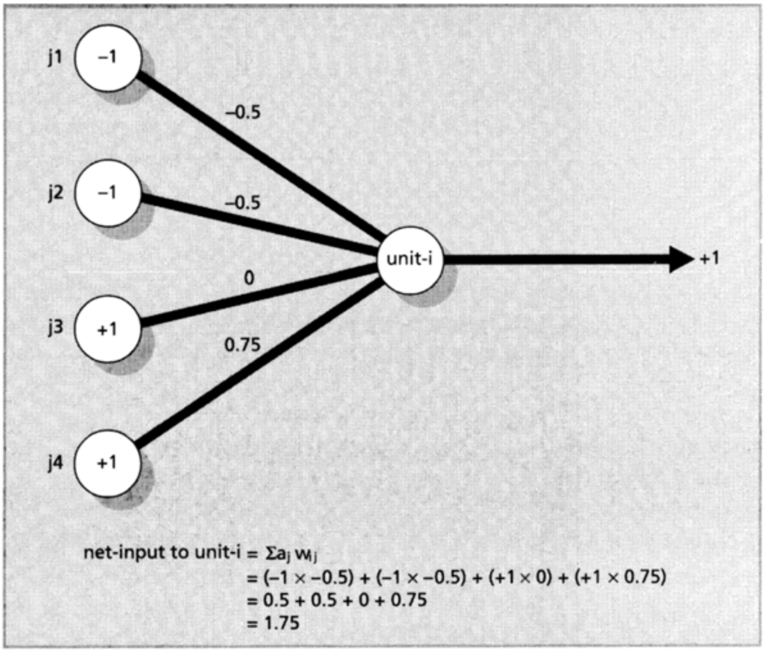

FIGURE 1.6

Diagram showing how the inputs from a number of units are combined to determine the overall input to unit-i. Unit-i has a threshold of 1; so if its net input exceeds 1 then it will respond with +1, but if the net input is less than 1 then it will respond with −1.

a number of modelling techniques. However, others have argued that connectionism represents an alternative to the information-processing paradigm (Smolensky, 1988; Smolensky, Legendre, & Miyata, 1993). Indeed, if one examines the fundamental tenets of the information-processing framework, then connectionist schemes violate one or two. For example, symbol manipulation of the sort found in production systems does not seem to occur in connectionist networks. We will return to the complex issues raised by connectionist networks later in the book.

COGNITIVE NEUROPSYCHOLOGY

Cognitive neuropsychology is concerned with the patterns of cognitive performance in brain-damaged patients. Those aspects of cognition that are intact or impaired are identified, with this information being of value for two main reasons. First, the cognitive performance of brain-damaged patients can often be explained by theories within cognitive psychology. Such theories specify the processes or mechanisms involved in normal cognitive functioning, and it should be possible in principle to account for many of the cognitive impairments of brain-damaged patients in terms of selective damage to some of those mechanisms.

Second, it may be possible to use information from brain-damaged patients to *reject* theories proposed by cognitive psychologists, and to propose new theories of normal cognitive functioning. According to Ellis and Young (1988, p. 4), a major aim of cognitive neuropsychology:

> is to draw conclusions about normal, intact cognitive processes from the patterns of impaired and intact capabilities seen in brain-injured patients…the cognitive neuropsychologist wishes to be in a position to assert that observed patterns of symptoms could not occur if the normal, intact cognitive system were not organised in a certain way.

The intention is that there should be bi-directional influences of cognitive psychology on cognitive neuropsychology, and of cognitive neuropsychology on cognitive psychology. Historically, the former influence was the greater one, but the latter has become more important.

Before discussing the cognitive neuropsychological approach in more detail, we will discuss a concrete example of cognitive neuropsychology in operation. Atkinson and Shiffrin (1968) argued that there is an important distinction between a short-term memory store and a long-term memory store, and that information enters into the long-term store through rehearsal and other processing activities in the short-term store (see Chapter 6). Relevant evidence was obtained by Shallice and Warrington (1970). They studied a brain-damaged patient, KF, who seemed to have severely impaired short-term memory, but essentially intact long-term memory.

The study of this patient served two important purposes. First, it provided evidence to support the theoretical distinction between two memory systems. Second, it pointed to a real deficiency in the theoretical model of Atkinson and Shiffrin (1968). If, as this model suggests, long-term learning and memory depend on the short-term memory system, then it is surprising that someone with a grossly deficient short-term memory system also has normal long-term memory.

The case of KF shows very clearly the potential power of cognitive neuropsychology. The study of this one patient provided strong evidence that the dominant theory of memory at the end of the 1960s was seriously deficient. This is no mean achievement for a study on one patient!

Cognitive neuropsychological evidence

How do cognitive neuropsychologists set about the task of understanding how the cognitive system functions? A crucial goal is the discovery of *dissociations,* which occur when a patient performs normally on one task but is impaired on a second task. In the case of KF, a dissociation was found between performance on short-term memory tasks and on long-term memory tasks. Such evidence can be used to argue that normal individuals possess at least two separate memory systems.

There is a potential problem in drawing sweeping conclusions from single dissociations. A patient may perform poorly on one task and well on a second task simply because the first task is more complex than the second, rather than because the first task involves specific skills that have been affected by brain damage. The solution to this problem is to look for double dissociations. A *double dissociation* between two tasks (1 and 2) is shown when one patient performs normally on task 1 and at an impaired level on task 2, and another patient performs normally on task 2 and at an impaired level on task 1. If a double dissociation can be shown, then the results cannot be explained in terms of one task being harder than the other.

In the case of short-term and long-term memory, such a double dissocation has been shown. KF had impaired short-term memory but intact long-term memory, whereas amnesic patients have severely deficient long-term memory but intact short-term memory (see Chapter 7). These findings suggest there are two distinct memory systems which can suffer damage separately from each other.

If brain damage were usually very limited in scope, and affected only a single cognitive process or mechanism, then cognitive neuropsychology would be a fairly simple enterprise. In fact, brain damage is often rather extensive, so that several cognitive systems are all impaired to a greater or lesser extent. This means that much ingenuity is needed to make sense of the tantalising glimpses of human cognition provided by brain-damaged patients.

Theoretical assumptions

Most cognitive neuropsychologists subscribe to the following assumptions (with the exception of the last one):

- The cognitive system exhibits *modularity,* i.e., there are several relatively independent cognitive processes or modules, each of which functions to some extent in isolation from the rest of the processing system; brain damage typically impairs only some of these modules.
- There is a meaningful relationship between the organisation of the physical brain and that of the mind; this assumption is known as *isomorphism.*
- Investigation of cognition in brain-damaged patients can tell us much about cognitive processes in normal individuals; this assumption is closely bound up with the other assumptions.
- Most patients can be categorised in terms of *syndromes,* each of which is based on co-occurring sets of symptoms.

Syndromes

The traditional approach within neuropsychology made much use of syndromes. It was claimed that certain sets of symptoms or impairments are usually found together, and each set of co-occurring symptoms was used to define a separate syndrome (e.g., amnesia; dyslexia). This syndrome-based approach allows us to impose some order on the numerous brain-damaged patients who have been studied by assigning them to a fairly small number of categories. It is also of use in identifying those areas of the brain mainly responsible for cognitive function such as language, because we can search for those parts of the brain damaged in all those patients having a given syndrome.

In spite of its uses, the syndrome-based approach has substantial problems. It exaggerates the similarities among different patients allegedly suffering from the same syndrome. In addition, those symptoms or impairments said to form a syndrome may be found in the same patients solely because the underlying cognitive processes are anatomically adjacent.

There have been attempts to propose more specific syndromes or categories based on our theoretical understanding of cognition. However, the discovery of new patients with unusual patterns of deficits, and the occurrence of theoretical advances, mean that the categorisation system is constantly changing. As Ellis (1987) pointed out, “a syndrome thought at time *t* to be due to damage to a single unitary module is bound to have fractionated by time *t*+2 years into a host of awkward subtypes.”

How should cognitive neuropsychologists react to these problems? Some cognitive neuro-psychologists (e.g., Parkin, 1996) argue that it makes sense to carry out group studies in which patients with the same syndrome are considered together. He introduced what he called the “significance implies homogeneity [uniformity] rule”. According to this rule, “if a group of subjects exhibits significant hetereogeneity [variability] then they will not be capable of generating statistically significant group differences” (Parkin, 1996, p. 16). The potential problem with this rule is that a group of patients can show a significant effect even though a majority of the individual patients fail to show the effect.

Ellis (1987) argued that cognitive neuropsychology should proceed on the basis of intensive single-case studies in which individual patients are studied on a wide range of tasks. An adequate theory of cognition should be as applicable to the individual case as to groups of individuals, and so single-case studies provide a perfectly adequate test of cognitive theories. The great advantage of this approach is that there is no need to make simplifying assumptions about which patients do and do not belong to the same syndrome.

Another argument for single-case studies is that it is often not possible to find a group of patients showing very similar cognitive deficits. As Shallice (1991, p. 432) pointed out, "as finer and finer aspects of the cognitive architecture are investigated in attempts to infer normal function, neuropsychology will be forced to resort more and more to single-case studies."

Ellis (1987) may have overstated the value of single-case studies. If our theoretical understanding of an area is rather limited, it may make sense to adopt the syndrome-based approach until the major theoretical issues have been clarified. Furthermore, many experimental cognitive psychologists disapprove of attaching great theoretical significance to findings from individuals who may not be representative even of brain-damaged patients. As Shallice (1991, p. 433) argued:

> A selective impairment found in a particular task in some patient could just reflect: the patient's idiosyncratic strategy, the greater difficulty of that task compared with the others, a premorbid lacuna [gap] in that patient, or the way a reorganised system but not the original normal system operates.

A reasonable compromise position is to carry out a number of single-case studies. If a theoretically crucial dissociation is found in a single patient, then there are various ways of interpreting the data. However, if the same dissociation is obtained in a number of individual patients, it is less likely that all the patients had atypical cognitive systems prior to brain damage, or that they have all made use of similar compensatory strategies.

Modularity

The whole enterprise of cognitive neuropsychology is based on the assumption that there are numerous *modules* or cognitive processors in the brain. These modules function relatively independently, so that damage to one module does not directly affect other modules. Modules are anatomically distinct, so that brain damage will often affect some modules while leaving others intact. Cognitive neuropsychology may help the discovery of these major building blocks of cognition. A double dissociation indicates that two tasks make use of different modules or cognitive processors, and so a series of double dissociations can be

Syndrome-based approach vs. single-case studies

syndrome-based approach	*Single-case studies*
Advantages	*Advantages*
Provides a means of imposing order and categorising patients. Allows identification of cognitive functions of brain areas. Useful while major theoretical issues remain to be clarified.	Avoids oversimplifying assumptions, No need to find groups of patients with very similar cognitive deficits.
Disadvantages	*Disadvantages*
Oversimplification based on theoretical assumptions. Exaggeration of similarities among patients.	Evidence lacks generalisability and can even be misleading.

used to provide a sketch-map of our modular cognitive system.

The notion of modularity was emphasised by Fodor (1983), who identified the following distinguishing features of modules:

- Informational encapsulation: each module functions independently from the functioning of other modules.
- Domain specificity: each module can process only one kind of input (e.g., words; faces).
- Mandatory or compulsory operation: the functioning of a module is not under any form of voluntary control.
- Innateness: modules are inborn.

Fodor's ideas have been influential. However, many psychologists have criticised mandatory operation and innateness as criteria for modularity. Some modules may operate automatically, but there is little evidence to suggest that they all do. It is implausible to assume the innateness of modules underlying skills such as reading and writing, as these are skills that the human race has developed only comparatively recently.

From the perspective of cognitive neuropsychologists, these criticisms do not pose any special problems. If the assumptions of information encapsulation and domain specificity remain tenable, then data from brain-damaged patients can continue to be used in the hunt for cognitive modules. This would still be the case even if it turned out that several modules or cognitive processors were neither mandatory nor innate.

It is not only cognitive neuropsychologists who subscribe to the notion of modularity. Most experimental cognitive psychologists, cognitive scientists, and cognitive neuroscientists also believe in modularity. The four groups differ mainly in terms of their preferred methods for showing modularity.

Isomorphism

Cognitive neuropsychologists assume there is a meaningful relationship between the way in which the brain is organised at a physical level and the way in which the mind and its cognitive modules are organised. This assumption has been called *isomorphism,* meaning that two things (e.g., brain and mind) have the same shape or form. Thus, it is expected that each module will have a different physical location within the brain. If this expectation is disconfirmed, then cognitive neuropsychology and cognitive neuroscience will become more complex enterprises.

An assumption that is related to isomorphism is that there is localisation of function, meaning that any specific function or process occurs in a given location within the brain (Figure 1.7). The notion of localisation of function seems to be in conflict with the connectionist account, according to which a process (e.g., activation of a concept) can be distributed over a wide area of the brain. There is as yet no definitive evidence to support one view over the other.

Evaluation

Are the various theoretical assumptions underlying cognitive neuropsychology correct? It is hard to tell. Modules do not actually "exist", but are convenient theoretical devices used to clarify our understanding. Therefore, the issue of whether the theoretical assumptions are valuable or not is probably best resolved by considering the extent to which cognitive neuropsychology is successful in increasing our knowledge of cognition. In other words, the proof of the pudding is in the eating. Farah (1994) argued that the evidence does not support what she termed the locality assumption, according to which damage to one module has only "local" effects. According to Farah (1994, p. 101), "The conclusion that the locality assumption may be false is a disheartening one. It undercuts much of the special appeal of neuropsychological architecture."

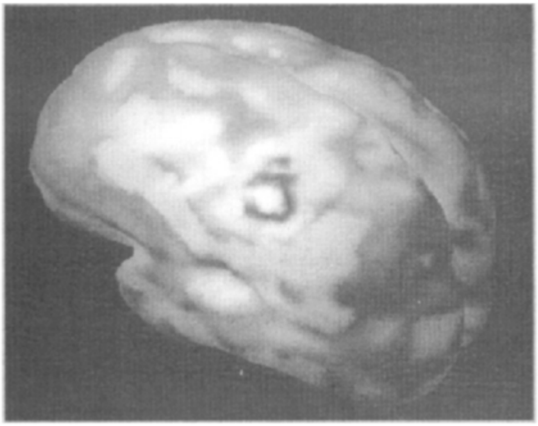

FIGURE 1.7

PET scans can be used to show localisation of function within the brain. This three-dimensional PET scan shows the metabolic activity within the brain during a hand exercise. The exercise involved moving the fingers of the right hand. The front of the brain is at the left. The most active area appears white; this is the motor cortex in the cerebral cortex where movement is coordinated. Photo credit: Montreal Neurological Institute/McGill University/CNRI/Science Photo Library.

One of the most serious problems with cognitive neuropsychology stems from the difficulty in carrying out group studies. This has led to the increasing use of single-case studies. Such studies are sometimes very revealing. However, they can provide misleading evidence if the patient had specific cognitive deficits prior to brain damage, or if he or she has developed unusual compensatory strategies to cope with the consequences of brain damage.

COGNITIVE NEUROSCIENCE

Some cognitive psychologists argue that we can understand cognition by relying on observations of people's performance on cognitive tasks and ignoring the neurophysiological processes occurring within the brain. For example, Baddeley (1997, p. 7) expressed some scepticism about the relevance of neurophysiological processes to the development of psychological theories:

> A theory giving a successful account of the neurochemical basis of long-term memory …would be unlikely to offer an equally elegant and economical account of the psychological characteristics of memory. While it may in principle one day be possible to map one theory onto the other, it will still be useful to have *both* a psychological and a physiological theory…Neurophysiology and neurochemistry are interesting and important areas, but at present they place relatively few constraints on psychological theories and models of human memory.

Why was Baddeley doubtful that neurophysiological evidence could contribute much to psychological understanding? The main reason was that psychologists and neurophysiologists tend to focus on different levels of analysis. In the same way that a carpenter does not need to know that wood consists mainly of

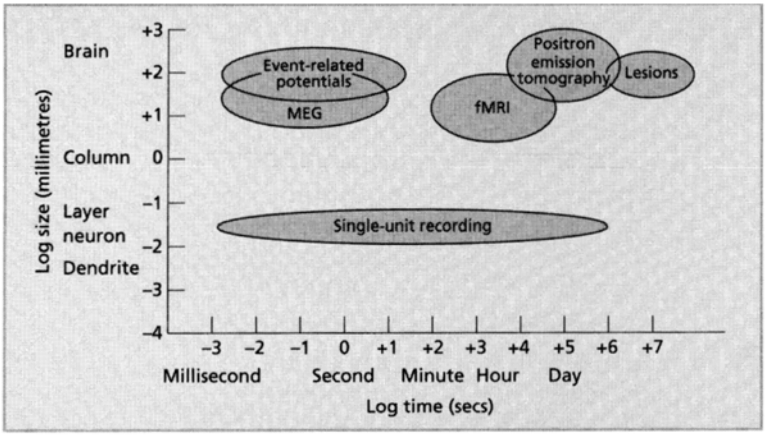

FIGURE 1.8

The spatial and temporal ranges of some techniques used to study brain functioning. Adapted from Churchland and Sejnowski (1991).

atoms moving around rapidly in space, so it is claimed that cognitive psychologists do not need to know the fine-grain neurophysiological workings of the brain.

A different position was advocated by Churchland and Sejnowski (1991, p. 17), who suggested:

> It would be convenient if we could understand the nature of cognition without understanding the nature of the brain itself. Unfortunately, it is difficult, if not impossible, to theorise effectively on these matters in the absence of neurobiological constraints. The primary reason is that computational space is consummately vast, and there are many conceivable solutions to the problems of how a cognitive operation could be accomplished. Neurobiological data provide essential constraints on computational theories, and they consequently provide an efficient means for narrowing the search space. Equally important, the data are also richly suggestive in hints concerning what might really be going on.

In line with these proposals, there are some psychological theories that are being fairly closely constrained by findings in the neurosciences (see Hummel & Holyoak, 1997, and Chapter 15).

Neurophysiologists have provided several kinds of valuable information about the brain's structure and functioning. In principle, it is possible to establish *where* in the brain certain cognitive processes occur, and *when* these processes occur. Such information can allow us to determine the order in which different parts of the brain become active when someone is performing a task. It also allows us to find out whether two tasks involve the same parts of the brain in the same way, or whether there are important differences. As we will see, this can be very important theoretically.

The various techniques for studying brain functioning differ in their spatial and temporal resolution (Churchland & Sejnowski, 1991). Some techniques provide information about the single-cell level, whereas others tell us about activity over much larger groups of cells. In similar fashion, some techniques provide information about brain activity on a millisecond-by-millisecond basis (which corresponds to the timescale for thinking), whereas others indicate brain activity only over much longer time periods such as minutes or hours.

Some of the main techniques will be discussed to give the reader some idea of the weapons available to cognitive neuroscientists. The spatial and temporal resolutions of some of these techniques are shown in Figure 1.8. High spatial and temporal resolutions are advantageous if a very detailed account of brain functioning is required, but low spatial and temporal resolutions can be more useful if a more general view of brain activity is required.

Single-unit recording

Single-unit recording is a fine-grain technique developed over 40 years ago to permit the study of single neurons. A micro-electrode about one 10,000th of a millimetre in diameter is inserted into the brain of an animal to obtain a record of extracellular potentials. A stereotaxic apparatus is used to fix the animal's position, and to provide the researcher with precise information about the location of the electrode in three-dimensional space. Single-unit recording is a very sensitive technique, as electrical charges of as little as one-millionth of a volt can be detected.

The best known application of this technique was by Hubel and Wiesel (1962, 1979). They used it with cats and monkeys to study the neurophysiology of basic visual processes. Hubel and Wiesel found there were simple and complex cells in the primary visual cortex, but there were many more complex cells. These two types of cells both respond maximally to straight-line stimuli in a particular orientation (see Chapter 4). The findings of Hubel and Wiesel were so clear-cut that they constrained several subsequent theories of visual perception, including that of Marr (1982; see Chapter 2).

Evaluation

The single-unit recording technique has the great value that it provides detailed information about brain functioning at the neuronal level, and is thus more fine-grain than other techniques (see Figure 1.8). Another advantage is that information about neuronal activity can be obtained over a very wide range of time periods from small fractions of a second up to several hours or days. A major limitation is that it is an invasive technique, and so would be unpleasant to use with humans. Another limitation is that it can only provide information about activity at the neuronal level, and so other techniques are needed to assess the functioning of larger areas of the cortex.

Event-related potentials (ERPs)

The electroencephalogram (EEG) is based on recordings of electrical brain activity measured at the surface of the scalp. Very small changes in electrical activity within the brain are picked up by scalp electrodes. These changes can be shown on the screen of a cathode-ray tube by means of an oscilloscope. A key problem with the EEG is that there tends to be so much spontaneous or background brain activity that it obscures the impact of stimulus processing on the EEG recording.

A solution to this problem is to present the same stimulus several times. After that, the segment of EEG following each stimulus is extracted and lined up with respect to the time of stimulus onset. These EEG segments are then simply averaged together to produce a single waveform. This method produces *event-related potentials* (ERPs) from EEG recordings, and allows us to distinguish genuine effects of stimulation from background brain activity.

ERPs are particularly useful for assessing the timing of certain cognitive processes. For example, some attention theorists have argued that attended and unattended stimuli are processed differently at an early stage of processing, whereas others have claimed that they are both analysed fully in a similar way (see

Chapter 5). Studies using ERPs have provided good evidence in favour of the former position. For example, Woldorff et al. (1993) found that ERPs were greater to attended than unattended auditory stimuli about 20–50 milliseconds after stimulus onset.

Evaluation

ERPs provide more detailed information about the time course of brain activity than do most other techniques, and they have many medical applications (e.g., diagnosis of multiple sclerosis). However, ERPs do not indicate with any precision which regions of the brain are most involved in processing. This is due in part to the fact that the presence of skull and brain tissue distorts the electrical fields emerging from the brain. Furthermore, ERPs are mainly of value when the stimuli are simple and the task involves basic processes (e.g., target detection) occurring at a certain time after stimulus onset. As a result of these constraints (and the necessity of presenting the same stimulus several times) it would not be feasible to study most complex forms of cognition (e.g., problem solving; reasoning) with the use of ERPs.

Positron emission tomography (PET)

Of all the new methods, the one that has attracted the most media interest is *positron emission tomography* or the PET scan. The technique is based on the detection of positrons, which are atomic particles emitted by some radioactive substances. Radioactively labelled water is injected into the body, and rapidly gathers in the brain's blood vessels. When part of the cortex becomes active, the labelled water moves rapidly to that place. A scanning device next measures the positrons emitted from the radioactive water. A computer then translates this information into pictures of the activity levels of different parts of the brain. It may sound dangerous to inject a radioactive substance into someone. However, only tiny amounts of radioactivity are involved.

Raichle (1994b) has described the typical way in which PET has been used by cognitive neuroscientists. It is based on a subtractive logic. Brain activity is assessed during an experimental task, and is also assessed during some control or baseline condition (e.g., before the task is presented). The brain activity during the control condition is then subtracted from that during the experimental task. It is assumed that this allows us to identify those parts of the brain that are active only during the performance of the task. This technique has been used in several studies designed to locate the parts of the brain most involved in *episodic memory,* which is long-term memory involving conscious recollection of the past (see Chapter 7). There is more activity in the right prefrontal cortex when participants are trying to retrieve episodic memories than when they are trying to retrieve other kinds of memories (see Wheeler, Stuss, & Tulving, 1997, for a review).

Evaluation

One of the major advantages of PET is that it has reasonable spatial resolution, in that any active area within the brain can be located to within about 3 or 4 millimetres. It is also a fairly versatile technique, in that it can be used to identify the brain areas involved in a wide range of different cognitive activities.

PET has several limitations. First, the temporal resolution is very poor. PET scans indicate the total amount of activity in each region of the brain over a period of 60 seconds or longer, and so cannot reveal the rapid changes in brain activity accompanying most cognitive processes. Second, PET provides only an indirect measure of neural activity. As Anderson, Holliday, Singh, and Harding (1996, p. 423) pointed out, "changes in regional cerebral blood flow, reflected by changes in the spatial distribution of intravenously administered positron emitted radioisotopes, are assumed to reflect changes in neural activity." This assumption may be more applicable at early stages of processing. Third, it is an invasive technique, because

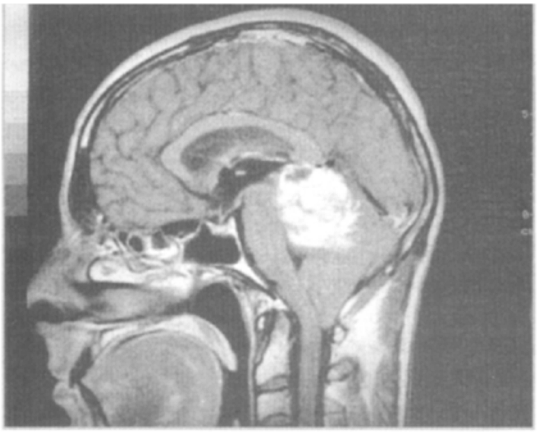

FIGURE 1.9

MRI scan showing a brain tumour. The tumour appears in bright contrast to the surrounding brain tissue. Photo credit: Simon Fraser/Neuroradiology Department, Newcastle General Hospital/Science Photo Library.

participants have to be injected with radioactively labelled water. Fourth, it can be hard to interpret the findings from use of the subtraction technique. For example, it may seem plausible to assume that those parts of the brain active during retrieval of episodic memories but not other kinds of memories are directly involved in episodic memory retrieval. However, the participants may have been more motivated to retrieve such memories than other memories, and so some of the brain activity may reflect the involvement of motivational rather than memory systems.

Magnetic resonance imaging (MRI and fMRI)

What happens in *magnetic resonance imaging* (MRI) is that radio waves are used to excite atoms in the brain. This produces magnetic changes which are detected by an 11-ton magnet surrounding the patient. These changes are then interpreted by a computer and turned into a very precise three-dimensional picture. MRI scans (Figure 1.9) can be used to detect very small brain tumours. MRI scans can be obtained from numerous different angles. However, they only tell us about the *structure* of the brain rather than about its *functions*.

The MRI technology has been applied to the measurement of brain activity to provide *functional MRI* (fMRI). Neural activity in the brain produces increased blood flow in the active areas, and there is oxygen and glucose within the blood. According to Raichle (1994a, p. 41), "the amount of oxygen carried by haemoglobin (the molecule that transports oxygen…) affects the magnetic properties of the haemoglobin… MRI can detect the functionally induced changes in blood oxygenation in the human brain." The approach based on fMRI provides three-dimensional images of the brain with areas of high activity clearly indicated. It is more useful than PET, because it provides more precise spatial information, and shows changes over shorter periods of time. However, it shares with PET a reliance on the subtraction technique in which brain activity during a control task or situation is subtracted from brain activity during the experimental task.

A study showing the usefulness of fMRI was reported by Tootell et al. (1995b). It involves the so-called waterfall illusion, in which lengthy viewing of a stimulus moving in one direction (e.g., a waterfall) is followed immediately by the illusion that stationary objects are moving in the opposite direction. There were two key findings. First, the gradual reduction in the size of the waterfall illusion over the first 60 seconds of observing the stationary stimulus was closely paralleled by the reduction in the area of activation observed in the fMRI. Second, most of the brain activity produced by the waterfall illusion was in V5, which is an area of the visual cortex known to be much involved in motion perception (see Chapter 2). Thus, the basic brain processes underlying the waterfall illusion are similar to those underlying normal motion perception.

Evaluation

Raichle (1994a, p. 350) argued that fMRI has several advantages over other techniques:

> The technique has no known biological risk except for the occasional subject who suffers claustrophobia in the scanner (the entire body must be inserted into a relatively narrow tube). MRI provides both anatomical and functional information, which permits an accurate anatomical identification of the regions of activation in each subject. The spatial resolution is quite good, approaching the 1–2 millimetre range.

One limitation with fMRI is that it provides only an *indirect* measure of neural activity. As Anderson et al. (1996, p. 423) pointed out, "With fMRI, neural activity is reflected by changes in the relative concentrations of oxygenated and deoxygenated haemoglobin in the vicinity of the activity." Another limitation is that it has poor temporal resolution of the order of several seconds, so we cannot track the time course of cognitive processes. A final limitation is that it relies on the subtraction technique, and this may not accurately assess brain activity directly involved in the experimental task.

Magneto-encephalography (MEG)

In recent years, a new technique known as *magneto-encephalography* or MEG has been developed. It involves using a superconducting quantum interference device (SQUID), which measures the magnetic fields produced by electrical brain activity. The evidence suggests that it can be regarded as "a direct measure of cortical neural activity" (Anderson et al., 1996, p. 423). It provides very accurate measurement of brain activity, in part because the skull is virtually transparent to magnetic fields. Thus, magnetic fields are little distorted by intervening tissue, which is an advantage over the electrical activity assessed by the EEG.

Anderson et al. used MEG in combination with MRI to study the properties of an area of the visual cortex known as V5 (see Chapter 2). They found with MEG that motion-contrast patterns produced large responses from V5, but that V5 did not seem to be responsive to colour. These data, in conjunction with previous findings from PET and fMRI studies, led Anderson et al. (1996, p. 429) to conclude that "these findings provide strong support for the hypothesis that a major function of human V5 is the rapid detection of objects moving relative to their background." In addition, Anderson et al. obtained evidence that V5 was active approximately 20 milliseconds after V1 (the primary visual cortex) in response to motion-contrast patterns. This is more valuable information than simply establishing that V1 and V5 are both active during this task, because it helps to clarify the *sequence* in which different brain areas contribute towards visual processing.

Evaluation

MEG possesses several valuable features. First, the magnetic signals reflect neural activity reasonably directly. In contrast, PET and fMRI signals reflect blood flow, which is assumed in turn to reflect neural activity. Second, MEG supplies fairly detailed information at the millisecond level about the time course of cognitive processes. This matters because it makes it possible to work out the sequence of activation in different areas of the cortex.

Techniques used by cognitive neuroscientists

Method	*Strengths*	*Weaknesses*
Single-unit recording	Fine-grain detail.	Invasive.
Information obtained over a wide range of time periods.	Only neuronal-level information is obtained.	
ERPs	Detailed information about the time course of brain activity.	lack precision in identifying specific areas of the brain. Can only be used to study basic cognitive processes.
PET	Active areas can be located to within 3–4 mm. Can identify a wide range of cognitive activities.	Cannot reveal rapid changes in brain activity. Provides only an indirect measure of neural activity. Findings f rom a subtraction technique can be hard to interpret.
MRl and fMRI	No known biological risk. Obtains accurate anatomical information. fMRl provides good information about timing.	Indirect measure of neural activity. Cannot track the time course of most cognitive processes.
MEG	Provides a reasonably direct measure of neural activity,	Irrelevant sources of magnetism may interfere with measurement.
Gives detailed information about the time course of cognitive processes.	Does not give accurate information about brain areas active at a given time.	

There are some major technical problems associated with the use of MEG. The magnetic field generated by the brain when thinking is about 100 million times weaker than the Earth's magnetic field, and a million times weaker than the magnetic fields around overhead power cables, and it is very hard to prevent irrelevant sources of magnetism from interfering with the measurement of brain activity. Superconductivity requires temperatures close to absolute zero, which means the SQUID has to be immersed in liquid helium at four degrees above the absolute zero of −273°C. However, these technical problems have been largely (or entirely) resolved. The major remaining disadvantage is that MEG does not provide structural or anatomical information. As a result, it is necessary to obtain an MRI as well as MEG data in order to locate the active brain areas.

Section summary

All the techniques used by cognitive neuro-scientists possess strengths and weaknesses. Thus, it is often desirable to use a number of different techniques to study any given aspect of human cognition. If similar findings are obtained from two techniques, this is known as converging evidence. Such evidence is of special value, because it suggests that the techniques are not providing distorted information. For example, studies using PET, fMRI, and MEG (e.g., Anderson et al., 1996; Tootell et al., 1995a, b) all indicate clearly that area V5 is much involved in motion perception.

It can also be of value to use two techniques differing in their particular strengths. For example, the ERP technique has good temporal resolution but poor spatial resolution, whereas the opposite is the case with fMRI. Their combined use offers the prospect of discovering the detailed time course *and* location of the processes involved in a cognitive task.

The techniques used within cognitive neuro-science are most useful when applied to areas of the brain that are organised in functionally discrete ways (S.Anderson, personal communication). For example, as we have seen, there is evidence that area V5 forms such an area for motion perception. It is considerably less clear that higher-order cognitive functions are organised in a similarly neat and tidy fashion. As a result, the various techniques discussed in this section may prove less informative when applied to such functions.

You may have got the impression that cognitive neuroscience consists mainly of various techniques for studying brain functioning. However, there is more than that to cognitive neuroscience. As Rugg (1997, p. 5) pointed out, “The distinctiveness [of cognitive neuroscience] arises from a lack of commitment to a single ‘level’ of explanation, and the resulting tendency for explanatory models to combine functional and physiological concepts.” Various examples of this explanatory approach are considered during the course of this book.

OUTLINE OF THIS BOOK

One problem with writing a textbook of cognitive psychology is that virtually all the processes and structures of the cognitive system are interdependent. Consider, for example, the case of a student *reading* a book to prepare for an examination. The student is *learning,* but there are several other processes going on as well. *Visual perception* is involved in the intake of information from the printed page, and there is *attention* to the content of the book (although attention may be captured by irrelevant stimuli). In order for the student to profit from the book, he or she must possess considerable *language skills,* and must also have rich *knowledge representations* that are relevant to the material in the book. There may be an element of *problem solving* in the student’s attempts to relate what is in the book to the possibly conflicting information he or she has learned elsewhere. Furthermore, what the student learns will depend on his or her *emotional state*. Finally, the acid test of whether the learning has been effective and has produced *long-term memory* comes during the examination itself, when the material contained in the book must be *retrieved*.

The words italicised in the previous paragraph indicate some of the main ingredients of human cognition, and form the basis of our coverage of cognitive psychology. In view of the interdependent functioning of all aspects of the cognitive system, there is an emphasis in this book on the ways in which each process (e.g., perception) depends on other processes and structures (e.g., attention; long-term memory; stored representations). This should aid the task of making sense of the complexities of the human cognitive system.

CHAPTER SUMMARY

- Cognitive psychology as a science. Cognitive psychology is unified by a common approach based on an analogy between the mind and the computer. This information-processing approach views the mind as a general-purpose, symbol-processing system of limited capacity. There are four main types of cognitive psychologists: experimental cognitive psychologists; cognitive scientists; cognitive neuropsychologists; and cognitive neuroscientists, who use various techniques to study brain functioning.
- Cognitive science. Cognitive scientists focus on computational models, in which theoretical assumptions have to be made explicit. These models are expressed in computer programs, which should produce the same outputs as people when given the same inputs. Three of the main types of computational model are semantic networks, production systems, and connectionist networks. Semantic networks consist of concepts, which are linked by various relations (e.g., is-similar-to). They are useful for modelling the structure of people's conceptual knowledge. Production systems are made up of productions in the form of "IF…THEN" rules. Connectionist networks differ from previous approaches in that they can "learn" from experience, for example, through the backward propagation of errors. Such networks often have several structures or layers (e.g., input units; intermediate or hidden units; and output units). Concepts are stored in a distributed manner.
- Cognitive neuropsychology. Cognitive neuropsychologists assume that the cognitive system is modular, that there is isomorphism between the organisation of the physical brain and the mind, and that the study of brain-damaged patients can tell us much about normal human cognition. The notion of syndromes has lost popularity, because syndromes typically exaggerate the similarity of the Symptoms shown by patients having allegedly the same condition. It can be hard to interpret the findings from brain-damaged patients for various reasons: patients may develop compensatory strategies after brain damage; the brain damage may affect several modules; patients may have had specific cognitive impairments *before* the brain damage.
- Cognitive neuroscience. Cognitive neuroscientists use various techniques for studying the brain, with these techniques varying to their spatial and temporal resolution. Important techniques include single-unit recording, event-related potentials, positron emission tomography, functional magnetic resonance imaging, and magneto-encephalography. Critics argue that neurophysiological findings am often at a different level of analysis from the one of most value to cognitive psychologists. In addition, such findings often fail to place significant constraints on psychological theorising.

FURTHER READING

- Ellis, R., & Humphreys, G. (1999). *Connectionist psychology; A text with readings*. Hove, UK: Psychology Press. Connectionism has become very influential within cognitive science, and this approach is discussed very thoroughly in this book.
- Gazzaniga, M.S., Ivry, R.B., & Mangun, G.R. (1998). *Cognitive neuroscience: The biology of the mind.* New York: W.W.Norton & Co. This is a comprehensive book in which the relevance of the cognitive neuroscience approach to the major areas of cognitive psychology is considered in detail.
- McLeod, P., Plunkett, K., & Rolls, E.T. (1998). *Introduction to connectionist modelling of cognitive processes.* Oxford; Oxford University Press. The principles and applications of connectionism are presented, and this book should even enable you to build your own connectionist models!
- Rugg, M.D. (1997). *Cognitive neuroscience.* Hove, UK: Psychology Press. Several experts discuss the ways in which cognitive neuroscience has benefited their area of research.
- Wilson, R.A., & Keil, F. (1999). *The MIT encyclopaedia of the cognitive sciences.* Cambridge, MA: MIT Press. This enormous book has extensive coverage by experts of computational intelligence, the neurosciences, and cognitive psychology.