idea contexts. Individual words and the phonetics of words was suppressed in order to be proficient

at high speed reading. Now I am slow and ponderous but still only 'see' words as ideas not sounds' (p. 406).

EB was 54 when he suffered his stroke so he had been reading in his abnormal fashion for over 30 years. It is not unreasonable to suppose, therefore, that his reading and related skills may have been organized abnormally. More specifically, his background suggests that he may have developed a reading and writing strategy that did not rely on phonology at all thus making preserved writing in the absence of speech a distinct possibility. This makes it difficult to assume that EB is a patient whose cognitive system was organized normally prior to brain damage. One could, of course, argue that EB's achievements are an indication that speaking and spelling are separable but the counter-argument here is that cognitive neuropsychology must explain what is typical not what may be possible.¹

Fortunately, other less controversial data now supports the separate existence of a spelling system (see chapter 8). Nonetheless one can see that if EB had been the only case of preserved spelling in the absence of speech, we might well have been misled. This emphasizes another difficulty for single case approaches – replication. In science a theory will not be accepted if the findings that generated it are not capable of replication. However, the radical approach seemingly fails to meet this basic tenet of scientific investigation. Imagine a situation in which one person suffers a brain lesion that results in a selective inability to read poetry as opposed to other written material and a second shows the reverse deficit. Within radical cognitive neuropsychology this would necessitate a distinction being made between the processing of poetry and other written material. More importantly the conclusions derived from these single cases would hold even if that pattern of impairment were never observed again. How could the single-case methodology ensure that the above dissociation was a genuine reflection of normal cognitive processes?

Supporters of the radical approach argue that testing flaws can be avoided by using extensive testing regimes in which converging evidence for a conclusion is derived from performance across a number of tasks. A more difficult problem is to combat inclusion of a patient whose data are atypical from the outset (see above). Care can be taken to avoid any patient who may have evidence of abnormal brain development, and the appearance of findings that are inconsistent with the developing theoret-

ical framework could also indicate that the data are mistaken in some way. A danger here, however, is that the inclusion or exclusion of cases becomes increasingly guided by theory rather than objective fact – a fact, incidentally, that can be as easily employed in studies based on group data. These conditions may reduce the risk of misleading information being incorporated into theorizing but one must still feel uneasy about dissociations based on the observation of just one or two patients. The answer here is to strengthen an idea by the collation of patients with similar deficits – something that appears to be done even by those most strongly associated with the single-case position.

At the time of writing both viewpoints seem firmly entrenched (see McCloskey, 1993 and Robertson et al., 1993) but, emerging from the criticisms of both sides, we can see how both approaches might be modified so as to become more acceptable. Group studies could benefit from stricter criteria for subject inclusion and the use of more stringent statistical procedures could reduce the inclusion of atypical group members. Proponents of the single-case approach would achieve greater support if they were to take some steps towards dealing with the issue of replication. In particular the development of some means of aggregating patients with similar deficits would dispel much of the current criticism.

In writing this book I was posed with a dilemma as to which viewpoint to adopt. Unlike Ellis and Young (1988), who chose the single-case viewpoint and thus mainly excluded group-based data, I have opted for a less radical approach and decided to use both single case and group-based data. This has inevitably exposed me to a much wider database and has the benefit of introducing some additional topics – particularly in the area of memory impairment.

Animal Studies

One of the more contentious areas of cognitive neuropsychology concerns its links with research on animals. Human neuropsychology is disadvantaged because the lesions suffered by patients lack precision in that they do not selectively target specific brain regions. Animal studies do not suffer from this drawback but a serious problem is whether the data from animal studies are relevant to understanding human cognition. This is a difficult issue, all the more so because of the ethical debate that continues.

Within neuropsychology it appears that animal studies seem most relevant in relation to 'lower' levels of cognition. In chapter 2, for example, we will see some overlap between human and animal data concerning the visual system. Similarly object recognition and attention

¹ An interesting case in point is the mnemonist S (Luria, 1968). This man had an incredible memory based largely on the use of imagery. No one doubts the power of the data but it does not figure in any current theory of normal memory function.

may also be areas in which animal work is informative about the organisation of human abilities. Memory is an area in which animal studies have been of varying influence. The work of Montreal group (see chapter 10) has shown some definite links between the role of the frontal lobes in the memory performance of non-human primates and humans. However, the classic work of O'Keefe and Nadel (1978), in which the rat hippocampus was specifically proposed as a cognitive map does not appear to generalize to human hippocampal function.

At the highest levels of human cognition: language, thought, executive function and consciousness; animal studies are likely to be of least relevance because here it is most likely that we are dealing with developments within the brain that are unique to humans.

The vastness of the animal literature along with my own doubts about how far one can generalize from animals to humans means that animal studies are not prominently featured. However, where animal research has provided some necessary clarification or enlightenment I have included it in my account.

Neuroimaging Techniques

Within the last 15 years there have been dramatic developments in neuroimaging techniques and these are having an increasing impact on neuropsychological research. A large number of patients now routinely receive computerized axial tomography (CAT) scanning. This is a technique based on X-rays which allows a three-dimensional picture of the brain to be created. CAT scans do serve many useful purposes but they are limited in terms of the detail they show – it is not unusual for patients with quite severe neurological difficulties to have a clear CAT scan.

A more recent and refined means of obtaining a structural picture of the brain is magnetic resonance imaging (MRI) sometimes known as nuclear magnetic imaging (NMR). This technique involves momentarily aligning atoms in a particular direction within a magnetic field using radio waves. When these waves are turned off the atoms produce a voltage (magnetic resonance) and this is used to construct an image. Plate 1.1 shows a CAT scan and an MRI scan and shows the greater degree of detail available in the latter. MRI scans are primarily used to identify which regions of the brain have been damaged. However, more recently, MRI techniques have been adapted in order to measure physiological changes in the brain, a technique known as functional magnetic resonance imaging (fMRI) (see Moonen, 1995, for a recent review of this technique and its applications).




Plate 1.1 Four different ways of imaging the brain: (a) a photograph, (b) CAT scan, (c) PET scan (normally different areas of activation would be shown in different colours), (d) MRI scan. From Posner & Raichle, *Images of Mind*. Copyright © 1994 Scientific American Library. Used with permission of W. H. Freeman.

Although use of functional MRI is increasing Positron Emission Tomography (PET) scanning is the most widely used means of detecting physiological abnormalities in the brain. PET involves injecting a patient with a water solution containing a positron emitting radioactive isotope. After a minute or so the radioactive isotope accumulates in the brain in direct proportion to the local blood flow – thus the greater the brain activity the greater the radioactive emission. In this way the extent of regional blood flow can be measured and those areas with greater or lesser activity identified. The results of PET scans are shown on images where brighter colours indicate higher levels of activity (see Plate 1.1).

Recently there has been increasing use of PET scanning as a means of investigating functional localisation in normal brains. The basic technique involves placing an individual into a PET scanning machine and then asking them to perform a particular cognitive task. By examining the different patterns of activation with different tasks it is possible to specify which brain structures are active during which tasks. Even greater

