

## “Shallow Draughts Intoxicate the Brain’’: Lessons from Cognitive Science for Cognitive Neuropsychology

Karalyn Patterson,<sup>a</sup> David C. Plaut<sup>b</sup>

<sup>a</sup>*MRC Cognition and Brain Sciences Unit*

<sup>b</sup>*Carnegie Mellon University*

Received 13 June 2008; received in revised form 8 October 2008; accepted 18 October 2008

---

### Abstract

This article presents a sobering view of the discipline of cognitive neuropsychology as practiced over the last three or four decades. Our judgment is that, although the study of abnormal cognition resulting from brain injury or disease in previously normal adults has produced a catalogue of fascinating and highly selective deficits, it has yielded relatively little advance in understanding how the brain accomplishes its cognitive business. We question the wisdom of the following three “choices” in mainstream cognitive neuropsychology: (a) single-case methodology, (b) dissociation between functions as the most important source of evidence, and (c) a central goal of diagramming the functional architecture of cognition rather than specifying how its components work. These choices may all stem from an excessive commitment to strict and fine-grained modularity. Although different brain regions are undoubtedly specialized for different functions, we argue that parallel and interactive processing is a better assumption about cognitive processing. The essential goal of specifying representations and processes can, we claim, be significantly assisted by computational modeling which, by its very nature, requires such specification.

*Keywords:* Cognitive neuropsychology; Connectionist modeling; Modularity; Dissociations; Interactivity; Single-case methodology; Learning

---

*A little learning is a dangerous thing;  
drink deep, or taste not the Pierian spring:<sup>1</sup>  
there shallow draughts intoxicate the brain,  
and drinking largely sobers us again.*

Alexander Pope, *An Essay on Criticism*, 1709

---

Correspondence should be sent to Karalyn Patterson, MRC Cognition and Brain Sciences Unit, Cambridge, UK. E-mail: karalyn.patterson@mrc-cbu.cam.ac.uk

## 1. Cognitive neuropsychology and cognitive science

This article concerns cognitive neuropsychology, the branch of cognitive neuroscience that takes, as its evidential base, patterns of preserved and disrupted cognitive abilities arising from brain injury or disease. We began our preparations for writing the paper by considering the relationship between cognitive neuropsychology, the topic of this article, and cognitive science, the topic of *topiCS*. In some sense, we thought, the former must be a subset or sub discipline of the latter, and thus one would expect a cognitive science journal to contain at least the occasional paper on cognitive neuropsychology. A quick survey of the contents of the last 4 years of *Cognitive Science* (surely *the* cognitive science journal), however, reveals not one article that could be identified as dealing with cognitive neuropsychological research. There are a few papers that touch on matters of significant concern in cognitive neuropsychology, such as the importance and interpretation of dissociations (of which more later); but not a single *Cognitive Science* paper in this 4-year period presents data or models of neuropsychological phenomena.

Given that the core of cognitive science is computational accounts of human cognition, might this situation have arisen because researchers with a computational bent have interested themselves exclusively in normal rather than impaired cognition? *Au contraire*: computational modelers, especially those with a PDP/connectionist perspective, have a long history of attention to disorders of cognition. Early on, this interest perhaps derived from two principal sources. First of all, many connectionist models employ distributed representations. These are considered to have a number of advantages (see Hinton, McClelland, & Rumelhart, 1986) of which one is relative robustness and graceful degradation in response to mild or moderate damage (Rumelhart & McClelland, 1986; Sejnowski, 1998). An obvious way to demonstrate this particular characteristic was to damage a working simulation and observe its behavior. Secondly, connectionist simulations are very often models in which the knowledge necessary to perform the task—represented by the weights on connections between units—is acquired through learning rather than specified by the modeler. In the context of this focus on learning, a number of models also addressed the question of *relearning*. That is, once a simulation had been trained to criterion on the cognitive task of interest, the researchers went on to investigate not only how it performed under damage but also how well and how quickly it relearned when given a period of retraining (Hinton & Sejnowski, 1986; Plaut, 1996; Sejnowski & Rosenberg, 1987).

Despite the importance of these two early reasons for exploring the effects of damage in computational models, they resulted mainly in demonstrations of basic principles rather than attempts to provide detailed theoretical accounts of cognitive disorders. From about the mid-1990s onward, however, a nontrivial amount of research has had the goal of simulating precise patterns of performance by real neurological patients. In other words, although the absence of such work from the *Cognitive Science* journal suggests that cognitive neuropsychology has a low profile in cognitive science, we argue that cognitive science, in the form of computational modeling, is having a major impact on cognitive neuropsychology, at least as practiced by neuropsychological researchers sympathetic to this approach. The remainder of this article consists of some speculation as to why cognitive neuropsychology is so

relatively absent from the field of cognitive science, followed by some exploration of what cognitive science can do for cognitive neuropsychology. In our view, it would not be too wild an exaggeration to suggest that—by teaching it to drink deeper—cognitive science may offer cognitive neuropsychology not just sobriety but a lifeline.

## 2. What is cognitive neuropsychology?

Cognitive neuropsychology uses patterns of cognitive disorders as evidence to derive inferences about cognitive function in the normal brain. Our review will be restricted to acquired disorders, that is cognitive impairments in adults for whom there is actual or presumptive evidence that the cognitive ability in question was normal prior to the event (e.g., stroke) or process (e.g., neurodegenerative disease) responsible for the impairment. Sometimes the main or partial goal is to understand the role of the specific brain regions or structures damaged. If damage to the right parietal lobe consistently produces hemispatial neglect to the left side of space, then a neuroanatomical account of visuospatial processing must include a role for this brain region. More often, however, especially in the approach dubbed “ultra-cognitive” neuropsychology (Shallice, 1988a), researchers have had little if any interest in neuroanatomy; their motivation is only to construct functional accounts of a cognitive ability. Nevertheless, they consider that valuable evidence about functional organization can be obtained by studying its disruption.

Cognitive neuropsychology dates back about one-and-a-half centuries, to the time when French and German neurologists like Paul Broca and Carl Wernicke began to engage in careful observation of their patients’ acquired language deficits and then to relate these to the abnormal brain regions observed in postmortem brain analyses. There are several good reviews of the history of this discipline and the approaches that it favored/disfavored during different periods (see especially Shallice, 1988a, 1988b). Here we will concentrate on the most recent flowering of cognitive neuropsychology, which came to prominence in the 1970s and spawned a journal devoted to the topic (called *Cognitive Neuropsychology*) in 1984. Activity in cognitive neuropsychology had, and to some extent still has, the following prototypical format. Find a neurological patient with a cognitive deficit acquired as a result of brain injury. Ideally, this deficit would be quite specific, though not so highly specific as to preclude general theoretical interest (jokes abound in the field about hypothetical case studies like “impaired digit span in a left-hander bilingual in French and Swahili”). The impairment should either be novel, that is, not previously reported in the literature, or of a cleaner, more dramatic form than previously published cases. Also, the deficit would ideally occur in a domain in which cognitive psychologists were busily engaged in research on its normal function. The neuropsychologist would then set about investigating the nature of the deficit, on the basis of both standardized tests and those of his/her own devising, with the goal of ruling out coincidental explanations and eventually trying to understand something new about the ability in question. For example, patients with deep dyslexia—one of the most intensively investigated disorders in the 1970–1980s—can read aloud many real words correctly but fail to read aloud even very simple nonsense words like *neg* or *dake*.

Researchers had to rule out boring explanations for this failure (e.g., that the patients could not *say* such unfamiliar combinations of phonemes) and could then move on to trying to understand the relevance of this dramatic failure for the processes involved in reading familiar words, using and developing theories about reading then available in the field of cognitive psychology.

### 3. Problems with cognitive neuropsychology

Three central aspects of the cognitive neuropsychology paradigm or approach require description here, as it is our hypothesis that these three may account in large part for the problems with the paradigm and for its low profile in cognitive science. The three are not entirely independent but, as they are all salient, we will describe them separately for emphasis. In possible order of importance from least to most, these three features may be labeled as follows: (a) an almost complete reliance on single-case studies; (b) an almost universal goal of establishing dissociations; (c) theorizing that consists almost entirely of descriptive, box-and-arrow diagrams with little, indeed sometimes no attempt to specify actual mechanisms of the putative processes.

#### 3.1. *How many cases?*

During the heyday of cognitive neuropsychology, the majority of publications were single-case studies. Part of the “justification” for this was pragmatic. Some interesting patterns of disorder are relatively rare, especially in the environment of a single research center. Furthermore, even if multiple cases were likely to exist around a given centre, many of the researchers engaged in the enterprise were experimental psychologists rather than clinicians, and considered themselves lucky when a neurologist or speech therapist referred even a single interesting case to study. Using results from a single case to draw conclusions about cognition in the general population does, however, rely on the assumption that cognitive abilities have the identical functional organization in all brains. This assumption was not always made explicit and/or defended, and unsurprisingly came in for some questioning if not outright criticism (e.g., Seidenberg, 1988). Perhaps the surprise is how relatively infrequent such criticism was. After all, in other subdisciplines of cognitive science and neuroscience where it is too time-consuming or costly to test multiple participants—including animal lesion studies, single-cell recordings in nonhuman primates, and human experimental psychophysics—the norm has always been to report on at least two or three subjects. Cognitive neuropsychology publications did occasionally stretch to a few patients with a similar disorder (e.g., Warrington, 1975; three cases of selective impairment of semantic memory), but this was definitely the exception rather than the rule. There was even the occasional claim that “valid inferences about the structure of normal cognitive systems from patterns of impaired performance are only possible for single-patient studies” (Caramazza & McCloskey, 1988, p. 519); but the fact is that even cognitive neuropsychologists who did not and would not adopt this extreme view mostly did do single-case research.

Another, perhaps slightly more honorable, justification for the reliance on single cases was a realization of inter-case variability, even between patients given the same syndrome classification. The belief seems to have been that averaging across results from such individuals would further muddy the already murky waters into which one was trying to peer (Ellis & Young, 1988; Shallice, 1979). This belief—based on the assumption that no two brain lesions will ever be identical—was the main reason for Caramazza and McCloskey's (1988) argument that single-case methodology was not only legitimate but also essential.

Even if one accepts the concern about the comparability of different lesions, there does seem to be a certain inconsistency between (a) an emphasis on variability between cases with the same putative deficit, and (b) the assumption that all brains are organized in the same way. This is not necessarily as sinister as it may sound, because variability in results might come from factors irrelevant to the actual function of interest. For example, if one patient thinks it is more important to be accurate and another that it is more important to respond, then testing these two might yield a considerable difference in the rates and types of errors even if the underlying deficit is identical. Nevertheless, single-case methodology is not the only possible response to this concern about nonidentical lesions; and recently at least some cognitive neuropsychologists have argued that a more trustworthy and informative source of evidence than case studies are case series, in which the same investigations are applied to a set of patients with the same core deficit (Patterson et al., 2006). Performance is presented for the individual cases as well as averaged across the group, enabling appreciation of both similarities and differences between the individual patients. Again, this is similar to the pattern of published data in neurophysiology or psychophysics.

A third, and perhaps the most important, justification for single-case methodology was an argument based on adequacy. If one accepts the assumption that “all brains/people do it the same way” (where “it” means some cognitive ability like reading aloud or auditory-verbal short-term memory or face recognition etc., and “the same way” means the same computational mechanisms), then surely one would only need to demonstrate a particular outcome in a single patient in order to generalize the conclusion to the rest of the population.

### *3.2. Dissociation: A legitimate goal and gold standard?*

The typical aim of a classical cognitive neuropsychology study was to find a dissociation: good performance on one task or type of material paired with poor performance on another, in the same patient. Needless to say, the two things shown to dissociate had to bear enough relationship to one another to make a contrast between them theoretically telling: A patient with a focal lesion in primary auditory cortex, unable to discriminate different speech sounds but with normal recognition of visually presented objects, would attract no interest.

The best treatment that we know of on types of dissociations and their interpretation is in Shallice (1988a). Suffice it to say here that the gold standard was always a double dissociation. Merely finding better performance on task or material-type A than B would always be open to the interpretation that A is just easier (or, in Shallice's terms, less resource-demanding) than B; but normal A paired with deficient B in one patient and normal B paired with deficient A in another was considered proof that these two abilities rely on separate

brain mechanisms. Here is an example. Some of the concepts that we know, such as *zebra* or *radiator*, are concrete, that is, they refer to tangible objects in the world; other concepts such as *idea* or *willingness* lack clear physical referents and are therefore considered to be abstract. Suppose that one finds a patient whose comprehension of words referring to concrete concepts is still within the normal range, but whose performance in comprehending abstract words is way below that of any age- and education-matched normal individual, and another case with the reverse pattern of success in comprehending these two classes of words. This is called a classical, or at least strong, double dissociation. By traditional cognitive neuropsychology logic, nothing more than these two cases will ever be needed to justify the conclusion that the brain has separate representations of, or separate mechanisms for processing of, concrete and abstract concepts. A hundred or a thousand patients showing an association, that is, correlated degrees of deficit on concrete and abstract words, would not offset even one double dissociation, because the brain regions supporting the two separate mechanisms might be located in such a way that the great majority of brain lesions would either damage or spare both (Coltheart, 2004).

The assumption that the functional organization of cognition can be unequivocally revealed by dissociation, however rarely that dissociation might be observed, has also not gone unquestioned; but such (published) concerns were infrequent and had little impact on the cognitive neuropsychology mainstream. We point interested readers to a thoughtful critique by Goldberg (1995) and cannot resist quoting from it, mainly because this excerpt expresses the point so well but also because it contains a misprint or typographical error with an unintended but amusing double entendre: “This is not to say that all instances of isolated strong dissociations are theoretically useless. This is to say, however, that they must be approached with a degree of weariness, pending the demonstration of their high prevalence in the presence of a particular lesion location, or/and converging evidence from other sources” (Goldberg, 1995, p. 195). The author surely meant to advocate a wary rather than a weary response, but the former does not preclude the latter. Another excellent critique of cognitive neuropsychology, of which we shall speak more in the next section, is by Seidenberg (1988). His cautionary note with respect to dissociation and single-case methodology was that, when a complex and adaptive system like the brain is damaged, it may still manage to perform a task but in a fashion not all that informative about normal function. Seidenberg’s summary of this problem ended with the following, also amusing, analogy: Do we learn much about the ways in which people normally travel to catch their flights at an airport by studying one man who—when there is a huge snowstorm blocking the roads—ditches his car and hires a helicopter? We should emphasize here that Seidenberg (1988) was not, and we most certainly are not, dismissing the potential to advance our understanding of human cognition by studying its disruption. The issue is not whether but how.

### 3.3. *The emptiness of the theorizing*

As already indicated, the third of the three worrying characteristics of the cognitive neuropsychological approach highlighted here is perhaps the most serious. It is also the one where, as we will attempt to demonstrate below, cognitive science may make the most

positive contribution. Here is how Seidenberg (1988) expressed this concern; we quote his words because we do not think that we can improve on them.

One of the main characteristics of the cognitive neuropsychological approach as it has evolved over the past few years...is that very little attention is devoted to specifying the kinds of knowledge representations and processing mechanisms involved. A concomitant characteristic is a lack of specificity as to how tasks...are performed. It is nonetheless assumed that one can draw valid inferences about the “functional architecture” despite the lack of this information. (p. 405)

In research of this kind, a deficit is attributed to a particular component because tests have been used to rule out involvement of other components. It is very difficult to accept this logic when the properties of the other components are themselves minimally known... The result is a superficial link between test and model, yielding very weak inferences about the locus of the deficit. Localising deficits within the functional architecture has nonetheless come to replace specification of processing mechanisms as the primary goal of this research. (p. 410)

These were not idle or empty concerns in Seidenberg’s (1988) hands; all of his criticisms were backed and illustrated by specific examples from neuropsychological research of that era.

It should be emphasized that the criticisms of cognitive neuropsychology with respect to dissociation and style of theorizing can also be leveled at much research employing functional brain imaging in normal individuals, which in the last decade or so has first joined and then overtaken cognitive neuropsychology as a source of evidence about brain–behavior relationships. The best functional imaging studies, like the best cognitive neuropsychology studies, are thoughtful and cautious in their conclusions. Thus, if the peak activation for a task involving face recognition is observed in brain region A but the peak activation for a task involving object recognition is observed in brain region B, not all researchers would conclude that recognition of faces and of objects are completely dissociable processes and that the former “resides” in area A and the latter in area B. As in cognitive neuropsychology, however, some imaging researchers have published just such simple and unjustifiable conclusions, rather than acknowledging (a) that regions A and B might both be involved, but differentially, in the two tasks; (b) that regions A and B might interact; and/or (c) that one might need to specify the representations and processes involved in these two tasks before such results have much value.

#### **4. The legacy of modularity?**

In our view, these three problematic aspects of cognitive neuropsychology as it has come to be practiced all stem from a strong assumption that the cognitive system has a fundamentally modular organization. Although there has been considerable debate over the specific properties that define a module (e.g., Coltheart, 1985; Fodor, 1983), the general idea is now commonplace: The cognitive system consists of a collection of separable components, each of which carries out a specific information-processing function that can be characterized independently of the functions of the other modules. Note that this claim is far more specific

than the mere presence of some degree of functional specialization in the brain—it stipulates distinct, informationally encapsulated and computationally idiosyncratic subsystems. Although multiple modules participate in any complex cognitive behavior, the degree of interaction among them is assumed to be very limited—sometimes restricted to a simple linear sequence with the final output of one module being passed on as input to the next. It is also generally assumed, at least within cognitive neuropsychology, that each module is separable not just functionally but neuroanatomically, making it possible for brain damage to impair the operation of one module while leaving the others largely intact. The modularity assumption has been so pervasive within cognitive neuropsychology that Ellis (1987) concluded “there can be no argument with the fact of modularity, only about its nature and extent” (p. 402). And for some cognitive neuropsychologists, as well as some functional imagers, it appears that the goal is to proliferate dissociations and modules into ever more fine-grained distinctions. Whereas we do not take issue with the general principle that different brain regions are specialized for different types of information and/or computation, we think that run-away module inflation (especially without concern for mechanism) is as dangerous as run-away economic inflation.

Given a starting assumption of many largely independent modules, the standard practice of cognitive neuropsychology makes perfect sense. Identifying behavioral dissociations among brain-damaged patients—particularly when the impairments are highly selective—is an appropriate means of delineating the modules and their functions. Moreover, single-case studies are justified because it seems reasonable to assume that neurologically intact individuals all have the same basic set of modules and connectivity—the same “functional architecture”—and differ only in the relative sophistication and efficiency of the modules themselves. Finally, it is natural to postpone a detailed specification of the internal operation of the modules until the overall architecture is firmly established. There is a sentence in the very first paragraph of Fodor’s influential book, *The Modularity of Mind*, that might seem precisely to recommend this “order of battle” (note that at this point in the book, Fodor was using the term “faculty” for module): “the best research strategy would seem to be divide and conquer: first study the intrinsic characteristics of each of the presumed faculties, then study the ways in which they interact” (Fodor, 1983, p. 1). Fodor’s very next sentence, however, says “Viewed from the faculty psychologist’s perspective, overt, observable behaviour is an interaction effect par excellence” (p. 1).

It is perhaps worth commenting that cognitive neuropsychology rarely if ever addressed the issue of how the modules got there in the first place. If cognitive neuropsychologists assumed that the modules were acquired on the basis of experience, then they probably thought that—as their goal was to analyze the consequences of damage to a largely steady-state adult brain—they did not need to, or at least did not have time to, address learning; it goes without saying that no researcher or even research team can work on everything. More often, perhaps following Chomsky, cognitive neuropsychologists probably assumed that the modules were innate: “...we take for granted that the organism does not learn to grow arms or to reach puberty.... When we turn to the mind and its products, the situation is not qualitatively different from what we find in the case of the body” (Chomsky, 1980, pp. 2–3).



Though it was not often made explicit, an assumption of innate rather than learned modules might be thought to buttress the assumption that different individuals have the same functional architecture; hence, studying a single case is sufficient.

Our concern is that cognitive neuropsychology has become increasingly detached from other areas within cognitive science and neuroscience in large part because the modularity assumption licenses a lack of consideration of representations and processes, and it is exactly in this respect that cognitive science can make a critical contribution to cognitive neuropsychology. To be clear, it is certainly possible to develop computationally explicit implementations of modular systems (e.g., Anderson & Lebiere, 1998; Coltheart, Curtis, Atkins, & Haller, 1993; Coltheart, Rastle, Perry, Langdon, & Ziegler, 2001; Kosslyn, Flynn, Amsterdam, & Wang, 1990; Newell, 1990; Simon, 1969), but doing so typically requires a substantial increase in the amount of communication and interaction among the modules. Essentially, more of the work is performed between modules rather than within them, calling into question the appropriateness of the modular organization in the first place. Perhaps it would be more fruitful to start with a theoretical framework grounded in interactivity and then explore the extent to which it can give rise not only to normal cognitive behavior but also to the types of selective deficits observed in neuropsychological research.

## 5. Distributed interactive processing

Suppose that, instead of modularity, one started with a very different set of theoretical assumptions, namely that cognitive processes fundamentally involve massively parallel interactions among distributed patterns of neural activity over multiple brain regions specialized for different types of information. What implications would this have for how one should study brain-damaged patients to obtain evidence about the organization and operation of the normal system? We will discuss this question from the perspective of the three features of cognitive neuropsychology identified above, once again taking the three separately though once again noting that they are not truly independent.

### 5.1. How many cases?

A recommendation to replace modularity with interactivity might seem orthogonal to the issue of whether all brains have the same functional architecture and therefore whether one should study single cases or case series. We think, however, that the relationship between these issues will become clear as we develop an argument that, except under special circumstances,  $n$  should be  $>1$  and as large as possible.

We should start by acknowledging that there *are* special circumstances. For example, the April 2008 issue of *Brain and Language* includes an elegant report (Davis et al., 2008) of a patient studied immediately poststroke who had hypoperfusion (low blood flow) essentially restricted to Broca's area in the posterior region of the left frontal lobe. At this stage, there was minimal infarction (death of brain tissue due to the loss of oxygen associated with low blood flow). Using modern stroke-treatment procedures, the patient was re-perfused as

quickly as possible. While the necessary investigations and preparations for this procedure were occurring, however, the patient performed language tests that revealed some selective and severe deficits, for example, in receptive and expressive syntactic processing. When the same tests were performed again following the successful re-perfusion, the patient had recovered these language abilities. The likelihood of accruing multiple patients suffering just these events in just these brain regions within one researcher's lifetime is very low indeed. This is therefore, in our judgment, an absolutely justifiable "case" for a single-case study. The important thing to note is the authors' conclusion: "...adequate function of Broca's area was necessary for a number of language comprehension and production tasks for this individual" (Davis et al., 2008, p. 56). The authors presumably expect this conclusion to generalize to other people, as do we, but they were careful not to assume the generalization.

The potentially crucial difference between single-case versus case-series designs has been highlighted in the context of explanations for various forms of acquired reading disorder. Several sets of researchers (e.g., Farah & Wallace, 1991; Patterson & Lambon Ralph, 1999; Price & Devlin, 2003) have suggested that, since reading is a relatively late-acquired skill both in human history and in the lifetime of any given human, there are probably no brain regions dedicated solely to any of the component abilities of reading. For example, reading involves recognizing/identifying letters (or characters in nonalphabetic writing systems) and employing orthographic representations of combinations of letters to compute other sorts of representations such as sounds and meanings of words. The argument is that the brain learns to read by expanding and adapting phylogenetically and ontogenetically older functions developed for visual processing and spoken language. This view predicts that lesions should never produce deficits in reading alone. The demands of reading may "stress" these component abilities more than other tasks do, such that reading impairments will be particularly prominent following damage to the brain areas supporting these functions; but careful testing should uncover difficulty with nonreading tasks that require the same functions. And indeed, if one tests patients with three different forms of acquired dyslexia, it appears that the cases of each type have associated deficits in nonreading abilities that make sense in the context of the nature of the reading impairment. (a) Patients with pure alexia (or alexia without agraphia), whose reading disorder is characterized by slow and error-prone letter identification, typically have deficits in nonreading tasks that require rapid and to-some-extent parallel visual identification (Cumming, Patterson, Verfaellie, & Graham, 2006; Mycroft, Berhmann, & Kay, unpublished data). (b) Patients with phonological dyslexia, who rely to an abnormal extent on word meaning in reading aloud and are poor at producing and/or judging information about the sounds of written letter strings, typically have deficits in phonological tasks that do not involve written-word stimuli (Crisp & Lambon Ralph, 2006; Rapcsak et al., 2008). (c) In patients with surface dyslexia, oral reading success is modulated by the predictability of the words' spelling-sound correspondences; these patients read *few*, *new* and *grew* aloud correctly but, given *sew*, they are most likely to say "sue" (Woollams, Lambon Ralph, Plaut, & Patterson, 2007). This kind of typicalization error turns out to be the characteristic of the patients' performance in other tasks not involving written words; asked to produce the past-tense forms of spoken

verbs, the patients correctly produce “blink” → “blinked” but also incorrectly produce “sink” → “singed.” Researchers who have studied this typicalization pattern across a number of tasks (e.g., Patterson et al., 2006; Rogers et al., 2004) have attributed the patients’ behavior in all of the tasks to degraded semantic representations.

Every paper referred to in the preceding paragraph presented data from a case series of patients with the form of reading disorder under consideration. Why are at least some researchers who used to feel comfortable with single-case research changing their strategy? One important aspect is that the papers just mentioned differ from standard cognitive neuropsychology not only in the number of patients tested but also in their goal. These studies were seeking not dissociation but theoretically meaningful association, and not modularity but interactivity. Beyond the simple fact that demonstrating a result in multiple different patients adds weight to the conclusion, it seems especially important, even crucial, to have this additional weight when the conclusion to be drawn is one so antithetical to the standard cognitive neuropsychology enterprise. A submitted single-case study arguing for a meaningful association between modules would be rejected on the grounds noted above in the section on dissociation, that is, that the single patient’s brain lesion just happened to affect both of the in-fact-separate but in this case “accidentally” associated abilities. This argument becomes a little harder to sustain if multiple cases reveal the association, though note—by traditional cognitive neuropsychology logic—not that much harder. As indicated earlier, by this logic, just one demonstration of a dissociation between two abilities can vitiate any number of observed associations.

So what is a cognitive neuropsychologist to do if, according to his/her theory, the specific ways that brain regions operate and interact yield clear predictions of associations as well as dissociations? Encouraged by recent work, we think that there are some sensible ways to proceed. (a) Make sure it is a good, well-specified theory—ideally, a computational model. (b) Make sure that the tests are sufficiently sensitive and appropriate. For example, patients with frontal lobe lesions can be disproportionately impaired in naming pictures of actions relative to pictures of objects; and in time-honored cognitive neuropsychology style, this “dissociation” has sometimes been interpreted as a selective deficit in verb processing. A recent study by D’Honnin and Pillon (2008), however, reveals that the disadvantage for action naming can be eliminated by asking the patient to name actions from video rather than from static pictures, indicating that the originally observed dissociation does not implicate verb processing per se. (c) Try to obtain converging evidence from a different research enterprise, such as experiments with normal individuals that include either functional imaging or measurement of individual differences. (d) Test as many patients as possible. (e) Scale the patients’ degree of deficit in ability Y as a function of the degree of deficit in ability X. Different patients with the same “syndrome” invariably differ in severity. The association argument about pure alexia says that these patients have difficulty in discriminating one letter from another visually similar one because they have difficulty in discriminating any visual stimulus from another one similar to it (such as checker-boards differing only in a few squares: Mycroft et al., unpublished data; Roberts, Lambon Ralph, & Caine, 2008). If this is correct, then the degree of reading deficit should correlate with the level of impairment in same-different matching of checker-boards. Of course, even this outcome will not

satisfy someone with a hard-line commitment to dissociations as the “...major tool of neuropsychological discovery and theory-building” (Goldberg, 1995, p. 195), because a lesion affecting two dissociable functions might have a bigger impact on both when it is a bigger lesion. Again, however, an accumulation of cases may help in this regard. A bigger lesion might also be expected to affect aspects of cognition other than X and Y; if an increasing deficit in X predicts an increasing deficit in Y (the two abilities argued to associate) but not in anything else, the association story gains credence.

There is another strong argument for case series. Genuine understanding of a cognitive function must include knowledge about its variability. As Seidenberg (1988) said, if we want to know about the behavior of getting to an airport, we need to determine how lots of different people do it, under both normal and abnormal conditions. Also, from the perspective of a non-Chomskyan belief that most cognitive functions are shaped as much by learning from experience as by innate structure, individual differences are to be expected, and ignored at our peril.

### 5.2. *Dissociation: A legitimate goal and gold standard?*

To a large extent, the reasons for and ways of dealing with this problem have been covered under the previous section on “how many cases?” We should emphasize again that, just as we are not “dissing” cognitive neuropsychology in general, we are not rejecting the demonstration of dissociations as a legitimate goal. Our point is merely that this should not be the only goal or the gold standard. Dividing up the brain’s functions into an ever larger catalogue of more specific modules on the basis of dissociations in single patients may seem like scientific progress to some, but not to us. Take the example of a comparison between reading single words aloud and naming pictures of objects whose names correspond to the stimulus words for oral reading, so that correct responses in the two tasks are perfectly matched for factors like length and phonological complexity. Several single-case studies have demonstrated the superiority of oral reading over naming for patients with poor ability to read pseudowords (such as *neg* or *dake*); the latter is treated as evidence that the advantage for reading over naming could not be attributed to general knowledge of how letters translate to sounds. Reflecting the traditional line of cognitive neuropsychology reasoning, some researchers have concluded either that (a) there must be a separate route linking whole printed words to their pronunciations (Coslett, 1991; Sartori, Masterson, & Job, 1987), or (b) there are separate output representations for object naming and word reading (Orpwood & Warrington, 1995). This is precisely the kind of strategy of inferring functional architecture without specifying mechanism—and without considering interactivity—that we think gets cognitive neuropsychology in trouble. A different and, in our view, more sensible interpretation of the oral reading advantage is that even a patient with inaccurate pseudoword reading might activate some appropriate information about a written word’s sound that can combine and interact with semantic information about the word, whereas object naming can only rely on semantic activation of phonology (Hillis & Caramazza, 1991, 1995; Plaut, McClelland, Seidenberg, & Patterson, 1996). As an alternative to postulating new and underspecified “routes” or modules, this kind of explanation based on graded degrees of deficit plus pervasive interactivity seems a more plausible and fruitful approach.

### 5.3. *The emptiness of the theorizing*

To anyone interested enough in cognitive science to be reading this paper in the first place, it will presumably be obvious why cognitive science is so germane to the solution of this third problem of cognitive neuropsychology. Computational modeling is all about specifying mechanism. In fact, unlike a descriptive box-and-arrow model, a computational model cannot be said to exist until its representations and processes are specified. If the model is constructed and “runs” and produces the target behavior, this is an unequivocal demonstration that this behavior is realizable with the processes and mechanisms embodied in the model; and—to take the argument a step further into the domain of cognitive neuropsychology—if a component of the model is subsequently damaged and the model then mimics the performance of real neurological patients, this is an unequivocal demonstration that the observed deficit could arise from this kind of damage to this kind of representation or mechanism. Such demonstrations do not, of course, guarantee that the model is a veridical reflection of how normal people perform this cognitive function well and/or how brain-injured people perform it poorly; and there are always questions about the dimensions on which a model’s success should be evaluated and how closely it needs to reproduce human behavior to qualify as a successful simulation (see Seidenberg & Plaut, 2006, for discussion). It is, however, an achievement of a truly different magnitude when human behavior can be related to a computational, that is, thoroughly specified, model as opposed to a box-and-arrow diagram informed in large part by researchers’ intuitions. If one thing about cognition is certain, it is surely this: our slow, serial, reasoning-style manner of *thinking about* how we might perform some cognitive function (such as pronouncing written words or retrieving memories of what happened last week) will not be a reasonable guide to the fast, parallel, interactive computations that our brains perform when engaged in these functions. Here again is someone else’s splendid expression of this point:

Human-designed systems have certain properties which ... are related to the limitations of our mental processes.... Because we are unable to talk or think about more than a very small number of processes taking place simultaneously, we isolate them into sub-assemblies so that each can be treated separately. We also like hierarchical systems to make explicit the flow of control. I used to think that these principles of modularity, rigorous sequentiality, and hierarchical control might underlie the structure and function of all elaborate systems. I now believe that while these principles may be at the heart of artificial engineering, natural engineering is different. Biological systems have processes which are more flexibly organised....(Brenner, 1997, pp. 38–39)

## 6. A way forward?

Some of the concerns that we have raised above are more properly thought of as issues of better and worse research design, and as such the same criticisms can be leveled at virtually any research discipline or approach, including functional imaging and even cognitive

science as well as cognitive neuropsychology. On the other hand, cognitive neuropsychology seems more culpable than average in some regards (or perhaps we are just more sensitive to its failings because we know it so well), and it is especially in these aspects where we think that some of the modeling approaches of cognitive science can help. This final section presents two examples that we hope will serve to illustrate the value of going beyond simple information-processing characterizations of cognitive processing to computationally explicit working implementations of the system, both in normal operation and following damage. The claim is not that the models offer correct or complete accounts of the phenomena, but rather that—by virtue of their computational explicitness—they advance theory and its relation to data in important ways. Both examples relate to neuropsychological findings that bear on the internal structure of semantic memory.

A longstanding debate within cognitive neuropsychology concerns whether semantic memory consists of multiple subsystems and, if so, whether these subsystems are organized by modality or by category (Caramazza, Hillis, Rapp, & Romani, 1990; Shallice, 1988b, 1993). Apparent evidence for category-based organization comes from the fact that some neurological patients have impairments to conceptual knowledge that are restricted largely to either living or nonliving things (Warrington & McCarthy, 1983, 1987; Warrington & Shallice, 1984). For example, patient JBR correctly identified the pictures of only 6% of living things but 90% of nonliving things. The reverse dissociation, although typically not as dramatic, has also been observed; for example, patient YOT was perfect at matching different pictures of the same type of animal but only 69% correct for humanmade objects. According to traditional cognitive neuropsychology logic, this double dissociation would imply distinct semantic subsystems for living and nonliving things; but Warrington and Shallice (1984) recognized that some aspects of the data were problematic for such a simple account. In particular, performance on certain subcategories fell on the “wrong” side of the living–nonliving divide. JBR, with a living-things deficit, was also impaired on metals, fabrics, musical instruments, and precious stones but unimpaired on body parts; YOT, with a nonliving impairment, showed the reverse pattern. Warrington and Shallice suggested that the observed dissociations could be explained if semantic knowledge was separated not by the category of the item but instead by the modality of the semantic attributes particularly germane to the items. The argument was that our knowledge of living things (and also the anomalous nonliving categories) relies more on sensory/perceptual information such as shape and color, whereas humanmade objects depend at least as much on functional aspects of knowledge such as how they are used. Although subsequent research has not yielded completely consistent or unequivocal evidence for this pattern (see, for example, Caramazza & Shelton, 1998), the Warrington and Shallice “sensory–functional” theory (SFT) remains highly influential.

Farah and McClelland (1991) developed a distributed connectionist implementation of the SFT in which both visual (picture) and verbal (word) inputs mapped onto semantic units that were organized into visual and functional subgroups, with connections both within and between the subgroups (see Fig. 1). Based on an analysis of dictionary definitions, the semantic representations of nonliving things were given approximately equal numbers of visual and functional semantic features, whereas those for living things were heavily weighted toward visual semantics. Damage to the visual-semantics component of the model

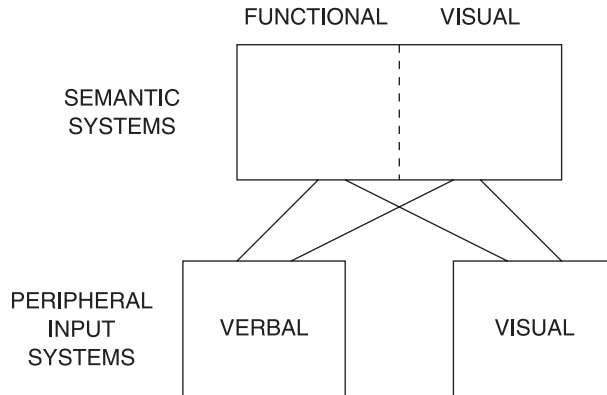


Fig. 1. General architecture of the Farah and McClelland's (1991) model. Each rectangle denotes a separate set of processing units. Functional and visual semantic units receive connections from each other as well as from visual and verbal input.

yielded far greater impairment on living than nonliving things, whereas damage to functional semantics yielded moderately more impairment on nonliving than living things, in accordance with the observed empirical patterns. In this respect, the Farah and McClelland's model can be viewed as a straightforward implementation of Warrington and Shallice's descriptive theory. The specific value of the implementation becomes clearer when considering an additional finding that appears inconsistent with the SFT; patients with a disproportionate deficit for living things often have impaired knowledge of *functional* as well as visual aspects of items in this category (see, e.g., Sartori & Job, 1988). In a strictly modular, box-and-arrow version of the theory, there is no reason to expect that damage to visual semantics would compromise functional semantic information. By contrast, in Farah and McClelland's implementation, visual and functional semantic units learn to interact and support each other, so that damage to one yields a mild but reliable deficit in the other. This is a clear example of the value of a fully explicit, computational instantiation of a theory for judging whether the theory is consistent with a particular empirical finding.

Farah and McClelland's (1991) model retains Warrington and Shallice's (1984) proposal of distinct semantic subsystems for visual and functional semantics, but it demonstrates non-obvious consequences of learning and of interactivity between "modules." The second example we consider challenges the very notion of distinct subsystems, instead arguing for graded modality-specific functional specialization within semantics.

A further source of evidence that has been used to argue for separate modality-based semantic systems comes from modality-specific naming disorders. For example, patients with optic aphasia (Lhermitte & Beauvois, 1973) are impaired at naming objects when presented visually, although they can demonstrate relatively intact recognition of the objects (e.g., by gesturing the correct use) and can name them when presented through other modalities (e.g., tactile input—feeling the object with eyes closed—or hearing a spoken description of the thing to be named). On the widely held assumption that object naming requires semantic mediation, this pattern of performance suggests a disconnection of visual

semantics (sufficient for recognition) from verbal semantics (required for naming) (Beauvois, 1982). Apart from being post hoc, this account is challenged by findings that not all visual naming is equally impaired; in particular, generating the names of actions in response to visually presented objects or depicted actions can be significantly better than visual object naming (Campbell & Manning, 1996; Teixeira Ferreira, Guisano, Ceccaldi, & Poncet, 1997). For example, when shown a set of 30 pictures of objects, patient AG (Campbell & Manning, 1996) was only 27% correct in answering “What is the name of this?” but 67% correct in answering “What can you do with this?”

Plaut (2002) developed an interactive connectionist simulation in which semantic representations develop under the pressure of learning to mediate between multiple input and output modalities in performing various tasks (see also Rogers et al., 2004). In particular, the model learned to map from visual or tactile input for each of 100 objects to the associated object or action name (e.g., “pen” vs. “write”) via a common set of intermediate (semantic) units. Consistent with the two-dimensional organization of neocortex, the semantic units were given topographically organized spatial positions between the input and output modalities such that some units were functionally closer than others to a particular modality (see Fig. 2). Learning in the network was biased to favor short over long connections, so that semantic units “near” a given modality became more important in accomplishing the mappings involving that modality. As a result, damage to connections between visual representations and regions of semantics near phonology disrupted visual object naming far more than either visual recognition or tactile naming—exactly the pattern




Fig. 2. Architecture of the Plaut (2002) model. Each gray square corresponds to a unit whose activity level is indicated by the size of the white square within it. The two-dimensional positions of units reflect their topographic organization, such that all sets of connections (indicated by arrows) except those from the Task units are subject to a learning bias that favors short over long connections.



observed in optic aphasia. Moreover, like the patients, the model was better at naming actions than objects from vision, due to the interactive support that action naming receives from action representations.

Thus, as in the Farah and McClelland's model, properties of interactivity that are apparent only through explicit simulation turn out to be critical in explaining otherwise puzzling aspects of the empirical findings. In addition, in this case, the simulation creates a type of organization that is unavailable when theorizing in terms of separable components. The semantic representations in the Plaut (2002) model are neither modality-independent nor modality-specific. Rather, they exhibit a graded degree of modality-specific organization, with regions near the inputs and outputs being more modality-specific and regions near the "center" of semantics being more amodal. This kind of organization seems at least consistent with what is known so far about the neuroanatomy of semantic memory (e.g., Bright, Moss, & Tyler, 2004; Gloor, 1997; Patterson et al., 2007), though we hasten to add that, at this stage, *so far* is not *very far*. As already expressed, our purpose in the current context is not to argue for the correctness of the model but rather to emphasize the value of interpreting neuropsychological deficits in terms of computationally explicit representations and processes.

## 7. Concluding comment

The phenomena of cognitive neuropsychology can be striking and dramatic. Deep dyslexic patients, once normal readers, may succeed in reading aloud the words *inn* and *bee* but fail on *in* and *be* (Gardner & Zurif, 1975). Patients with pure alexia, also once competent readers, may easily identify a word when asked, "What does C, H, U, R, C, H spell?" but fail to identify the same word when given it in written form (Patterson & Kay, 1982; note that identification from print is definitely easier than identification from oral spelling for normal adults). A patient with semantic dementia who goes out every morning to buy his regular newspaper so that he can do the Sudoku puzzle in the paper utterly failed to recognize a Sudoku puzzle when given one from a puzzle book in the clinic (Lambon Ralph & Patterson, 2008). These dramatic and selective deficits must be informative about the organization of cognitive processes in the brain. It is perhaps understandable that, when a fascination with such phenomena first became fairly widespread, the researchers thus engaged found it sufficiently thirst-quenching to take shallow draughts, and draw shallow conclusions, from this very deep spring of information. Drinking deeper is much harder work; but if cognitive neuropsychology is to provide any kind of genuine understanding of the massively parallel and interactive functioning of the brain, that harder work will be necessary.

## Note

1. For readers who do not know, as we did not, the *Pierian spring* apparently refers to the classical tradition that the Muses were born in the Pieria region of northern Greece.

## References

- Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*. Hillsdale NJ: Erlbaum.
- Beauvois, M.-F. (1982). Optic aphasia: A process of interaction between vision and language. *Proceedings of the Royal Society of London, Series B*, 298, 35–47.
- Brenner, S. (1997). *Loose ends*. London: Current Biology Ltd.
- Bright, P., Moss, H., & Tyler, L. K. (2004). Unitary vs. multiple semantics: PET studies of word and picture processing. *Brain and Language*, 89, 417–432.
- Campbell, R., & Manning, L. (1996). Optic aphasia: A case with spared action naming and associated disorders. *Brain and Language*, 53, 183–221.
- Caramazza, A., Hillis, A. E., Rapp, B. C., & Romani, C. (1990). The multiple semantics hypothesis: Multiple confusions? *Cognitive Neuropsychology*, 7, 161–189.
- Caramazza, A., & McCloskey, M. (1988). The case for single-patient studies. *Cognitive Neuropsychology*, 5, 517–528.
- Caramazza, A., & Shelton, J. R. (1998). Domain-specific knowledge systems in the brain: The animate-inanimate distinction. *Journal of Cognitive Neuroscience*, 10, 1–34.
- Chomsky, N. (1980). Rules and representations. *Behavioral and Brain Sciences*, 3, 1–15.
- Coltheart, M. (1985). Cognitive neuropsychology and the study of reading. In M. I. Posner & O. S. M. Marin (Eds.), *Attention and performance XI* (pp. 3–37). Hillsdale, NJ: Erlbaum.
- Coltheart, M. (2004). Are there lexicons? *Quarterly Journal of Experimental Psychology*, 57A, 1153–1171.
- Coltheart, M., Curtis, B., Atkins, P., & Haller, M. (1993). Models of reading aloud: Dual-route and parallel-distributed-processing approaches. *Psychological Review*, 100, 589–608.
- Coltheart, M., Rastle, K., Perry, C., Langdon, R., & Ziegler, J. (2001). DRC: A dual route cascaded model of visual word recognition and reading aloud. *Psychological Review*, 108, 204–256.
- Coslett, H. B. (1991). Read but not write “idea”: Evidence for a third reading mechanism. *Brain and Language*, 40, 425–443.
- Crisp, J., & Lambon Ralph, M. A. (2006). Unlocking the nature of the phonological-deep dyslexia continuum: The keys to reading aloud are in phonology and semantics. *Journal of Cognitive Neuroscience*, 18, 348–362.
- Cumming, T., Patterson, K., Verfaellie, M., & Graham, K. S. (2006). One bird with two stones: Letter-by-letter reading in pure alexia and semantic dementia. *Cognitive Neuropsychology*, 23, 1130–1161.
- D’Honnin, P., & Pillon, A. (2008). Verb comprehension and naming in frontotemporal degeneration: The role of the static depiction of actions. *Cortex*, 44, 834–847.
- Davis, C., Kleinman, J. T., Newhart, M., Gingis, L., Pawlak, M., & Hillis, A. E. (2008). Speech and language functions that require a functioning Broca’s area. *Brain and Language*, 105, 50–58.
- Ellis, A. W. (1987). Intimations of modularity, or, the modelarity of mind: Doing cognitive neuropsychology without syndromes. In M. Coltheart, G. Sartori & R. Job (Eds.), *The cognitive neuropsychology of language* (pp. 397–408). Hillsdale, NJ: Erlbaum.
- Ellis, A. W., & Young, A. W. (1988). *Human cognitive neuropsychology*. Hove and London: Lawrence Erlbaum.
- Farah, M. J., & McClelland, J. L. (1991). A computational model of semantic memory impairment: Modality-specificity and emergent category-specificity. *Journal of Experimental Psychology: General*, 120, 339–357.
- Farah, M. J., & Wallace, M. A. (1991). Pure alexia as a visual impairment: A reconsideration. *Cognitive Neuropsychology*, 8, 313–334.
- Fodor, J. A. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.
- Gardner, H., & Zurif, E. (1975). BEE but not BE: Oral reading of single words in aphasia and alexia. *Neuropsychologia*, 13, 181–190.
- Gloor, P. (1997). *The temporal lobe and limbic system*. New York: Oxford University Press.
- Goldberg, E. (1995). Rise and fall of modular orthodoxy. *Journal of Clinical and Experimental Neuropsychology*, 17, 193–208.

- Hillis, A. E., & Caramazza, A. (1991). Mechanisms for accessing lexical representations for output: Evidence from a category-specific semantic deficit. *Brain and Language*, *40*, 106–144.
- Hillis, A. E., & Caramazza, A. (1995). Converging evidence for the interaction of semantic and sublexical phonological information in accessing lexical representations for spoken output. *Cognitive Neuropsychology*, *12*, 187–227.
- Hinton, G. E., McClelland, J. L., & Rumelhart, D. R. (1986). Distributed representations. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing* (Vol. 1, pp. 77–109). Cambridge, MA: MIT Press.
- Hinton, G. E., & Sejnowski, T. J. (1986). Learning and relearning in Boltzmann machines. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing* (Vol. 1, pp. 282–317). Cambridge, MA: MIT Press.
- Kosslyn, S. M., Flynn, R. A., Amsterdam, J. B., & Wang, G. (1990). Components of high-level vision: A cognitive neuroscience analysis and accounts of neurological syndromes. *Cognition*, *34*, 203–277.
- Lambon Ralph, M. A., & Patterson, K. (2008). Generalization and differentiation in semantic memory: Insights from semantic dementia. *Annals of the New York Academy of Sciences*, *1124*, 61–76.
- Lhermitte, F., & Beauvois, M.-F. (1973). A visual-speech disconnection syndrome: Report of a case with optic aphasia, agnosic alexia and colour agnosia. *Brain*, *96*, 695–714.
- Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: MIT Press.
- Orpwood, L., & Warrington, E. K. (1995). Word specific impairments in naming and spelling but not reading. *Cortex*, *31*, 239–265.
- Patterson, K., & Kay, J. (1982). Letter-by-letter reading: Psychological descriptions of a neurological syndrome. *Quarterly Journal of Experimental Psychology*, *34*, 411–441.
- Patterson, K., & Lambon Ralph, M. A. (1999). Selective disorders of reading? *Current Opinion in Neurobiology*, *9*, 235–239.
- Patterson, K., Lambon Ralph, M. A., Jefferies, E., Woollams, A., Jones, R., Hodges, J. R., & Rogers, T. T. (2006). “Pre-semantic” cognition in semantic dementia: Six deficits in search of an explanation. *Journal of Cognitive Neuroscience*, *18*, 169–183.
- Patterson, K., Nestor, P. J., & Rogers, T. T. (2007). Where do you know what you know? The representation of semantic knowledge in the human brain. *Nature Reviews Neuroscience*, *8*, 976–987.
- Plaut, D. C. (1996). Relearning after damage in connectionist networks: Toward a theory of rehabilitation. *Brain and Language*, *52*, 25–82.
- Plaut, D. C. (2002). Graded modality-specific specialization in semantics: A computational account of optic aphasia. *Cognitive Neuropsychology*, *19*, 603–639.
- Plaut, D. C., McClelland, J. L., Seidenberg, M. S., & Patterson, K. (1996). Understanding normal and impaired word reading: Computational principles in quasi-regular domains. *Psychological Review*, *103*, 56–115.
- Price, C. J., & Devlin, J. T. (2003). The myth of the visual word form area. *NeuroImage*, *19*, 473–481.
- Rapcsak, S. Z., Beeson, P. M., Henry, M. L., Leyden, A., Kim, E., Rising, K., Andersen, S., & Hyesuk Cho, M. S. (2008). Phonological dyslexia and dysgraphia: Cognitive mechanisms and neural substrates. *Cortex*, forthcoming.
- Roberts, D., Lambon Ralph, M. A., & Caine, D. (2008). More than meets the eye: Does letter-by-letter reading derive from visual or from reading-specific problems? *Paper presented to the British Neuropsychological Society Meeting*, London.
- Rogers, T. T., Lambon Ralph, M. A., Garrard, P., Bozeat, S., McClelland, J. L., Hodges, J. R., & Patterson, K. (2004). Structure and deterioration of semantic memory: A neuropsychological and computational investigation. *Psychological Review*, *11*, 205–235.
- Rumelhart, D. E., & McClelland, J. L. (1986). Distributed representations. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing* (Vol. 1, pp. 110–146). Cambridge, MA: MIT 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.
- Sartori, G., Masterson, J., & Job, R. (1987). Direct route reading and the locus of lexical decision. In M. Coltheart, G. Sartori & R. Job (Eds.), *Cognitive neuropsychology of language* (pp. 59–77). London: Erlbaum.

- Seidenberg, M. S. (1988). Cognitive neuropsychology and language: The state of the art. *Cognitive Neuropsychology*, 5, 403–426.
- Seidenberg, M. S., & Plaut, D. C. (2006). Progress in understanding word reading: Data fitting vs. theory building. In S. Andrews (Ed.), *From inkmarks to ideas: Current issues in lexical processing* (pp. 25–49). Hove: Psychology Press.
- Sejnowski, T. J. (1998). Memory and neural networks. In P. Fara & K. Patterson (Eds.), *Memory* (pp. 162–181). Cambridge, England: Cambridge University Press.
- Sejnowski, T. J., & Rosenberg, C. R. (1987). Parallel networks that learn to pronounce English text. *Complex Systems*, 1, 145–168.
- Shallice, T. (1979). Case-study approach in neuropsychological research. *Journal of Clinical Neuropsychology*, 1, 183–211.
- Shallice, T. (1988a). *From neuropsychology to mental structure*. Cambridge, England: Cambridge University Press.
- Shallice, T. (1988b). Specialisation within the semantic system. *Cognitive Neuropsychology*, 5, 133–142.
- Shallice, T. (1993). Multiple semantics: Whose confusions? *Cognitive Neuropsychology*, 10, 251–261.
- Simon, H. A. (1969). *The sciences of the artificial*. Cambridge, MA: MIT Press.
- Teixeira Ferreira, C., Guisano, B., Ceccaldi, M., & Poncet, M. (1997). Optic aphasia: Evidence of the contribution of different neural systems to object and action naming. *Cortex*, 33, 499–513.
- Warrington, E. K. (1975). Selective impairment of semantic memory. *Quarterly Journal of Experimental Psychology*, 27, 635–657.
- Warrington, E. K., & McCarthy, R. (1983). Category-specific access dysphasia. *Brain*, 106, 859–878.
- Warrington, E. K., & McCarthy, R. (1987). Categories of knowledge: Further fractionation and an attempted integration. *Brain*, 110, 1273–1296.
- Warrington, E. K., & Shallice, T. (1984). Category specific semantic impairments. *Brain*, 107, 829–854.
- Woollams, A. M., Lambon Ralph, M. A., Plaut, D. C., & Patterson, K. (2007). SD-Squared: On the association between semantic dementia and surface dyslexia. *Psychological Review*, 114, 316–339.