Modelling the Selection of Routine Action: Exploring the Criticality of Parameter Values

Richard Cooper
Department of Psychology
Birkbeck College
University of London
Malet St., London, WC1E 7HX
R.Cooper@psyc.bbk.ac.uk

Tim Shallice
Department of Psychology
University College
University of London
Gower St., London, WC1E 6BT
T.Shallice@psychol.ucl.ac.uk

Abstract

Several authors have distinguished automatic behaviour of routine or well-learnt action sequences from controlled behaviour of novel actions. In this paper we present an interactive activation model of routine action selection based on the Contention Scheduling theory of Norman & Shallice (1986). The model, developed in the specific domain of coffee preparation, provides a good account of normal behaviour in a complex yet routine task. In addition, we report lesioning studies which show breakdown of action selection qualitatively similar to that seen in a variety of neurological patients (action disorganisation syndrome, utilisation behaviour, and Parkinson's disease). These lesioning studies are based on the systematic variation of critical system parameters. Such parameters, which are implicit in all interactive activation models, raise complex methodological issues relating to the criticality of their values. We address these issues by reporting results of a detailed exploration of the parameter space.

Introduction

Within the psychological literature on action there is a frequent distinction between two qualitatively different types of action. Schneider & Shiffrin (1977), for example, distinguish between automatic and controlled action, whereas Norman & Shallice (1986) distinguish between routine and non-routine action. The distinction essentially concerns the degree to which an action complex can be carried out in the absence of intentional control, and is seen most clearly within dual-task situations. Thus, it is frequently possible to carry out a routine or well-practiced task (e.g., driving) whilst simultaneously performing an attentionally more demanding task (e.g., having a complex conversation), with little or no interference between the tasks. When the two tasks are both attentionally demanding, however, (e.g., negotiating a new set of roadworks whilst having a complex conversation) simultaneous performance becomes very difficult.

Norman & Shallice (1986) and Shallice (1988) argue that distinct cognitive systems are responsible for the control of routine and non-routine action, and that the different properties of these two forms of action result from differences in the systems underlying their performance.

According to the Norman & Shallice theory, routine action is controlled by selection mechanisms operating on an activation-based network of action schemas, with the selected schema controlling action. They term this system Contention Scheduling (CS). Non-routine action, on the other hand, makes use of an attentionally demanding central process, the Supervisory Attentional System (SAS). In non-routine situations the SAS controls behaviour by modulating the activations of schemas in CS, but in routine situations CS may function autonomously.

Routine and non-routine action are further differentiated by the types of errors which characterise the domains. Error types in routine action include: omission and anticipation errors (i.e., neglecting to perform a sub-action, or performing one sub-action too early, such as adding boiling water to an empty teapot, having neglected to first add tea); capture errors and utilisation behaviour (in which behaviour is "captured" by the environment and diverted from some initially intended action to some other environmentally relevant behaviour); object and place substitutions (in which an appropriate action is performed with inappropriate arguments (e.g., having prepared a mug of coffee, putting the coffee, instead of the milk, in the fridge); and perseverations (i.e., performing one sub-action multiple times, even after the sub-action's goal has been achieved). These types of error, or lapses, are not observed in the controlled execution of non-routine action.

The domain of action differs from a number of other cognitive domains which have been modelled in that it necessarily involves sequentiality. Domains such as visual processing, object recognition, and reading, can be viewed in terms of essentially static mappings. Action, in contrast, requires that issues of sequentiality be addressed, both in terms of constraints imposed by task structure (some actions, e.g., stirring the coffee, may only be possible once certain preconditions have been satisfied, e.g., picking up a spoon or stirrer) and physical limitations imposed by resource constraints (e.g., the number of limbs available).

One consequence of the inherent sequentiality of the action domain is that many standard modelling techniques (e.g., multi-layer feedforward networks) are unsuitable. Data on lapses in routine action (as described above) also argue against approaches in which each action acts as a cue for its following action so yielding an action sequence (cf. Houghton & Hartley, 1996), such as recurrent network approaches (e.g., Jordon, 1986; Elman, 1990). Indeed, action lapses in normals and patients, such as those described above, were a fundamental consideration in the development of the CS theory in terms of interactive activation. The model reported here is therefore based on a number of interactive activation networks (cf. Mc-

Clelland & Rumelhart, 1981), in which action is controlled via an activation-based selection mechanism. The model is reported in more detail elsewhere (Cooper, Shallice, & Farringdon, 1995). In this paper we 1) present an updated version of the model; 2) provide new results of lesioning studies; and 3) provide a detailed analysis of the model's parameter space.

The Model

The central component of the CS theory is the schema network. In formal terms, this is a directed acyclic graph in which nodes correspond either to action schemas (partially order sequences of actions) or goals. Edges from the schema nodes point to goal nodes, and *vice versa*, so the network may be thought of as consisting of alternate "layers" of schema and goal nodes. The schema nodes pointed to by a goal node correspond to distinct methods of achieving that goal, and the goal nodes pointed to by a schema node represent the subgoals which must be achieved in order to execute the schema. Thus, the goal of preparing a mug of coffee might be achieved in a variety of ways (e.g., either by preparing instant coffee or by preparing percolated coffee), and the schema for preparing instant coffee will include a number of subgoals (e.g., boil water, add coffee grinds, add cream, etc.). It is assumed that in reality a person's schema network will contain nodes for all of that person's routine activities, and that the routinisation of activities involves the addition of new nodes to the network.

Each schema node has a state and a numerical activation value (which, in the simulations described here, can range from 0.0 to 1.0). The state is either selected or not selected, and affects the flow of activation throughout the schema network. When selected, schemas pass activation to their subschemas (i.e., those schemas which may achieve the schema's subgoals). This tends to excite subschemas, causing them to become selected, and hence to activate their subschemas in turn. This top-down flow of activation is tempered by three further activation sources. Firstly, all schemas have triggering conditions. These are conditions which, when satisfied by the external world, cause the schema's activation to increase. Thus, the triggering conditions for a pick-up schema might include the existence of a suitable pickup-able object within reach. The existence of such an object would then tend to activate the pick-up schema. Evidence for this source of excitation comes from capture errors, in which behaviour can be diverted by aspects of the environment into familiar, but unintended, action sequences (e.g., William James' (James, 1890) example of changing into his pyjamas when moments before he had been intending to dress for dinner).

The two remaining sources of excitation/inhibition act to stabilise the network and prevent all schemas from becoming simultaneously active. Lateral inhibition acts between schemas which either correspond to alternate means of achieving the same goal (e.g., preparing instant versus percolated coffee), or which share resource or subschema requirements (which is assumed to amount to sharing a common subgoal). Lateral inhibition tends to inhibit activation values, but acts differentially across competing schemas, such that more highly activated schemas are inhibited less. Acting against lateral inhibition is self activation, which is a further excitatory source of activation applying to all schemas. In the absence of top-down and environmental activation sources, lateral inhibition and self activation act to yield a network that is in an unstable equilibrium. Top-down and environmental activation force the system from this state into local minima in which at most one schema from each competing subset is highly active.

As noted above, the schema/goal network is layered. Nodes at the bottom layer correspond to primitive actions. These are assumed to correspond to low-level units of action which are non-decomposable at the cognitive level. For the purposes of the current simulations, this level includes schemas for, for example, pick-up (though we recognise that such schemas could be further decomposed into rotate-forearm, grasp, etc.). Selection of a primitive action triggers the corresponding motor action, which will generally result in the schema's goal being achieved.

When a schema's goal is achieved, the schema is temporarily inhibited. This reduces the lateral inhibition on competing schema, pushing the system out of its local minimum and allowing competing schemas the opportunity to control behaviour.

Schema selection is threshold based. If a schema's activation exceeds the selection threshold (0.6 on a scale of 0.0 to 1.0 in the simulations reported here) the schema is selected (unless a competing schema is even more active). Selected schemas remain selected until their activation is exceeded by that of a competitor (even if their activation falls below the selection threshold).

The most critical part of the model that remains to be described concerns the nature of object representations and the relationship between these and the schema network. The original verbal descriptions of CS did not specify how objects were selected by schemas in order to produce complete action specifications (e.g., if pick-up is selected, how does the effector system know what should be picked up?). It was found necessary to extend the CS theory in order to produce a complete simulation.

We therefore assume that, parallel to the schema network, there exists a network of object representations. These representations compete through the effects of lateral inhibition and self activation. When a primitive schema is selected, the objects corresponding to the most active appropriate representations are used to fill the argument roles of the corresponding primitive action.

The triggering of schemas by the external world is mediated by the activation values of object representations (such that highly active representations are more effective triggers). Object representations are also activated by schemas in which the corresponding objects may participate. Thus, the representation of a cup on a table will tend to excite the pick-up schema, and *vice versa*. In the absence of sufficient top-down activation, such positive feedback loops can lead to utilisation behaviour and capture errors.

The above system would be sufficient if all schemas had at most one argument. However, some schemas require multiple arguments (e.g., empty *the spoon* into *the coffee mug*). Such schemas require multiple objects to be active at the moment of argument selection. In the current implementation we solve this problem by introducing multiple activations for object representations. Each representation has a separate activation for each functional role that the corresponding object can serve. Thus, a mug containing a partially made cup of coffee can act as either a source or a target for the movement of coffee-making ingredients. Given this, a successful attempt at emptying a specific spoon into a specific coffee mug will require that, when the empty schema is selected, the spoon is active as an implement and the mug is active as a target.

Processing within the system, as in most interactive activation networks, is cyclic. On each cycle the activations of all schema nodes are recalculated, based on their activation at the beginning of the cycle (modulated by a persistence/decay factor) and their net excitation or inhibition (which is blurred through the addition of normally distributed random noise). The selection mechanism then adjusts the state of all nodes based on the new activation values.

The precise behaviour of the complete simulation is determined by eight parameters which specify the levels of, and relationships between, the various activation sources. These are:

S/L A factor which controls the relative proportions of self activation and lateral inhibition within all interactive activation networks.

I/E(s) A factor which controls the relative proportions of internal or top-down activation and external or environmental activation of all schema nodes in the schema interactive activation network.

I/E(o) A factor used to scale the external or schema-based activation of all object representation nodes in the object interactive activation network.

Comp/NonComp A factor which controls the relative proportions of competitive (i.e., self activation + lateral inhibition) and non-competitive (i.e., top-down or schema-based + environmental) activations which contribute to a node's net excitation/inhibition.

Persistence The degree to which all activations persist from one cycle to the next.

Noise The standard deviation of normally distributed random noise added to all net influences.

Selection Threshold The activation which must be exceeded before a schema can be selected.

Rest Activation The activation level to which schema nodes and object representations with no net input tend.

Each of these parameters may range in value from 0.0 to 1.0.

The large number of parameters is due to our insistence on being explicit about all numerical factors which may influence the simulation's behaviour. We discuss our methodology for exploring this parameter space below.

General Behaviour

The CS model as described here is intended to model action control in complex but well-learned tasks, such as cleaning one's teeth, starting a car and dressing. The complexity of such activities makes them relatively unattractive for investigation by standard experimental psychology methods and, apart from natural history studies of lapses in their execution (e.g., Reason, 1979, 1984), we know of no such investigations. It is therefore not possible to provide a detailed quantitative comparison of the model's behaviour with that of subjects on some appropriate task. There are, however, detailed neuropsychological studies relevant to the domain. The most detailed of these have been carried out on the Action Disorganisation Syndrome (Schwartz, Reed, Montgomery, Palmer, & Mayer, 1991; Schwartz, Mayer, Fitzpatrick-De Salme, & Montgomery, 1993). These studies have examined the behaviour of patients performing routine daily activities (such as preparing their breakfast from the objects provided on a hospital breakfast tray). We have therefore attempted to simulate one such task, that of preparing the morning coffee, with evaluation being based on the ability to simulate normal behaviour and, after lesioning, the ability to produce the types of errors shown by patients performing the task.

Normal Behaviour

With appropriate values for the parameters, the system is able to correctly perform the complex task of coffee preparation. This task involves adding coffee granules, sugar and milk to an existing mug of hot water. A transcript of the behaviour is shown in figure 1. In this figure, schemas are shown in italics, with selection being indicated by a '+' prefix and deselection indicated by a '−' prefix. Primitive actions are shown in roman font. The level of indentation shows the degree of embedding of subschemas within parent schemas.

Note that this protocol was generated with a situation containing a number of coffee and sugar sources. One effect of random noise is to yield different choices of which sugar and coffee sources to use.

Lesioning Studies

In order to evaluate the behaviour of the model after damage, we conducted a number of simulations in which key parameters were varied. We suggest that parameters may be related to neurophysiological damage or disruption in two possible ways. Firstly, certain parameters may reflect neurotransmitter concentration or sensitivity. For example, we believe there to be a relationship between the effectiveness of the dopamine system and the S/L parameter. Secondly, parameters based on relative proportions of excitation between different subsystems (e.g., I/E(s) and I/E(o)) may reflect relative connectivity within and between those subsystems. It can be argued that variation of these parameters corresponds to physical damage to such connectivity.

Our lesioning studies have focussed on the variation of two parameters. Firstly, we have investigated the behaviour of the model whilst systematically varying the I/E(s) parameter over its entire range (from 0.00 to 1.00, at 0.01 intervals, with 20 simulations at each value). This parameter was considered in order to test the hypothesis that action disorganisation can be understood in terms of insufficient top-down or intentional

```
+coffee
  +coffee-from-pack
    +pick-up-schema
    Picking up coffee-packet1 with left-hand
    -pick-up-schema
    +tear-schema
    Tearing coffee-packet1 (with right-hand)
    -tear-schema
    +pour-schema
    Pouring coffee-packet1 into coffee-mug1
    -pour-schema
    +put-down-schema
    Putting coffee-packet1 down
    -put-down-schema
  -coffee-from-pack
  +sugar-from-bowl
    +pick-up-schema
    Picking up spoon1 with left-hand
    -pick-up-schema
    +dip-spoon-schema
    Dipping spoon1 into sugar-bowl1
    -dip-spoon-schema
    +empty-spoon-schema
    Emptying spoon1 into coffee-mug1
    -empty-spoon-schema
    +put-down-schema
    Putting spoon1 down
    -put-down-schema
  -sugar-from-bowl
  +milk-from-carton
    +pick-up-schema
    Picking up milk-carton1 with right-hand
    -pick-up-schema
    +tear-schema
    Tearing milk-carton1 (with left-hand)
    -tear-schema
    +pour-schema
    Pouring milk-carton1 into coffee-mug1
    -pour-schema
    +put-down-schema
    Putting milk-carton1 down
    -put-down-schema
  -milk-from-carton
-coffee
```

Figure 1: Action selection in the coffee preparation domain

control within the CS system (cf. Schwartz et al., 1991). Decreasing the parameter from the value required for normal behaviour (approximately 0.95), is equivalent to decreasing the top-down control within the schema subsystem whilst simultaneously increasing the environmental influence within the subsystem. As the parameter was varied, qualitatively different forms of behaviour were observed. Occasional utilisation and object substitution errors arose with very slight reductions (with I/E(s) ≈ 0.90). With moderate reductions (I/E(s) ≈ 0.60), utilisation errors dominated, and with large reductions (I/E(s) ≈ 0.20), utilisation errors were mixed with perseverative errors. These results provide some support for the initial hypothesis, in that all these forms of behaviour were observed to varying degrees in the studies reported by Schwartz *et al.* (1991, 1993). There are difficulties, however, in interpreting the data. In particular, many attempted action sequences include physically impossible subsequences (e.g., attempting to pour the contents from a container without first opening the container, or attempting to stir the coffee without first picking up a stirring implement). It is not clear that such errors do not invoke supervisory attention. If this is the case, then it will be necessary to model at least some SAS functions before the behaviour of the model after error can be compared with that of patients.

A further set of lesioning studies examined the dependence of the model's behaviour on the parameter governing the ratio of self activation and lateral inhibition (S/L). These studies were motivated by the arguments of Robbins & Sahakian (1983) that activation of the striatal dopamine system corresponds to increased activation of schemas within the CS model. These arguments suggest that decreasing S/L should result in behaviour similar to that shown by Parkinson's disease patients, in whom striatal dopamine is known to be deficient (Robbins, 1991). One feature of the behaviour of Parkinson's patients is that action initiation is typically greatly slowed, but once action has been initiated, it may proceed relatively normally. This behaviour was indeed exhibited by the simulation. The coffee preparation task was performed accurately over a wide range of the S/L parameter's values (from 0.20 to 0.65), but at lower values within that range action initiation was disproportionately slowed, with the onset of action being delayed by as much as the time required for the "normal" simulation to complete the entire task.

Parameter Criticality

As noted above, the current model has a large number of parameters. In order to fully understand the model's behaviour, and avoid charges of parameter fitting, we have conducted extensive explorations of the parameter space.

The first issue to be addressed concerns finding a set of parameter values that yields qualitatively normal behaviour. The model is sufficiently well-behaved to allow a systematic approach to finding such a parameter set. The first stage is to fix the parameters which govern global network behaviour. The relevant parameters are persistence and rest activation, and the relevant global behaviour is smoothness of activation profiles. This is done by setting noise to be very low and the selection threshold very high. The various weighting parameters (S/L, Comp/NonComp, I/E(s), and I/E(o)) can be set to any value in their range (zero to one) as alterations to these parameters do not affect the gross activation flow throughout the network.[1] Rest activation and persistence can now be counterbalanced in order to achieve smooth activation profiles. This amounts to searching a two-dimensional space with one degree of freedom, and presents no difficulties.

[1] These parameters were engineered precisely to have this property: altering any one of them will affect the relative weighting of activation sources, but not the total activation flow over a suitably long time interval.

Once smooth activation profiles have been achieved, attention can be focussed on competitive effects. With the Comp/NonComp parameter set to 1.0 (so that there are only competitive influences within the networks), the S/L parameter can be adjusted to yield appropriate competitive behaviour. Similarly, with Comp/NonComp set to 0.0, the parameters governing non-competitive activation flow (I/E(s) and I/E(o)) can be adjusted. I/E(s) must be set so as to get appropriate action selection (with no utilisation behaviour, but with some environmental triggering). I/E(o) must be set to ensure that appropriate arguments are active when low-level schemas are selected. Comp/NonComp can then be adjusted to yield a suitable balance between the competitive and non-competitive activation sources.

This approach is not guaranteed to yield a configuration of parameters leading to error-free behaviour, and adjustments may be required at each step to maintain the properties obtained at the previous step, but it does provide a reliable means of producing a point in the eight-dimensional parameter space that is sufficiently close to a well-behaved parameter setting to allow fine tuning. Four rather different parameter settings found by the above approach, all of which yield well-formed (i.e., virtually error free) behaviour, are given in table 1. The first two columns of this table show how modifications to persistence can be counterbalanced by adjustments to rest activation, independently of other parameter values, whilst the last two columns demonstrate how the other parameters can be adjusted once persistence and rest activation have been set.

S/L	0.50	0.50	0.60	0.60
I/E(s)	0.95	0.95	0.95	0.95
I/E(o)	0.10	0.10	0.07	0.05
Comp/NonComp	0.65	0.65	0.80	0.80
Persistence	0.80	0.88	0.80	0.88
Noise	10^{-3}	10^{-3}	10^{-4}	10^{-4}
Sel. Threshold	0.60	0.60	0.70	0.50
Rest Activation	0.10	0.05	0.10	0.05

Table 1: Four parameter configurations yielding well-formed behaviour

A substantive test of parameter criticality consists of showing that behaviour remains qualitatively normal in a sizable region surrounding some "standard" parameter values. In order to investigate parameter criticality in this sense we performed a number of simulations in which each of the eight parameters was varied whilst all other parameters were held fixed at standard values.[2] This yielded, for each parameter, a range of stable behaviour, as shown in table 2.

Table 2 shows that considerable variation in some parameters (most notably S/L and Comp/NonComp) may be tolerated without qualitatively affecting performance of the task. In order to further investigate parameter criticality we extended this investigation by varying a range of parameters simultaneously. This was intended to examine stability over a larger, multidimensional region of the parameter space. A range of values was chosen for each of five parameters (S/L (3 values), I/E(s) (3 values), Comp/NonComp (5 values), Noise (3 values) and Selection Threshold (6 values)). 20 simulations were conducted at each point on the five-dimensional grid defined by these parameter values (a total of $20 \times 3 \times 3 \times 5 \times 3 \times 6 = 16200$ simulations). The resulting action sequences were automatically categorised as correct (if they each comprised a sequence of 12 error-free actions leading to "correctly" prepared coffee similar to that in figure 1) or incorrect.[3] The outcome of this analysis is summarised in table 3, in which each entry consists of a parameter value and the total percentage of correct attempts at the task for that value, whilst the other four parameters range over all their possible values in the parameter subspace. Thus, the first entry states that 60% of attempts at coffee preparation were successful when S/L was fixed at 0.50 and the other four parameters varied over all tested values.

Parameter	Standard	Stable Range
S/L	0.50	0.50 – 0.65
I/E(s)	0.95	0.94 – 0.96
I/E(o)	0.10	0.06 – 0.14
Comp/NonComp	0.65	0.50 – 0.65
Persistence	0.80	0.79 – 0.81
Noise	10^{-3}	$\frac{1}{\infty} - 10^{-3}$
Sel. Threshold	0.60	0.50 – 0.95
Rest Activation	0.10	0.09 – 0.11

Table 2: Parameters and their values for coffee preparation

Table 3 shows that behaviour was disproportionately bad in two conditions: when Comp/NonComp was very low (only 24% correct at 0.45); and when noise was high (only 26% correct at 0.0025). A further analysis was therefore carried out to factor out the effects of these parameters. Table 4 shows that when noise is very low, and Comp/NonComp is between 0.55 and 0.65, the model is stable in 94% of the region of parameter space investigated. We believe this demonstrates remarkable stability across a significant region of the parameter space.

		Comp/NonComp				
		0.45	0.50	0.55	0.60	0.65
Noise	0.0001	37%	78%	94%	93%	94%
	0.0005	30%	69%	83%	72%	78%
	0.0025	6%	33%	54%	31%	7%

Table 4: The effects of Comp/NonComp and Noise

[2] The standard values were those given in the first column of table 1. No attempt was made to generate standard parameter values which were particularly robust to change.

[3] Note the strictness of this criterion: just one action slip in the 240 actions required to prepare 20 mugs of coffee resulted in parameter settings being classed as incorrect.

Parameter	Value ⇒ % Correct					
S/L	0.50 ⇒ 60%	0.55 ⇒ 60%	0.60 ⇒ 51%			
I/E(s)	0.94 ⇒ 60%	0.95 ⇒ 57%	0.96 ⇒ 55%			
Comp/NonComp	0.45 ⇒ 24%	0.50 ⇒ 57%	0.55 ⇒ 77%	0.60 ⇒ 65%	0.65 ⇒ 57%	
Noise	0.0001 ⇒ 79%	0.0005 ⇒ 66%	0.0025 ⇒ 26%			
Sel. Threshold	0.40 ⇒ 30%	0.50 ⇒ 50%	0.60 ⇒ 64%	0.70 ⇒ 75%	0.80 ⇒ 72%	0.90 ⇒ 53%

Table 3: The effects of parameter variation

General Discussion

The CS model as presented here is a refinement of a previously presented model (Cooper et al., 1995). The most important difference concerns the way parameters are incorporated into the model. In particular, the current model employs single L/S, Comp/NonComp, and persistence parameters across all interactive activation networks, reducing the parameter space from twelve to eight dimensions. This reduction was achieved by normalising all sources of excitation and inhibition, so that the activations within all networks are within similar bounds. The reduction in parameters enabled the detailed examination of the parameter space presented above.

The model is currently being developed in two directions. Firstly, we are attempting to further explore the relationship between the S/L parameter and Parkinson's disease by replicating quantitative data (Malapani, Pillon, Dubois, & Agid, 1994) on a choice reaction time task in normals and Parkinson's patients. Secondly, we are integrating some supervisory functions (e.g., task switching and error recovery) into the model in order to account for behaviour in more complex tasks currently being investigated by Schwartz *et al.*

Acknowledgements

We are grateful to David Glasspool, Michael Montgomery and Jonny Farringdon for assistance with the model described here. We are also grateful to Stephanie Warrick and two anonymous reviewers for comments on an earlier draft of this paper. Initial development of the model was supported by the UK Joint Council Initiative in Cognitive Science and Human Computer Interaction, grant #9212530.

References

Cooper, R., Shallice, T., & Farringdon, J. (1995). Symbolic and continuous processes in the automatic selection of actions. In Hallam, J. (Ed.), *Hybrid Problems, Hybrid Solutions*, pp. 27–37. IOS Press, Amsterdam.

Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, *14*(2), 179–211.

Houghton, G., & Hartley, T. (1996). Parallel Models of Serial Behaviour: Lashley Revisited. *PSYCHE*, *2*(25).

James, W. (1890). *The Principles of Psychology*. Holt, New York, NY.

Jordon, M. I. (1986). Serial order: A parallel distributed processing approach. Tech. rep. 8604, Institute for Cognitive Science, University of California, San Diego, CA.

Malapani, C., Pillon, B., Dubois, B., & Agid, Y. (1994). Impaired simultaneous cognitive task performance in Parkinson's disease: A dopamine-related dysfunction. *Neurology*, *44*, 319–326.

McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: part 1. An account of basic findings. *Psychological Review*, *88*(5), 375–407.

Norman, D. A., & Shallice, T. (1986). Attention to action: Willed and automatic control of behavior. In Davidson, R., Schwartz, G., & Shapiro, D. (Eds.), *Consciousness and Self Regulation: Advances in Research and Theory, Volume 4*, pp. 1–18. Plenum, New York, NY.

Reason, J. T. (1979). Actions not as planned: The price of automatization. In Underwood, G., & Stevens, R. (Eds.), *Aspects of Consciousness*, chap. 4, pp. 67–89. Academic Press, London, UK.

Reason, J. T. (1984). Lapses of Attention in Everyday Life. In Parasuraman, W., & Davies, R. (Eds.), *Varieties of Attention*, chap. 14, pp. 515–549. Academic Press, Orlando, FL.

Robbins, T. W. (1991). Cognitive deficits in Schizophrenia and Parkinson's Disease: Neural basis and the role of Dopamine. In Willner, P., & Scheel-Krüger, J. (Eds.), *The Mesolimbic Dopamine System: From Motivation to Action*, pp. 497–528. John Wiley & Sons, New York, NY.

Robbins, T. W., & Sahakian, B. J. (1983). Behavioural effects of psychomotor stimulant drugs: Clinical and neuropsychological implications. In Creese, I. (Ed.), *Stimulants: Neurochemical, Behavioural, and Clinical Perspectives*, pp. 301–338. Raven Press, New York, NY.

Schneider, W., & Shiffrin, R. M. (1977). Controlled and automatic human information processing. I: Detection, search and attention. *Psychological Review*, *84*, 1–66.

Schwartz, M. F., Mayer, N. H., Fitzpatrick-De Salme, E. J., & Montgomery, M. W. (1993). Cognitive theory and the study of everyday action disorders after brain damage. *Journal of Head Trauma Rehabilitation*, *8*, 59–72.

Schwartz, M. F., Reed, E. S., Montgomery, M. W., Palmer, C., & Mayer, N. H. (1991). The quantitative description of action disorganisation after brain damage: A Case Study. *Cognitive Neuropsychology*, *8*(5), 381–414.

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