DOES COGNITIVE NEUROPSYCHOLOGY HAVE A FUTURE?

Reflections stimulated by Rapp, B. (Ed.) (2001). *The handbook of cognitive neuropsychology: What deficits reveal about the human mind*. Philadelphia: Psychology Press.

Trevor A. Harley
University of Dundee, UK

INTRODUCTION

Any book called *The Handbook of Something* (compared with a book with a less ambitious title, such as *Explorations in something* or *Some current research on something*) sets itself a lofty and difficult target. The title suggests that the book is a definitive work: a comprehensive picture of the state of the art of modern cognitive neuropsychology. Furthermore, "handbook" implies that it will tell the reader how to do something, or reveal how something works—in this case, according to the subtitle, it will tell the reader what deficits reveal about the human mind, and by implication a great deal about how the mind works. How well does Rapp's *Handbook of Cognitive Neuropsychology* (hereafter simply *HCN*) match up to these noble aims?

There are at least two methods of assessing any handbook's success. The first method is the pragmatic one of comparing the material covered by the book with the actual content of the subject as it is currently practised. Given the prominence of the editor of the *HCN* and its contributors in their field, such a comparison will be particularly illuminating because it compares what researchers are actually doing with what they think—or at least, what they claim—that they are doing. The second method is to compare the contents of a handbook against the espoused theoretical and methodological aims of the subject. This comparison is interesting because it will reveal any divergence between the current practice of a subject and its theoretical goals and methodological foundations. Perhaps what cognitive neuropsychologists say they are doing, and what they are actually doing, are two very different things. Furthermore, any divergence might call for a fundamental reassessment of those goals and foundations.

COGNITIVE NEUROPSYCHOLOGY: A SNAPSHOT

To compare the contents of the *HCN* with how cognitive neuropsychology is currently practised,

---

3

we need a snapshot of the subject. I have taken such a picture by surveying articles published in the journal *Cognitive Neuropsychology* over the last 4 years. The reasons for this decision should be obvious: This journal publishes top-quality research by top-quality researchers in and only in the area of cognitive neuropsychology. *Cognitive Neuropsychology* defines the subject as it is practised. Four years provides a reasonably sized sample that just overlaps with the contents covered by the *HCN* (it cites many papers from the late 1990s): Obviously both journal and book suffer some publication lag, although it is probably larger for a book. The last 4 years of *Cognitive Neuropsychology* cover Volumes 15–18 (1998–2001).

How does the *HCN* survey the area? It contains 24 chapters and a useful glossary. Table 1 shows the content of the *HCN*, divided up into subject areas. These areas are mainly each covered by an individual chapter, although in a couple of cases, where it is appropriate, I have combined chapters into a superordinate category. I have classified each of the articles in *Cognitive Neuropsychology* published during the census period using the same scheme. Some articles cover more than one topic, and a few cover three.

Clearly there are some significant differences between what is presented in the book and what is published in the journal. Compared with what the *HCN* might lead us to believe, researchers are obsessed disproportionately with reading and semantics. Hence the *HCN* is more *comprehensive* than the subject it sets out to portray, but less *representative*. Such a conclusion is not so much a criticism of the book as of the area, but anyone reading the *HCN* will gain a distorted view of what modern cognitive neuropsychology is about. In practice it is a narrow subject: Large expanses of human behaviour remain essentially unexplored.

There are several reasons why this narrowness might have arisen. First, researchers understandably prefer more tractable problems. For example, in the psycholinguistics literature there is vastly more research on normal comprehension than normal production because it is easier to carry out experiments on comprehension. Similarly, reading deficits are easier to study than deficits in musical

Table 1. *The number of chapters and articles on particular topics[a] in cognitive neuropsychology in Rapp's (2001)* Handbook of Cognitive Neuropsychology *(* HCN, *column 2), all articles in the journal* Cognitive Neuropsychology *(column 3, labelled* CN*), and all articles in* Cognitive Neuropsychology *excluding reviews, those in special issues, and replies (column 4,* CN *base)*

Topic	HCN	CN	CN (base)
Theoretical foundations	3	1	0
Object recognition & perception	1	12	8
Faces	1	23	4
Spatial cognition	1	5	5
Attention	2	7	7
Lexical architecture	1	1	1
Morphology	1	0	0
Reading	1	25	16
Writing and spelling	1	5	5
Spoken word production	1	11	10
Bilingualism	1	0	0
Sentence comprehension	1	2	2
Sentence production	1	0	0
Memory	2	21	4
Semantics	1	16	13
Time perception	1	0	0
Mathematics	1	4	2
Music	1	1	1
Action	2	14	7

[a] In a few cases, topics are combined where appropriate (e.g., the theoretical and review chapters, and attention and consciousness, which are the two chapters of the attention section).

ability. It is more convenient, and easier to control extraneous variables, if you use lists of monosyllabic words rather than multisyllabic ones for your materials. Tractability is related to the second problem, which is that researchers are necessarily limited by the availability of interesting patients. "Interesting" here means novel, theoretically important, demonstrating clear-cut deficits, or any combination of these. There is also a skill in spotting theoretically interesting behaviour, and in having the ability to test and to have access to patients with appropriate aetiologies. Nevertheless, there is a large degree of luck in that the right patient has to turn up in the right place at the right time. This factor introduces a systematic bias into cognitive neuropsychology. Third, many researchers are naturally attracted to grand topics—the best example of which is

semantics. Finally, after a while, topics take on a life of their own. Perceptual and functional features have come to trouble a generation.

Narrowness is revealed not just in the topics that fascinate neuropsychologists, but also within those topics. This specialisation in some areas—and neglect of others—is reflected by the *HCN*. For example, Farah's chapter on consciousness is, perhaps necessarily, exclusively about forms of visual perception without awareness, and does not touch upon many of the topics that most psychologists and philosophers would say are relevant to consciousness (e.g., qualia, and why they exist; why there is a problem about consciousness at all; materialism; epiphenomenalism—although it is unclear how neuropsychology could address these issues). Studies of reading often focus disproportionately on monosyllabic words. Mathematical cognition focuses on number recognition, arithmetical facts, and simple calculation. Of course, this narrowness is largely motivated by the tractability issue raised above, but the result is that the range of human behaviour studied is even further restricted.

Although bilingualism and second language acquisition are rarely studied by neuropsychologists, a range of languages are covered in the survey period in *Cognitive Neuropsychology*: one Chinese, probably four French, one German, one Hebrew, and probably six Italian case studies, with the rest presumably English (although annoyingly, this is not always explicitly stated, and the first language of the patient is not always obvious). Such multiculturalism is good to see, but often just reflects the nationality of the researchers. It has, of course, long been realised that different languages impose different processing constraints (e.g., on parsing; see Cuetos & Mitchell, 1988), and offer different opportunities for researchers to study the mind (Japanese *kanji* and *kana* perhaps being the most famous example). Nevertheless, English remains the language of choice of neuropsychologists, and much might be lost by this narrowness.

In summary, cognitive neuropsychology is in practice a very narrow, specialist subject. If you doubt this, look through an introductory psychology text to see what impact it has had.

THE THEORETICAL GOALS AND METHODOLOGICAL FOUNDATIONS OF HUMAN COGNITIVE NEUROPSYCHOLOGY

Who first put the cognitive into cognitive neuropsychology? Although the general approach can be traced back to the nineteenth century (Shallice, 1988), and the techniques that are now recognised to be those of cognitive neuropsychology were first applied at least in the 1970s (e.g., Marshall & Newcombe, 1973), the explicit promulgation of the goals and methods of the subject can be attributed to several pioneers (e.g., Caramazza, 1984a, 1986; Coltheart, 1985; Ellis & Young, 1988; Shallice, 1988).

It is worth reconsidering cognitive neuropsychology's early manifesto. There were two fundamental and mirror-image (bordering on circular) aims (Coltheart, 1985; Ellis & Young, 1988):

1. To explain the patterns of impaired and intact performance seen in people with brain damage in terms of damage to one or more components of a model of normal cognitive functioning.

2. To draw conclusions about normal, intact cognitive processes from the patterns of performance observed in people with brain damage. Drawing such conclusions necessitates belief in the *transparency assumption*, whereby it is assumed that the cognitive system of a brain-damaged person is the same as that of a normal participant except for the local modification caused by the brain damage (Caramazza, 1986).

Furthermore, most (but by no means all) of the pioneers argued that traditional group-based experiments using brain-damaged participants are not informative as single-case studies, and may indeed give misleading results. There were several reasons for this concern. For example, it may be difficult to define any theoretically meaningful group, averaging across participants may conceal interesting differences among them, and it is illegitimate to extrapolate from a group of

brain-damaged participants to normal participants[1] (see footnote 2 for references). Hence the proper unit of analysis was deemed to be the single-case study. This methodological stricture constitutes a third defining characteristic of cognitive neuropsychology.

3. To make inferences about the processes underlying human behaviour using single-case studies.

In the early days there was some debate about the role of the brain in explaining behaviour and about the importance of the localisation of function. Positions taken by researchers varied from those who argued that the localisation of psychological function in the brain is of no interest (Morton, 1984), to those who believed that localisation is of interest but is distinct from the evaluation of the model as a cognitive theory (Ellis & Young, 1988), or that at the time it was of limited usefulness (Shallice, 1988). In each case, the role of the brain was downgraded compared with traditional neurology. Hence we can identify a fourth, more controversial, characteristic of the cognitive neuropsychology programme.

4. Knowing where a psychological process happens in the brain is, at best, presently of limited psychological interest, and indeed it may tell you nothing.

Taken together, these four points identify the original research programme of cognitive neuropsychology. Although all but the fourth may seem relatively uncontroversial (or at least, acceptable) now, in the 1980s they were sufficiently radical to generate some vehement debates, particularly on single-case versus group studies[2].

Shallice (1988) christened the approach that rejects group studies and localisation data as inappropriate sources of evidence for theorising about the mind *ultra-cognitive neuropsychology*. He himself withdrew from the brink of condoning the behaviour of ultra-cognitive neuropsychologists, arguing that in principle some group studies might provide useful data, and that anatomical considerations should not be rejected as irrelevant, although both approaches are beset with practical and theoretical differences that limit their usefulness, at least at that time. It is worth noting, however, that the theorising in Shallice's book is based on very few group studies, and hardly any localisation data.

Methodological debate in the 1980s centred on single-case studies, modularity, and inference from brain-damaged people to non-brain-damaged people (essentially the *subtractivity assumption*— that brain-damaged behaviour is normal behaviour minus the components that have been damaged). Cognitive neuropsychologists would say that the proof of the pudding is in the eating: Get on with the job, and see what we have discovered. And one only has to look at the *HCN* to say that we appear to have learned an awful lot. Nevertheless, criticisms of cognitive neuropsychology persist in the literature, now focusing particularly on the *seriality assumption*. The seriality assumption states that cognitive processing is best described in terms of serial processing, modular information processing models of the sort encapsulated in "box diagrams." Given the contemporary emphasis on the parallel nature of cognitive processing, driven particularly by the advent of connectionism, it is clear that these models and this assumption are still, at best, controversial. Parallel processing systems are less easy to explore using brain-damage data than serial, modular flowchart systems (Kennedy, 1996), so traditional cognitive neuropsychology might be backing the wrong horse. It must be noted, however, that the move to computational cognitive

---

[1] Although, as Shallice (1988) points out, taking Caramazza's (1986) argument to its conclusion means that we should be wary of drawing conclusions about normal functions from group studies with normal participants!

[2] One could fill a paper with citations, but at the risk of boring older readers, for examples see Badecker and Caramazza, 1985, 1986; J. Bub and Bub, 1988; Caplan, 1986, 1988, 1991; Caramazza, 1984a, 1986; McCloskey, 1993; Robertson, Knight, Rafal, and Shimamura, 1993; Schwartz, 1984; Sokol, McCloskey, Cohen, and Aliminosa, 1991; Zurif, Gardner, and Brownell, 1989; Zurif, Swinney, and Fodor, 1991; see Bates, McDonald, MacWhinney, and Appelbaum, 1991, versus Caramazza, 1991, for a particularly vituperative debate.

neuropsychology—particularly connectionist cognitive neuropsychology—makes it easier to explore the neuropsychological processes of parallel systems. The use of connectionist cognitive neuropsychology also means that we focus more than we might otherwise do on what is happening inside the boxes, and have to make explicit representational and processing assumptions, as advocated by Seidenberg (1988). The end result is that neuropsychologists are no longer susceptible to the accusation that lesion-correlation studies might be restricted to uncovering the functional architecture (Saffran, 1982). Nevertheless, in spite of Seidenberg's exhortation, and the clear advantages of computational modelling, discovering the functional architecture is still the most pressing concern among cognitive neuropsychologists, and connectionism is widely, and incorrectly, disregarded—a point to which I will return.

Early articles in the journal *Cognitive Neuropsychology*, with a few exceptions, followed the manifesto of ultra-cognitive neuropsychology. They reported single-case studies, mentioned little localisation data, and made little of what they did mention, and contained few neuroimaging photographs (although of course these were also less readily available).

Recent years have seen something of a change in the types of article published in *Cognitive Neuropsychology*. Some of the changes may reflect changes in editorial policy, but there is no reason to assume other than that they reflect changing interests and attitudes by researchers. In my survey, not all the papers in *Cognitive Neuropsychology* report single-case studies. (I include as single-case studies papers that describe or compare more than one person, but that provide detailed background, analysis, and interpretation of each individual.) Seven papers report what are best described as traditional group experiments, with up to 108 participants. Of these seven group experiments,

three are studies of normal participants (that nevertheless clearly cast light on neuropsychological issues). Furthermore, 26 primary articles report detailed localisation data, in that they contain a picture of a brain scan of some sort. Ten articles have as a major focus the localisation of functions in the brain, or the establishing of neural pathways. Finally, for obvious reasons, early articles contain little mention of computational modelling, while more recent issues contain connectionist models and evaluate data against models.

Hence there has been a methodological and theoretical shift in the practice of cognitive neuropsychology. The early methodological assumptions have been in practice abandoned (perhaps revealing the danger of a programmatic approach to science—researchers just ignore the methodological strictures if what is forbidden appears to work). Of most theoretical interest are the new emphasis on localisation, and the growth of connectionist modelling. I will look at each of these in more detail.

SHOULD THERE BE ANY "NEURO" IN COGNITIVE NEUROPSYCHOLOGY?

I have only a vague understanding of how my car works. I do know where the engine, battery, oil reservoir, and windscreen liquid reservoir all are; I even know where some parts of the engine are. Knowing all this, however, is of no use in understanding *how* the car works. Other people have used the analogies of calculators, cars, chocolate factories, computers, radios, television sets, vacuum cleaners, or even what smoke coming out of a factory tells you about what is being made inside, to the brain[3]. All these authors make the point that knowing where the bits of a thing are doesn't tell you much about how the thing works. Sceptics of

---

[3] The sources of these analogies are as follows: cars (Marin, Saffran, & Schwartz, 1976), chocolate factories (Coltheart, in press), computers (many, including Johnson-Laird, 1983), multiplication devices (calculators, we might call them now; Deutsch, 1960), radio and television sets (Gregory, 1961), vacuum cleaners (Marin et al., 1976). I would be interested to hear of any other analogies I might have missed.

neuropsychology also make the point that randomly smashing a car (or car, calculator, or whatever) with a hammer will not tell us anything interesting about how the car (or whatever) works[4]. On the other hand, some (and an increasing number) of neuropsychologists appear to believe that the (neural) localisation of (psychological) function is of more than neurological interest: They believe that it can inform psychological models.

Clearly the identification of distinct neural structures in the brain in itself can tell us nothing about cognition. Studying the brain can only make a contribution to understanding the mind when we can relate the two in some way. Modern imaging techniques promise a way of mapping this relationship. It is useful to distinguish between two types of imaging: That which tells us where and that which tells us when. "Where" imaging includes CAT, PET, and NMR scans and emphasises the location of processes in the brain. "When" imaging uses ERP techniques to explore the time course of processing. The emphasis is relative. "Where" imaging can build up a picture of the supposed location of processing over time, although at present temporal resolution is not very good, and "when" imaging depends on the location of electrodes. "Where" imaging is more prevalent, but "when" might be more useful.

Imaging the localisation of function of components of anything—cars, brains—can be described at four levels.

1. The "tokenism" level. A glossy magazine might print a picture of the inside of a car because it looks good and is fun to do. Sometimes brain images appear to be included just because they are available. They add nothing to the argument of the paper. Imaging devices are fun to use, are expensive, and require big grant money to obtain and maintain. No doubt all this adds to their popularity.

2. The "where things are" level. For example, the engine is on the left side of engine compartment. Even this level of description, of course, assumes that one can identify the engine: We have to know beforehand what the engine does. After tokenism, this is the most elementary and least interesting type of localisation knowledge.

3. The "relative location" level. For example, the engine is above the axle. This level is more interesting, because it might tell us how things are connected. It still suffers from the same problems as level 2, and in addition is subject to a new assumption, *proximity*. This assumption states that adjacent regions of a complex machine do related tasks. While plausible some of the time, this assumption is clearly often wrong (e.g., the engine and wheels are some way away yet the point of the engine is to drive the wheels; the engine and windscreen reservoir might be adjacent, yet they are as functionally unrelated as can be found in a car). Similarly we cannot assume that because regions of the brain are near to each other, they do similar things.

4. The "functional" level. For example, the engine drives the axles, which in turn are connected to and drive the wheels. Car manuals provide clear diagrams showing where the components are, what each does, and how they cohere to create a machine that fulfils a particular function. The diagrams explain why components with particular functions are in the places that they are. The problem here is that you already need to know mostly what each component does before you can make sense of it. In addition, the absolute location of components is rarely critical to the workings of the system: what is more important is the pattern of connectivity.

Recent issues of *Cognitive Neuropsychology* are full of brain scans. There has been an explosion of interest in neuroimaging of different sorts as the technology becomes more widespread and more accurate. (Of course, it isn't just the journal

---

[4] The point that random damage to things tells you nothing has been disputed. Marin et al. (1976) point out that random damage to the steering system of a car might reveal that the steering mechanism operates independently of the engine; they also point out that optimistically smashing things up is the preferred method of investigation in nuclear physics.

Cognitive Neuropsychology that does this; as Uttal, 2001, points out, recently most articles in *Science* that have anything to do with psychology report imaging data.) An examination of the articles in *Cognitive Neuropsychology* that include brain scans, looking at the extent to which the images contribute to our understanding of cognitive processes, suggests that most of these papers use imaging at the tokenism and where-things-are levels.

This growth in neuroimagery has generated a minor side industry in the growth of criticism about the techniques and what they tell us. Uttal (2001) has derogatorily dubbed this emphasis on imaging "the new phrenology." Of course, this is not to deny that there have been spectacular advances in the techniques themselves, or that they are of exceptional use in patient diagnosis and treatment. The debate is about the extent to which knowing about where things happen in the brain is informative for cognitive models. There are three types of criticism of neuroimaging.

The first type of criticism concerns the relationship between brain activity and the techniques used in imaging (e.g., McCloskey, in the *HCN*, discusses whether blood oxygenation really does indicate greater cognitive activity). I find this the least interesting type of criticism. Researchers will soon provide a definitive rebuttal of these problems, if they have not already done so.

The second type of criticism concerns statistical and methodological issues, particularly about the logic of the subtraction method for identifying where particular processes take place. Neuroimages tell us where the parts of the brain are that become active while the participant is engaged in certain tasks. Assumptions are made about the subtraction method that enable us to deduce what processes are happening in that part of the brain. In the subtraction method, the pattern of brain activation generated by a task that employs the cognitive component of interest is compared to that generated by a baseline task that activates all but the component of interest. The assumptions of this method (e.g., that we can identify cognitive components of tasks, that we can find appropriate baseline tasks, and that the additional components of tasks do not modify the other components) might not

always be correct, with the consequence that psychological processes are incorrectly identified with neural locations (see, for discussion of problems with this methodology, D. Bub, 2000; Poeppel, 1996; Pulvermuller, 1999; Uttal, 2001; Van Orden & Paap, 1997). These are serious concerns that may not be answerable. It is possible that a conjunction method, which examines the neural correlates of two or more distinct tasks that share only the processing difference of interest, may be superior (Price & Friston, 1997). Nevertheless, even the conjunction method makes some assumptions (e.g., about how tasks are decomposed into their components).

The third type of criticism is most interesting because in principle it vitiates the whole enterprise of looking at localisation in the brain. This criticism poses the problem that brain functions are not localised in brain structures in any meaningful way; or, if they are, we might not be able to deduce the relationship. It is worth looking at this point, which has a lengthy tradition, in more detail. It was indeed first applied to the lesion-correlation methods of neuropsychology. Among the first to make this point were Deutsch (1960), who argued that one could not hope to find any part in a calculator (to rephrase slightly) that corresponded to multiplication by taking it to pieces, and Gregory (1961), who argued that we would never learn how a television set works by pulling bits out of one and seeing what happens.

Other critics have sprung to the defence of ultra-cognitive neuropsychology—and indeed for a neuropsychology-less psychology. One of the most strident campaigners for ultra-cognitive neuropsychology has been John Morton. He has consistently argued that a cognitive level of explanation of behaviour is necessary. Adequate accounts of psychological processes must involve purely psychological constructs. Without a full functional architecture and cognitive theory, the biological level (e.g., information about localisation of function) will tell us nothing about this cognitive level. In particular, one cannot use localisation data to prove or disprove psychological data unless one can show that there is a one-to-one correspondence between brain location and psychological process

(e.g., Mehler, Morton, & Jusczyk, 1984; Morton, 1984; Morton & Frith, 1995). Mehler et al. give the following example. Suppose we find that all grammatical functions were affected by the planum temporale, and that the planum temporale was concerned only with grammar. Thus we might localise the grammar module in the planum temporale. Now a new researcher discovers that the planum temporale is also involved in listening to music. We would be most unlikely to conclude, on the basis of the localisation data, that listening to music involves grammatical computation. Rather, we would probably attempt to redefine the type of computation that we think the planum temporale does—for example, we might conclude that it plays a general role in sequencing. The point is that it is the cognitive analysis that constrains the mapping between the brain location and psychological process.

Others have also questioned whether clearly definable neurological structures can be straightforwardly related to cognitive modules. Farah (1994) argued that the locality assumption—the idea that brain damage to one component of the functional architecture has just local effects—breaks down if the brain is an interactive structure. For example, in an interactive system, damage to the semantic system has consequences for the visual processing and lexical retrieval systems. Interaction of this type, however, enables a connectionist model of naming to explain the co-occurrence of visual, semantic, and lexical phenomena in the breakdown of object naming in dementia (e.g., Tippett & Farah, 1994). Van Orden, Pennington, and Stone (2001) question the logic behind research based on the search for double dissociations. They argue that attempts to postulate modules that can be doubly dissociated are at best dubious. They argue that it is difficult to identify modules, and in particular modules that correspond to neurological structures, because the process is theory-laden. They discuss the example of the apparent double dissociation between lexical and nonlexical reading routes: No "pure" patients have been found who unambiguously demonstrate this dissociation. They also argue that as there is no theory-free method of determining whether any particular person with

deep dyslexia, for example, is an instance of a pure case, there is no theory-free way to decide whether deep dyslexia implies the existence of a lexical module. Coltheart (in press) also questions whether functional modules need correspond to anatomical modules. He argues that cognitive neuroscience is dependent on cognitive neuropsychology in that one can only meaningfully image cognitive processes when we already have a detailed account of what these processes do. We can only properly make sense of neuroimaging data when we localise previously identified cognitive primitives. Furthermore, Coltheart argues, data from imaging studies cannot be used to falsify theories of cognition. We would not reject a dual-route model of reading just because we cannot find two distinct neural structures corresponding to each of the routes: The functional architecture might be mapped onto the anatomical architecture in some complex, nonmodular way. (And we might add that therefore the reverse must also be true: Researchers would not reject a single-route model if we appeared to be able to identify two anatomical routes.) Finally, an additional problem for attempts to localise function is that brain areas that play no functional role in a process may nevertheless become activated when it is taking place (see Farah & Aguirre, 1999). McCloskey, in his concluding chapter, discusses results suggesting that the V1 region and the LGN (lateral geniculate nucleus) become activated during mental imagery. These results might suggest that these regions, and the visual processes with which they are associated, play a functional role in imagery; on the other hand, their activation might just be a side effect—by feedback from connections from higher visual areas. On the other hand, McCloskey argues, if lesion data were to show a consistent relationship between damage to these regions and disruption of imagery processes, then we would be likely to conclude that those regions play an essential functional role. Yet cognitive neuropsychologists are now in danger of giving more prominence to "where" imaging data than to lesion-function correlation data.

The essence of these arguments is that distinct neurological structures need not correspond to

functional modules—indeed, there might not be any modules. To be able to decide whether there are and whether there is any correspondence, you need to have a complete theory of cognition before you begin interpreting images. Hence imaging can, in principle, add nothing new. There is a level of psychological theorising—the cognitive level—that can only be studied at this level, and information from lower levels will tell us nothing about what happens at the cognitive level. In summary, we need a theory of cognition before we can properly understand what is happening at the lower levels.

These arguments are not, however, at a level of logical proof. Therefore I can see three related rejoinders to these accusations. The first rejoinder is that imaging will piggyback on experimental cognitive psychology and neuropsychology, eventually getting to the point where it can stand on its own feet. Eventually we might be able to give a complete mapping between brain structures and cognitive processes, and be able to trace all possible feedback routes. Once again, however, imaging techniques might only be able to do this when we have a complete theory of cognition. The second rejoinder is that we might not be able to dispense with the cognitive level, but that does not mean that the biological level might not be related to it in an interesting way. The third is that as resolution and techniques improve, imagers and neuroscientists will look at more and more specific parts of the brain, and be able to relate those to cognitive processes. As I think of dedicated teams of scientists examining small groups of neurons, I think: rather them than me. Furthermore, the arguments above are irrespective of scale: Whatever the size of the neural structures that we are examining, we will still first need a theory of cognition in order to be able to make sense of them.

The contributors to the *HCN* are just as bemused and divided by this reductionist shift as many of us. Most of the chapters are in the traditional semi-ultra-cognitive mode, describing single-case studies, and with a good sprinkling of boxology diagrams. A few mention anatomy and neuroimaging and there are two brain pictures (in the chapter on memory), although the final impression is that this information contributes little to the overall cognitive story. That is, much of the localisation information is at the tokenism level.

Many of the contributors (and presumably many psychologists) appear to find themselves somewhere between two extremes: that localisation data in principle tell us nothing about how the mind works, and that they are necessary. (In fact, I cannot find any reference saying that localisation data are necessary, perhaps because it is such an implausible claim: A computational device could always be physically realised in some different way.) One can generate thought experiments in which localisation data might adjudicate between alternative hypotheses: Consider, for example, the alternative serial-discrete and cascading-interactive models of lexicalisation (e.g., Dell, 1986; Harley, 1984; Levelt, Roelofs, & Meyer, 1999). Assume that all the methodological problems have been ironed out. Suppose that we could unambiguously localise the part of the brain responsible for semantic encoding and the part responsible for phonological encoding. (It might of course prove impossible even to get to this point.) Then, further assuming we could image the activation of those areas within tens of milliseconds, we carry out the imaging experiment and find that the semantic and phonological areas light up at the same time. Would we then conclude that lexicalisation must be interactive? I doubt that such a demonstration would change the minds of the serialists. First, we would have needed an unambiguous demonstration that the areas involved do just what is thought, and no more. Second, consider what we might find if the serial model were in fact the correct one. The semantic area would light up, and then the phonological area. But what happens when the semantic level is passing information to the phonological area? Then both levels might be active at the same time. Or we might see activation travel from the semantic area, which then fades back into inactivity, along some neural pathway to the phonological area, which then lights up. This approach is implausible. It also ignores complications such as lemmas, the structure of the phonological system, other intervening physical levels that might be of no psychological interest, and so on.

I don't think either camp would resign on the basis of such data.

Consideration of a study by Ainsworth-Darnell, Shulman, and Boland (1998) provides another example of how unlikely it is that researchers will change their minds solely on the basis of localisation data. They concluded (on the basis of clear and convincing ERP data) that different parts of the brain become involved when syntactic and semantic anomalies are present, and therefore that these processes are represented separately. As far as I can tell, no one working in parsing has changed their mind about the modularity of the parsing system on the basis of such data. Instead, researchers have a drawn a distinction between representational and processing modularity. Even though different types of information are stored in different parts of the brain, they might still all be active at the earliest stages of processing (Pickering, 1999).

In the chapter on face recognition, De Haan mentions studies such as that of Clark, Keil, Maisog, Courtney, Ungerleider, and Haxby (1996) showing that the brain area involved in face recognition overlaps with that involved in non-face object recognition, suggesting that at some level the same visual primitives are computed for both types. But it would be extraordinary if there were not overlap at some level. The point here is that psychologists are even less likely to change their minds about psychological level issues on the basis of localisation data than they are on the basis of any other data.

In the final chapter McCloskey considers the possibility that, with the advent of effective functional neuroimaging, traditional lesion-study cognitive neuropsychology has outlived its usefulness. Studies of brain-damaged patients possess several disadvantages: Lesions are often large and usually straddle functionally distinct brain regions, and post-trauma reorganisation complicates interpretation. On the other hand, imaging studies also have their disadvantages (see above). He concludes that imaging and lesion studies should not be seen as competitors, but as complementary approaches. The world remains to be convinced. So far we have learned much from lesion studies, but little from imaging ones.

THE ROLE OF COMPUTATIONAL MODELLING IN COGNITIVE NEUROPSYCHOLOGY

There is now an emphasis on computational modelling that was almost completely absent at the genesis of the discipline. This change is a good one because it means that researchers now focus more on the processes involved. In addition, as noted above, connectionist models address the parallel nature of many psychological processes. Connectionism also appears to begin to bridge the gap between mind and brain, in that complex processes are related to low-level, neural-like processes.

I counted the number of papers on connectionism that were published in *Cognitive Neuropsychology* within the last 5 years. Of the 71 primary articles, just 6 focus on computational modelling, or attempt to fit experimental data to computational models. Hence, according to my index, the average cognitive neuropsychologist is nearly four times as interested in faces as in computational modelling. This lack of interest in connectionism is also reflected in the *HCN*. For example, the Hinton and Shallice (1991) model of semantics, reading, and deep dyslexia is amazingly not mentioned once in the whole handbook. There is also very little on the "triangle model" of reading and its account of phonological dyslexia (Farah, Stowe, & Levinson, 1996; Harm & Seidenberg, 1999; Plaut, McClelland, Seidenberg, & Patterson, 1996; Seidenberg & McClelland, 1989).

Why is connectionism so widely ignored? There are probably many reasons. Those of us who teach neuropsychology know that connectionism is not taught in proportion to its impact (or at least, in proportion to what its impact should be). It might just be too *difficult* for many people. Perhaps connectionism is ignored because most researchers consider its impact to be exaggerated and that journal editors are wrong to give it such prominence. Another reason is the worry that many have: That connectionist models are too powerful—they can explain everything, and can be difficult to falsify. (Indeed, Coltheart notes in the first chapter of the *HCN* that the same accusation has

also been—unjustly—made of cognitive neuro-psychology. Furthermore, Plaut and McClelland (2000) have recently argued that distributed connectionist models are *more* rigorous in respect of falsifiability than other types of model. Connectionist models discover representations, without additional assumptions being incorporated directly into the models. Models should also be judged on what they cannot do, as well as what they can do.) Of course, there has been much debate about the extent to which explicit rules are necessary in particular domains of language processing, and debate about the extent to which particular models adequately describe experimental data. Finally, computational modellers have neglected (or at least, not yet got round to) even larger tracts of human behaviour than cognitive neuropsychologists. Nevertheless, it is striking that so much of an important and controversial topic can be largely ignored. It is also wrong to ignore an approach that, more than any other over the last few years, has at last focused on the actual processes involved in cognition.

AGREEMENT IN COGNITIVE NEUROPSYCHOLOGY

A final concern that the *HCN* brings home is about the lack of agreement among cognitive neuro-psychologists. A *Handbook of how my car works* would tell me just that: about my car's components, what they do, how they work, and how they contribute to how my car works overall. The *HCN* describes a set of theories about how the mind works. While a description of any subject will describe only theories, what is quite remarkable about those described in the *HCN* is the extent to which they conflict. Furthermore, the conflict between theories is often at a high level: To what extent does the mind use symbolic rather than subsymbolic processing? How modular is it? How closely tied are psychological processes to neural pathways? How many routes are involved in any one process? and so on. Here are a couple of examples from the *HCN*. Shelton and Caramazza,

in their chapter on the organisation of semantic memory, argue for a domain-specific knowledge hypothesis that views knowledge as being organised into broad domains deriving from specialised neural mechanisms, against the otherwise prevalent modality-specific, sensory-functional theory. Nickels's chapter reflects the dominant view in studies based on normal and brain participants, and computational modelling, that there is a stage of lemma access in speech production; Caramazza (1997) argues convincingly against the existence of such a stage. There is even disagreement about what commonly used terms mean: As Nickels notes in her chapter on spoken word production, the words "semantics" and "concepts" are both used to refer to general preverbal aspects of knowledge and to lexically specific aspects of meaning. To these examples one can add: How many routes are involved in reading? Is there a general phonological deficit underlying phonological dyslexia? Is speech production an interactive process? How many phonological buffers are there? and so on. While debate and controversy are signs of a healthy, developing subject, one can have too much of a good thing. Although any particular description of a theory sounds sensible, overall the *HCN* leaves me in a turmoil of confusion. I want to know the answers, but by the end of the *HCN* (and reading the last 4 years of *Cognitive Neuropsychology*) I didn't always know what the questions were. For example, the concept of "words" in speech production seems to be different from the concept of "words" in reading. Arguments about the role of lemmas have dominated research on language production, but not language comprehension.

Given the lack of consensus, each contributor naturally brings their own bias to their chapter. For example, Shelton and Caramazza unsurprisingly—but controversially—conclude that there is little support for the modality-specific and the perceptual-functional knowledge hypotheses, instead arguing that semantic memory is primarily organised into content-specific semantic domains. Nickels, in her even-minded review of spoken-word production, remains sceptical that interactive models can account for all the necessary normal and neuropsychological phenomena.

HARLEY

While the complexity of the subject contributes to this lack of agreement, the type of models favoured by many cognitive neuropsychologists is still the worst culprit. In a box-and-arrow diagram an impairment can arise from damage either to information or processes contained in the box, or from difficulty accessing the box, or from both (see Shallice, 1988; Warrington & Shallice, 1979). This problem is particularly manifest in theories about the semantic system(s). Computational models here possess a great advantage: The processes involved are made explicit, and the models are testable (again, see Seidenberg, 1988, for his widely cited and widely ignored discussion of this point).

Hence the *HCN* inadvertently does an excellent job of conveying the lack of agreement endemic to contemporary cognitive neuropsychology. The key question is whether there is more controversy and less agreement in cognitive neuropsychology than in other comparable subjects. I suspect that there is. After all, if we cannot agree on whether knowing or not knowing where something happens in the brain is informative for psychological models, why should we expect to agree on less fundamental matters?

CONCLUSIONS

In summary, Rapp's *Handbook of Cognitive Neuropsychology* is a fine—if expensive—book. It deserves particular praise for its attempt to show cognitive neuropsychologists that there is a world outside the processing of isolated words. Nevertheless, the *HCN* fails to convince because of the failings of the subject. The *HCN* is biased in its coverage, it is biased in its theoretical orientation, it does not supply the answers, it largely ignores the most powerful tool available (computational modelling), and the current research that it describes deviates from the goals and methods of the programme. Of course, these failings are mostly not the fault of the editor or contributors. Indeed, they do an excellent job of a very difficult task.

Perhaps the most striking thing about the *HCN* is that it might have been published at the beginning of the end of the life of the subject. It's a pity, because it had only just reached puberty, and was starting to get on just fine with its computational cousin but, already partly suffocated by its own narrowness, and weakened by internecine squabbling, it was eventually murdered by the toys.

REFERENCES

Ainsworth-Darnell, K., Shulman, H. G., & Boland, J. E. (1998). Dissociating brain responses to syntactic and semantic anomalies: Evidence from event-related potentials. *Journal of Memory and Language*, *38*, 112–130.

Badecker, W., & Caramazza, A. (1985). On considerations of method and theory governing the use of clinical categories in neurolinguistics and cognitive neuropsychology: The case against agrammatism. *Cognition*, *20*, 97–125.

Badecker, W., & Caramazza, A. (1986). A final brief in the case against agrammatism: The role of theory in the selection of data. *Cognition*, *24*, 277–282.

Bates, E., McDonald, J., MacWhinney, B., & Appelbaum, M. (1991). A maximum likelihood procedure for the analysis of group and individual data in aphasia research. *Brain and Language*, *40*, 231–265.

Bub, D. (2000). Methodological issues confronting PET and fMRI studies of cognitive function: With special reference to Frackowiak et al. (1997) "Human brain function". *Cognitive Neuropsychology*, *17*, 467–484.

Bub, J., & Bub, D. (1988). On the methodology of single-case studies in cognitive neuropsychology. *Cognitive Neuropsychology*, *5*, 565–582.

Caplan, D. (1986). In defense of agrammatism. *Cognition*, *24*, 263–276.

Caplan, D. (1988). On the role of group studies in neuropsychological and psychopathological research. *Cognitive Neuropsychology*, *5*, 535–548.

Caplan, D. (1991). Agrammatism is a theoretically coherent aphasic category. *Brain and Language*, *40*, 274–281.

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

Caramazza, A. (1984b). On drawing inferences about the structure of normal cognitive systems from the analysis of patterns of impaired performance: The case for single-patient studies. *Brain and Cognition*, *5*, 41–66.

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

Caramazza, A. (1985). On drawing inferences about the structure of normal cognitive systems from the analysis of patterns of impaired performance: The case for single-patient studies. *Brain and Cognition*, *5*, 41–66.

Caramazza, A. (1986). On drawing inferences about the structure of normal cognitive systems from the analysis of patterns of impaired performance: The case for single-patient studies. *Brain and Cognition*, *5*, 41–66.

Caramazza, A. (1991). Data, statistics, and theory: A comment on Bates, McDonald, MacWhinney, and Applebaum's "A maximum likelihood procedure for the analysis of group and individual data in aphasia research". *Brain and Language*, *41*, 43–51.

Caramazza, A. (1997). How many levels of processing are there in lexical access? *Cognitive Neuropsychology*, *14*, 177–208.

Clark, V. P., Keil, K., Maisog, J. M., Courtney, S., Ungerleider, L. G., & Haxby, J. V. (1996). Functional magnetic resonance imaging of human visual cortex during face matching: A comparison positron emission tomography. *Neuroimage*, *4*, 1–15.

Coltheart, M. (1985). Cognitive neuropsychology. In M. Posner & O. S. M. Marin (Eds.), *Attention and performance*, *Vol. XI* (pp. 3–37). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Coltheart, M. (1996). Phonological dyslexia: Past and future issues. *Cognitive Neuropsychology*, *13*, 749–762.

Coltheart, M. (in press). Cognitive neuropsychology. In J. Wixted (Ed.), *Handbook of experimental psychology: Vol. 4. Methodology* (3rd ed.). New York: John Wiley.

Cuetos, F., & Mitchell, D. C. (1988). Cross-linguistic differences in parsing: Restrictions on the use of the late closure strategy in Spanish. *Cognition*, *30*, 73–105.

Dell, G. S. (1986). A spreading-activation theory of retrieval in sentence production. *Psychological Review*, *93*, 283–321.

Deutsch, J. A. (1960). *The structural basis of behaviour*. Cambridge: Cambridge University Press.

Ellis, A. W., & Young, A. W. (1988). *Human cognitive neuropsychology*. Hove, UK: Lawrence Erlbaum Associates Ltd.

Farah, M. J. (1994). Neuropsychological inference with an interactive brain: A critique of the "locality" assumption. *Behavioral and Brian Sciences*, *77*, 43–104.

Farah, M. J., & Aguirre, G. K. (1999). Imaging visual recognition: PET and fMRI studies of the functional anatomy of human visual recognition. *Trends in Cognitive Sciences*, *3*, 179–186.

Farah, M. J., Stowe, R. M., & Levinson, L. L. (1996). Phonological dyslexia: Loss of a reading-specific component of the cognitive architecture? *Cognitive Neuropsychology*, *13*, 849–868.

Gregory, R. L. (1961). The brain as an engineering problem. In W. H. Thorpe & O. L. Zangwill (Eds.), *Current problems in animal behaviour* (pp. 307–330). Cambridge: Cambridge University Press.

Harley, T. A. (1984). A critique of top-down independent levels models of speech production: Evidence from non-plan-internal speech errors. *Cognitive Science*, *8*, 191–219.

Harm, M. W., & Seidenberg, M. S. (1999). Phonology, reading acquisition, and dyslexia: Insights from connectionist models. *Psychological Review*, *106*, 491–528.

Hinton, G. E., & Shallice, T. (1991). Lesioning an attractor network: Investigations of acquired dyslexia. *Psychological Review*, *99*, 74–95.

Johnson-Laird, P. (1983). *Mental models*. Cambridge: Cambridge University Press.

Kennedy, M. (1996). Cognitive neuropsychology and aphasia: A critical analysis. In B. Dodd, R. Campbell, & L. Worrall (Eds.), *Evaluating theories of language: Evidence from disordered communication* (pp. 161–183). London: Whurr Publishers.

Levelt, W. J. M., Roelofs, A., & Meyer, A. S. (1999). A theory of lexical access in speech production. *Behavioral and Brain Sciences*, *22*, 1–75.

Marin, O. S. M., Saffran, E. S., & Schwartz, M. F. (1976). Dissociations of language in aphasia: Implications for normal function. *Annals of the New York Academy of Sciences*, *280*, 868–884.

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

McCloskey, M. (1993). Theory and evidence in cognitive neuropsychology: A "radical" response to Robertson, Knight, Rafal, and Shimamura (1993). *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *19*, 718–734.

Mehler, J., Morton, J., & Jusczyk, P. W. (1984). On reducing language to biology. *Cognitive Neuropsychology*, *1*, 83–116.

Morton, J. (1984). Brain-based and non-brain based models of language. In D. Caplan, A. R. Lecours, & A. Smith (Eds.), *Biological perspectives in language* (pp. 40–64). Cambridge, MA: MIT Press.

Morton, J., & Frith, U. (1995). Causal modeling: A structural approach to developmental psychology. In

D. Cicchetti & D. J. Cohen (Eds.), *Developmental psychopathology: Vol. 1. Theory and methods* (pp. 357–390). New York: John Wiley.

Pickering, M. J. (1999). Sentence comprehension. In S. Garrod & M. J. Pickering (Eds.), *Language processing* (pp. 123–153). Hove, UK: Psychology Press.

Plaut, D. C., & McClelland, J. L. (2000). Stipulating versus discovering representations. *Behavioral and Brain Sciences*, *23*, 489–491.

Plaut, D. C., McClelland, J. L., Seidenberg, M. S., & Patterson, K. E. (1996). Understanding normal and impaired word reading: Computational principles in quasi-regular domains. *Psychological Review*, *103*, 56–115.

Poeppel, D. (1996). A critical review of PET studies of phonological processing. *Brain and Language*, *55*, 317–351.

Price, C. J., & Friston, K. J. (1997). Cognitive conjunction: A new approach to brain activation experiments. *Neuroimage*, *5*, 271–270.

Pulvermuller, F. (1999). Words in the brain's language. *Behavioral and Brain Sciences*, *22*, 253–336.

Rapp, B. (Ed.). (2001). *The handbook of cognitive neuropsychology: What deficits reveal about the human mind*. Philadelphia: Psychology Press.

Robertson, L. C., Knight, R. T., Rafal, R., & Shimamura, A. P. (1993). Cognitive neuropsychology is more than single-case studies. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *19*, 710–717.

Saffran, E. (1982). Neuropsychological approaches to the study of language. *British Journal of Psychology*, *73*, 317–337.

Schwartz, M. F. (1984). What the classical aphasia categories can't do for us and why. *Brain and Language*, *21*, 3–8.

Seidenberg, M. S. (1988). Cognitive neuropsychology and language: The state of the art. *Cognitive Neuropsychology*, *5*, 403–426.

Seidenberg, M. S., & McClelland, J. L. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review*, *96*, 523–586.

Shallice, T. (1988). *From neuropsychology to mental structure*. Cambridge: Cambridge University Press.

Sokol, S. M., McCloskey, M., Cohen, N. J., & Aliminosa, D. (1991). Cognitive representations and processes in arithmetic: Inferences from the performance of brain-damaged subjects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *17*, 355–376.

Tippett, L. J., & Farah, M. J. (1994). A computational model of naming in Alzheimer's disease: Unitary or multiple impairments? *Neuropsychology*, *8*, 1–11.

Uttal, W. R. (2001). *The new phrenology: The limits of localizing cognitive processes in the brain*. Cambridge, MA: MIT Press.

Van Orden, G. C., & Paap, K. R. (1997). Functional neuroimages fail to discover pieces of mind in parts of the brain. *Philosophy of Science Proceedings*, *64*, S85–S94.

Van Orden, G. C., Pennington, B. F., & Stone, G. O. (2001). What do double dissociations prove? *Cognitive Science*, *25,* 111–172.

Warrington, E., & Shallice, T. (1979). Semantic access dyslexia. *Brain*, *102*, 43–63.

Zurif, E., Gardner, H., & Brownell, H. H. (1989). The case against the case against group studies. *Brain and Cognition*, *10*, 237–255.

Zurif, E., Swinney, D., & Fodor, J. A. (1991). An evaluation of assumptions underlying the single-patient-only position in neuropsychological research: A reply. *Brain and Cognition*, *16*, 198–210.