

## Specific impairments of planning

BY T. SHALLICE

*M.R.C. Applied Psychology Unit, 15 Chaucer Road, Cambridge, CB2 2EF, U.K.  
and National Hospital, Queen Square, London WC1N 3BG, U.K.*

An information-processing model is outlined that predicts that performance on non-routine tasks can be impaired independently of performance on routine tasks. The model is related to views on frontal lobe functions, particularly those of Luria. Two methods of obtaining more rigorous tests of the model are discussed. One makes use of ideas from artificial intelligence to derive a task heavily loaded on planning abilities. A group of patients with left anterior lesions has a specific deficit on the task. Subsidiary investigations support the inference that this is a planning impairment.

### INTRODUCTION

Over the last 10 years a very fruitful interaction has grown up between a number of areas in cognitive psychology and corresponding areas in neuropsychology. One field that would seem ripe for such a development is that of executive functions and their disorders. Within neuropsychology, Luria's (1966) concept of a system specialized for the programming, regulation and verification of activity is widely used clinically to explain certain types of specific disorder, particularly some that can arise from lesions involving the frontal lobes (see, for example, Lhermitte *et al.* 1972; Walsh 1978). Moreover, within cognitive psychology there have been many discussions of seemingly related topics such as control processes, attention centres and central executive working memories.

However, there has been relatively little input from cognitive psychology to the understanding of high-level cognitive disorders. One reason has been that theories within cognitive psychology have tended to contain at most a single selection or general executive component, as in the very influential theory of Shiffrin & Schneider (1977). As I consider models that have only a single selection or general executive component insufficiently powerful to help to explain high-level cognitive disorders, I shall first introduce a more complex model developed by D. Norman and myself (Norman & Shallice 1980). I shall then discuss the neuropsychological evidence relevant to it.

Norman and I adopted a position common in psychology (see, for example, Miller *et al.* 1960) that at one level both cognition and action depend upon the 'running' of highly specialized routine programs (or productions or schemas), each of which will produce a specific output for a certain range of inputs. In our usage the basic unit is the 'schema' (short for an action schema or thought schema), a unit that can control a specific overlearned action or skill such as drinking from a container, doing long division, making breakfast, or finding one's way home from work. Schemas are held to be activated in various ways, for instance from 'triggers' released by perception and from the output of other schemas. For instance, I recently went into a room in the Unit where I work and found myself making a pulling movement with my left hand in the air. Somewhat puzzled, I realized that the cord that controlled

the light switch had been removed, but that the act of passing through the door and the dim light in the room together with the requirements of other higher-level schemas for light had triggered a specific schema that existed for activating the light switch in that particular room.

Since schemas can be activated totally independently of each other by different aspects of the situation, there is nothing to stop many schemas being activated at the same time. The novel aspect of our model is that the critical process of selection of the small subset of schemas that are to be 'run' at any time is held to involve not one but two qualitatively distinct processes: one, contention scheduling, involved in both routine and non-routine selection and the other, the Supervisory Attentional System, only in the latter (see figure 1).




FIGURE 1. A simplified version of the Norman & Shallice (1980) model representing the flow of control information. The lines with arrows represent activating input, the cross-hatched lines represent the primarily mutually inhibitory function of contention scheduling. The term 'effector system' refers to specific purpose-processing units involved in schema operation for both action and thought schemas. In the latter case schema operation involves placing information in short-term stores that can activate the trigger data base.

The major motivation for assuming two types of selection process comes from artificial intelligence. Different problem-solving programs, and even different levels within the same program, approach the issue of selection of which basic operation is to be executed at any one time in different ways.

One type of approach occurs in simulations of problem-solving employing production systems, a method that is in frequent use today. Production systems are composed of a potentially large number of separate, highly specific units – productions – each of which can become a candidate for operation if its particular, possibly quite complex, condition is satisfied (see Newell & Simon 1972). For the present purposes the relation between a production and its eliciting condition is analogous to the schema–trigger relation. In the operation of production system programs an analogous situation can occur to that described for our model, in that a number of productions can have their conditions satisfied at the same time. Modern programs therefore contain as part of the functional architecture of the system a conflict resolution device, which selects from among the initial selection according to simple criteria such as the characteristic importance of that production being selected or how recently the selecting data have arrived or how specific are the conditions for selection (see McDermott & Forgy 1978). Selection by this process is crude but fast.

In our model, a component called 'contention scheduling' has this function of ensuring *in a routine way* the efficient use of the limited effector and cognitive resources, given that competition to use these resources for many different purposes exists. It needs to select a highly restricted number of compatible schemas so that they can control such resources as they require until their goal is achieved, unless a much higher priority schema is triggered. Schema operation occurs when activation in a schema control unit reaches threshold, selection being ensured by mutual inhibition between the units by an amount depending upon their incompatibility. (The process is related to that discussed in Shallice (1972).) After selection of a schema a (probably somewhat reduced) inhibitory effect needs to be maintained by it until its goal is achieved, even if certain of its perceptual or short-term memory triggers are no longer present. In addition, schema-produced activation of its component (subordinate) schemas needs to be maintained so that they can be selected if appropriate triggers occur.

The essence of difficult problems is, however, that even though at one level the solution uses routine operations, which operations are used and in what order is not routine. Selection by using the 'strongest' triggers, as in contention scheduling, would not by itself be adequate. A second current approach that is used in the simulation of the solution of such problems is for the program to distinguish between procedures used to tackle problems that it 'knows' how to answer and those that it uses where the first type of procedure unexpectedly fails or where no solution procedure is 'known'. Problems are divided into the routine and non-routine and for the latter the program incorporates a general programming or planning component that can in principle be applied to any problem domain where it also has specialized knowledge (see Boden 1977, ch. 12). Selection by using this general programming component now becomes slow and flexible instead of fast, routine and unchanging. In the model, non-routine selection is held to involve the *biasing* (not the replacement) of the operation of contention scheduling by additional activation of appropriate schemas from another mechanism – the Supervisory Attentional System (S.A.S.) – which contains the general programming or planning systems that can operate on schemas in every domain.

That the process involved in non-routine selection may indeed require an additional system not involved in routine selection is supported by certain failures in the normal selection of action, 'action lapses', particularly the sort recently studied by Reason (1979) and Norman (1981) called 'capture errors'. Capture errors are best illustrated by an example: one of Reason's subjects described how, when passing through his back porch on the way to get his car out, he stopped to put on his Wellington boots and gardening jacket as if to work in the garden.

Consider what would happen on the model if one were to carry out a routine task that does not require continuous monitoring and activation from the S.A.S. Its component schemas can normally be selected by using contention scheduling alone, so the S.A.S. could be directed toward activating some non-competing schema and still the component schemas in the routine action could be satisfactorily selected by contention scheduling alone. Occasionally, though, a schema that is in fact incorrect could become more strongly activated in contention scheduling and capture the effector systems. The S.A.S. being directed elsewhere, would not immediately monitor this, so one has a 'capture error'.

Findings from a pioneering diary study of Reason (1982) provide further support for interpretation of this type of error in terms of the model. Subjects rate lapses as occurring when they are 'pre-occupied' and 'distracted', which would correspond to no activation being

received for the intended action from the S.A.S. which is instead activating a third non-competing schema. Also, both captured and capturing actions are rated as occurring 'very often' and being 'automatic'. This would mean that they could, as required on the model, be carried out controlled by contention scheduling alone. Moreover, the captured and capturing actions were rated as having very similar stimulus characteristics, so performance of the captured action would have triggered the activation of the incorrect schema. All these factors maximize the chance of an incorrect schema being more activated in contention scheduling than the correct one – and so a lapse occurring.

#### NEUROLOGICAL CORRESPONDENCES

What would this model predict about high-level cognitive impairment? In particular, what should happen with a specific deficit to the S.A.S.? The performance of routine tasks should not be affected even if they required considerable special-purpose processing resources. However, there would be a difficulty in coping with novelty or in planned initiative.

As the functioning of the S.A.S. within the model can be seen as a specification in an information-processing framework of Luria's (1966) unit for the programming, regulation and verification of activity, it is hardly surprising that the predicted impairment fits well with the classical view of frontal lobe dysfunction. Take, for instance, the position Goldstein held after World War I on such difficulties as summarized by Rylander (1939, p. 20): disturbed attention, increased distractibility, a difficulty in grasping the whole of a complicated state of affairs... well able to work along old routine lines... (but)... cannot learn to master new types of task, in new situations... at a loss'.

Unfortunately, although some single case studies and some group studies in which frontal lobe patients are compared with normal controls support this type of characterization, group studies in which the groups are defined by the site of lesion frequently do not produce a deficit in patients with anterior lesions when compared with those having posterior lesions (see, for example, Reitan 1964). Indeed, an unpublished study undertaken by Warrington, Oldfield and myself of performance of patients with localized lesions used a battery of 10 tests that were simplified versions of ones where from the literature one might expect a frontal deficit, including three of Luria's and word fluency (Milner 1964), the Stroop (Perret 1974) and Cognitive Estimates (Shallice & Evans 1978). Yet even on these tests only one produced a significant effect of anterior-posterior location in the basic analysis. This was a category sorting task, a type of task that has quite often given frontal deficits (see, for example, Milner 1964) – in our case the simplest possible variety, Weigl's Sorting Task (Weigl 1941).

Why so many tasks do not give specific deficits with frontal lesions, I shall discuss briefly later. Whatever the reason, one implication is that the existence of a high-level general-purpose programming unit is not yet clearly established from neuropsychological evidence.

The present theory, though, has an advantage over older related theories, which may help to obtain more clear-cut evidence. On older theories what tasks would require this sort of high-level cognitive system can only be intuitively specified. The present model allows more rigorous inferences to be made in two ways, the more indirect by using the properties of contention scheduling, the more direct by using the properties of the S.A.S.

## USING THE PROPERTIES OF CONTENTION SCHEDULING

If the S.A.S. is inoperative, this will leave the organism operating under the control of contention scheduling alone. The behaviour of such an organism should then be predictable from the properties of contention scheduling. Two of its characteristics are relevant. First, as it is a 'routine' structure its parameters change slowly. If then the environmental situation is such that a trigger is present that strongly activates a schema, it will not normally be possible to prevent the schema being selected. Response perseveration should occur.

By contrast, consider the situation where there are no strong environmental triggers. In the intact organism the appropriate schema would need to receive additional input from the S.A.S. to ensure its selection. However, if the S.A.S. is absent, possibly the activation level of all schemas will not reach threshold, or random fluctuations could cause another inappropriate schema to become dominant. This possibility is increased as recent inputs have a heavier weighting. Conflict resolution in production systems – the model for contention scheduling – often utilizes this procedure (see McDermott & Forgy 1978), and in contention scheduling it takes place through trace decay in the trigger data base. So, when there is no S.A.S., the capture of contention scheduling by a schema triggered by any new input is especially likely to occur. One should therefore, as a result of the lesion, observe distractibility in some situations but in others response perseveration.

The combination of distractibility and an inability to shift from making the well learned response to a stimulus has long been known to occur with frontal lobe lesions in animals. Pribram (1973) gives a number of examples of increased distractibility in such animals. Moreover, the classical deficit in delayed response tasks that such animals show, together with their hyperactivity, can be reduced to a very considerable extent by minimizing distractions such as by putting the animal in the dark. Complementarily they have difficulty in extinction situations and in discrimination reversal (see Fuster (1980) for a review of the relevant literature).

In humans this combination of deficits is less well established. Abnormal perseveration, of action and thought schemas, not merely percepts, is very well demonstrated in the Wisconsin Card-Sorting Task (Milner 1964; Nelson 1976), and Milner showed that with dorsolateral frontal lesions this can happen in patients with normal scores on W.A.I.S. I.Q. Distractibility is less well established experimentally but frequently mentioned in clinical reports (see, for example, Rylander 1939).

## EXTRAPOLATION FROM ARTIFICIAL INTELLIGENCE

Any strong support for the idea that specific impairments of an S.A.S. system can be observed requires, however, some principled way of selecting a task that would make strong use of its functions. The obvious method is to attempt to derive predictions from the field in which the theory originated, namely artificial intelligence. In theory one should develop a simulation of a task that makes heavy use of a general-purpose planning system, show that the behaviour of normal subjects conforms to the simulation and show that specific deficits on this task can occur. Here for the first two requirements I shall rely on extrapolations from tasks in the same problem domain.

One problem domain much used in artificial intelligence simulation of problem solving is that of look-ahead puzzles in which a set of stacks of blocks has to be constructed from a starting

configuration in series of individual moves; the Tower of Hanoi is one example. The difficulty in such problems lies in determining the appropriate order of simple moves rather than in any inherently spatial factors. It is an appropriate domain for planning programs as the goal is achieved by decomposition into subgoals. Yet this decomposition can lead to special planning difficulties through the tackling of the subgoals in the wrong order, through a pre-requisite for the solution of a subgoal being missing or even through the problem being such that any purely linear sequence of achieving subgoals could not work (see Sussman 1975). It is also an appropriate domain for the present purposes, as most people would not appear to have many special-purpose subroutines available for it. It is therefore plausible that they too would have to have recourse to a general programming unit if they have one.




FIGURE 2. Three subproblems of the Tower of London test. The initial position is the same for all.

Sussman's (1975) seminal problem-solving program HACKER, which incorporates a general-purpose programming unit, was tested by using this type of problem. Moreover, Anzai & Simon (1979) have simulated the detailed protocol of an individual subject's repeated attempts to solve the Tower of Hanoi. This simulation too incorporates a high-level learning mechanism 'completely independent of the particular task and fully applicable to other problem-solving environments in which heuristic search occurs'. The problem domain therefore satisfies the first two requirements.

A particular subclass of problems in this domain was designed that were tractable for normal subjects and whose parameters could be varied to produce a graded difficulty test. In each problem three beads, one red, one green and one blue, have to be moved from a starting configuration on three sticks of unequal length to a target position in a minimum number of moves. Four problems are 2 or 3 moves deep, four are 4 moves deep and four are 5 moves deep. In the easier problems (e.g. problem no. 2 (see figure 2)) the planning resources required are minor; thus a strategy of moving directly to target pegs can work. In the medium level of difficulty a means-ends analysis or error correction procedure is required, since the simple strategy of working forward by moving directly to the target-peg could only succeed if a less obvious first move were tried (B  $\rightarrow$  peg 1 in no. 6 is non-dominant in contention scheduling, as the first move in the previous problem no. 5 is R  $\rightarrow$  peg 1 or peg 2). In the harder problems, quite complex planning difficulties can occur. Thus in no. 10 achieving the subgoal of reaching the target for the red bead can hinder the subgoal of achieving it for the blue; indeed this problem is very similar to one that led HACKER into a bug it could not cope with simply at one stage of its development (see Sussman 1975, p. 106).

In a collaborative experiment with R. McCarthy I tested 61 patients from the National Hospital with unilateral localized lesions of various aetiologies and 20 control subjects on the task. The patients were allocated to anterior or posterior groups by using a procedure easy to

apply with the computed tomography (c.t.) scan. Patients were placed in an anterior group if their lesion involved the frontal lobe and more than half of its mass was anterior of the bisector of the line joining the nasion to the inion (anterior and posterior skull sites used in the National Hospital c.t. scan procedure); otherwise they were placed in a posterior group. The mean number of problems solved at the first attempt in less than 60 seconds by each group of patients is shown in figure 3. For the patient groups there was a significant interaction between hemisphere and anterior-posterior location with comparisons between the left hemisphere groups and between the anterior groups both giving a significant deficit for the left anterior group. An analysis with a scoring system allowing corrections and taking into account solution time gave similar results. There was therefore a specific deficit on the task for the left anterior group.




FIGURE 3. Tower of London test: percentage correct at the first attempt for the four lesion groups and the control group.

The left anterior group was not impulsive. The time from when the card with the target position was presented to the first move of the beads was measured for each problem. Analysis of these initial planning times irrespective of the success of the solution attempt gave a significant interaction between hemisphere and anterior-posterior location within the lesion groups; the left anterior group was the slowest of all.

One of the many major methodological problems with group studies is the patient selection procedure even in an 'unselected' consecutive series. Based as it is on a clinical process, in no way does it provide any sort of random sample of possible lesion sites. Severity differences can therefore exist between groups. Moreover, as all psychological tests have multiple components, a deficit on a task does not necessarily implicate any particular component. However, as expected, none of the other tasks used for control purposes (baseline tests) produced an anterior deficit. More important, covarying the results of the critical task against those of other tests strengthened the theoretical inferences.

The Tower of London task, as we call the planning task, appears to have many processing components in common with the W.A.I.S. subtest, Block Design, both being visual-motor tasks with the possibility of verbal mediation. However, when the analysis of the present results

is repeated covaried against Block Design, the form of the interaction is unchanged. It cannot therefore be caused by the components that they utilize in the same fashion. This conclusion is reinforced by the results of individual patients where a double dissociation exists (this is only true for the number solved measure of the Tower of London test). Two left anterior patients scored below any normal subject on the Tower of London (25 and 42 %) but well on Block Design (scaled scores 10 and 17), while two right posterior patients scored well on the Tower of London (67 and 84 %) but below any control on Block Design (scaled scores 5 and 4). Therefore the two tasks appear to be differentially loaded on two factors, most plausibly a planning and a spatial processing one. In particular, the Tower of London appears to have a major planning component.

However, before accepting this interpretation two other possibilities need to be considered. One other difference that exists between the two tasks is that the Tower of London task has the greater short-term storage load. However, covarying the results against span again leaves the interaction unaffected. Moreover, on anatomical grounds this suggestion seems unlikely: a deficit of short-term storage for verbal or visual information would be expected with posterior rather than anterior lesions (see Walsh (1978) for review).

The most plausible alternative to a planning explanation, given that the deficit is observed in left anterior but not right anterior patients, is that some verbal process is being interfered with. This could be just verbal mediation or the process that Luria called the 'regulating function of speech'. As an indirect test of the hypothesis that the left anterior deficit was attributable to some speech-based process, R. McCarthy and I, using normal subjects, investigated the effect on the task of articulatory suppression. The subjects continually repeated 'ABCDEFGH' to themselves while doing it. Now suppressing, for instance, increased the subjects' error rate from 1 to 12 % on a task that requires the direction of action by inner speech – a modified version of the De Renzi & Vignolo (1962) Token Test. Moreover, Baddeley *et al.* (1981) have found a pattern of deficits by using articulatory suppression, which is compatible with its affecting at a later time information being retained with the help of inner speech. Therefore if the left anterior deficit is due to a dysfunction of some speech-based process, then articulatory suppression would be expected to affect performance of the task. In fact there was no support for this hypothesis as the effect was totally insignificant.

A more direct test of the planning hypothesis would require a more detailed analysis of the results. A *post hoc* analysis of this sort in general supports a planning interpretation. In particular, a 'planless' procedure of working forward by using the *a priori* most obvious moves until the beads can be moved directly to their target positions will work for questions 1–5, but not after that. The left anterior group shows a very sharp drop in performance after question 5, but the other patients do not at all. Moreover, this effect is present on the analysis of the nature of the first move, so it is not merely a failure to remember or implement a plan completely.

## DISCUSSION

The Tower of London Task was constructed so that satisfactory performance would require the use of a general programming system, such as the Supervisory Attentional System, if such a system exists. A particular group of patients, the left anterior group, was found to be specifically impaired on the task. Moreover, alternative interpretations of the deficit were not supported in subsidiary investigations. In particular the pattern of results on the Tower of London



test is very different from those on baseline tests, none of which produce a specifically anterior deficit, and in particular from that on Block Design, a test that must otherwise overlap very much in the processing resources it requires. The findings therefore support the idea that a general programming system, such as the S.A.S., is necessary for satisfactory performance of non-routine tasks, such as those that involve planning or the overcoming of previously dominant associations; but that it is not necessary for 'routine' tasks, where the previously learned relations between triggers and schemas are sufficient to ensure appropriate schema selection in contention scheduling.

Two questions are posed by this interpretation. First, this characterization of baseline tests as being 'routine' and not requiring programming resources may well be true of a task like span, but is it not an oversimplification for a task like Block Design? Deficits on Block Design have been found in certain patients with frontal lobe lesions (Lhermitte *et al.* 1972). However, the question is resolved if, in accordance with their findings, one assumes that Block Design requires both spatial processing resources and programming resources, with the former being the limiting factor in many right posterior patients and the latter the crucial deficit in occasional anterior patients. The different pattern of results obtained between Block Design and the Tower of London tasks still needs to be explained in terms of two types of processing resource – most plausibly programming and spatial – but the difference becomes one of degree, the Tower of London requiring the more programming resources and the less spatial resources.

The requirement for specific processing resources in many tasks designed to assess S.A.S. functions relates also to the second issue. This is a much more serious problem than the first. If the S.A.S. exists and is impaired in left anterior lesions, why in group studies where the groups are anatomically defined are specific deficits not more often found in groups with anterior lesions by comparison with an appropriate control group, namely patients with more posterior lesions? Worse still, it may be that replication poses an especial problem in such studies. The replication failures that we surprisingly obtained are not unique; thus considering word fluency, the classic finding is of a specific left frontal deficit (Milner 1964), but Ramier & Hécaen (1970) and Perret (1974) found a right frontal deficit too, while Newcombe (1969) and Coughlan (1976) found no specific frontal deficit!

In my view the methodological problems, well recognized but rarely discussed, that are always involved in neuropsychological group studies are, if anything, greater for detecting potential anterior deficits. In particular this is so for group composition, which depends upon the clinical population available and on complex interactions between lesion sites and extent, aetiology and the patient's suitability for testing (Teuber, in the discussion of Milner (1964), makes a relevant analysis of replication difficulties on Wisconsin Card-Sorting).

Yet another factor is that cognitive tasks derived from the psychometric tradition typically contain questions that are gently graded in difficulty. This may have the effect of 'shaping' appropriate trigger-schema relations; on the difficult items the strategy is partly made routine.

The possibility of obtaining tasks showing specific impairment of the S.A.S. by using group studies is also reduced by a more positive factor. Damage to the S.A.S. would theoretically not necessarily have to vary only in degree. There could be qualitatively different types of impairment. So the performance of tasks that require certain of its components need not be impaired by a lesion that affected other components.

The plausibility of this type of fractionation of the system is increased by considering the analogy between the S.A.S. and the corresponding part of Sussman's (1975) program HACKER.

This contains a number of functionally distinct units, of which, for instance, the impairment of one (the 'Critics Gallery') could lead to impulsiveness, and of another (the 'Bug Identifier') to perseveration. However, the identification of such hypothetical syndromes by means of group studies, especially given the grave problems with this methodology, would probably prove extremely difficult. Instead a shift will probably be required to the methodology already proving more powerful for the investigation of impairments of specific resources, namely the case study approach (see Shallice 1979).

Considerable methodological difficulties therefore remain in the neuropsychology of planning, while conceptually the field remains rather primitive. For instance the critical relation between planning and language has not been discussed here at all. However, the overall plausibility of a general programming system, S.A.S., which can be separately impaired, has been considerably strengthened by these results. Yet it might appear that this theoretical position merely echoes classical views on impairments of higher mental function. If one approached the subject with poetic licence, the S.A.S. could almost be characterized as 'the special workshop of the thinking processes', which is how Burdorch described the frontal lobes in 1819 (see Rylander 1939, p. 14)! However, the tortuous history of the investigation of such disorders since then (see Rylander 1939; Teuber 1964; Walsh 1978) should reinforce the need to have such concepts properly grounded theoretically. Links with cognitive psychology and even artificial intelligence should provide the necessary foundations.

I am grateful to Elizabeth Warrington for providing facilities for the research to be undertaken and for many helpful discussions, to Roz McCarthy for her assistance over many aspects of the paper and to Marie-France Beauvois for her advice on the first draft of the paper.

#### REFERENCES

- Anzai, Y. & Simon, H. A. 1979 The theory of learning by doing. *Psychol. Rev.* **86**, 124-140.
- Baddeley, A. D., Eldridge, M. & Lewis, V. J. 1981 The role of subvocalization in reading. *Q. Jl exp. Psychol.* (In the press).
- Boden, M. 1977 *Artificial intelligence and natural man*. New York: Basic Books.
- Coughlan, A. 1976 An investigation of selected verbal skills and verbal memory following left hemisphere lesions. Ph.D. thesis, University of London.
- De Renzi, E. & Vignolo, L. A. 1962 The Token Test: a sensitive test to detect receptive disturbances in aphasia. *Brain* **85**, 665-678.
- Fuster, J. M. 1980 *The prefrontal cortex*. New York: Raven.
- Lhermitte, F., Derouesné, J. & Signoret, J. L. 1972 Analyse neuropsychologique du syndrome frontal. *Revue neurol.* **127**, 415-440.
- Luria, A. R. 1966 *Higher cortical functions in man*. London: Tavistock.
- McDermott, J. & Forgy, C. 1978 Production system conflict resolution strategies. In *Pattern-directed inference systems* (ed. D. A. Waterman & F. Hayes-Roth), pp. 177-199. New York: Academic Press.
- Miller, G. A., Galanter, E. H. & Pribram, K. H. 1960 *Plans and the structure of behavior*. New York: McGraw-Hill.
- Milner, B. 1964 Some effects of frontal lobectomy in man. In *The frontal granular cortex and behavior* (ed. J. M. Warren & K. Akert), pp. 313-334. New York: McGraw-Hill.
- Nelson, H. 1976 A modified card sorting test sensitive to frontal lobe defects. *Neuropsychologia* **12**, 313-324.
- Newcombe, F. 1969 *Missile wounds of the brain*. London: Oxford University Press.
- Newell, A. & Simon, H. A. 1972 *Human problem solving*. Englewood Cliffs, N.J.: Prentice-Hall.
- Norman, D. A. 1981 Categorization of action slips. *Psychol. Rev.* **88**, 1-15.
- Norman, D. A. & Shallice, T. 1980 Attention to action: willed and automatic control of behavior. *Center for Human Information Processing Technical Report* no. 99.
- Perret, E. 1974 The left frontal lobe in man and the suppression of habitual responses in verbal categorical behavior. *Neuropsychologia* **12**, 323-330.
- Pribram, K. H. 1973 The primate frontal cortex - executive of the brain. In *Psychophysiology of the frontal lobes* (ed. K. H. Pribram & A. R. Luria), pp. 293-314. New York: Academic Press.

- Ramier, A. M. & Hécaen, H. 1970 Rôle respectif des attientes frontales et de la latéralisation lésionnelle dans les déficits de la 'fluence verbale'. *Revue neurol.* **123**, 17-22.
- Reason, J. T. 1979 Action not as planned. In *Aspects of consciousness* (ed. G. Underwood & R. Stevens), pp. 67-89. London: Academic Press.
- Reason, J. T. 1982 Lapses of attention. In *Varieties of attention* (ed. R. Parasuraman, R. Davies & J. Beatty). New York: Academic Press. (In the press.)
- Reitan, R. 1964 Psychological deficits resulting from cerebral lesions in man. In *The frontal granular cortex and behavior* (ed. J. M. Warren & K. Akert), pp. 295-312. New York: McGraw-Hill.
- Rylander, G. 1939 Personality changes after operations on the frontal lobes. *Acta psychiat. neurol. Suppl.* no. 30.
- Shallice, T. 1972 Dual functions of consciousness. *Psychol. Rev.* **79**, 383-393.
- Shallice, T. 1979 Case study approach in neuropsychological research. *J. clin. Neuropsychol.* **1**, 183-211.
- Shallice, T. & Evans, M. E. 1978 The involvement of the frontal lobes in cognitive estimation. *Cortex* **14**, 294-303.
- Shallice, T. & McCarthy, R. 1982 (In preparation.)
- Shiffrin, R. M. & Schneider, W. 1977 Controlled and automatic human information processing. II. Perceptual learning, automatic attending, and a general theory. *Psychol. Rev.* **84**, 127-190.
- Sussman, G. J. 1975 *A computational model of skill acquisition*. New York: American Elsevier.
- Teuber, H-L. 1964 The riddle of frontal lobe function in man. In *The frontal granular cortex and behaviour* (ed. J. M. Warren & K. Akert), pp. 410-444. New York: McGraw-Hill.
- Walsh, K. W. 1978 *Neuropsychology: a clinical approach*. Edinburgh: Churchill Livingstone.
- Weigl, E. 1941 On the psychology of the so-called process of abstraction. *J. abnorm. soc. Psychol.* **36**, 3-33.