The Handbook of Cognitive Neuropsychology

What Deficits Reveal About the Human Mind

Edited by

Brenda Rapp
Johns Hopkins University

2001

PSYCHOLOGY PRESS
Taylor & Francis Group

1

Assumptions and Methods in Cognitive Neuropsychology

Max Coltheart

WHAT IS COGNITIVE NEUROPSYCHOLOGY?

Cognitive neuropsychology is a branch of cognitive psychology.

Cognitive psychology is that science which seeks to learn more about the nature of the mental processes responsible for our ability to perform such basic cognitive activities as understanding and producing language, recognizing objects and people, storing information in memory and subsequently being able to retrieve it, acting intelligently upon the physical world, and so on. Cognitive psychologists also interest themselves in higher-level cognitive processes such as reasoning and problem-solving, the formation of beliefs about the world and about other people and how we assess the plausibility of such beliefs once they have occurred to us, and the cognitive nature of social interaction.

The most common way of investigating cognition is to study people with normal cognitive abilities as they are performing some particular cognitive task; the cognitive psychologist develops a theory about the nature of the mental processes used for performing this task, makes deductions from this theory about what results should be observed in some experiment involving that kind of cognition, carries out such an experiment, and considers whether the data confirm or falsify the theory.

Another way of doing cognitive psychology is to collect data relevant to the theory from people who have disorders of cognition. Any such disorder might be *acquired* or it might be *developmental*. If someone in whom a particular cognitive function had been normal suffers damage to the brain which specifically impairs that form of cognition that person has an acquired disorder of cognition. In contrast, someone who has not been able to acquire the cognitive ability to a normal degree has a developmental disorder of cognition. These disorders of cognition are relevant to theories about normal cognition, because such theories make predictions not only about the results of experiments with cognitively intact people, but also about ways in which brain damage can impair cognition, and about ways in which the acquisition of a cognitive ability can go awry. Thus, data from people with acquired or developmental disorders of cognition can confirm or falsify theories about cognition; this way of doing cognitive psychology is cognitive neuropsychology.

For example, suppose one believed that our ability to recognize all kinds of visual stimuli—objects, faces, and the printed word—used a single common visual recognition system. This theory predicts that if in a particular person brain damage has impaired the ability to recog-

nize visually-presented objects, such a person will also have an impairment in the ability to recognize faces and an impairment in the ability to recognize visually-presented words. If experimental investigations of this person reveal that face recognition and visual word recognition are both intact, that falsifies the original theory about visual recognition. The same would be true if one found, in a person who had no history of brain damage, that as he grew up he had acquired the ability to recognize faces and printed words but was never able to acquire the ability to recognize visually-presented objects.

Cognitive neuropsychology is not a kind of neuropsychology (even though it is highly relevant to neuropsychology) because, to put the matter in a nutshell, cognitive neuropsychology is about the mind, while neuropsychology is about the brain. Consider the hypothetical patient described above, who after brain damage could no longer recognize objects but could still recognize faces and visually-presented words. Presumably that means that there must be a region of the brain that is needed if we are to recognize objects, but not needed for recognizing faces or printed words. One might use brain-imaging methods with this patient to try to discover where in the brain this object-recognition system is located. However, whether one succeeded or failed in the attempt to locate such a region in the patient's brain is simply irrelevant to the conclusion drawn about the cognitive organization of visual recognition on the basis of the patient's performance. The theory that objects, faces, and visually-presented words are all recognized by a single common visual recognition system is still falsified by the finding of a person in whom brain damage affected only object recognition, regardless of what can or can't be discovered from studying the damaged brain of such a person.

Cognitive neuropsychology is also not necessarily concerned with treatment or rehabilitation, even though it might be relevant, even highly relevant, to them. Suppose one became convinced that all of the kinds of acquired or developmental cognitive disorders studied by cognitive neuropsychologists were completely untreatable—convinced, that is, that a child with a particular developmental disorder will never be able to acquire that cognitive ability, no matter what treatment is used, and that a brain-damaged patient who has lost a particular cognitive ability will never be able to regain it, no matter what treatment is used. Even if all of this were true, that would have no implications at all for the practice of cognitive neuropsychology. Data from people with cognitive disorders could still be used to confirm or falsify theories about normal cognition. On the other hand, it might be true one day—though it isn't true at present—that our theories about how normal cognitive processes operate have become so detailed that they make predictions about whether particular treatment techniques will or won't be effective. If that day comes, then treatment studies will be relevant to theories about normal cognition, and cognitive neuropsychology will be directly concerned with issues involving treatment.

Of course, there is already a way in which treatment studies can yield cognitive-neuropsychological evidence, since sometimes theories make predictions about the outcomes of treatment studies. If whether or not function Y depends upon function X is a question of relevance to a theory, and if a patient has an impairment of function X, and if successful treatment of function X is achieved, then one should observe an improvement in function Y. Failure to observe this would count as evidence against the theory.

Finally, although cognitive neuropsychology is distinct from clinical neuropsychology, it is important for clinical neuropsychology in a number of ways. For example, highly sophisticated methods of cognitive assessment can be developed if one uses as a starting point a detailed theory of the relevant cognitive system. The PALPA (Psycholinguistic Assessments of Language Processing in Aphasia) battery for assessing disorders of language (Kay, Lesser, & Coltheart, 1992) and the BORB (Birmingham Object Recognition Battery) for assessing disorders of visual perception and visual recognition (Riddoch & Humphreys, 1993) are good examples of such theory-based assessment methods. Cognitive-neuropsychological assessment methods such as these allow the clinician to pinpoint precisely which aspects of cognition are

intact and which are impaired in particular patients, and this provides a specific focus for treatment and rehabilitation.

THE CASE OF AC

Investigations of the patient AC by Coltheart et al. (1998) provide a useful example that provides a more detailed illustration of just how cognitive neuropsychologists seek to learn more about the nature of normal cognition from investigations of people with impairments of cognition.

AC was a 67-year-old man at the time of these investigations, and had formerly been employed as a clerical worker with the public railway system in New South Wales, Australia. He had a history of cardiovascular and cerebrovascular disease, and a CT scan performed four days after he had suffered a stroke revealed a recent lesion in the territory of the left middle cerebral artery, plus a number of older small lesions in both cerebral hemispheres indicative of earlier more minor cerebrovascular incidents.

Assessment of AC's linguistic and other cognitive abilities revealed some impairments that were surprisingly severe given the relatively minor damage evident in the CT scan. For example, AC's reading ability was almost completely abolished; he could not even judge that *A* and *a* are two forms of the same letter while *A* and *e* are not, a very elementary reading task. This was not because of some impairment of vision, since he could judge that *A* and *A* are the same while *A* and *E* are not, and could copy such letters correctly, indicating that he could see perfectly well. And since he had been a clerk, it is not likely that his severely impaired ability to read would have been present prior to his stroke. AC's writing was also almost completely abolished; he could write almost nothing to dictation except his name and address, and when shown an uppercase letter and asked to write its lowercase form next to it, he could not do so. This was not a motor problem, because he was entirely capable of copying an uppercase letter in its uppercase form, as mentioned above.

AC's ability to generate spoken words, in conversation or when trying to name a picture, was also severely affected; he knew what he wanted to say, but could rarely find the actual word he wanted to produce. He was also very poor at copying pictures and at drawing to dictation.

Studying his reading, his writing, or his drawing could well have provided interesting new information about how these cognitive activities are normally carried out. However, a conversation with AC involving his knowledge about the properties of objects—in particular, animals—led to a series of studies of his knowledge of word meanings. This conversation was prompted by a paper with the title "The Oyster with Four Legs" (Sartori & Job, 1988) which had appeared just before we first met AC. The conversation went as follows:

MC: How many legs does an oyster have?
AC: A few.
MC: I see. What about an ant?
AC: Some.
MC: A caterpillar?
AC: No legs.
MC: What about a snake?
AC: None.
MC: And a seagull?
AC: Four legs.

This inability to provide the required information about animals was formally tested by choosing a set of 18 animals, 9 possessing legs and 9 not, and asking AC to respond "Yes" (i.e., it has legs) or "No" (it hasn't) when the names of these animals were spoken to him. Chance performance would be 9 out of 18; AC scored 10 out of 18.

Why did he have such difficulty performing this task? One possibility is a category-specific loss of semantic information where the affected category is animals: This was the case with the patient described by Sartori and Job (1988). That was easy to test: A set of 20 inanimate objects, 10 with legs (e.g., chair, table, sofa, etc.) and 10 without legs, was chosen. When asked to indicate whether or not these objects had legs, in response to the spoken names of the objects, AC again failed to perform better than chance, scoring 11 out of 20 correct. So the difficulty here was not with loss of semantic information specifically concerning animals.

Perhaps it had something to do with the word or the concept "leg"? No, because AC was also at chance when the task was to respond to the question "Does it have a tail?", scoring 12 out of 22 correct and on a subsequent occasion, 11 out of 22 correct with the same items.

Was it just knowledge of the *parts* of objects that AC had lost? No, because he was unable to make judgements based just on the overall shapes of objects ("Is it round or not?" yielded 15 out of 28 correct) or just on their color ("Is it colored or just black and white?" yielded 12 out of 20 correct).

Well, perhaps what he had lost was information about the *perceptual properties* of objects; so he was then tested for his knowledge of nonperceptual information about objects. Asked to classify animals as Australian or not, he did very well (18 out of 20 correct; so he knew that a kangaroo is Australian and an elephant is not, despite his inability to report how many legs each has). Asked to classify animals as dangerous or not, he also did very well (19 out of 20 correct), and he also succeeded with the question "Do people usually eat this animal?" (23 out of 24 correct) and "Does this animal live in water or not?" (18 out of 20 correct). These four tests all require access to nonperceptual knowledge about objects, and AC performed all four very well.

His failures on the legs, tails, shape, and color tests could be because he had lost all *perceptual* knowledge about objects; or was it just specifically *visual* knowledge which had gone? That was tested by asking for perceptual but nonvisual knowledge. On a test of auditory perceptual knowledge (i.e., "Does it make a sound?"), he scored 24 out of 26 correct; and on a test of olfactory perceptual knowledge (i.e., "Does it have a smell?"), he scored 19 out of 20 correct.

These results support the following conclusion: What happened to AC is that he lost all information about the visual properties of objects, while still retaining knowledge about their nonperceptual properties, and also knowledge about perceptual properties which are not visual.

What does this tell us about how knowledge of objects is represented mentally? If everything we know about objects is represented in a single object knowledge system, it is hard to see how one particular form of knowledge (e.g., visual knowledge) could be lost while other forms of knowledge are still essentially intact. We are therefore led to the view that there is a system of knowledge about what objects look like which is quite separate from other stores of knowledge about other kinds of properties of objects.

Now, it is possible that there are just two systems of object knowledge, one which tells us what objects look like and another which tells us everything else about them (i.e., whether an object has a smell, lives in Australia, makes a noise, lives in water, etc.). But it doesn't seem very plausible to argue that all the forms of perceptual information are stored together with nonperceptual information, except for visual information, which has its own proprietary store. What seems more plausible, given that we want to say that there's a system just for visual knowledge of objects, is that there is also a system just for auditory knowledge about objects, and another one just for olfactory knowledge about objects—indeed, a separate system of perceptual object knowledge for each of our senses, plus a system of nonperceptual object knowledge. If so, we should find patients who know that a radio has no legs and a table does, but can't tell you which of them makes a noise (these hypothetical patients have lost their system of auditory knowledge about objects but still have their system of visual knowledge), patients who know that gasoline isn't normally drunk and vodka is, but can't tell you which of these has a smell (loss of olfactory object knowledge), and so on.

A key point here is that, although we have been led to this theory about how object knowledge is stored by studies of someone with impaired object knowledge, the theory is a theory about how such knowledge is represented in the intact person—it is a theory about all of us, including about AC prior to his brain damage. That, of course, is the essence of cognitive neuropsychology: building a theory about normal cognition from a study of abnormal cognition.

Like any other methodology, cognitive neuropsychology makes certain assumptions—that is, there are certain things which need to be true, or at least approximately true, if studies of people with impaired cognition are to be informative about the nature of intact cognitive systems. Those assumptions are discussed below, but before they can be discussed, it is necessary first to say something about a concept that is central for those assumptions—the concept of *modularity*.

MODULARITY

Let's suppose the cognitive task we are interested in is the very general task of understanding visual stimuli, whether they be objects or faces or printed words. One hypothesis that might be offered here is that such understanding involves three processing stages: detection of visual features, followed by visual recognition, followed by access to stored meanings. That hypothesis is represented in diagrammatic form in Figure 1.1.

What's being proposed here is that there is one single system containing representations of each of the stimuli we have learned to recognize by sight, whether these be objects or faces or visually-presented words, and a separate single system containing information about the meanings associated with each of these stimuli. The representation of a particular stimulus in the visual recognition system would be linked to the feature detectors for just those visual features which that stimulus possesses, so that when just those features are present in the visual stimulus, the representation of that stimulus in the visual recognition system would be maximally activated. This would lead in turn to the activation of the meaning of that stimulus in the semantic system. The visual recognition system does not contain any information about meanings; its domain is visual structure, not semantics.

According to this theory, then, the system we use for understanding what we see has three components. Such components are typically depicted by cognitive neuropsychologists as boxes in box-and-arrow diagrams like that of Figure 1.1. We need some term to refer to these components, and a commonly used term is *module*. A theory like that represented in Figure 1.1 is referred to as a modular theory, since it claims that the system we use for understanding what we see is composed of three modules. So it is necessary to say something about what's meant by the term module.

The most useful discussion of what this term means is to be found in Fodor (1983), even though what Fodor said about modularity has been widely misunderstood. Fodor approached this issue by listing a number of properties which are *characteristic* of modules. These properties included the following:

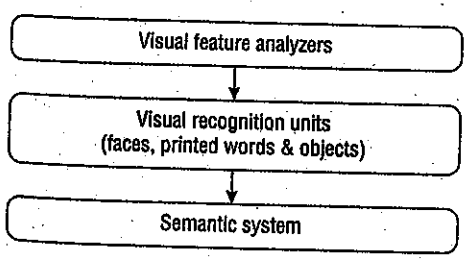

Figure 1.1. A schematic representation of a possible theory of object recognition.

- domain specific
- innately specified
- informationally encapsulated
- fast
- hardwired (neurally specific)
- autonomous
- not assembled

Misunderstandings of Fodor have mainly arisen via the belief that what he was offering was intended as a definition of modularity—that is, that he was proposing a list of properties all of which were necessary conditions for applying the term. But Fodor was very clear that this was not his intention: "I am not, in any strict sense, in the business of defining my terms. . . . So what I propose to do instead of defining 'modular' is to associate the notion of modularity with a pattern of answers to such questions as 1–5" (Fodor, 1983, p. 37; Here the five questions he was referring to were questions about possible features of modularity, questions such as "Is the system innately specified?" or "Is the system domain-specific?")

And again: "Given that a system has any of the properties in question, then the likelihood is considerable that it has all the rest. . . . However, I doubt that a claim that strong could be empirically sustained, since it is reasonably easy to think of psychological processes that are fast but not encapsulated, or involuntary but not innate, and so forth. The present contention, in any event, is relatively modest: it's that if a psychological system has most of the modularity properties, then it is very likely to have all of them" (Fodor, 1983, p. 137).

So, for example, Fodor would not *require* that a system be innate for it to be called a module; he is simply suggesting that most systems that deserve to be called modules turn out to be innate. But the system that skilled readers use for reading, whilst certainly not innate, is fast, neurally-specific, domain-specific, and autonomous (Fodor means "automatic" by this term), and so is modular to some interesting extent—compare: "One would thus expect—what anyhow seems to be desirable—that the notion of modularity ought to admit of degrees. The notion of modularity that I have in mind certainly does. When I speak of a cognitive system as modular, I shall therefore always mean 'to some interesting extent'" (Fodor, 1983, p. 37).

Elsewhere (Coltheart, 1999) I have suggested a concept of modularity that departs very slightly from Fodor's, but appears more useful, at least for cognitive psychologists and cognitive neuropsychologists, and which I am adopting throughout this chapter:

> Fodor was not concerned with defining modularity, nor with specifying any properties that a cognitive system must necessarily possess for it to be considered modular. I am suggesting that one can be more ambitious than this, by defining 'module' as 'a cognitive system whose application is domain specific'; here domain-specificity is a necessary condition for the applicability of the term 'modular'. Now it is necessary to say something about what might be meant by 'domain-specific'. I mean that a cognitive system is domain-specific if it only responds to stimuli of a particular class: so that to say that there is a domain-specific face recognition module is to say that there is a cognitive system which responds when its input is a face, but does not respond when its input is, say, a written word, or a visually-presented object, or someone's voice. There's no circularity here, since the claim that there is a face recognition module does not derive merely from the existence of faces as a stimulus class that is conceptually distinguishable from other stimulus classes such as written words, objects or voices. The claim is derived from empirical observations, as follows. Suppose first that we were entertaining the idea that there was a single module for visual recognition which accepts as inputs faces, objects and printed words. Then we noticed that in the neuropsychological literature there were reports of patients with impaired visual word recognition but retained face recognition and of patients with impaired visual word recognition but retained visual object recognition. We also noticed reports of patients with impaired visual object recognition but retained face recognition [1] and patients with impaired visual object recognition but retained visual word recognition. Finally, we also noticed reports of patients with impaired face recognition but retained visual object recognition and of patients with impaired face recognition but retained visual object recognition. This collection of results refutes our original idea that there is a cognitive module whose domain is the recognition of all forms of visual stimuli. Instead, there are three separate modules: a face recognition module, a visual object recognition module, and a visual word recognition module [a claim that I believe is in fact correct]. The domain-specificities are obvious here: visual objects, faces, printed words." (Coltheart, 1999, pp. 118–119)

Note that this conception of domain specificity does not require that there be some kind of gatekeeper which decides which stimuli should be admitted to a module and which should not. The retina does not respond to auditory stimuli (its domain is light in a certain wavelength range) but no one would suppose that this is because there is some mechanism which actively blocks sound waves from reaching the retina. Sound waves pass through the whole head; but only some structures in the head are activated by these waves. In the same way, visual information from print would reach the face recognition module, but that module would not contain any structures that would be activated by this particular kind of visual information.

Here, then, a module is an information-processing mechanism which is domain-specific (i.e., which is responsible for processing just certain kinds of information for certain specific purposes). Modules can themselves be modular in organization; as we've seen from the example of AC discussed earlier, the semantics module itself is made up of a number of different modules.

So is Mike McCloskey's laser printer, discussed in Chapter 24. In that chapter McCloskey reports data from studies of that printer showing that an overload of memory affects the printing of graphics but not the simultaneous printing of text. From this dissociation he inferred with extreme caution that inside his laser printer " . . . text is somehow treated differently from graphics" (this volume, p. 594); a less cautious but perfectly reasonable inference is that the printer contains two distinct modules, one a memory for graphic information and the other a memory for text information. This proposal about the printer's functional architecture might be refuted by further investigations of the printer's behavior, but given the data reported in Chapter 24 the proposal is not only reasonable but plausible.

Here, "functional architecture" means a description of an information-processing system in terms of what its modules are and what the pathways of information flow between these modules are. Any information-processing system, be it a person or a printer, has a functional architecture—and also a hardware architecture, this being a neural architecture in the case of persons.

Having said what I mean by the terms "module" and "functional architecture," I can now return to the issue of the assumptions of cognitive neuropsychology.

THE ASSUMPTIONS OF COGNITIVE NEUROPSYCHOLOGY

Firstly, let's be clear what "assumption" means here. It does not mean "here are some things which cognitive neuropsychologists assert to be true of people." It means something quite different: "here are some things which need to be true of people, at least to a good approximation, for cognitive neuropsychology to be able to succeed in its aims." There are four major assumptions, as follows.

Assumption 1: Functional Modularity

Since a fundamental aim of cognitive neuropsychology is to discover functional architectures of cognitive systems, and since a functional architecture is a configuration of modules, cognitive neuropsychology will get nowhere except when cognitive systems actually are configurations of modules. Fodor (1983) is interesting again here. In that book he offers an account of the mind

as consisting of two systems. There's the system of input modules and there's the central system. The input modules are responsible for encoding and recognizing perceptual inputs. The central system is responsible for such "higher-level" cognitive processes as the fixation of belief. The input-module system is, of course, modular, but the central system is nonmodular. What follows from that, according to Fodor, is that, because cognition which involves the central system is nonmodular, it is not amenable to scientific study. Indeed, he offered this as Fodor's First Law of the Nonexistence of Cognitive Science. If he is correct that some forms of cognition depend upon cognitive systems that are not functionally modular, then these forms of cognition cannot be successfully studied by cognitive neuropsychology. It remains to be seen whether this is the case.

Assumption 2: Anatomical Modularity

Even if cognitive systems are actually configurations of modules, it does not follow that any of these modules is realized in some specific and relatively small region of the brain. It could be that the neural tissue for each individual functional module is very widely spread throughout the brain. In other words, it is possible for there to be functional modularity but no anatomical modularity. If so, almost any form of brain damage must affect very many—even all—modules. In that case, cognitive neuropsychology would get nowhere because the functional modularity of cognition would not manifest itself in the performance of brain-damaged patients—for example, even if the face recognition system and the spoken-word recognition system were two quite distinct functional modules, one would never see a patient who could recognize faces but not spoken words, nor a patient who could recognize spoken words but not faces.

Assumption 3: Uniformity of Functional Architecture Across People

Even if cognitive systems are actually configurations of modules (i.e., there is functional modularity), and even if individual functional modules generally are realized in restricted brain regions (i.e., there is anatomical modularity), cognitive neuropsychology would still not get anywhere if different individuals had different functional architectures for the same cognitive domain. That's because it would not be possible to make inferences about the functional architecture of the cognitive systems of people in general from the data of a single patient, however well-justified the inference was from that patient's data to claims about that particular patient's functional architecture. Of course, this assumption is not peculiar to cognitive neuropsychology; it is widespread throughout the whole of cognitive psychology. Thus if this assumption is false, that's not just bad news for cognitive neuropsychology; it is bad news for all of cognitive psychology.

Assumption 4: Subtractivity

Cognitive neuropsychology treats the functional architecture of an impaired cognitive system as the functional architecture of the intact cognitive system with one or more of its components damaged or deleted. The assumption here is that brain damage can impair or delete existing boxes or arrows in the system, but cannot introduce new ones: that is, it can subtract from the system, but cannot add to it. If this were not so, studying impaired systems could not tell us about normal systems. This is not at all to say that brain-damaged patients do not adopt new strategies to cope with their impairments, strategies that intact individuals do not use. Such compensatory strategies could consist of using the normal functional architecture in an abnormal way. For example, in patients with pure alexia (letter-by-letter reading), a condition discussed in Chapter 10, visual word identification is typically achieved by spelling out the letters of a word one-by-one, and identifying the whole word after having spelled out all of its

letters. That is not how intact readers recognize words. But intact readers have a letter-naming capability, and they are also capable of recognizing a word if it is spelled aloud to them. Hence the procedures for word recognition used by the pure-alexic patient have not been introduced to the language system by the brain damage; these procedures are present in the intact system, but are not normally used to assist visual word identification.

If one or more of these four assumptions is untrue, what are the implications for cognitive neuropsychology? The worst-case scenario is happy oblivion. The cognitive neuropsychologist's data appear to make sense, and interesting theories about functional architectures are suggested by them; little does the cognitive neuropsychologist know that this happy state of affairs has only arisen because a false assumption has been assumed to be true. Here, all is illusion, but there is no way for the cognitive neuropsychologist to know this. Yet, a little thought shows how unlikely this scenario is. Suppose for example that Assumptions 1, 2, and 4 are true but Assumption 3 is false. This will mean that patients A and B will, before their brain damage, have possessed functionally modular cognitive architectures which were also anatomically modular, and that their brain damage has subtracted from, but not added anything to, these architectures; but the architectures of each patient were, and so still are, very different. This will be immediately apparent: Patient A's data will strongly suggest a certain cognitive architecture, so will patient B's, but the two architectures thus suggested will be different ones. Thus, the cognitive neuropsychologist will not be happily oblivious here: It will be obvious that something is deeply wrong if different patients yield quite different ideas about what the functional architecture is.

What if Assumption 1 is false, or Assumption 2? In both cases, the consequence will be that we simply won't see patients with highly selective cognitive disorders, patients such as AC (whose semantic impairment was just for visual information). But we do—and there are many other examples of such highly selective cognitive disorders in other chapters of this book. If just Assumption 1 is false, then we won't see such cases because the cognitive system does not have the modular organization that would cause them. If just Assumption 2 is wrong, then we won't see such cases because the brain does not have the degree of anatomical modularity that would allow brain damage to produce such disorders.

Finally, Assumption 4: If the other assumptions were all correct but this one were false, that would soon be evident, because the functional architectures inferred from studies of brain-damaged patients could not be successfully applied to the explanation of results obtained by cognitive psychologists studying people without brain damage.

In sum, then, if any of these four assumptions were false, that would have soon become apparent, because cognitive-neuropsychological research would soon have run into severe difficulties; thus there is a kind of fail-safe mechanism here. From this follows an even more important point: If one considers that cognitive neuropsychological research over the past 30 years or so has yielded much coherent evidence about cognition that has led to proposals about cognitive architectures which are also useful in explaining data from studies of people with intact cognitive systems, that conclusion constitutes evidence that all four of these quite fundamental assumptions about mind and brain are in fact correct.

INFERENCE FROM DATA TO THEORY IN COGNITIVE NEUROPSYCHOLOGY

Cognitive neuropsychology seeks to interpret cognitive disorders as selective impairments of functional architectures. Suppose, for example, you have a particular theory about how visually-presented objects are recognized, and you express the functional architecture proposed by this theory in the form of a modular box-and-arrow diagram. Then, whenever you come across a patient with any kind if impairment of visual object recognition, you will try to determine which boxes and arrows of the diagram are still intact in the patient's visual recognition

system and which have been damaged. If you can demonstrate that a certain pattern of damage to the functional architecture you've proposed would cause that damaged system to show just the same symptoms that the patient is showing, then you have achieved two things: you've obtained empirical support for the theory of visual object recognition, and you've gained an insight into exactly what is wrong with this patient that's much deeper than simply saying, "The patient has an impairment of visual object recognition." If on the other hand you can demonstrate that there is *no* pattern of damage to the boxes and arrows which would cause the system to behave like the patient, then you have also achieved something: you have shown that the theory is false. The data from the patient have refuted the theory.

It is important here to distinguish between three different kinds of data that brain-damaged patients may yield, these kinds being associations of deficits, dissociations of deficits, and double dissociations of deficits.

Associations, Dissociations, and Double Dissociations

With associations the patient is impaired on task X and also on task Y. For example, he or she is poor at understanding printed words and poor at understanding spoken words. Here the two impairments are said to be associated, because both are present in the same patient

With dissociations the patient is impaired on task X but responds normally on task Y. For example, he or she is poor at understanding printed words but normal at understanding spoken words. Here the two impairments are said to be dissociated, because one is present and the other is absent in the same patient

With double dissociations we typically need two patients: patient A who is impaired on task X but normal on task Y, and patient B who is normal on task X but impaired on task Y. Both patients show a dissociation; when the two dissociations occur in opposite directions, this is called a double dissociation.

One can still speak of dissociations between two tasks even if performance is impaired on both tasks. If a patient is impaired at both task A and task B, but is significantly more impaired on the second task than on the first, that can be treated as a dissociation; and if a second patient is observed who is also impaired at both task A and task B, but is significantly more impaired on the first task than on the second, that can be treated as a double dissociation. A more detailed and technical analysis of associations, dissociations, and double dissociations may be found in Chapter 10 of Shallice (1988).

It is possible to have double dissociations within a single patient. For example, Rapp, Benzing, and Caramazza (1995) described a patient who was worse with nouns than with verbs when producing them in speech, but worse with verbs than nouns when producing them in print.

How can we use data like these to learn more about cognition?

We might interpret the data as evidence that there is a single cognitive module responsible for word comprehension, and that module is damaged. If so, one would expect that both spoken words and printed words would suffer.

A problem here is that this interpretation could be correct, but an alternative is obvious: There are separate cognitive systems for the comprehension of speech and print, but these systems are so close together in the brain that it is unlikely, or even impossible, for only one to be damaged. Take, for example, Gerstmann's syndrome, which is diagnosed when four symptoms are present: acalculia, finger agnosia, right–left disorientation, and dysgraphia. It is conceivable that this association of four impairments occurs through damage to one single cognitive system that we use for finding square roots, finding our way around the environment, knowing which finger is the thumb, and writing; but this is hardly a plausible cognitive theory. A far more plausible interpretation for such a syndrome is that these tasks involve a variety of separate cognitive systems but since these systems are located so close together in the brain (in this case, in the left parietal cortex), any damage to the region on which one of the

cognitive systems depends is highly likely to have also affected the regions upon which the other cognitive systems depend, thereby producing impairments in the other cognitive tasks.

Although this challenge to the use of association data in cognitive neuropsychology is a standard one, it is by no means so strong as to completely eliminate associations from consideration. Suppose, for example, one were studying a patient who had both impaired ability to read aloud and also impaired spelling ability. Detailed investigation of his reading showed that he could read aloud pronounceable nonwords such as *choe* or *plood* successfully, but he often made errors with real words which disobey standard spelling-to-sound rules, exception words such as *shoe* or *flood*; and his errors here were "regularization errors," such that he read these words according to the rules, pronouncing *shoe* to rhyme with *toe* and *flood* to rhyme with *mood*. This form of acquired dyslexia is known as surface dyslexia, and is frequently interpreted in relation to a dual-route theory of reading aloud.

In such a theory, one way to read aloud is by applying letter-to-sound rules to the printed letter string. This will succeed with nonwords (and also with regular words, which obey standard spelling-to-sound rules) but will fail with exception words, producing regularization errors for such words. The other way to read aloud is via a lexicon of known words: the letter string is recognized as an entry in an orthographic input lexicon and that allows retrieval of its pronunciation from a phonological output lexicon. This will succeed with all words, regular or exception, but will fail for all nonwords, since nonwords are not represented in the orthographic input lexicon (nor in the phonological output lexicon). Given this theory, if brain damage impairs the orthographic input lexicon but spares the nonlexical (rule-based) method of reading, the patient will be able to read nonwords successfully, but will make regularization errors with exception words—that is, will exhibit surface dyslexia.

Exactly the same kind of theory can be applied to the task of spelling. In a dual-route theory of spelling, there's a rule-based procedure (applying sound-to-spelling rules) which will succeed for nonwords and for words which obey such rules, but which will yield regularization errors in spelling words which don't obey such rules; so *flood* would be spelled *flud* and *shoe* as *shoo* by this procedure. And there's a second procedure for spelling: looking up the item to be spelled in an orthographic output lexicon. This will succeed with all words, regular or exception, but will fail for all nonwords, since nonwords are not represented in the orthographic output lexicon. Given this theory, if brain damage impairs the orthographic output lexicon but spares the nonlexical (rule-based) method of spelling, such a patient will be able to spell nonwords successfully, but will make regularization spelling errors with exception words—that is, will exhibit the form of acquired dyslexia known as surface dysgraphia.

What has all this got to do with associations? Well, suppose that we see a patient who exhibits both surface dyslexia (e.g., reads *flood* to rhyme with *mood*) and surface dysgraphia (e.g., spells *flood* as *flud*). It is very tempting to conclude that this shows that it is wrong to distinguish between an orthographic input lexicon and an orthographic output lexicon. One could interpret such a patient much more economically by supposing that there is only one orthographic lexicon, used both for reading and for spelling. A single lesion, just of this orthographic lexicon, would produce both surface dyslexia and surface dysgraphia. If instead one wished to distinguish between an orthographic input lexicon and an orthographic output lexicon, then one would have to propose that any patient like this has two separate lesions.

But a defender of the two-orthographic-lexicon theory could reply: Yes, many patients with surface dyslexia also show surface dysgraphia, but that's not because there is only one orthographic lexicon; instead, there are indeed separate input and output orthographic lexicons, but they are neuroanatomically close, and so when one is affected by brain damage the other is also very likely to be affected too. This association of deficits is not telling us anything about the functional architecture of the reading and spelling systems, the defender would say; it is just an accident of the way in which the components of this system are represented in the brain.

But now suppose that an investigator showed that the irregular words that a surface dyslexic and surface dysgraphic patient could not read were exactly those that he could not spell, and all of the irregular words he could read correctly he could also spell correctly. What are we to conclude from this? Either there are distinct orthographic input and output lexicons, and this item-specific correspondence between reading and spelling is just a huge coincidence; or, far more plausibly, the item-specific correspondence between reading and spelling here is telling us that there is just a single orthographic lexicon use both for reading and for writing. Some of the entries in this lexicon have been damaged; for all words which are irregular and whose entries have been damaged, both reading and spelling will have to be nonlexical, and therefore erroneous. This is an example of a point made in more detail in McCloskey's chapter, namely, that cognitive-neuropsychological data are not limited to gross level-of-performance measures such as percent correct. More sophisticated analyses such as the analysis of consistency of performance in the above example can overcome the limitations of inferences from associations of deficits.

Dissociation. Suppose we observe a patient who is poor at understanding printed words but good at understanding spoken words. What might we infer from this about the cognitive psychology of language? We might interpret the data as evidence that there are separate cognitive modules responsible for written word comprehension and for spoken word comprehension, and the first of these modules is damaged whilst the second one is still intact.

In this case the problem is that this interpretation could be correct, but an alternative is obvious: there might be just one cognitive module responsible for word comprehension, and comprehension of written words might be harder for it than comprehension of spoken words. So when it is partially damaged, printed words will suffer more than spoken words.

This challenge to the use of single dissociations in cognitive neuropsychology is referred to as the *resource artifact* objection by McCloskey (Chapter 24, this volume), and he points out that in many situations it is so implausible a challenge that it is not a serious problem for attempts to make inferences from single-dissociation data. One example he uses is the patient who can produce picture names in writing but not in speaking. To rebut the claim that this result implies that there are separate orthographic and phonological output lexicons, one would have to argue that spoken picture naming is somehow intrinsically harder than written picture naming. Such a claim would seem to have no justification at all.

Double dissociation, suppose we observe the following two patients: Patient A is impaired on comprehending printed words, but normal at comprehending spoken words; patient B is normal on comprehending printed words but impaired on comprehending spoken words.

What might we infer from this about the cognitive psychology of language? We might interpret the data as evidence that there are separate cognitive modules responsible for spoken word comprehension and written word comprehension. One patient has damage to the first module, the other to the second. That reasoning is not subject to the objections to associations and single dissociations; that's one reason for the popularity of the double dissociation method. For example, the claim that the dissociation found in patient A arose because there's just a single language comprehension system which finds printed-word comprehension harder than spoken-word comprehension is refuted by the data from patient B, for whom printed-word comprehension is easier than spoken-word comprehension

However, the inference that there are separate cognitive modules responsible for spoken word comprehension and written word comprehension is not the correct one to draw from the data of patient A and patient B, or at least it is too vaguely stated, because of the ambiguity of the term "separate." Suppose, as surely is the case, that spoken-word comprehension is achieved by a cognitive system which consists of a number of different components (modules); and suppose this is also true of written-word comprehension. If we say that there are separate cognitive modules responsible for spoken word comprehension and written word comprehension, "separate" might be taken to mean "completely distinct" (i.e., having no modules in

common). It would not be reasonable to infer, from the double dissociation seen with patients A and B, that the spoken-word comprehension and written-word comprehension systems have no modules in common.

What is reasonable to infer is the conclusion that the two systems are not identical–and, in particular, that there is *at least one* module that is part of the spoken-word comprehension system but not part of the written-word comprehension system (this module is damaged in patient A but intact in patient B), and that there is *at least one* module that is part of the written-word comprehension system but not part of the spoken-word comprehension system (this module is damaged in patient B but intact in patient A).

TWO COMMON MISUNDERSTANDINGS ABOUT COGNITIVE NEUROPSYCHOLOGY AND DOUBLE DISSOCIATIONS

There are two misunderstandings that are sufficiently common that I once heard both made in conference talks at the same conference. Because they are so common, it is worth identifying them and explaining just why they are mistaken.

Speaker A at this conference argued that the use of double dissociation data to make inferences about the structure of cognitive systems is always based on a simple logical fallacy, known since Aristotle: the fallacy of affirmation of the consequent. The argument from double dissociation, he said, always runs thus:

1. If cognitive architecture A were true, then one would see double dissociation B;
2. Double dissociation B has been observed;
3. Therefore, it follows that cognitive architecture A must be true.

This is indeed fallacious. But it isn't the way cognitive neuropsychologists argue. Instead, they argue thus:

1. I have observed double dissociation B.
2. Let me think . . . what's a cognitive architecture which, when damaged in two different ways, would produce double dissociation B?
3. Well, here's one: cognitive architecture A.
4. So I hypothesize that the architecture of the cognitive system is architecture A.

Cognitive neuropsychologists never claim that certain data logically *require* a particular architecture, because, of course, it is never the case that data logically require a particular theory. A theory is a reasonable interpretation of data, not a logical consequence of data.

Speaker B at this meeting argued that inference based on double dissociation evidence is always fatally flawed because it is always the case that there might be some cognitive architecture different from the one proposed that is also compatible with the observed double dissociation data.

This is indeed true. But it isn't just true of the double dissociation approach. It is true of all the methods psychology uses when building theory from data. Indeed, it is true of all science, not just psychology. Consider, for example, Mendel's postulation of the gene to explain the data he had collected concerning the characteristics of successive generations of sweet peas. He didn't say, nor could he have, that this was the only possible explanation of his data.

Nor would Gell-Mann have said that the only possible explanation of the data he considered was the proposal that quarks exist. Observations from experimental physics indicated the existence of some 100 particles in the atom's nucleus. Gell-Mann proposed that all of those particles, including the neutron and proton, are composed of fundamental building blocks that he named "quarks." The quarks are permanently confined by forces coming from the exchange of "gluons." He and others later constructed the quantum field theory of quarks and gluons,

called "quantum chromodynamics," which seems to account for all the nuclear particles and their strong interactions.

Mendel did not claim that the only logically possible explanation of the sweet-pea data was that genes exist; Gell-Mann did not claim that the only logically possible explanation of the particle-physics data was that quarks exist; and a cognitive neuropsychologist would not claim that the only logically possible explanation of a particular double dissociation is that a certain cognitive architecture exists.

Despite the errors made by the two above mentioned speakers, these kinds of objections do crop up quite a lot as a response to work based on double dissociation evidence. A cognitive neuropsychologist might describe data documenting a double dissociation, then describe a cognitive architecture which, when damaged in two different ways, would yield this double dissociation, and conclude by saying: "So I am proposing that this is the cognitive architecture of the system." To which someone replies, appearing to believe that a devastating criticism is being made: "That doesn't have to be the case. There could well be another quite different architecture that would also yield the same double dissociation."

There are two effective responses that the cognitive neuropsychologist can make here. The first is just to say "I agree." A puzzled silence will follow.

The second is to say "Tell me what this other architecture is." If the critic cannot do so, the criticism loses all force. If the critic can do so, then we have two different hypotheses about what the functional architecture is, both of which can explain the data; so then we need new empirical work which will allow us to decide which hypothesis is to be preferred—for example, new data which are compatible with one of the hypotheses and inconsistent with the other.

SYNDROMES, THE SINGLE-CASE-STUDY APPROACH, AND COGNITIVE NEUROPSYCHOLOGY

A syndrome is a collection of symptoms which often co-occur in the same individual. As mentioned above, the symptoms finger agnosia, dyscalculia, left–right disorientation, and impaired spelling with intact reading co-occur sufficiently often that this conjunction of symptoms has a name: Gerstmann's syndrome. One can ask: Why do these symptoms co-occur? One possible answer is that there is some cognitive module which plays a role in finger identification, calculation, left-right orientation, and spelling. Damage to that module would therefore result in all four symptoms; and in that case, studying patients with Gerstmann's syndrome would provide information about the functional architecture of some cognitive system. However, there are no currently-proposed theories of any cognitive processing system which contain a module that plays a part in the performance of all four of these tasks, and it is very hard to conceive of what such a theory (and module) would be like. Hence, we are led to seek another answer to the question of why these symptoms co-occur, an answer in terms of neural rather than functional architecture; namely, that these tasks depend on cognitive modules which are functionally distinct but realized adjacently in the brain (in the left parietal cortex, as it happens). In that case, studying Gerstmann's syndrome can tell us nothing about functional architectures of cognition.

Cognitive neuropsychologists believe that the kind of argument made in the previous paragraph regarding Gerstmann's syndrome applies in general to all syndromes. So they do not engage in the investigation of syndromes, and the reason they don't is because studying a syndrome will not yield information about functional architectures of cognition. Consider the syndrome known as Broca's aphasia, the typical symptoms of which are halting, effortful, nonfluent speech, impoverished grammatical structure in sentence production, especial difficulty with grammatical function words; and an especial difficulty with prefixes and suffixes. If all of these symptoms arose from a single impairment of the language-production system, we would need to propose a functional architecture for that system which contained one element (box or arrow) which, when impaired, would cause the system to show every one of these symptoms. This very idea was proposed by Kean (1977) and by Zurif and Caramazza (1976). Kean proposed that all of the symptoms of Broca's aphasia arise because of damage to a system needed for processing morphemes that do not carry stress. Zurif and Caramazza proposed that all of the symptoms of Broca's aphasia arise because of damage to a syntactic parsing module of the language system.

Subsequent studies, however, showed not only that neither of these theories of Broca's aphasia are correct, but also that no theory which offers an explanation for all of the symptoms of Broca's aphasia in terms of damage to a single component of the language-processing system could be correct. That is because, even though these symptoms commonly co-occur, they do not invariably co-occur; that is, they dissociate. Symptoms which dissociate cannot of course have a single common cause.

That is not to say that the study of syndromes has no part at all to play in cognitive neuropsychology. To see how this is so, consider the modern history of the cognitive-neuropsychological study of acquired dyslexia. This began 25 years ago with the seminal paper of Marshall and Newcombe (1973). In that paper the authors made two important contributions: They defined three syndromes of acquired dyslexia (surface dyslexia, deep dyslexia, and visual dyslexia), and in relation to a theory of normal skilled reading they interpreted each syndrome as arising from a particular pattern of preserved and impaired components of that theory. This revolutionized the study of acquired dyslexia. Detailed studies of many cases of deep dyslexia (for a review see Coltheart, Patterson, & Marshall, 1980) and of surface dyslexia (for a review see Patterson, Marshall, & Coltheart, 1985) followed; and other syndromes of acquired dyslexia were defined and investigated, such as phonological dyslexia (Beauvois & Derousne, 1979) and letter-by-letter reading (Patterson & Kay, 1982).

But then what happened next in the cognitive-neuropsychological study of acquired dyslexia was that syndrome-oriented research was gradually abandoned. The motivation for abandoning the syndrome approach was as follows:

> The concept of the syndrome has been a useful one in developing work relating (acquired) dyslexic syndromes to theories about reading. However, its usefulness is likely to be short-lived. The reason is that, if a dyslexic syndrome is a specific pattern of preservations and impairments of reading abilities . . . and if a modular theory of reading is appropriate, it follows that there are many different possible dyslexic syndromes. Any unique pattern of impairments to the boxes and arrows of (the theory) will produce a unique syndrome; since (the theory) has enough boxes and arrows to produce a large number of different unique patterns of impairments, it generates a large number of different syndromes. (Coltheart, 1984a, p. 370)

Following this logic (and see also Marshall, 1984), Howard and Franklin (1988) pointed out that, since their theory of language-processing contained 27 components (boxes or arrows), the number of possible syndromes according to this theory is $2^{27}-1$, which is a large number (2,220,075). If there are several million possible syndromes, the concept of syndrome will not be a useful one. Reasons for abandoning the concept of the syndrome were further spelled out by Caramazza (1984), Coltheart (1987), and Ellis (1987).

But if so, how does one generalize the findings from one patient to findings from another, if the two patients are not being treated as representative examples of the same syndrome?

> The generalizations do not take the form of claiming that there exists a single syndrome which many patients exhibit. Instead these generalizations take the form of claiming that there exists a single theory of the relevant cognitive system which can offer interpretations of the various sets of symptoms exhibited by various different patients. (Coltheart, 1984b, p. 6)

Hence,

> Even if every (dyslexic) patient exhibited a unique reading disorder, it might still be possible to interpret every patient's behaviour in the context of a single theory for reading. The assumption that a single theory should be applicable to all patients allows each new patient to be an appropriate source of data for testing the theory; and this permits one to generalize from previous to future patients even if one has rejected the policy of thinking in terms of syndromes. (Coltheart, 1984a, p. 371)

The moral is: When one is dealing with a domain of cognition about which nothing cognitive-neuropsychological is currently known, it is likely to be profitable to begin by seeking to define syndromes within that domain, just as Marshall and Newcombe (1973) did for reading. This will be a useful initial ground-clearing exercise which will provide the cognitive neuropsychologist with ideas about what kinds of distinctions are of relevance to this domain (words vs. nonwords and regular words vs. irregular words, for example). Such ideas can be used to develop a modular model of the functional architecture of the system responsible for performance in this cognitive domain. That's the point at which the syndrome approach has outlived its usefulness. From that point on, the job is to use data from individual patients, not groups of patients, to test that model and any other models which are also formulated.

COMMON CONCERNS

In this section I will discuss four worries about cognitive neuropsychology that people often articulate.

"Aren't these kinds of cases very rare?"

Certainly. But why is this a problem? What seems to be behind this question is the suggestion that, because a particular pattern of cognitive impairment is very rarely seen in the neuropsychology clinic, it is somehow unrepresentative of the population in general—that is, conclusions about how cognitive processing systems are organized in all of us that are reached on the basis of rarely-seen patterns of impairment are unsafe. However, such conclusions are safe if we make the following assumption: that the architecture of cognition is constant across people. This is the assumption of uniformity of functional architectures which was discussed earlier, where it was pointed out that *all* of cognitive psychology, not just cognitive neuropsychology, relies upon this assumption.

Naturally there will be individual differences in the contents and perhaps even the efficiencies of particular cognitive modules; all that is being assumed is that normal individuals who have been exposed to the same environments will possess the same set of cognitive modules.

"The case you reported is not a pure case; he had more than one impairment of the cognitive system in which you are interested"

One criticism levelled at the work of Marshall and Newcombe (1973) was that their cases were not pure examples of the reading syndromes they were defining. Take, for example, the dual-route theory of reading discussed earlier in this chapter. As was pointed out earlier, a person with an impairment affecting just of the orthographic lexicon in such a theory would be normal at nonword reading while being poor at reading irregular words: that is surface dyslexia. But the first surface dyslexics described in the literature (by Marshall & Newcombe, 1973) also had some impairment of nonword reading, though this was far less marked than their impairment of exception word reading. There are two ways of responding to this.

The first approach is to take the impurity of these cases as evidence against the theory (since the theory predicts that pure cases should exist) and even as evidence against the usefulness of the syndrome approach.

The second approach is to take the view that a brain lesion will in most cases affect a number of different cognitive processing systems (those whose anatomical loci are close together) even when the systems themselves are functionally distinct. Hence it will be rare to find patients whose lesion affects the lexical reading route but completely spares the nonlexical reading route. But it should be possible to find such patients, given enough patients, and patience.

This second approach turned out to be the correct one. Eventually, surface dyslexic patients with very impaired irregular word reading but normal nonword reading were found (Bub, Cancelliere, & Kertesz, 1985; McCarthy & Warrington, 1986)—that is, pure cases of surface dyslexia.

"If every patient is unique, how can you replicate your results?"

Consider what replication means in cognitive psychology. Suppose a cognitive psychologist reports, say, that high-frequency words yield shorter latencies than low-frequency words in a reading-aloud experiment. Another cognitive psychologist, interested in this finding, wants to be sure that it is true in general, rather than being confined just to the subjects, or just to the stimuli, used by the first investigator. So the second investigator chooses comparable but new sets of high-frequency and low-frequency words, and a comparable but new set of subjects, and carries out a reading-aloud experiment with these.

Now suppose a cognitive neuropsychologist were to report a patient with a severe impairment of semantic knowledge who was nevertheless normal at picture and object naming. This is something that has never been reported, so that, if it ever is, it will certainly immediately attract the attention of other cognitive neuropsychologists. The reason for this is that all current theories about how pictures and objects are named propose that the *only* route from a picture or object to its name is via its semantic representation. If that is so, every patient with an impairment of the semantic system must also have an impairment of picture and object naming. So the discovery of a patient with a semantic impairment but no impairment in naming pictures or objects would refute all existing theories of how pictures are named.

Any scientist would want to replicate a finding which is inconsistent with all relevant existing theories. So it is essential to consider what might count as replication in this context. The analogue here to replication as it is done in cognitive psychology might be thought of as: choosing a new but comparable patient and a new but comparable set of pictures or objects, and carrying out a naming experiment with these. But what is meant here by "comparable patient"? Certainly not just another patient with a semantic impairment. If one such patient were chosen, and proved completely unable to name pictures or objects, this would not be taken as a failure to replicate the study with the first patient, because the conclusion drawn from that first study was *not* that all patients with semantic impairments will have intact picture and object naming. It was that *some* patients with semantic impairments can have intact picture and object naming.

In fact, replication in this sense may sometimes be impossible in cognitive neuropsychology. Let's imagine that the relevant cognitive system is such that there is a pathway from pictures and objects to their names which bypasses semantics (that's why it is possible to name pictures and objects normally even when the semantic system is impaired). The neural realizations in the brain of this pathway and of the semantic system might be such that it is almost impossible for damage to the semantic system to spare the picture/object-to-name pathway. It happened with the first patient referred to here but may never happen again, in which case the result is literally unreplicable, no matter how genuine.

The way out here is to appeal to the assumption of uniformity of functional architectu across people. If one accepts that it has been demonstrated that this first patient possess pathway from pictures/objects to names that bypasses semantics, and if the uniformity a sumption is made, then it follows that all people have such a pathway, even if no subsequen patient is ever seen who is normal at picture and object naming but has a semantic impai ment. Anyone who wanted to challenge the view that there is a pathway from pictures/obje to names that bypasses semantics in all people would therefore either have to abandon th uniformity assumption or seek to show that there was something unreliable about the study o which this claim was based.

And here there's a different type of replicability which is relevant—within-patient rather tha across-patient replicability. If a cognitive neuropsychologist noticed that, in a set of words give to a patient to read, words of Romance origin were read much better than words of Germani origin, where the words were not chosen with that variable in mind, it would be poor science to take this as evidence that the reading system's performance is etymologically sensitive. At a minimum, one would select two sets of words, one set for each type of origin, matched on all the variables that are known to affect reading by people with acquired dyslexia, and administer these to the patient. If the remarkable etymological effect still occurred, a prudent investigator would anticipate scepticism and seek to head it off in advance by further studies seeking to replicate the effect with different word sets. While it is certainly true that spectacular findings of this kind in cognitive neuropsychology may not always have been adequately supported in this way, there is nothing about the discipline of cognitive neuropsychology which prevents such within-patient replications from being carried out, and they often are.

"Aren't these kinds of theories too powerful? Aren't they able to explain anything?"

This objection rests upon an ambiguity of the concept of "explaining everything." A theory which can explain *all logically possible data* is too powerful, because there is no possible observation which could refute it. But a theory which can explain *all so-far-observed data* is not too powerful; on the contrary, it is exactly what we want. Cognitive neuropsychologists grow their box-and-arrow diagrams because each new box or arrow is compelled by some new datum; to put this another way, if any box or any arrow were removed, some piece of data would no longer have an explanation. Thus the boxes and arrows are not added by whim; they are demanded by data.

And, in any case, the kinds of theories current in cognitive neuropsychology cannot explain all logically possible data. For example, the dual-route theory of reading aloud discussed elsewhere in this chapter would be instantly refuted by the discovery of a patient who is more successful at reading nonwords than regular words, because there is simply no way of damaging any component or components of that theory which would result in an impaired system that read nonwords better than regular words. And as discussed above, the theory of picture and object naming according to which mediation via semantics is obligatory, a theory which is currently accepted by all, would be instantly refuted by the discovery of a patient with intact picture and object naming but with a semantic impairment.

SUMMARY

Modular theorizing is widespread in contemporary cognitive psychology. In very many cognitive domains, cognitive psychologists develop theories about the system via which the cognitive activity in question is performed by proposing a modularly-organized functional architecture for that system. Such cognitive psychologists generally assume not only modularity, but also

architectural uniformity: that is, they assume that any such architecture is uniform across all the people they might study (apart from merely quantitative individual differences, such as differences in vocabulary size, or differences in capacity of working memory, for example).

Any theory of this kind can be effectively tested by carrying out studies of people with acquired or developmental disorders of the relevant domain of cognition (that is, can be effectively investigated by cognitive neuropsychology), provided that a third and a fourth assumption are made. It needs to be assumed that the functional modules are also anatomically modular (allowing them to be susceptible to dissociations after brain damage), and it needs to be assumed that brain damage or abnormal development cannot add new modules, or new pathways of communication between modules: Abnormality must only take the form of impairing or deleting modules or pathways of the normal system. Unless both assumptions were at least approximately correct, nothing coherent would have emerged from cognitive-neuropsychological research—which is by no means the case, as the other chapters in this volume so clearly demonstrate.

REFERENCES

Beauvois, M. F., & Derouesné, G. (1979). Phonological alexia: three dissociations. *Journal of Neurology, Neurosurgery, and Psychiatry, 42*, 115–124.

Bub, D., Cancelliere, A., & Kertesz, A. (1985). Whole word and analytic translation of spelling-to-sound in a non-semantic reader. In K. E. Patterson, M. Coltheart, & J. C. Marshall (Eds.), *Surface dyslexia: Neuropsychological and cognitive studies of phonological reading* (pp. 15-34), London: Lawrence Erlbaum.

Caramazza, A. (1984). The logic of neuropsychological research and the problem of patient classification in aphasia. *Brain & Language, 21*, 9–20.

Coltheart, M. (1984a). Theoretical analysis and practical assessment of reading disorders. In Cornoldi, C. (Ed.), *Aspects of Reading and Dyslexia*. Padua, Italy: Cleup.

Coltheart, M. (1984b). Acquired dyslexias and normal reading. In Malatesha, R. N., & Whitaker, H. A. (Eds.), *Dyslexia: A Global Issue*. The Hague: Martinus Nijhoff.

Coltheart, M. (1987). Functional architecture of the language-processing system. In Coltheart, M., Sartori, G., & Job, R. (Eds.), *The cognitive neuropsychology of language*. London: Lawrence Erlbaum.

Coltheart, M. (1999). Modularity and cognition. *Trends in Cognitive Sciences, 3*, 115–120.

Coltheart, M., Inglis, L., Cupples, L., Michie, P., Bates, A. & Budd, B. (1998). A semantic subsystem specific to the storage of information about visual attributes of animate and inanimate objects. *Neurocase, 4*, 353–370.

Coltheart, M., Patterson, K., & Marshall, J. C. (Eds.) (1980). *Deep dyslexia*. London: Routledge & Kegan Paul.

Ellis, A. W. (1987). Intimations of modularity, or, the modelarity of mind: Doing cognitive neuropsychology without syndromes. In Coltheart, M., Sartori, G. & Job, R. (Eds), *The cognitive neuropsychology of language*. London: Lawrence Erlbaum.

Fodor, J. A. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.

Howard, D., & Franklin, S. (1988). *Missing the meaning? A cognitive neuropsychological study of the processing of words by an aphasic patient*. Cambridge, MA: MIT Press.

Kay, J., Lesser, R., & Coltheart, M. (1992). *PALPA: Psycholinguistic Assessments of Language Processing in Aphasia*. Hove, UK: Lawrence Erlbaum.

Kean, M-L. (1977). The linguistic interpretation of aphasic syndromes. *Cognition, 5*, 9–46

Marshall, J. C. (1984). Towards a rational taxonomy of the acquired dyslexias. In Malatesha, R. N. & Whitaker, H. A. (Eds.), *Dyslexia: A Global Issue*. The Hague: Marinus Nijhoff.

Marshall, J. C., & Newcombe, F. (1973). Patterns of paralexia: A psycholinguistic approach. *Journal of Psycholinguistic Research, 2*, 175–199.

McCarthy, R., & Warrington, E. K. (1986). Phonological reading: Phenomena and paradoxes. *Cortex, 22*, 359–380.

Patterson, K. E., & Kay, J. (1982). Letter-by-letter reading: Psychological descriptions of a neurological syndrome. *Quarterly Journal of Experimental Psychology, 34A*, 411–441

Patterson, K. E., Marshall, J. C., & Coltheart, M. (Eds.) (1985). *Surface dyslexia: Cognitive and neuropsychological studies of phonological reading*. Hove, UK: Lawrence Erlbaum.

Rapp, E,, Benzing, L., & Caramazza, A. (1995). The modality-specific representation of grammatical category. Paper presented at the 36th Annual Meeting of the Psychonomic Society, Los Angeles, CA

Riddoch, M. J., & Humphreys, G. W. (1993). *BORB: Birmingham Object Recognition Battery*. Hove, UK: Psychology Press.

Sartori, G., & Job, R. (1988). The oyster with four legs: A neuropsychological study on the interaction of visual and semantic information. *Cognitive Neuropsychology, 5*, 105–132

Shallice, T. (1988). *From Neuropsychology to Mental Structure*. Cambridge, UK: Cambridge University Press

Zurif, E., & Caramazza, A. (1976). Psycholinguistic structures in aphasia: Studies in syntax and semantics. In T. Whitaker & H. A. Whitaker (Eds), *Studies in Neurolinguistics, Vol 2*. New York: Academic Press.