

CONTENTION SCHEDULING AND THE CONTROL OF ROUTINE ACTIVITIES

Richard Cooper

Birkbeck College and University College London, UK

Tim Shallice

University College London, UK and SISSA, Trieste, Italy

The control of routine action is a complex process subject both to minor lapses in normals and to more severe breakdown following certain forms of neurological damage. A number of recent empirical studies (e.g. Humphreys & Ford, 1998; Schwartz et al., 1991, 1995, 1998) have examined the details of breakdown in certain classes of patient, and attempted to relate the findings to existing psychological theory. This paper complements those studies by presenting a computational model of the selection of routine actions based on competitive activation within a hierarchically organised network of action schemas (cf. Norman & Shallice, 1980, 1986). Simulations are reported which demonstrate that the model is capable of organised sequential action selection in a complex naturalistic domain. It is further demonstrated that, after lesioning, the model exhibits behaviour qualitatively equivalent to that observed by Schwartz et al., in their action disorganisation syndrome patients.

INTRODUCTION

Lapses in action are a common occurrence in everyday life. Reason (1979), for example, cites behaviours such as perceptual confusions (e.g. putting shaving cream, instead of toothpaste, on a toothbrush), insertions (e.g. turning on a light on entering a room, despite it being daytime) and omissions (e.g. failing to add tea to a teapot before adding water, and only realising the error when the resulting "tea" is clear). In a diary study involving 35 normal participants, he found an average of just under 1 such slip reported by each participant per day. Many neurological patients also show problems

with the control of action. In some cases the errors made by such patients may be seen as exaggerated forms (in both quantity and quality) of the lapses seen in normals. Thus, Schwartz et al. (1991) report a patient who, on various occasions, added butter and oatmeal to a mug of hot water whilst ostensibly preparing coffee.

In fact, the control of action is subject to a range of neurological impairments. These include:

1. *The grasp reflex*: An inability to inhibit an explicitly forbidden simple response such as grasping when the open palm of the hand is touched (cf. De Renzi & Barbieri, 1992);

Requests for reprints should be addressed to Dr Richard Cooper, Department of Psychology, Birkbeck College, Malet Street, London WC1E 7HX, UK (Tel: +44 171 631 6211; Fax: +44 171 631 6312; Email: R.Cooper@psyc.bbk.ac.uk).

We are grateful to Jonathan Farrington, Michael Montgomery, Myrna Schwartz, and Stephanie Warrick for assistance with the implementation and development of the model presented here. We are also grateful for the continuing support of John Fox, and for the highly constructive comments of all of the above colleagues, as well as David Glasspool, David Plaut, and two anonymous reviewers, on earlier drafts of this manuscript. Initial research was supported by a grant from the UK Joint Research Council Initiative in Cognitive Science and Human Computer Interaction, reference #9212530, to Tim Shallice and John Fox. Preparation of the final manuscript was supported in part by a grant from the US National Institutes of Health to Myrna Schwartz (#R01 NS31824-05).

2. *Anarchic hand syndrome*: Where one hand carries out unintended actions which block or undo intended actions performed with the other hand (cf. Della Sala, Marchetti, & Spinnler, 1991; Goldberg, Mayer, & Togliani, 1981);

3. *Utilisation behaviour*: Where the patient picks up and uses objects when they are clearly not appropriate to the task at hand (cf. Brazzelli, Colombo, Della Sala, & Spinnler, 1994; Hashimoto, Yoshida, & Tanaka, 1995; Lhermitte, 1983; Shallice, Burgess, Schon, & Baxter, 1989);

4. *Ideational apraxia (in the sense of Poeck & Lehmkuhl, 1980)*: An impairment of the ability to perform actions appropriate to objects, when multiple sub-actions have to be made, but in which individual actions may be correctly imitated;¹

5. *Action disorganisation syndrome*: Where the patient's goal-directed action is generally disorganised, with frequent errors, including action omissions, utilisations, and argument substitutions (cf. Duncan, 1986; Humphreys & Forde, 1998; Schwartz et al., 1991, 1995, 1998);

6. *Bradykinesia*: A significant slowing of willed initiation of action sequences as seen in Parkinson's disease patients (cf. Malapani, Pillon, Dubois, & Agid, 1994; Owen et al., 1992); and

7. *Stereotypy*: A tendency toward repeated or stereotyped action occurring with amphetamine psychosis (cf. Lyons & Robbins, 1975).

This list is not complete, and the precise relationship between some of the impairments is unclear (e.g. between ideational apraxia and action disorganisation syndrome), but lapses and acquired disorders of human action such as these pose a problem for any account of the control of action. Clearly, for such an account to be viable it must not only explain normal behaviour but also address lapses and breakdowns of that behaviour. One influential account of the control of action which purports to do this is that of Norman and Shallice (1980, 1986). The original account of the theory

was stated verbally. All subsequent accounts (e.g. Shallice, 1982, 1988; Shallice & Burgess, 1996) have similarly been stated in only verbal terms. As a consequence it has not been possible to test properly either that the Norman and Shallice theory provides an adequate account of the control of normal behaviour, or that it provides a viable account of behaviour following neurological damage. This paper addresses this failing by presenting a computational model of the control of action based on the theory developed by Norman and Shallice. It has two main functions:

1. to show that the Norman and Shallice framework is a possible candidate for the organisation and control of action selection in an environment of multiple objects and when the organism has multiple competences which it can utilise; and

2. to show that, when damaged, the system exhibits behaviours similar to those observable in neurological patients.

The remainder of the paper begins by presenting the theoretical and empirical background behind the work. This includes a discussion of the "level" of action with which we are concerned, a review of the principal findings relating to the breakdown of behaviour at that level (both in normals and subsequent to certain forms of neurological damage), and a brief presentation of the Norman and Shallice theory. This is followed by a detailed presentation of the computational model. The behaviour of the model, both in normal and abnormal modes of functioning, is then discussed, and correspondences between abnormal model behaviour and some neurological patient behaviour (specifically behaviour of patients exhibiting action disorganisation syndrome) are highlighted. The paper concludes with a consideration of the relationship between the current model and other models of sequential behaviour and a speculative exploration of some further neurological disorders potentially explicable by the model.

¹ It is likely that this disorder can be subdivided, since some patients have frontal lesion sites quite different from the most frequent left temporoparietal junction localisation (De Renzi & Lucchelli, 1988). The former would be the more relevant subgroup here. Ideational apraxia will not be considered further in this paper.

THEORETICAL AND EMPIRICAL BACKGROUND

Three Levels of Action

The organisation of human action can be analysed on a variety of different levels. At lower level execution of an action must be understood in terms of the biomechanics of the required movements and the physical properties of any targets (e.g. McLeod & Dienes, 1996; Rosenbaum, Vaughan, Barnes, Marchak, & Slotta, 1990; Viviani, 1990). However, the variety of our individual motor skills and the way that each has its own specific characteristics apparently not continuous with those of any other skill has led to the use, since the time of Head (1926), of concepts such as schema; thus Schmidt (1975) proposed that a specific type of action is controlled by a *motor response schema*. Motor response schemas were held to be formed by abstracting over movements of the "same general type" (Schmidt, 1975, p. 235) a relationship between the movements' initial conditions, response specifications, sensory consequences, and response outcome. Schmidt (1975, p. 235) argued that such schemas were "more important to the subject than [...] any of the stored instances, which [...] are forgotten more quickly over time than is the schema". The approach is compatible with the idea that the lowest level at which individual actions are specified is in terms of mixtures of discrete sensory-motor mappings (Wolpert & Ghahramani, 1997).

The idea of motor schemas has been elaborated by Arbib and his colleagues (Arbib, 1985; Iberall & Arbib, 1990; Jeannerod, Arbib, Rizzolatti, & Sakata, 1995), who argue that individual subcomponents of a larger skill unit are each represented by lower-level schemas. So during prehension, the sub-actions for "reach," "preshape," "enclose," "rotate forearm," and that for "selecting numbers of fingers" would each be potentially elicitable by object affordances. At this level of organisation the individual sub-components of the action need to be carried out at precisely specified times and to be coordinated with each other. They are also closely specified with respect to the detailed physics both of the

environment and of the effector system. However, while the individual subactions are designed to be carried out in parallel with particular other subactions, arbitrary combination is not in general possible.

For much higher levels of action control, concepts such as *scripts* (Schank & Abelson, 1977) and *Memory Organisation Packets* (MOPs) (Schank, 1982) have been proposed to represent the organisation of well-learned activities such as going to a restaurant or visiting a doctor's surgery. Here the sub-units can vary in the time required to carry them out and they can at times be separated by other unrelated activities—for example, one can go into a shop to buy something whilst on the way to the doctor's surgery. When a MOP is activated, only one of its subactions is generally carried out at a time, but the subactions can on occasions be combined with other activities, so one can read whilst waiting at the doctor's surgery. Moreover, the detailed physics of the effector system and of the environment are not relevant. The size of the steps at the doctor's surgery or the shape of the door handle do not affect the basic execution of the MOP. Even the nature of how individual actions are specified is not crucial. One can stand or sit in the waiting room. Similarly one can pay the bill at a restaurant by cash, credit card, or cheque—each requiring quite distinct types of movement—without affecting the overall organisation of the MOP.

Between these two extremes lies the organisation of the components of a whole range of well-learned activities—making breakfast, cleaning one's teeth, starting a car, dressing, and so on. In these situations parallel execution of subactions only rarely occurs (although it may be possible and the sub-actions themselves (e.g. *stir the coffee*) appear cognitively to be represented as discrete units. In these respects the organisation resembles that of the doctor-visit domain. However, like the reaching domain the local physics of the environment and the specific timing are relevant, but on a grosser grain, and interruption, while entirely feasible, is to be avoided as it leads to disorganisation of the action. Thus, in toothbrushing, although one can delay or do something else between putting

toothpaste on the toothbrush and brushing one's teeth, it is not generally advisable.

This intermediate level of organisation has sufficient specificity (see Table 1) to be considered a distinct domain in the organisation of action. The key conceptual differences with the lower level are that the selection of which subaction is to be carried out is the critical issue, and this process may be subject to voluntary control. Key conceptual differences with respect to the higher level include the need to represent the local physics of the subactions and the requirement to provide an account (at a single level) for *all* behaviours occurring over a particular period. Thus, whereas in the higher domain one must extract from the set of goals being tackled over a (generally longer) period of time just those behaviours that are necessary for tackling one of these goals, in the intermediate domain all behaviours are prototypically nested under a single higher-level goal.

Lapses and Errors in Intermediate Domain Action Selection

The intermediate domain of action is further differentiated from the higher and lower levels by the forms of breakdown to which it is subject, both in normals and following neurological damage.

Lapses in Normal Action Selection

The complexity of the intermediate domain of action makes it relatively unattractive for investigation by the methods of standard human experimental psychology. Indeed, we know of no such studies. However, natural history studies of action lapses relevant to the domain have been carried out. Reason (1979, 1984, 1990) has conducted a series of diary studies in which normal volunteers were required to note any "actions not as planned" over periods of several weeks. The systematic nature of many of the observed slips and lapses allowed Reason to develop a system of error categorisation. Norman (1981) analysed a further set of action slips and developed a related categorisation. The full range of errors identified by these authors is too large to consider here. However, a number of error

types are particularly relevant to the intermediate domain of action, including:

1. *Capture*: An unintended action sequence performed in place of an intended action sequence that is, nevertheless, appropriate given the environmental cues (e.g. putting on gardening boots upon entering the garage, instead of getting the car out as intended);

2. *Omission*: An action sequence in which one step or subtask is not performed, despite the lack of any intention to omit the step or subtask (e.g. failing to add tea to a teapot before adding water when making a pot of tea);

3. *Anticipation*: An action sequence in which one step or subtask is performed earlier in the sequence than usual, despite the lack of any intention to alter the usual ordering (e.g. when filling a bucket from a tap, putting a lid on the bucket before turning off the tap);

4. *Perseveration*: The unintentional repetition of a step or subtask (e.g. adding excessive teaspoonsful of sugar to coffee when distracted by an interesting conversation or event); and

5. *Object substitution*: An intended action carried out with an unintended object (e.g. applying shaving cream instead of toothpaste to a toothbrush).

These categories are neither disjoint nor definitive, and there can be difficulties in using them to classify certain action sequences. In particular, omission and anticipation errors are not necessarily clear-cut categories—an anticipation can be analysed in terms of omitting a step or subtask prior to the anticipatory error. Nevertheless, the error categories are important because they arise in "normal" functioning of the undamaged action selection system. That is, the undamaged action selection system is susceptible to a variety of errors. Any viable theory of intermediate domain action selection must account for this susceptibility, as well as the occurrence of the more extreme forms and patterns of error seen in neurological patients.

Action Disorganisation Syndrome

Earlier we noted a number of disorders of action control arising after neurological damage. Our

Table 1. *Levels of Action and Their Properties*

	<i>Local Physics of Environment and Effectors Relevant</i>	<i>Control of Timing Essential</i>	<i>Parallel Execution of Subactions</i>	<i>Subactions Combinable with Other Activities</i>	<i>Interruptions Possible Between Subactions</i>	<i>Subactions Cognitively Represented as Discrete Units</i>	<i>Selection of Specific Subactions and Arguments Critical</i>	<i>Subject to "lapses" to be done</i>	<i>Subject to Voluntary Control of Subactions</i>
Lowest level (e.g. reaching)	At a detailed	At a very fine level	Required grain	No	No	No	No (assumed)	No	No
Intermediate level (e.g. toothbrushing)	Only at a gross level (e.g. whether objects are reachable)	At a much grosser level (i.e. before the environment changes much)	Not generally required	Not generally required	In general, yes	Yes	Yes	Yes	Yes
Highest level (e.g. visiting doctor)	No	No	No	Often yes	Generally yes	Yes	Not necessarily (a variety of implementations may be possible)	Yes	Yes

focus in this paper is on just one of these—action disorganisation syndrome. A number of quantitative studies of this syndrome exist (e.g. Humphreys & Forde, 1998; Schwartz et al., 1991, 1995, 1998). In particular, Schwartz and colleagues (1991, 1995) have carried out a number of detailed studies in which patients were videotaped every day on occasions when they were carrying out activities such as making up their breakfast from the containers available on a hospital breakfast tray or brushing their teeth. For instance, Schwartz et al. (1991) videotaped 28 sessions of HH—a patient with a bilateral (predominantly right) medial frontal lesion—eating breakfast. In the investigation every episode during which he prepared his coffee was analysed. Coffee preparation is a task with a number of levels of control. Thus, preparing coffee involves, among other things, adding milk to the partially made coffee, and adding milk involves picking up a milk container, opening it, pouring the appropriate amount of milk from the container into the coffee mug, and then closing and/or discarding the container.

Typically HH had his main course centred on the part of the tray closest to him and around it were oatmeal, juice, a pint container of milk, hot water in a covered mug, utensils, an empty cup and a variety of condiments including instant coffee, sugar, and cream, all in single-portion containers. Acts carried out by HH were labelled by an action coding scheme in terms of an event hierarchy and four basic operations. The basic operations used were:

1. MOVE (x) TO (location) VIA (instrument) BY (manner);
2. ALTER (x) TO (status) VIA (instrument) BY (manner);
3. TAKE (x) (i.e. to take possession of object x);
4. GIVE (x) (i.e. relinquish possession of object x).

Manner was used to describe the actual movement by which the action was accomplished (e.g. tearing, pouring). In the event hierarchy, basic operations were grouped into larger units in terms of the way they achieved clear subgoals in the task of preparing and consuming breakfast. A crux action was

defined to be a basic operation which was invariably essential to the completion of a subgoal.

A number of measures of disorganisation were used. One related to independent basic operations, those which are not cruxes and which are not part of achieving the subgoal corresponding to the next crux action. As HH recovered the rate of such independent operations gradually declined from 80% to 20% of his actions. The second measure involved place or object substitutions, namely errors where the action was appropriate (in that it was not an independent) but one of its arguments was incorrectly filled. Many such substitutions occurred (place $n = 42$; object $n = 15$). Virtually everything on the tray was subject to misuse. Thus object substitutions involved putting in the mug (for coffee) oatmeal (nine times), butter (twice), salt (twice) and orange juice (twice), yet the patient was not agnostic. Schwartz et al. (1995) provide a similar analysis of JK, a somewhat related patient, and further quantitative studies of action disorganisation syndrome are reported by Schwartz et al. (1998) and Humphreys and Forde (1998).

Contention Scheduling and the Supervisory System

Over 15 years ago, Norman and Shallice suggested that the intermediate domain of action selection is characterisable in terms of an activation-trigger-schema framework (Norman & Shallice, 1980, 1986; see also Norman, 1981; Shallice, 1988). In particular, it was argued that selection of an action depended upon its corresponding schema being activated above threshold, with activation being received (horizontally) from “triggers” and (vertically) from so-called “source” schemas. Otherwise, schemas were in lateral inhibitory competition, with the amount of competition between any pair of schemas depending upon the degree of overlap in their effector system requirements. Selection of a schema allowed it to send activation to its “component” schemas and/or to control whatever part of the effector system it required.

The Norman and Shallice approach, termed contention scheduling, has a number of advantages:

1. It allows volitional top-down control of action selection to be combined with precise environmental triggering at the appropriate time.

2. It clearly complements the schema theory of Arbib and colleagues for the lower level of the selection of subactions within a basic action such as grasping.

3. It makes the prediction that the primary bottleneck in dual task performance should lie at the level of the selection of the appropriate action, which fits empirical research on psychological refractory periods (Pashler, 1994).² Moreover, as lateral inhibitory effects are held to depend on the degree of overlap in their effector system requirements, variations in the degree of overlap in dual task situations due to differences in response requirements of the tasks can be explained (McLeod, 1977; McLeod & Posner, 1984).

4. It can provide a basic system that a second system—the supervisory system—may modulate, and so contribute to a theoretical account of higher-level executive operations, and in particular of their disorders following frontal lobe lesions (Shallice & Burgess, 1996).

5. It can provide, under inadequate operation, an account of abnormal or inappropriate selection of actions analogous to the types of errors identified earlier.

The final point has two distinct empirical correspondences. First, it is held to correspond to the situation in normal participants when they are distracted, as in this situation the modulating system—the supervisory system—has its operations determined by some other task. Indeed, Reason (1984) has shown that the median rating given by participants for their mental state when making an action lapse is 6 on a 1–7 degree of distraction scale. Other aspects of Reason's findings fit with the theory's account. Both capturing and captured actions are well-learned; there should therefore be schemas organised to control the actions. Moreover, they are carried out in physically similar situations; the car-

rying out of one will generally occur in situations when the other would therefore be automatically triggered.

The second type of empirical correspondence concerns neurological patients with a range of difficulties either in the selection of appropriate actions or in the inhibition of inappropriate ones, as described earlier. Such patients typically have lesions involving the medial surfaces of the frontal lobes. The behaviour of these patients, which can be loosely described as distractible (Shallice, 1982, 1988), intuitively seems to correspond to damage within a system of the above sort. Indeed, distractibility, utilisation behaviour and the action disorganisation syndrome have all been characterised as resulting from the loss of top-down activation of the schema hierarchy, either because of loss of supervisory system control (Shallice, 1982; Shallice et al., 1989), or within a single system where no contention scheduling/supervisory system distinction is made (Schwartz et al., 1991, 1995).

Notwithstanding the above advantages, the original theorising about contention scheduling was subject to four main disadvantages:

1. The model was not implemented, so there was no guarantee that the accounts offered of various empirical phenomena were actually valid.

2. Contention scheduling—the system responsible for routine selection of action—was held to operate in the intact adult human modulated by a second system—the supervisory system—held to be responsible for the organisation of nonroutine (novel) behaviours. However, action lapses, distractibility, and utilisation behaviour, which can all manifest themselves in routine activities (indeed action lapses typically occur in such situations), were viewed as a consequence of loss of supervisory system modulation of contention scheduling (Shallice, 1982; Shallice et al. 1989). As Schwartz et al. (1991) pointed out this was unprincipled and action lapses could be viewed instead as the loss of

² Pashler claims his results to be in conflict with contention scheduling's predictions, but this is because he characterises it as predicting an impairment in producing a response rather than in selecting an action schema.

top-down volitional control *within* the routine action control system itself.

3. Although an earlier paper with a related perspective (Shallice, 1972) had argued that a function of trigger (then selection) input was to (consciously) set the arguments—that is, the specific goal—of a selected schema, there was no consideration of argument specification in the contention scheduling model. Yet errors of argument selection are a major form of action lapse (cf. Norman, 1981; Reason, 1979, 1984, 1990).

4. The model was concerned entirely with the hierarchical organisation of subcomponent activities and had no separable representation of goals, these being held to be represented at the level of the supervisory system. By contrast, Schwartz et al. (1991) argue that goals also need to be clearly represented within the mechanisms responsible for the hierarchical control of routine actions.

This paper therefore describes an implementation of the contention scheduling model and reports the behaviour of that model both under normal functioning and following lesioning.

THE MODEL

Few, if any, psychological theories are sufficiently well specified as to allow direct and complete implementation (cf. Cooper, Fox, Farrington, & Shallice, 1996), and the development of a computational model of action selection from the psychological model presented by Norman and Shallice (1986) has required significant elaboration of the original description. Some elaborations are clearly implementational, in that they are necessary for the complete specification of an executable computational model (e.g. the specification of mathematical equations governing the effects of excitation and inhibition on schema activation). Others, however, correspond to substantive theoretical extensions or elaborations. In the description below we attempt to isolate all assumptions of the computational model and indicate their status with respect to the theory/implementation distinction. Theoretical

assumptions are further subdivided into core and peripheral assumptions (corresponding respectively to Lakatos' [1970] hard-core and protective belt). We hold a strong theoretical commitment to core assumptions (CA), and believe them to be critical in producing the simulated behaviour. We hold less theoretical commitment to peripheral assumptions (PA), and regard them as being somewhat flexible and open to modification. In general, alternate theoretically acceptable mechanisms exist for the processes which these assumptions define, and an examination of such mechanisms remains for future work. We have no theoretical commitment to implementational assumptions (IA), which are intended only for computational completeness, and hold that the behavioural characteristics of interest are largely independent of these assumptions.

Functional Subcomponents and Their Interactions

Figure 1 illustrates the main functional subcomponents of our implementation of contention scheduling. At the heart of the model is the schema network. Nodes in this network correspond to the action schemas described earlier. Each node has an activation value and activations interact through a variety of excitatory and inhibitory mechanisms. Two other networks, which also operate according to the principles of interactive activation, serve to model object representations and resource requirements. In addition, a selection process oversees the schema network, interfacing with the object representation, resource, and motor systems. Schema selection occurs when the activation of a schema's node exceeds a threshold. Selection of a high-level schema acts to modify the flow of activation within the schema network such that component schemas of the selected schema receive additional excitation (thus increasing their likelihood of selection). Low-level schemas correspond to discrete actions. Selection at this level leads first to the assignment of object representations and resources to the corresponding action and then to the execution of that action.

Schema-related Structures and Processes

Schema-related structures and processes are defined by 10 core assumptions (CA1–CA10), 7 peripheral assumptions (PA1–PA7), and 5 implementational assumptions (IA1–IA5).

The Organisation of Schemas and Goals

It is assumed that the fundamental unit of organised behaviour within the intermediate domain is the schema. Schemas correspond to abstractions over goal-directed segments of action, and are assumed to exist at various levels. Thus schemas might exist for relatively low-level actions such as opening jars or higher-level sequences such as spreading butter/jam on bread/toast or preparing instant coffee.

Schemas are goal directed: Each schema has a goal that it achieves. Formally, a goal is a condition that may or may not be satisfied by the world. The goal is satisfied when the condition is true of the world. Schemas are effectively *methods* for achieving goals. Thus, a goal might be that a particular mug of coffee is sweet. A schema for achieving this goal might comprise: picking up a spoon, dipping the spoon into some sugar contained within a sugar-bowl, filling the spoon with sugar, transporting it and then tipping it into the mug. The same goal might be achieved by any number of other schemas, such as one involving sugar cubes. The schema that is most appropriate in any particular situation will depend on a variety of factors, including the objects available in the environment and individual preferences.

Although implicit in Norman and Shallice's (1986) account, previous descriptions of contention scheduling have not developed the relationship between schemas and goals, and the contention scheduling system has been portrayed simply as a hierarchy of schemas. Schwartz et al. (1991), however, argue that in order to capture the complexity of routine action it is necessary to represent goal information within the schema network, and this is the approach adopted here (see also Duncan, 1986).

CA1: Schemas are goal directed.

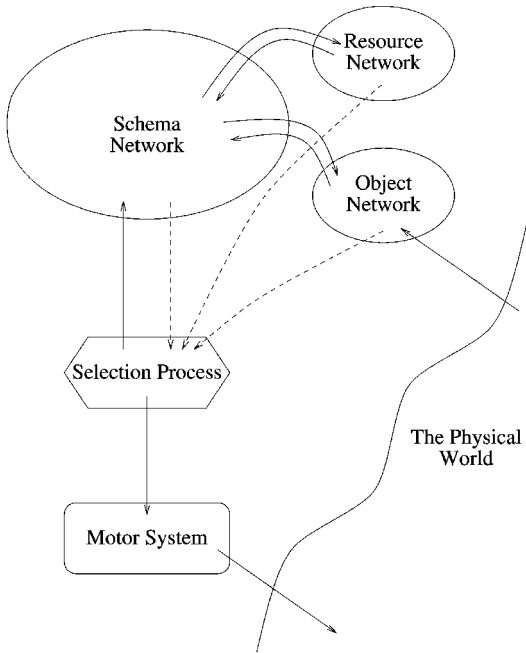


Fig. 1. The principal components of the contention scheduling implementation.

The model is hybrid in the sense that it incorporates both continuous variables and discrete gating of activation by symbolic flags. Continuous variable modes of operating were an intrinsic part of the original informal model (cf. Norman & Shallice, 1980, 1986), and they more easily represent changes in the weighting of several different sources of information (as in the combination of activation from intentional control and environmental triggering, or in processes like lateral inhibition). On the other hand, some-or-none effects of selection on lower level schemas and in particular all-or-none processes like argument-filling appear to require discrete representations.

The underlying execution model for the simulation is synchronous and cyclic. Thus, on each processing cycle all activation values (for schemas, object representations, and resources) are updated, schema selections and deselections are made, and, if appropriate, discrete actions are effected. These operations are effectively performed in parallel.

The components of schemas are also properly understood as subgoals (rather than subschemas). In the schema for sweetening coffee by adding sugar from a sugar bowl, for example, the first component of the schema is not the schema for “pick up teaspoon,” but the subgoal “teaspoon in hand.” This subgoal might be achieved by picking up an appropriate spoon from the kitchen table (a “basic level” schema with no components), or by opening the cutlery drawer, removing a teaspoon, and then closing the drawer (a higher-level schema with three components).

CA2: Schemas consist of a partially ordered set of subgoals.

Despite the importance of goals, they do not form a separate network or hierarchy. Rather, they are represented within the schema network as “layers” that mediate schema/subschema relations. Figure 2, for example, shows some of the schema/goal relationships used in the simulations reported later in this paper. The nodes beneath a schema node (i.e. the schema node's components) correspond to the subgoals which comprise that schema, and the nodes beneath a goal node (i.e. the goal node's components) correspond to schemas which may achieve that goal. The left-to-right ordering of subschemas in the figure is not indicative of any “hard-wired” ordering constraints. Rather, subgoal information concerning ordering relations (in the form of subgoal preconditions) is attached to a subgoal on a case by case basis, as discussed later.

Within the schema hierarchy the components of goal nodes should be understood disjunctively—any possible component schema may be employed to achieve a goal—whereas the components of schema nodes should be understood conjunctively—all subgoals of a schema must be achieved in order to complete the schema. This is indicated in Figure 2 by the presence of arcs joining conjunctive branches and the absence of arcs joining disjunctive branches. The structuring yields what is referred to within Artificial Intelligence as an and/or tree (see, e.g. Charniak & McDermott, 1985). Such trees are commonly used to represent goal hierarchies in Artificial Intelligence applications.

The Schema Network

The schema network is task dependent in that different tasks will require networks made up of different schemas and goals. The network shown in Figure 2 is a subpart of that employed in the simulations presented next.

As noted earlier, each schema has a corresponding activation value. These activation values vary with time (according to equations given in Appendix A) and determine the role of their corresponding schemas in controlling behaviour. In brief, if a schema's activation exceeds a predefined threshold then behaviour will be determined by that schema.

CA3: Schemas have an associated activation value. This value is a real number that varies over time.

Schema activation may be affected by the presence (or absence) of appropriate triggering situations in the environment. Thus, the presence of small portable objects in the environment may excite the schema for picking up such objects. Schemas may also be excited by “top-down” influences from higher-level schemas. Thus, a schema for *add sugar from sugar bowl to coffee mug* may excite a lower-level schema for *pick up teaspoon*. Top-down influences may also originate from the Supervisory System. (Indeed, it is postulated that the direct excitation or inhibition of schema activation values is the only way in which the Supervisory System may affect behaviour.) Standard interactive activation processes of competition operating over schema activation values provide two further sources of influence: a lateral influence that ensures that schemas that compete for resources inhibit each other, and a self influence that generally works to excite schemas, partially countering the lateral influence. The net influence on schema activation at any time is given by a weighted sum of these four sources (plus normally distributed pseudorandom noise). The precise effect of this net influence on a schema's activation is given by Equation 13 in Appendix A. This effect is dependent on a schema's current activation (which has a tendency to persist even with zero net influence) and its rest activation (the activation level to which it will eventually fall with zero net influence). If a schema is inhibited

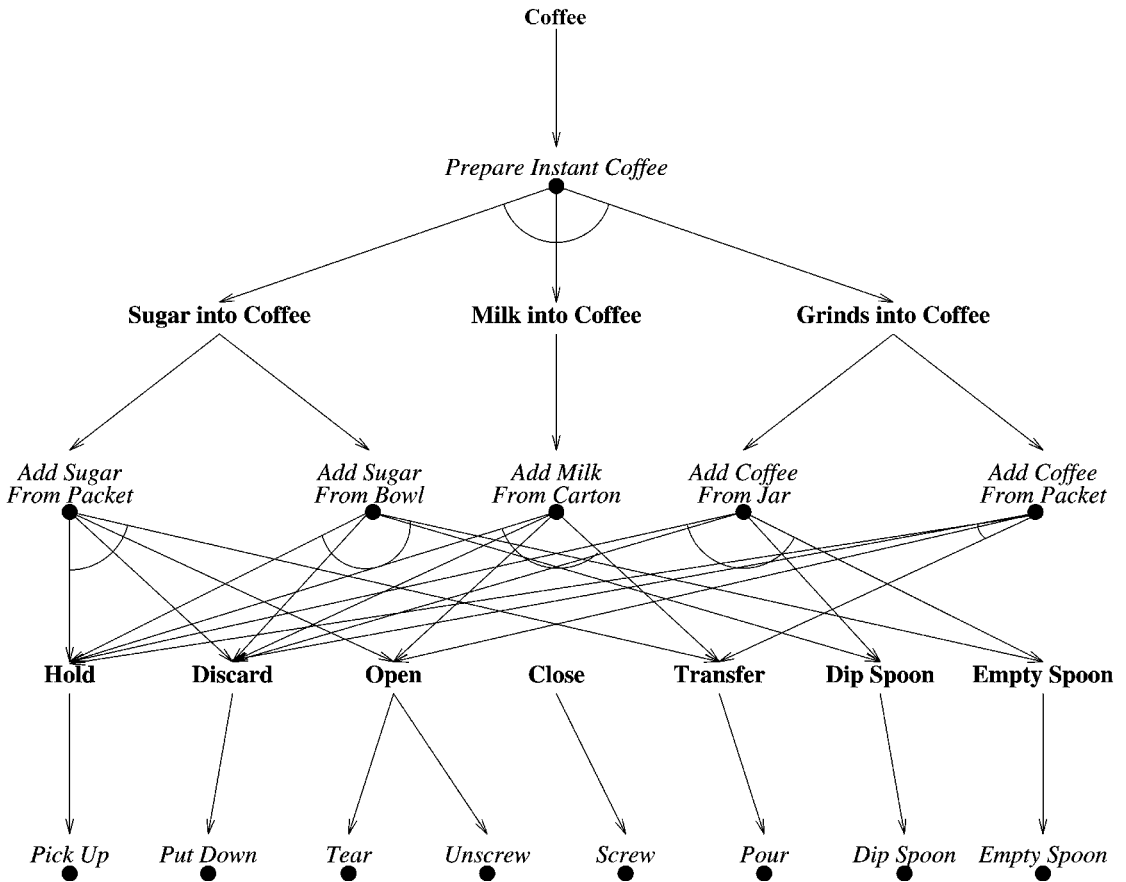


Fig. 2. Schema/goal organisation in the coffee preparation domain. Schemas are indicated by italic type and goals by bold type.

(i.e. its net influence is negative), its activation may be pushed below its resting activation.

CA4: Schema activations are influenced by five factors: top-down influence, environmental influence, self-influence, lateral influence, and random noise.

IA1: The net influence on a schema's activation is a weighted sum of the influences cited in CA4.

CA5: In the absence of any influence, schema activations tend to persist (i.e. if a schema is highly active and all influences are removed from it its activation will slowly decay to its resting value). In addition the resting value of schemas with zero net influence is greater than the minimum activation value. Schema activations may be inhibited to below this resting value.

In the simulations reported here activation values range from 0 to 1, and the selection threshold is .6. The criticality of these parameter values for appropriate behaviour is discussed later.

The notion of competition between schemas is central to the use of interactive activation within the schema network. Two nodes compete if (a) they are alternate means of achieving the same goal or (b) they share at least one subgoal. The idea here is that if two schemas both achieve the same goal then only one of those schemas needs to be selected (and effected), and if two schemas share a subgoal then they are likely also to share resource requirements. Recall that, unlike schemas, goals do not have activation values associated with them. The prime purpose of goals within the network is in

establishing competitive relationships between schema nodes.

PA1: Schemas compete if they are alternate means of achieving the same goal or if they share one or more subgoals. This competition is effected by a "lateral influence" on the activations of competing schemas.

Lateral inhibitory influences are commonly employed within interactive activation networks to ensure that only one element from a set is highly active at any one time. The most active element strongly inhibits its competitors, but the low activity of the competitors means that any inhibition on the most active element due to their activation is small. In the current implementation the lateral influence of one schema on another is proportional to the difference between the first schema's activation and rest activation. Consequently, if a schema's activation is above rest activation then its lateral influence on competitors will be inhibitory, but if a schema's activation is below rest activation then its lateral influence on competitors will be excitatory. Schemas at rest do not contribute to the lateral influence.

IA2: The degree of lateral influence of schema A on schema B (assuming A and B compete) is proportional to the difference between schema A's activation and rest activation. The total lateral influence on a schema is the sum of the lateral influences from all of its competitors.

Lateral influence and selfinfluence are opposing sources of activation which, in the absence of top-down and environmental influence, leave the network in an unstable equilibrium. Top-down and environmental influence destabilise this equilibrium and allow the network to move towards a stable state in which, in the normally functioning system, at most one schema from any competitive set exceeds the selection threshold. This schema is then selected, altering the network dynamics and causing the network to move toward a different stable state.

IA3: The selfinfluence of a schema on its own activation is directly proportional to the schema's activation. (This influence is generally excitatory, but see IA5.)

A number of factors affect the top-down influence of source schema activations on the activations of component schemas. These are discussed in detail in the sections following on Schema Selection and Deselection and Goal Achievement. The influence is generally zero, but if a source schema becomes sufficiently active (and is then selected) the effect can be excitatory (and proportional to the activation of the source schema).

The environmental influence on schemas, by contrast, is rarely zero and can be either excitatory or inhibitory. As discussed in the section on Object Representations and Schema/Object Interactions, we assume that the internal representations of objects have associated activation values similar to those of schemas. These value are used in calculating the degree to which an object representation triggers a schema. A relatively large number of highly salient objects associated with a schema will yield a strong influence on that schema, whereas fewer, less salient objects will have a weak influence on a schema, and objects whose representations are inhibited below their resting values may have an inhibitory affect on related schemas.

CA6: Schemas have associated "triggering conditions." The environmental influence on a schema is dependent on the extent to which its triggering conditions are satisfied by the (system's representation of the) current situation.

Assumption CA6 leaves open the question of how a schema's triggering situations are determined. As a general principle, we take schemas corresponding to discrete actions, which we refer to as basic-level schemas, to be triggered by situations which satisfy their preconditions. Thus the schema for *pick-up* is triggered by situations in which there is a free hand and a small portable object within reach of that hand. Higher-level schemas are activated solely by the presence of the objects involved in their component subschemas (so the schema for *prepare instant coffee* is activated by all coffee-related objects).

PA2: The triggering conditions of basic-level schemas are the preconditions of the corresponding

actions. Higher-level schemas are triggered by the presence of all objects relevant to those schemas.

a selected schema's source schema changes, the selected schema is deselected.

Schema Selection and Deselection

In addition to the earlier assumptions concerning schema activation, it is assumed that schemas may be in one of two states: selected or unselected. When a schema is selected it may excite its component schemas (i.e. those schemas that achieve its subgoals). Unselected schemas have no effect on the activation of their component schemas.

CA7: Schemas have a state, which may be either selected or unselected. If a schema is selected it may pass excitation to its component schemas. (This excitation is the top-down influence referred to in CA4.)

LA4: The amount of excitation passed from selected source to component schemas is directly proportional to the activation of the source schema and inversely proportional to the number of subgoals of the source schema.

Basic-level schemas, those at the lowest level of the schema network, have no subgoals and hence no component schemas. They correspond to discrete actions. Selection of a basic-level schema triggers the execution of its corresponding action.

CA8: When a basic-level schema is selected it triggers execution of its corresponding action.

A schema becomes selected when its activation exceeds a selection threshold. In the interest of simplicity this threshold is, in the current implementation, fixed and constant for all schemas, but Norman and Shallice (1986) suggest that the threshold may be different for different schemas, and that as a schema becomes well learnt its selection threshold may decrease.

CA9: When a schema's activation exceeds the selection threshold its state changes to selected.

Selected schemas remain selected until either their activation falls below that of a competitor, or their source schema's selection status changes (i.e. their source schema becomes selected or unselected).

CA10: When a selected schema's activation falls below that of one of its competitors, or when the state of

The net effect of selection is to bias competition in the schema network in favour of schemas with selected source schemas. Thus, if several schemas are competing, and one has a selected source schema, then that component schema will generally win the competition due to the additional excitation it receives from its source schema. The winning component schema will then be selected. Under normal functioning, component schemas will therefore be selected within the scope of their source schemas. This selection will continue to the lowest level of the schema network, whereupon discrete actions will be triggered.

Goal Achievement

Norman and Shallice (1986) suggest several conditions that may lead to a selected schema ceasing to operate (i.e. becoming deselected), including the satisfaction of the selected schema's goal. In this section we detail the mechanisms for goal monitoring assumed within the contention scheduling implementation. These mechanisms are deliberately naïve. Mechanisms for more sophisticated goal monitoring and error recovery undoubtedly exist. We consider such mechanisms to belong within the Supervisory System (cf. Shallice & Burgess, 1996).

We assume that, within contention scheduling, a schema's subgoals are "ticked off" as they are achieved by the system. For example, *add sugar from bowl to coffee mug* may have four subgoals: hold suitable implement, transfer sugar to implement, transfer sugar from implement to coffee mug, and discard implement. Performing the discrete action of picking up a spoon achieves the first of these subgoals and hence results in the achieved subgoal being ticked off from the source schema's subgoal list.

When a schema is selected its subgoal list is initialised to include all of that schema's subgoals. When a component schema is deselected, the component schema's goal is removed from its source schema's subgoal list.

PA3: When a schema's state changes from selected to unselected, all subgoals on the schema's subgoal list are marked as unachieved.

PA4: When a component schema is deselected, its goal is marked as achieved on the subgoal list of the component schema's selected source schema (i.e. it is ticked off).

From PA4, deselection at one level in the schema network results in goal achievement at the superordinate level. Deselection at the lower level results from inhibition of the lower-level schema (which leads to the standard mechanisms of deselection, discussed earlier, being invoked). This inhibition is controlled by goal achievement at the lower level.

PA5: When all subgoals of a selected schema have been achieved the schema is inhibited. This inhibition remains in force as long as the schema is selected.

In the current implementation the inhibition referred to in PA5 is achieved by temporarily reversing the nature of self influence on the relevant schema (i.e. by switching it from an excitatory to an inhibitory influence). This reversal causes a refractory effect on those schemas that correspond to achieved goals. Self influence is restored to an excitatory influence once the schema has been deselected.

IA5: If all subgoals of a selected schema are achieved then the effect of self influence on that schema is inhibitory. Otherwise it is excitatory.

The earlier assumptions mean that goal achievement substantially alters the dynamics of activation flow in the schema network. Under normal functioning the net effect of the assumptions will be that any schema corresponding to an achieved goal will be deselected, allowing a competing schema to be selected, and sequential behaviour to proceed.

The implementation of goal achievement described here differs from the notion of goal *satisfaction* described by Norman and Shallice (1986), in which satisfaction of a goal by external sources would lead to deselection of the corresponding schema. The assumption implicit in our implementation, that monitoring of goal achievement may be

based purely on proprioceptive feedback, is undoubtedly a simplification. The above assumptions relating to goal achievement may therefore prove to be inadequate in more complex or changing environments.

Serial Ordering of Schemas

At least three types of ordering constraint are apparent in intermediate domain action. Firstly, there are necessary constraints that arise from the physics of the environment. When adding sugar from a packet to a mug of coffee, for example, it is necessary first to open the packet before the sugar can be poured into the mug. Second, there are constraints that must be satisfied in order to complete the task successfully. Again, when adding sugar from a packet to a mug of coffee, the task requires that the sugar be poured from the packet before the packet is discarded, even though both actions are physically possible once the sugar packet has been opened. Third, variation in serial order may arise from individual preferences when ordering is arbitrary. Thus, when preparing instant coffee most, but not all, individuals will add coffee grinds to the hot water before adding the milk and sugar. Mechanisms to allow all of these sources of variation in serial order are included in the model described here.

The first form of necessary constraint, embodied in PA6, arises in the current implementation from the way in which the representation of the environment interacts with the schema network.

PA6: The environmental influence on basic-level schemas is limited to those schemas corresponding to actions that are physically possible given the (system's representation of the) current state of the environment.

One consequence of PA6 is that, for example, the *pick-up* schema is only triggered by situations in which there is an object suitable for picking up and a free hand available to perform the action. PA6 does not necessarily prevent schemas corresponding to physically impossible actions from being selected (as such schemas may still be excited above the selection threshold by other sources of activation), but the normally functioning system is biased against such selections.

Ordering constraints of the second and third variety, necessary task-specific constraints and arbitrary ordering preferences, are enforced through the explicit marking of preconditions on schema subgoals. These preconditions affect the flow of top-down excitation within the network.

PA7: The top-down influence on component schemas by selected source schemas is gated by goal and precondition achievement. That is, top-down excitation only flows to schemas whose goal has not been achieved but whose preconditions have been achieved.

Within the simulations reported here, necessary ordering constraints are only required at the lowest level, and these are all enforced by a general principle requiring that the principal goal (or crux, in the terminology of Schwartz et al. [1991]) of all lowest-level schemas must be achieved before the final "discard" goal is activated. For example, when adding sugar from a sugar packet to the coffee mug, the sugar must be poured from the packet before the packet is discarded.

The effect of precondition gating is to negatively bias goals whose preconditions have not been achieved, but the bias may be over-ridden by other factors. The mechanism does not therefore strictly enforce any ordering on a selected schema's subgoals.

Assumptions PA6 and PA7 lead to a relatively robust system that, under normal functioning, effects appropriate schema ordering. The assumptions are considered to be peripheral because other mechanisms for controlling serial order are consistent with the theory developed by Norman and Shallice (1980, 1986). Ordering constraints could, for example, be realised in terms of serial position varying activation gradients, as in the typing model of Rumelhart and Norman (1982) or the connectionist serial order models of speech production (Houghton, 1990) or short-term memory (Burgess & Hitch, 1992).

Object Representations and Schema/Object Interactions

Object representations and their interaction with the schema network are defined by one core

assumption (CA11), five peripheral assumptions (PA8-PA12), and one implementational assumption (IA6).

The Object Network

As noted earlier, the internal representations of objects, like schemas, have associated numerical activation values. These values serve two purposes. First, they are used in calculating the environmental influence on schemas: Objects with highly active representations tend to trigger relevant schemas more than equivalent objects with less active representations, and objects whose representations are inhibited below the resting activation inhibit relevant schemas. The second purpose of object activations concerns argument selection. When an action is to be executed, it is necessary to set its arguments (i.e. the objects to which the action is to be applied, such as a spoon in the case of a *pick-up* action). In the current implementation this assignment of arguments to actions takes the most active representation of an object that is appropriate to the action at the time when the action is to be executed.

In order for this process of argument selection to function correctly, however, it is necessary to allow the activation of an object representation to be dependent on the intended function of the object. For example, an action such as *pour* might take two arguments: a source container and a target container. In this case it is clearly necessary to distinguish the arguments: Each cannot simply be the most active representation of a container in the current environment. For this reason, an object representation's activation needs to be not just a single number, but a vector whose components represent the activation of the object representation with respect to different potential functional roles. In this way an object representation can simultaneously have high activation with respect to one purpose or function and low activation with respect to another. In the coffee preparation domain discussed later, three functional roles are employed: implement, source, and target.

PA8: The different functions that an object may serve are represented by separate activation values, and

competition between object representations operates within these functions.

Although structure in the form of part-whole relationships (as in McClelland & Rumelhart's [1981] model of letter perception) undoubtedly exists within the domain of object representations, such structure is considered beyond the scope of the current implementation. A consequence of this lack of structure is that within the (implementation of the) object representation domain there is no parallel to the top-down influence present in the schema domain. Most other aspects of the schema domain, however, are mirrored in the object domain. In particular, the effect of objects on schemas is mirrored by an effect of schemas on objects. Schemas influence objects involved in the situations that trigger them, with the degree of excitation or inhibition being dependent upon the difference from rest of the activation of the schema. This allows the development of positive feedback loops between the schema and object representation domains, whereby schemas excite (or inhibit) objects which then, by participating in triggering situations, excite (or inhibit) schemas, and so on. Capture errors may arise if this feedback is not appropriately controlled.

PA9: Object node activations are influenced by (at least) a lateral influence, a self influence, an influence from schema nodes, and random noise.

IA6: The net influence on an object representation's activation is a weighted sum of the influences cited in PA9.

PA10: The influence of a schema's activation on that of an object representation (for a particular function) is dependent on the extent to which the object representation is employed, serving that function, in the triggering conditions of the schema.

Recall that when a schema's activation is below rest the schema inhibits the representations of objects which it may use. Similarly, object representations below rest inhibit schemas which may use them. The feedback loop can therefore lead to negative priming (Tipper, 1985), in which the activations of (currently irrelevant) object representations

are pushed significantly below their rest value. We return to the psychological evidence for such priming when we examine the detailed functioning of the model.

Competition (i.e. lateral influence and self influence) also acts within the domain of object representations. Although competition would appear to prohibit the representation of multiple simultaneously active objects, the use of separate activations for different functions allows that several objects may be highly active at once, provided that they are active for different functions. As noted earlier, this appears to be necessary for argument selection within schemas that involve multiple arguments. It also simplifies argument selection during complex action sequences which require that many objects be active (for different purposes) for the duration of the action sequence.

The lateral influence on an object representation for a particular function is given by the sum of the relative differences from rest of the activations of all other object representations for that function, so competition operates between objects with respect to each function. As a consequence, at any one moment within the coffee preparation domain, at most one object will be an active implement, at most one object will be an active source, and at most one object will be an active target.

PA11: Object representations compete within functional domains. This competition is effected by a lateral influence on the activations of competing object representations, and a self influence on all object representation activations.

The self influence on an object representation for a particular function is proportional to the activation of that object representation for that function. Within the object representation domain self influence is always excitatory. Because there is no analogue in the object representation domain to schema selection there is no equivalent to the switching of self influence between excitatory and inhibitory modes within the domain.

A final issue relating to object representations concerns the creation and deletion of elements within the object representation network. In the simulations reported here, it is assumed that all

objects remain present throughout the entire task. Mechanisms for introducing and removing object representations are therefore not considered.

Argument Selection

As noted earlier basic level schemas correspond directly to discrete actions (i.e. actions whose sub-structure is not represented within the contention scheduling system). When a basic-level schema is selected the motor system is invoked to carry out the corresponding discrete action.

A single action can generally be applied to a variety of objects, frequently using different resources. Thus, a *pick-up* action might be effected with either hand, and might be applied to any number of appropriately small objects. Before an action can be carried out its argument and resource requirements need to be fully specified.

Each discrete action specifies a (possibly empty) set of object argument roles and a (possibly empty) set of resource argument roles (cf. Schwartz et al., 1991). Each argument role has a symbolic selection restriction (Katz & Fodor, 1963). For example, the *pour* action requires (i) a specific resource—a hand holding an open container, and (ii) a specific object—a target to pour into. Selection restrictions are used to ensure that only appropriate arguments fill the argument roles of actions.

CA11: Discrete actions, which correspond to basic-level schemas, specify selection restrictions on the objects and resources to which they may be applied.

When an action is selected, each resource argument role is mapped onto the most active appropriate resource. Object argument roles are similarly assigned, except that such roles also specify a function or purpose for their argument, and the function specifies the domain of object representation activation values that should be considered in the argument specification process. The function is generally inherited from the source schema leading to the component schema that triggered the action. Hence, if the *pick-up* schema is selected, the source schema that led to its selection will provide a context in which to interpret the argument selection (specifying, for example, if an implement [such as a

spoon] or a source [such as a sugar packet] should be picked up). Default functions are employed when, due to abnormal functioning, low-level schemas are selected in the absence of higher-level control.

PA12: Objects and resources are allocated to the argument roles of discrete actions according to selection restrictions marked on the argument roles of the action and the activation of nodes in the object representation and resource networks.

The use of selection restrictions on argument selection prevents attempts at physically impossible actions such as pouring the contents of a container into itself. However, domain-independent principles for specifying such conditions are unclear. If a selection restriction is too strong, action will be blocked when no appropriate arguments can be assigned to the argument roles (but when, for example, inappropriate arguments may exist). If the restrictions are too weak, impossible actions will be attempted. The simulations conducted here have used relatively strong selection restrictions (thus blocking actions when no appropriate arguments can be determined).

The Resource Network

The interaction of resource requirements with schema selection and the allocation of resources to selected schemas are defined by three peripheral assumptions (PA13–PA15) and one implementational assumption (IA7).

In general, schemas require allocation of appropriate resources for their successful execution. A low-level *pick-up* schema, for example, requires use of a hand to effect the corresponding action. Other schemas might require use of cognitive resources such as language processing mechanisms.

The problem of resource allocation is analogous to that of the selection of objects for actions. We therefore suggest that resources compete for activation in a network analogous to the object network, and that resource allocation follows the same principles as the allocation of objects to the argument roles of actions. Thus, when a basic-level schema is selected, resources are allocated to its corresponding action according to selection restrictions (speci-

fied by the action) and the relative activation of the resources, with the most active appropriate resource(s) being assigned to the resource argument role(s) of the action.

PA13: Resources participate in an interactive activation network, parallelling that of object representations.

PA14: Discrete actions specify the resources which they require, and when an action is executed the most active appropriate resources are allocated to it.

We assume that the parallels between the object representation and resource networks also extend to the reciprocal activation of resources by schemas and schemas by resources. Thus, the activation of both resources and object representations contribute to the activation of schemas which may use them, and, in turn, schemas excite the resources that they require. As with object representations, basic-level schemas are linked directly to the resources they require, and higher-level schemas are linked to all those resources to which their component schemas are linked. This allows high-level resource allocation considerations (e.g. what resources will be required by other subschemas within the current schema) to be combined with low-level resource allocation considerations (e.g. what resources are available now). The net effect of this is, for example, to select the appropriate hand when picking up an object when either hand can reach the object, but when one particular hand is required for a later operation within the schema.

PA15: If a resource can be utilised by a schema (or one of its components) then it activates that schema and vice versa.

IA7: The excitation of a schema on a utilisable resource (and vice versa) is proportional to the difference between the schema's activation and rest activation.

Cognitive resources are not employed in the simulations reported here, and no attempt has been made to include them in the resource network. Instead, the network implemented here consists of just two nodes, corresponding to the two hands.

Given this, the competitive effects of lateral influence and self influence within the resource network have been omitted from the implementation.

Parameters of the Model

A total of eight parameters govern the flow of activation within and between the various networks of the model. Two parameters control general aspects of network dynamics: rest activation (the activation level to which activations in all domains tend in the absence of any net input) and persistence (the degree to which activation values persist over time with zero net input). These parameters work together to yield smooth activation profiles throughout processing. In the simulations reported later, they take the values .10 and .80 respectively, but an increase in one can typically be countered by a decrease in the other to maintain smooth activation profiles (Cooper & Shallice, 1997).

A third parameter controls the standard deviation of normally distributed random noise that is added to the net influence in all domains before the effect of that influence on activation values is calculated. Some noise is essential if competitive processes are to separate otherwise equal competitors, and in order to provide some variation in behaviour. Random variability is also a characteristic of the nervous system. If the noise level is too high, however, it may over-ride competitive processes, leading to spurious action selection. Normally distributed noise with standard deviation of 10^{-3} was used in the simulations reported later.

Three balance parameters control the contribution of the various activation influences to the net input within all networks. In the most general case, nodes of a network have four inputs (ignoring noise): self influence, lateral influence, an internal influence (top-down influence in the case of schemas), and an external influence (the influence from the environment in the case of schemas and the influence from schemas in the case of object representations and resources). The first two of these relate specifically to competitive process. The parameter Self:Lateral controls the relative proportion of self influence and lateral influence in the final influence on a node. The second two inputs to

a node are non-competitive influences, and the Internal:External parameter controls the proportion of each of these. A third balance parameter controls the proportion of competitive and non-competitive influences contributing to the total excitation or inhibition of a node. (See Appendix A for further details.)

The parameters are specified as proportions of total activation from the various influences primarily because informal investigations revealed that the behaviour of the system was frequently determined by ratios of activation sources. For example, qualitatively similar behaviour was observed over a range of weightings of activation influences when the ratio of contributions from lateral influence and self influence were fixed. Such invariants can, by realising them in terms of parameters, be more easily be fixed whilst other aspects of the system are varied.

Within the schema network, a further parameter, the selection threshold, controls schema selection. This parameter specifies the activation level above which a schema node must be excited in order to be selected. When competition is functioning appropriately the model is relatively insensitive to the value of this parameter, but if the selection threshold is extremely high (above .95), schema selection tends to fail because schemas cannot become sufficiently active (unless schema persistence is also increased), and if the threshold is too low (below .50), schemas tend to be selected before competitive processes have achieved their purpose, and so spurious selection of schemas may occur. In this situation schema selection is rapid but generally inappropriate. The simulations reported below were conducted with a selection threshold of .60.

In the normally functioning system we assume that those parameters which are not specific to a particular domain of the model (the rest activation, persistence, Competitive:Non-competitive and Self:Lateral parameters) have the same value in all relevant domains. This assumption, which was not present in earlier work (cf. Cooper, Shallice, & Farrington, 1995), substantially reduces the parameter space and constrains the range of behaviours of which the normally functioning model is capable.

BEHAVIOUR OF THE MODEL

The model of action selection that we have presented applies to routine or well-learned actions within the intermediate domain. In evaluating the model we are primarily concerned with three issues: the ability of the model under “normal circumstances” to produce well-structured action sequences in complex hierarchically structured tasks; the susceptibility of the normally functioning model to the action lapses observed in normals and described earlier in the section on Lapses in Normal Action Selection; and the ability of the model, when damaged, to yield behaviour qualitatively similar to that described by Schwartz et al. (1991, 1995, 1998) as action disorganisation syndrome.

Applying the Model in a Specific Situation: The Coffee Preparation Domain

The model, and the theory on which it is based, is independent of any particular task—the same theory is held to account for all routine behaviours within the intermediate domain, ranging from dressing and grooming to routine aspects of driving. In order to examine these issues we must, however, focus on a specific task. The task we have adopted is based on one of the situations in which Schwartz et al. (1991) observed a neurological patient exhibiting extensive action disorganisation, namely the eating of an (institutional) breakfast, and in particular that part of the breakfast that involves the preparation of coffee. This situation was selected for four reasons:

1. The activity has been analysed by Schwartz et al. as involving a five-level hierarchy of action control (see Fig. 3). Thus the organisation of an appropriate sequence of sub-actions presents a considerable challenge to a controller that does not operate using purely symbolic procedures such as a production system.
2. The variety of sub-actions involved means that there are substantial opportunities for capture errors to occur.

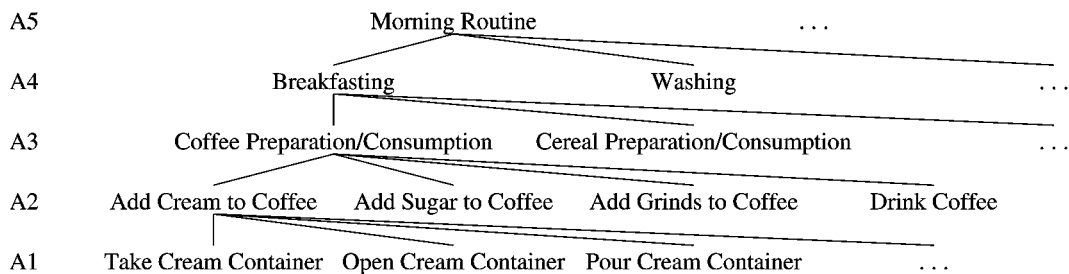


Fig. 3. Schwartz et al.'s hierarchy of control.

3. The number and variety of objects in the environment is sufficiently small to be computationally tractable but still allows the possibility of much inappropriate triggering and argument selection.

4. Behaviour of neurological patients on the task has been examined by Schwartz et al., and it is therefore possible to compare the behaviour of the model following lesioning with that of neurological patients.

Certain simplifications have been made to the Schwartz et al. situation for implementational reasons. The environment in which the patient was placed is modelled by a breakfast tray on which are positioned 13 objects. The tray is represented by an 8 × 4 grid. Objects have features representing contents (for packets and containers), state (open or closed), and position. Position is specified in terms of the coordinate grid, ranging from - 4 (left) to + 4 (right) in the first coordinate and 0 (towards participant) to + 4 (away from participant) in the second. Table 2 lists the simulated objects on the tray with their features.

Similar simplifications have been made within the schema domain: the schema repertoire of the system has been reduced to 22 (comprising those given earlier in Fig. 2 together with a number of “distractor” schemas such as *Eat from Spoon* and *Drink from Container*). The discrete actions of the coffee preparation world, together with their selection restrictions on resources and arguments, are shown in Table 3.

Normal Functioning in the Coffee Preparation Domain

Given the schema hierarchy and object representations described earlier, the model is indeed capable of producing hierarchically structured action and argument selection. Figure 4 shows the schemas selected and actions effected by the system in a typical run. Twelve actions are performed in total. The first four achieve the goal of adding coffee grinds to the mug, by picking up the packet of coffee grinds, tearing the packet, pouring its contents into the mug, and discarding the packet. In the next phrase, sugar is added from the sugar bowl to the mug with the aid of a spoon. Finally the milk carton is opened and a dollop of milk poured into the mug.

Schema Selection

Of particular interest within the transcript is the order of schema selection and the assignment of arguments to argument roles after an action has been selected. Although serial ordering at the low-

Table 2. The Contents of the Breakfast Tray

Object	Contents	State	Position
Sugar packet	2 measures of sugar	Closed	(- 2, + 2)
Salt packet	1 measure of salt	Closed	(- 2, + 2)
Coffee packet	1 measure of coffee	Closed	(- 2, + 2)
Sugar bowl	10 measures of sugar	Open	(- 3, + 2)
Cereal bowl	20 measures of oatmeal	Closed	(0, + 1)
Toast plate	4 pieces of toast	—	(+ 3, + 2)
Juice glass	5 measures of juice	Open	(+ 2, + 2)
Coffee mug	15 measures of hot water	Open	(- 2, + 2)
Coffee jar	50 measures of coffee	Open	(- 3, + 3)
Coffee jar lid	—	—	(- 2, + 2)
Milk carton	50 measures of milk	Closed	(+ 3, - 3)
Spoon	Empty	—	(+ 1, + 2)
Stirrer	—	—	(- 2, + 2)

Table 3. *Primitive Coffee Preparation Actions*

<i>Action</i>	<i>Resources</i>	<i>Arguments</i>
Pick Up	Empty hand	Object within reach
Put Down	Hand holding object	—
Dip Spoon	Hand holding empty spoon	Non-empty open container
Empty Spoon	Hand holding non-empty spoon	Open container
Eat From Spoon	Hand holding non-empty spoon	—
Stir	Hand holding stirrer	Non-empty open container
Pour	Hand holding non-empty open container	Open container
Drink	Hand holding non-empty open container	—
Tear	Hand holding packet + empty hand	—
Screw Open	Hand holding jar + empty hand	—
Screw Closed	Hand holding jar + hand holding screw lid	—
Swap	Hand holding object + empty hand	—

est level is partially determined by the ordering constraints specified in terms of symbolic preconditions, the order of schema selection is in practice largely driven by the environment. Once *Add Sugar from Bowl* has been selected, for example, the only subgoal of this schema that will be possible given the normal state of the environment is to pick up an implement. Schemas that may achieve this goal therefore receive excitation from the environment. All other sugaring actions (dipping the implement into a source, emptying the implement into a target, and setting the implement down) receive no environmental excitation. Under these conditions, the *Pick Up* schema will win out over its competitors (the other sugaring actions) because it will be receiving more net excitation than any of those competitors. Once an object has been picked up (hopefully a spoon), symbolic preconditions prevent top-down excitation from immediately triggering *Put Down*. Instead, environmental triggering biases selection toward *Dip Spoon*, which in due course is selected and then inhibited. Similar considerations then lead to the selection and inhibition of *Empty Spoon*. Finally, completion of *Empty Spoon* fulfils the preconditions of *Put Down*, and allows top-down excitation to trigger that schema.

The activation profiles of the various schemas throughout the run are shown in Fig. 5. Schema selection (as indicated by the cycle numbers) can be seen to correspond to the peaks in the activation profiles. Hierarchical grouping of schemas can also

be seen in this figure, with the activation peaks of selected schemas bracketing those of their subschemas.

Argument Selection

Argument selection is somewhat more complex. Consider the case of picking up the spoon when *Add Sugar from Bowl* is selected. When this schema is selected within *Prepare Instant Coffee*, the representation of the coffee mug will generally be highly active as a target due to excitation from *Prepare Instant Coffee*. *Add Sugar from Bowl* will activate the representation of the sugar bowl as a source, and the representation of the spoon as an implement. When *Pick Up* is selected, *Add Sugar from Bowl* specifies that it is an implement (and not a source or target) that should be picked up, and so the spoon, as the object whose representation is most highly active as an implement, is assigned the role of the object to pick up. This contrasts with the case of *Add Sugar from Packet*, where it is a source, and not an implement, that should be picked up. The different requirements of these two cases illustrate the rationale behind the use of distinct activation values for different object functions within the domain of object representations.

Figure 6 shows the activation profiles of the object representations (with respect to the function of “source”) throughout a trial of the coffee preparation task. The graph is based on the same trial as

6	+ <i>Prepare Instant Coffee</i>
34	+ <i>Add Coffee from Packet</i>
61	+ <i>Pick Up</i>
61	Picking up coffee packet with left hand
66	- <i>Pick Up</i>
90	+ <i>Tear</i>
90	Tearing coffee packet (with right hand)
94	- <i>Tear</i>
103	+ <i>Pour</i>
103	Pouring coffee packet into coffee mug
110	- <i>Pour</i>
127	+ <i>Put Down</i>
127	Putting down coffee packet
137	- <i>Put Down</i>
148	- <i>Add Coffee from Packet</i>
198	+ <i>Add Sugar from Bowl</i>
222	+ <i>Pick Up</i>
222	Picking up spoon with left hand
226	- <i>Pick Up</i>
242	+ <i>Dip Spoon</i>
242	Dipping spoon into sugar bowl
245	- <i>Dip Spoon</i>
249	+ <i>Empty Spoon</i>
249	Emptying spoon into coffee mug
255	- <i>Empty Spoon</i>
268	+ <i>Put Down</i>
268	Putting down spoon
277	- <i>Put Down</i>
285	- <i>Add Sugar from Bowl</i>
295	+ <i>Add Milk from Carton</i>
327	+ <i>Pick Up</i>
327	Picking up milk carton with right hand
332	- <i>Pick Up</i>
356	+ <i>Tear</i>
356	+ Tearing milk carton (with left hand)
360	- <i>Tear</i>
369	+ <i>Pour</i>
369	Pouring milk carton into coffee mug
376	- <i>Pour</i>
388	+ <i>Put Down</i>
388	Putting down milk carton
396	- <i>Put Down</i>
397	Trial complete

Fig. 4. Action selection in the coffee preparation domain. Schema names are presented in italic type, with selection indicated by "+" prefixes and deselection indicated by "-" prefixes. The numbers in the left column indicate the processing cycle on which the corresponding event occurred.

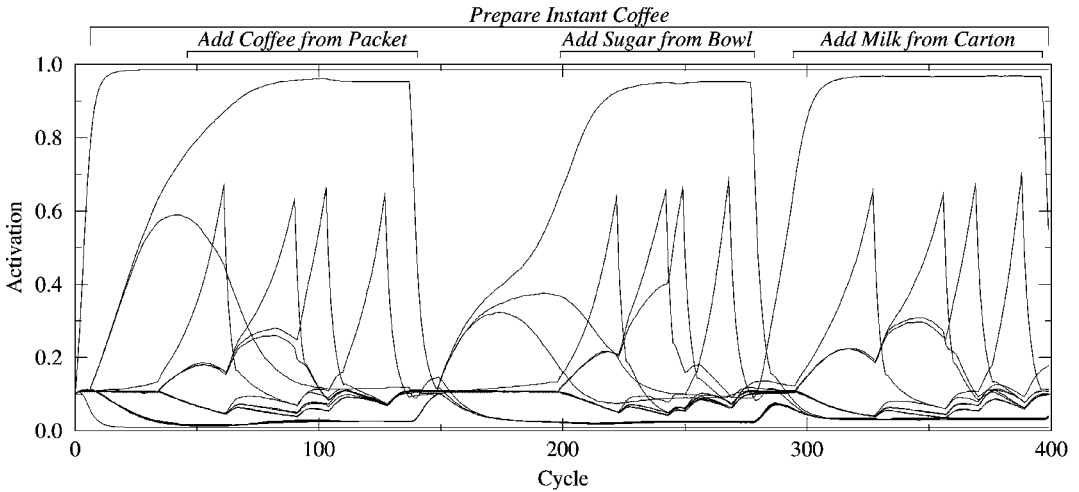


Fig. 5. Schema activation profiles in the coffee preparation domain.

that in Fig. 5. It can be seen that object representations cluster into three groups based on their activations. The least active representations correspond to those objects that are present throughout the duration of the task but that are entirely irrelevant. These include a salt packet, a cereal bowl, and a glass of orange juice. Of particular interest is the fact that the activations of these representations are below the rest activation of .10. Representations with intermediate levels of activation correspond to objects that are relevant to the task but not in use.

This level includes the various sources of sugar, coffee, and milk that might be used in preparing the coffee. The most active representations correspond to objects that are selected as arguments. The three peaks in the graph correspond (from left to right) to the representations of a coffee packet, a sugar bowl, and a milk carton. Activation of these representations begins at the intermediate level, peaks when the object is being used, and drops back to the intermediate level when the schema involving that object is completed (or to the lowest level if the

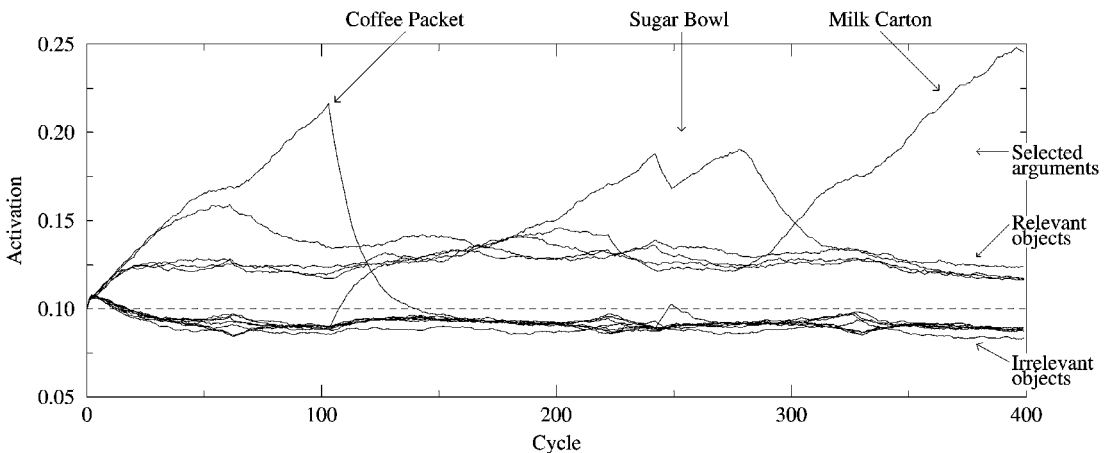


Fig. 6. Object representation profiles in the coffee preparation domain (see text).

object is no longer relevant to the task, such as the empty coffee packet at cycle 103 in this example).³

Psychological studies of priming provide some evidence for different levels of activation of object representations. A number of studies (e.g. Tipper, 1985; Tipper, Weaver, Kirkpatrick, & Lewis, 1991) have found that, under certain conditions, response times can be slower in a primed condition than in the control condition. Such negative priming occurs when a stimulus shares some features with a distractor. It is taken to indicate the inhibition (below some resting level) of distracting information. Negative priming provides strong support for the use of rest activation within our domain of object representations and for the observed inhibition of representations of irrelevant objects below this rest value.

There is also empirical support for the use of separate activation levels for distinct object purposes. Tipper, Weaver, and Houghton (1994) compared distractor inhibition (i.e. inhibition of irrelevant stimuli) in situations that differed with respect to their behavioural goals. Tipper et al. found that only those features of distractor representations that were most closely associated semantically with the particular action to be performed were inhibited. Other features of distractors could remain active and could even facilitate subsequent behaviour. This is precisely the behaviour expected from a system in which object representations have separate activation levels for distinct purposes: An object can be inhibited with respect to one purpose but remain active with respect to another. Thus, in the coffee preparation domain, although there is only one active target throughout the entire task (the coffee mug), several different objects are, at different times, active as sources as the task progresses.

Resource Selection

Resource selection also presents problems. Only two resources are available in the coffee preparation

task—the left hand and the right hand. In the transcript, the left hand is used to pick up the coffee packet and the spoon, but the right hand is used to pick up the milk carton. This is because the milk carton is on the right side of the simulated breakfast tray, whereas the other objects are on the left. Environmental activation of effectors takes account of reaching constraints so that appropriate effectors are active when action selection occurs.

Effects of Noise

Random noise within the schema, object, and resource networks leads to variability in the model's behaviour. This variability manifests itself in three distinct ways: variability in the total number of cycles to complete the task; variability in the order in which schemas that are not subject to ordering constraints are selected; and variability in the specific schemas selected for a goal when several schemas are equally applicable. Thus, in the coffee preparation task, the model will on some occasions add sugar before adding milk to the coffee mug, and on other occasions add milk first. In addition, when adding sugar, the model will on some occasions add sugar from a sugar bowl, and on other occasions add sugar from a sugar packet.

However, random noise is not added to the model solely to yield variability in behaviour. It is necessary to prevent competitive ties within the activation networks. Without noise, two competing schemas could, in principle, have the same activation and receive the same net excitation. Self influence and lateral influence between such tied schemas cancel each other out. At best, such tied schemas slow down competitive processes. At worst, they block competitive processes completely. Noise prevents such situations from arising and generally ensures that competition proceeds at a reasonable rate.

³The line on the graph that rises at approximately cycle 110 from the lowest level to the intermediate level corresponds to the coffee mug, which at cycle 103 becomes a relevant source of coffee preparation materials (in virtue of the addition of coffee grinds to the mug at this cycle). At the same time, the coffee packet drops to the lowest level as, being empty, it is no longer relevant to the coffee preparation task.

Dependence upon Parameter Values

Normal functioning of the model as reported here is dependent upon appropriate settings of the numerical parameters described in the section on Parameters of the Model. Elsewhere (Cooper & Shallice, 1997) we have shown that an earlier version of the model is robust with respect to a variety of parameter manipulations. In brief, a variety of different parameter configurations were found to yield appropriately structured error-free behaviour similar to that reported here. Although the differences between the earlier version of the model and the current version affect the precise ranges of parameter values over which the results hold, they do not alter the general pattern of those results.

Action Lapses in the Coffee Preparation Domain

The model described above is robust in the sense that, for significant areas of the parameter space, it produces error-free goal-directed action sequences. However, as noted in the section on Theoretical and Empirical Background, a variety of action lapses arise in the routine action of neurologically intact individuals. In order to account for these behaviours we assume that factors such as stress, fatigue, operation in an implicitly dual task situation, and failure to monitor the world can affect the values of system parameters, thus impairing the functioning of the system and ultimately resulting in action lapses. Most types of lapse may occur under a variety of conditions. The following sections illustrate some of the circumstances within the model under which each type of lapse described in the section on Lapses in Normal Action Selection may occur.

Capture Errors

Schemas are triggered by situations in the (representation of the) environment which are compatible with their execution. This triggering is normally insufficient by itself to lead to selection of the schema. However, if the environmental influence on schemas is too strong (or equivalently the top-down influence is too weak) then schemas at the lowest level may be selected purely through envi-

ronmental triggering, leading to toying behaviours. In the model this corresponds to situations in which the Internal:External parameter within the schema network (henceforth referred to as Internal:External,) is too low. The toying behaviours correspond to utilisation behaviour or low-level capture errors.

Higher-level capture errors require that top-down excitation is intact in order to structure the lower-level behaviours. Capture errors at such levels may still arise, however, if competition is inappropriately resolved, as when, for example, self activation in the schema network is unnecessarily high. An example of this behaviour is shown in Fig. 7(a).

Omission and Anticipatory Errors

There are at least two ways in which omission errors may arise within an activation-based action selection system such as the one described here. First, schemas corresponding to the omitted actions may simply not be activated to threshold, and hence not be selected. This occurs in the current model if self activation in the schema network is too low or if top-down and environmental influences are underweighted. Second, schemas might be selected, but their execution might be prohibited by, for example, the inability to select appropriate arguments or resources for the corresponding actions. The former possibility may lead to the omission of an entire subtask. The latter generally arises in conjunction with anticipation errors, where an action that should be performed later in a sequence is attempted before its preconditions have been satisfied. Figure 7b shows an error of the second variety—omission of the crux action when adding milk to the coffee. This error arose with reduced influence from top-down and environmental sources.

Perseverative Errors

Perseverative errors typically arise when competitive processes break down. If self activation in the schema network is too great, for example, or equivalently if lateral inhibition is insufficient, schemas may fail to be deselected at the appropriate time. The perseverative behaviour in Fig. 7(c) arose in precisely these circumstances.

- | | |
|----|--------------------------------------|
| 60 | Picking up sugar bowl with left hand |
| 62 | Drinking from sugar bowl |
| 63 | Putting down sugar bowl |
- a. A capture error: drinking (from the sugar bowl!) before the coffee is made
- | | |
|-----|--|
| 464 | Picking up milk carton with right hand |
| 504 | Tearing milk carton (with left hand) |
| 538 | Putting down milk carton |
- b. An omission error: failing to pour milk into the coffee
- | | |
|-----|---|
| 189 | Picking up spoon with left hand |
| 204 | Dipping spoon into sugar bowl |
| 212 | Emptying spoon into coffee mug |
| 225 | Putting down spoon |
| 252 | Picking up sugar packet with left hand |
| 272 | Tearing sugar packet (with right hand) |
| 283 | Pouring sugar packet into coffee mug |
| 299 | Putting down sugar packet |
| 314 | Picking up sugar packet with right hand |
| 343 | Pouring sugar packet into coffee mug |
| 377 | Putting down sugar packet |
- c. Perseverative errors whilst adding sugar
- | | |
|-----|---------------------------------------|
| 227 | Picking up cereal bowl with left hand |
| 271 | Pouring cereal bowl into coffee mug |
| 299 | Putting down cereal bowl |
- d. An object substitution error: using the cereal bowl as a source of coffee ingredients

Fig. 7. *Errors in preparing coffee.*

Apparent perseveration may also arise from perseverative object substitution errors. If an object remains active after it has been used, it may be selected for subsequent actions. Many action subsequences in the coffee preparation domain involve the same four basic actions (pick up, tear, pour, put down) and such re-use of an object can therefore appear to be perseveration.

The model in its current form does not exhibit so-called recurrent perseveration (Sandson & Albert, 1984), in which a subtask is repeated after

one or more intervening subtasks. This form of perseveration would seem to require a less rigid mechanism for recording goal achievement. At present once a subgoal is achieved it is ticked off from its source schema's subgoal list. The subgoal list is not reinitialised until the source schema is deselected. Recurrent perseveration could arise if, for example, this ticking-off process was subject to error or decay.

Object Substitution Errors

The correct selection of arguments depends upon a schema's correct arguments having the most active representations when the schema is selected. Object or place substitution errors arise when this condition is not met. The condition may fail when, for example, noise is high and/or the excitation of object representations by schemas is insufficient. The example in Fig. 7(d) (using the cereal bowl in place of a coffee-related substance for the source in one of the coffee subtasks) arose when both manipulations were made.

Abnormal Functioning in the Coffee Preparation Domain: Action Disorganisation Syndrome

As noted in the introduction, a primary objective of the current work is to show that, when lesioned, the model exhibits behaviours similar to those observable in neurological patients. In this section we report the effects of one form of impairment to the functioning of the model, namely modification of parameter values beyond the range yielding normal behaviour within the coffee preparation domain.

Rationale and Method

The approach adopted in this paper to accounting for neuropsychological phenomena is rather different from the approaches typically used within connectionist modelling to relate model behaviour to patient behaviour. Essentially two methods have been used—removal of connections or units and addition of noise (see, e.g. Hinton & Shallice, 1991). Here we have followed the approach pioneered by Martin, Dell, Saffran, and Schwartz (1994) for interactive activation models in their

work on modelling disorders of speech production. This is to consider the behaviour of the model in parameter space outside of the region that gives rise to normal performance.

In order to establish parameter modification as a plausible approach to the modelling of damage it is necessary to consider the possible relationships between parameter values and neurological correlates that may sustain damage. One can envisage at least two such relationships. First, a parameter may be taken to correlate with the level of a neurotransmitter. (Later we speculate that such a relationship may hold between the Self:Lateral parameter and the neurotransmitter dopamine.) In such cases, modifying the parameter's value would correspond to modifying the neurotransmitter concentration (or perhaps to modifying receptivity to the neurotransmitter) at the neurological level. Second, other parameters may be seen as reflecting connectivity patterns between functional subcomponents of the system. For example, the Internal:External parameter may be related to the ratio of the connectivity from the two activation pathways to schema nodes. Modifying such parameters corresponds to modifying the connectivity ratio. Such modification might arise, for instance, through partial ablation of either pathway.

Modelling impairment in terms of variation to a parameter is also consistent with the theoretical account given by Schwartz et al. (1991) of the action disorganisation of their patient HH. They argued that his errors arose primarily from the weakening of top-down control over action selection. A weakening of top-down control in the model corresponds to a decrease in the Internal:External parameter. We thus focus our discussion of behaviour in the coffee preparation domain following lesioning by examining behavioural patterns arising from variation of this parameter's value.

A number of simulations were performed in order to test both the model and Schwartz et al.'s hypothesis. Fifty trials of the coffee preparation task were simulated at each point in the parameter space as the value of the Internal:External parameter ranged from .00 to 1.00 in increments of .01. Any trial that was not complete (i.e. in which the

goal of preparing coffee had not been achieved) within 2500 cycles—more than 5 times the normal task completion time—was aborted. Transcripts, consisting of the sequence of actions attempted, were collected for all trials and scored using two (largely independent) scoring systems.

Scoring

A major difficulty within the domain of action is the characterisation of error types. The complexity of behaviour is such that there is no simple way in which to characterise behaviour, and therefore to test the predictions of different models. The first systematic treatment that allowed a provisional characterisation was the Action Coding System of Schwartz et al. (1991). This coding system was developed in order to provide a quantitative description of disorganisation (as opposed to error) in routine action. Within the system an action sequence is described in terms of the number of basic actions performed and the proportion of those actions that are independent (roughly, those actions that are neither crux actions nor actions contributing to the next crux action). The system also yields two further dependent measures: the number of crux and noncrux errors. Crux errors comprise object/place substitutions and, in the case of coffee preparation, drinking anticipations. Non-crux errors comprise omissions, instrument substitutions, and faulty action execution. We refer to this coding system as ACS₁.

ACS₁ was used to score the action transcripts produced by the simulations. (The results are presented later.) ACS₁ was adopted because of its use by Schwartz et al. (1991) in describing the action disorganisation of HH whilst preparing his morning coffee (i.e. the task on which the current simulations are based). Use of ACS₁ therefore allows a direct comparison of the model's behaviour with that of HH. Schwartz et al. (1995) also employ ACS₁ to analyse the disorganised behaviour of a related patient, JK, during a super-ordinate task, breakfasting. This provides a further reference point for the model.

ACS₁ is a coarse-grained coding system. It does not distinguish between different types of error, and so patients with qualitatively different error

patterns can yield similar quantitative descriptions. In recent work Schwartz et al. (1998) have adopted a more detailed coding system. We refer to this system, which focuses on quantifying the occurrences of specific types of errors, as ACS₂. Within ACS₂ errors are coded as omissions, object substitutions, sequence errors (including anticipations and perseverations), and action additions. A variety of other error types are also recognised, but the rates of these errors are generally very low. ACS₂ allows a much more detailed analysis of behaviour. Its relevance to the current work centres on its use by Schwartz et al. (1998) to analyse the disorganised action of a group of 30 closed head injury (CHI) patients. We have therefore further analysed the simulation's behaviour using a scoring system based on ACS₂.

The variant of ACS₂ that we have used (referred to here as ACS₂) differs from that developed by Schwartz et al. (1998) in two ways. Both of these differences concern the treatment of omission errors.

First, within ACS₂ omissions errors are broken into two subtypes: step and subtask omissions. A subtask omission occurs when a complete subtask is omitted from an action sequence. In the case of coffee preparation this may occur if, for example, there is no evidence of the subtask of adding milk being performed. Performance of any operation on the milk carton would be taken as evidence of inclusion of the subtask. A step omission occurs when a single action (e.g. pouring the milk into the coffee) is omitted from a sequence. The division of omissions into these two subtypes is intended to prevent a single high-level (i.e. subtask) omission from artificially inflating the omission score. In our domain a single such error is equivalent to the omission of four basic steps. It would not be surprising if, without this distinction, omissions were to dominate the error analysis.

A second, more significant, difference between ACS₁ and ACS₂ concerns the treatment of certain sequence errors, specifically Schwartz et al.'s (1998) category of anticipation-omission errors. Schwartz et al. distinguish between pure omissions (e.g. failing to use cream in coffee) and anticipation-omissions (e.g. closing a lunch-box without packing it).

The latter are considered to be sequence errors, and excluded from the count of omission errors. The distinction receives some support from the significantly different prevalence of the error types. However, in our opinion it is generally a difficult distinction to maintain. The examples cited here are a case in point. We have therefore not attempted to distinguish the two forms of error, and have included anticipation-omission errors in our category of step omissions.

Results

As noted earlier, simulations were performed with the Internal:External_i parameter varying over its entire range. When the value of the parameter was below .35 the behaviour of the model was entirely unstructured. Several hundred actions were performed at some parameter values, few were performed at others. In all cases all actions were independent in the ACS₁ sense. Within this region triggering activation from the environment is excessive, leading to behaviour that is driven entirely by the environment. Typically this consists almost exclusively of "toying" behaviour—repeatedly picking up and putting down one object (e.g. a spoon or a sugar packet). We do not consider this behaviour in detail, but note in passing its similarity to behaviour described in a utilisation behaviour patient by Shallice et al. (1989).

Figure 8 shows the ACS₁ analysis of the model's behaviour for values of the Internal:External_i parameter between .30 and 1.00. Two principal areas of qualitatively different behaviour are apparent. In the first area, when the parameter's value is less than .57, there is clear breakdown of behaviour. Even when the parameter's value is fixed (but within this region), large variations in behaviour between trials occur. These variations are evident in both the number of actions performed and the proportion of those actions that are independent. As the parameter's value increases toward 0.50 (corresponding to increasing top-down influence and decreasing environmental influence), however, structure begins to emerge. This is apparent from the general decrease in the proportion of independent actions occurring at and above values of .35. Closer examination of the transcripts reveals that,

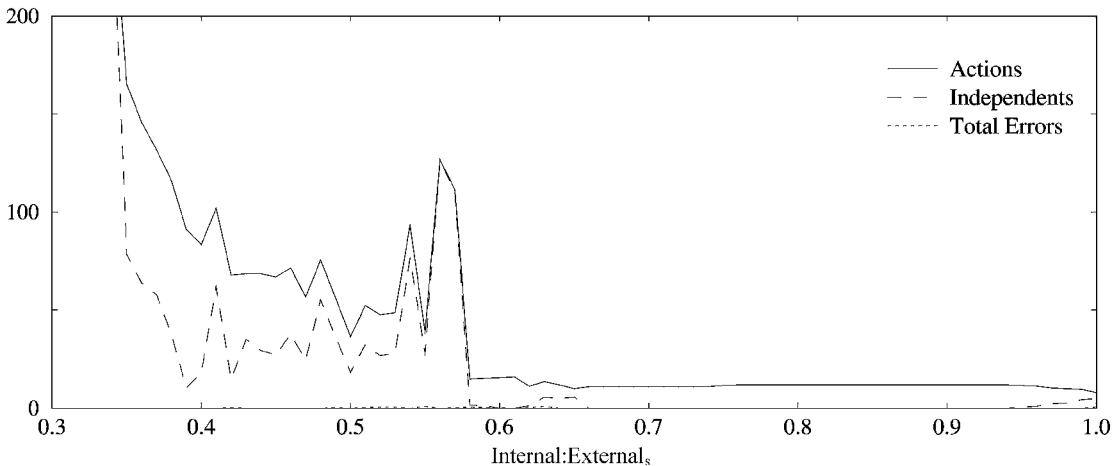


Fig. 8. Analysis of the model's action disorganisation using ACS.

within this region (between .35 and .50), one or more subtasks of coffee preparation (e.g. adding coffee grinds to the coffee mug) are generally performed appropriately, and that the remaining actions are independent toying actions. As the parameter's value continues to increase toward .57, however, the apparent structure again breaks down, and at approximately .57 all actions are independent. In the second principal area, when the parameter's value is greater than .57, behaviour is generally structured. Between 10 and 15 actions are performed on each trial and there are few independent actions. Behaviour, at least as measured by ACS₁, appears normal.

Figure 8 also shows the total number of errors, using the ACS₁ definition of error (i.e. crux plus noncrux errors). Few of the actions attempted by the model fall within this definition of "error" (which includes object, instrument, and place substitutions, drinking anticipations, omission of single actions, and faulty execution of an action, but excludes perseverations and omission of subtasks: see Schwartz et al., 1991, p. 396). Errors that do occur generally involve omission of a single action. Such errors occur sporadically when the parameter's value is between .35 and .65.

The ACS₂ analysis of the model's behaviour is shown in Fig. 9. The two graphs (i.e. Fig. 8 and 9), which are based on the same raw data, demonstrate

that there are major differences between the scoring systems. At lower values of the parameter, behaviour is dominated by sequence errors. Indeed, the vast majority of actions performed when the parameter's value is below .65 are sequence errors of the perseverative variety. Omission errors of both types are also common. When the parameter's value is below .50, approximately two step and two subtask omissions occur, but as the parameter's value increases towards .65 step omissions become more frequent and subtask omissions become correspondingly less frequent.

For higher values of the Internal:External_s parameter, use of ACS₂ reveals significant breakdowns of behaviour not apparent from ACS₁. Thus, Fig. 9 shows that error-free behaviour occurs only when the value of Internal:External_s is between .75 and .95. The region immediately to the left of this, between .65 and .75, is of particular interest. ACS₁ is insensitive to the action disorganisation in this region. ACS₂, however, characterises the behaviours as including one subtask omission and several sequence errors. In fact, analysis of the internal state of the model during performance of the task reveals that the errors in this region stem from a single perseverative object substitution error: The wrong object is selected during one subtask (e.g. the sugar packet is picked up in place of the milk carton when intending to add milk), and then used in place of the correct object. The actions within the subtask are then counted as sequence errors and the subtask

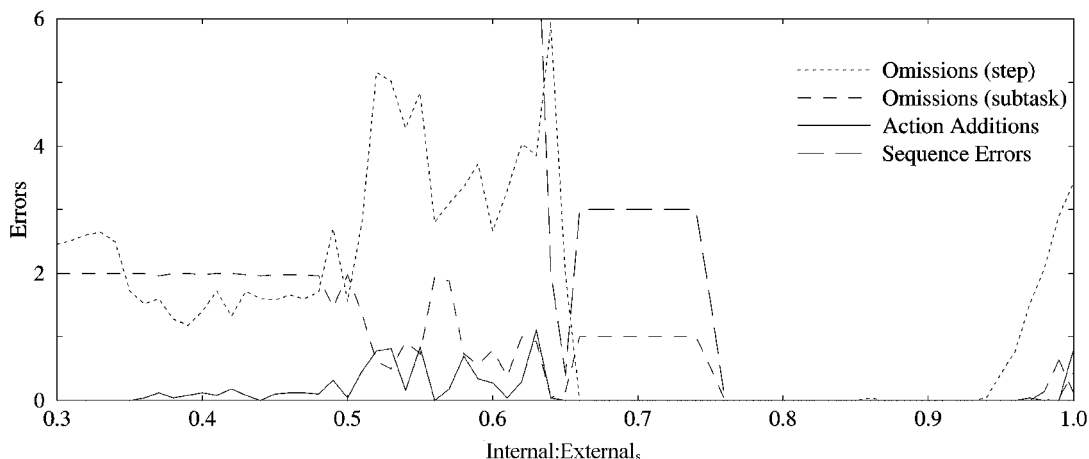


Fig. 9. Analysis of the model's errors using ACS₁. Note that sequence errors when Internal:External_s is below .63 are off the scale. The overall pattern of results is described in the Results section. Relations between the errors produced by the model and by neurological patients are considered in the Discussion.

itself is considered to have been omitted. Although the classification of these errors by ACS₂ may appear dubious, it is not possible to distinguish perseverative object substitution errors from perseverative sequence errors without referring to the internal state of the model during execution of the erroneous sequence. This is a general problem for empirical analyses of this type of behaviour.

Also shown in Fig. 9 are action additions. These occur at low rates throughout much of the parameter's range. Not shown on the graph are object substitution errors. No errors of this type were observed within the region of parameter space shown. Object substitution errors are one of the few forms of error included within the ACS₁ definition of error. Their absence is curious, but consistent with Fig. 8. We discuss possible reasons for the lack of these errors below.

Discussion

We have identified three principal data sets with which the model's behaviour following lesioning may be compared: the ACS₁ analysis of patient HH (Schwartz et al. 1991), the ACS₁ analysis of patient JK (Schwartz et al. 1995), and the ACS₂ analysis of a group of CHI patients (Schwartz et al., 1998).

Comparison of the model's behaviour, as analysed in Fig. 8, with the ACS₁ analysis of the behaviour of HH reveals that performance of the lesioned model shares several characteristics with that of the patient. Thus, HH's performance varied radically from day to day. The number of actions performed ranged from 10 to 40 (controls generally produced approximately 32 actions over the duration of the task). Similar variation, even with fixed parameter values, was observed between trials of the model. Furthermore, a substantial proportion of HH's actions (generally 20% to 50%) were independent. This pattern looks most similar to that produced by the model when the value of the Internal:External_s parameter is between .60 and .62, some way below the range of normal behaviour. This is consistent with Schwartz et al.'s (1991) hypothesis that action disorganisation behaviour results from a weakening of top-down control within a hierarchically organised activation-based action selection system. The pattern of errors produced by the model, however, differs from that produced by HH. Between 10% and 20% of HH's actions were errors (including significant numbers of both crux and noncrux errors), but the ACS₁ analysis of the model's behaviour reveals few errors of either sort. Of particular concern is the virtual lack of object substitution errors, which were common in HH's behaviour.

There are several possible explanations for this difference. HH performed the coffee preparation task within the context of preparing and eating his institutional breakfast. This raises the possibility of interference from other active breakfasting schemas and object representations on HH's coffee preparation behaviour. Indeed, Schwartz et al. (1991) report interleaving of other breakfasting tasks during HH's coffee preparation. Recall that when the value of the Internal:External parameter is between .65 and .75 the model produces a combination of subtask and omission errors that actually stem from a perseverative object substitution. The presence of such errors in the model's behaviour suggests that, if the representations of other objects relevant to breakfasting were also active (due to coffee preparation taking place in the context of breakfasting), then it is highly likely that they too would be substituted into coffee preparation schemas, yielding higher rates of object substitution errors.

An alternate explanation for the difference between HH's error patterns and that produced by the model is that HH's deficit may have affected more than just the Internal:External parameter. As noted earlier, the model produces genuine object substitution errors under certain forms of lesioning. Such errors arise most frequently with weakening of the excitation of object representations by schemas. It is therefore possible that HH's behaviour results from two deficits: a deficit in top-down control and a deficit in the excitation of object representations.

A further aspect of HH's behaviour noted by Schwartz et al. (1991) concerns his tendency to perseverate on certain tasks. Although perseveration was not common within the task of coffee preparation, HH produced numerous perseverations whilst engaged in a further routine task, that of toothbrushing. These perseverations were clearly strongly influenced by the environment (e.g. the presence of a tap with a continuous supply of water). Although perseverative errors are only indirectly recorded by ACS₁ (via the count of the total number of basic actions performed in a task), an informal analysis of the model's transcripts shows that, when the value of the Internal:External parameter is less than .75, many of the unnecessary

actions performed by the model are either utilisation errors or perseverative errors.

We turn now to JK (Schwartz et al., 1995), whose action disorganisation was similar in many ways to that of HH. Object and place substitution errors were again the predominant error type, with perseverative errors being common in the toothbrushing domain. Examination of JK's action scripts reveals two points of particular interest: the scripts include instances both of the toying behaviour exhibited by the model (picking up and putting down an object without using it) and of perseverative object substitutions similar to that discussed earlier. These behaviours provide further support for the model.

ACS₁ thus reveals several qualitative similarities between the behaviour of HH and JK and that of the lesioned model. However, as discussed earlier ACS₁ provides only a coarse characterisation of action disorganisation. The use by Schwartz et al. (1998) of ACS₂ to provide a more detailed characterisation of action disorganisation in the behaviour of 30 CHI patients performing a variety of everyday tasks (including coffee preparation), and under several conditions, provides more detailed data with which the model's performance may be compared. One of the principal findings of Schwartz et al. (1998) was that the predominant error type in their more severe patients was that of omission, and not object/place substitution, as in HH and JK. Perseverative errors were also relatively rare. Several factors may account for these differences. The CHI group was heterogeneous. Some CHI patients in the study did, like HH and JK, produce mainly object/place substitution errors. The conditions under which the CHI patients were tested were also more controlled than those in which HH and JK were observed. These factors may have led to different error patterns. In any case, the incidence of omission errors in the CHI group is a highly significant finding to which the model should speak.

Direct comparison of the model's behaviour as characterised by ACS₂ and the behaviour of Schwartz et al.'s (1998) CHI patients is limited by several factors. First, although coffee preparation was one task employed in Schwartz et al.'s testing, individual data for patients on this task alone is not

available. Second, it is not strictly appropriate to directly compare the model's mean behaviour at any one parameter value with group data. Third, as noted above differences exist between some aspects of the scoring systems. Despite these caveats, broad comparisons between the model's behaviour and that of Schwartz et al.'s (1998) CHI patients are still instructive.

Recall that normal functioning occurs in the model when the Internal:External_i parameter's value is between approximately .75 and .95. The comparison with HH and JK suggests that the behaviour of action disorganisation patients may be modelled by reducing this parameter's value to between .60 and .75. From Fig. 9 it can be seen that at lower values in this range, corresponding presumably to patients with more severe difficulties, omission errors are common. Some of these omissions are clearly anticipatory, and would therefore be counted as sequence errors by Schwartz et al. (1998). Others, however, are more clearly "pure" omissions, as observed in the CHI patients. The model also produces many sequence errors. As noted earlier, these are primarily perseverative errors. Sequence errors also form a large percentage of patient errors. The main difference between the model's behaviour and that of the CHI patients is again the model's relative lack of object/place substitution errors. The reasons suggested previously for this difference are again relevant. In sum, the qualitative fit between the model and patient behaviour is strong.

Recently, Humphreys and Forde (1998) have reported two further patients, HG and FK, showing disorganisation of routine behaviour. Two findings from these patients are of particular significance to the current work. First, although both patients show broadly similar behaviour to that of HH and JK, a detailed analysis of their perseverative errors revealed two distinct forms of perseveration, and a dissociation between these forms in the two patients. HG tended to perseverate by performing a single action (e.g. cutting sellotape in order to wrap a gift) over and over. This form of perseveration has been termed continuous perseveration by Sandson and Albert (1984). FK's perseverations, in contrast, tended to involve the

repetition of actions that had been successfully completed earlier in the task (e.g. adding milk to a cup of tea). Sandson and Albert term this recurrent perseveration. Although Schwartz et al. (1991) observed both forms of perseveration in HH's attempts at tooth-brushing, the dissociation noted by Humphreys and Forde suggests that these forms of perseveration may correspond to different forms of breakdown. This is consistent with the model. The perseverative errors produced by the model when the Internal:External_i parameter is varied are all of the continuous type. Recurrent perseverations would appear to involve the breakdown of goal-monitoring mechanisms. We have not investigated this in the current work.

In summary, the basic errors that the model produces—omissions, perseverations, additions, and substitutions—are all ones that neurological patients produce. More critically, reduction in the Internal:External_i parameter from between .75 and .95, where it reproduces normal behaviour, to the region of .60 to .75 produces many qualitative similarities to the behaviour of the two patients whose errors of action disorganisation have been most closely studied (patients HH and JK of Schwartz et al., 1991, 1995). Moreover, they support Schwartz et al.'s theoretical position that a reduction in top-down control leads to the action disorganisation syndrome. Despite the qualitative similarities between the simulation and the patient data, however, there are also qualitative differences. In particular, rates of object/place substitution errors are low, and recurrent perseveration has not been observed.

GENERAL DISCUSSION

The ability of the model to account for appropriately sequenced action in normals and its success in accounting for action disorganisation behaviour in neurological patients is encouraging. We now turn to two issues arising from the model and its development, namely similarities between this and related models, and further neurological deficits into which the model may provide insight.

Related Models

The dynamic and sequential nature of the action selection domain, combined with the hierarchical nature of action schemas and the forms of errors seen in normals and neurologically impaired individuals, pose a set of problems for any computational account of action selection. Take, for instance, simple recurrent network models (e.g. Elman, 1990; Jordon, 1986), which have been advocated in other sequential domains. Whilst Elman's work on syntax has shown that networks with feedback from intermediate layers are able to learn context representations that are sensitive to hierarchical structuring (Elman, 1993), we know of no work in which such networks have been shown to be able to account for errors of the type observed in the action domain (specifically omission and other ordering errors). The principal difficulty in obtaining such errors within recurrent networks appears to arise from the lack of any separate representation of hierarchical relations (i.e. source/component schema relationships) and order information (i.e. the relative ordering of component schemas within a single source schema). It is thus difficult for order information to be disrupted without disruption to hierarchical relations.

Notwithstanding these points, the model presented here has similarities with a number of models current in the literature. The basic mechanism, interactive activation, was first implemented by McClelland and Rumelhart (1981) in their model of letter and word perception. However, the application of an interactive activation approach to the action domain, where the dynamic nature of action is central, rather than the perception domain, where stimuli are typically regarded as static, means that there are of necessity a number of key differences.

In their model, McClelland and Rumelhart were not concerned with sequences of outputs—each input-output mapping was treated entirely independently of all others. For a given input, activation values in their model tended toward a single stable state, and this state was taken to represent the result of the perceptual process. One difference, therefore, between the model developed here and that of McClelland and Rumelhart concerns the

production of response sequences. This includes the incorporation of mechanisms for changing the dynamics of the network (through schema selection and deselection) and mechanisms for keeping track of behavioural goals.

A second group of models to which the contention scheduling model is related are the so-called competitive queueing models (Houghton, 1990), in turn linked to the typing model of Rumelhart and Norman (1982). Houghton and colleagues have developed a series of models based on interactive activation within the domain of spelling (Glasspool, 1998; Houghton, Glasspool, & Shallice, 1994; Shallice, Glasspool, & Houghton, 1995) and in the domain of speech production (Hartley & Houghton, 1996). Unlike the letter/word recognition domain of McClelland and Rumelhart, this domain shares with the domains discussed here the need to produce a number of responses in series. Notions of selection and inhibition after selection are common to contention scheduling and these models. The models differ from the contention scheduling model in several respects, however. First, the Houghton et al. models include only limited hierarchical organisation. They have a number of levels (typically two, corresponding to words and letters), which are fixed by the model architecture. The action selection domain demands much greater hierarchical flexibility. Thus, although the schema network employed within the coffee preparation task has a fixed structure with three levels, this is a consequence of the task domain, not a limitation arising from any architectural mechanism. Cooper (1998) uses an alternate schema network with fewer levels to model a visual and auditory reaction time task, and alternate schema networks with more levels could in principle be developed for tasks with greater hierarchical complexity. Second, argument and resource selection are not required in the spelling domain. The mechanisms that effect these processes are necessary to account for error classes involving incorrect or inefficient argument or resource selection, and are central to the contention scheduling model. Third, in the contention scheduling model, serial ordering of subgoals within a schema is based on gating of activation flow by

symbolic preconditions. This contrasts with the use of a varying context signal to control serial order as employed in the spelling models (but see Hartley & Houghton, 1996). Finally, no attempt has been made within the contention scheduling model to address the question of learning within the action selection domain. Learning is an integral part of the Houghton et al. spelling models.

Of potentially even greater relevance to the current work is the model developed by Maes (1989). This is another interactive activation model of action selection. The origins of the model, however, are in Artificial Intelligence, rather than cognitive psychology. Nodes in the model correspond to various competences, and activation flows between the nodes according to standard interactive activation principles. Like the contention scheduling model, conflicting nodes inhibit each other, and symbolic preconditions govern sequential behaviour. However, in contrast to our model, nodes in the model of Maes actively excite their precondition nodes. In addition, there is no argument selection or hierarchical structuring on those nodes: Activation flows directly between all nodes, which correspond to base-level actions with bound arguments.

It should be clear from this discussion that there are substantial differences between the contention scheduling model presented here and other models in the literature. These arise primarily from requirements of the domain of action selection and the theoretical claim of hierarchical structuring of schemas within this domain.

Further Applications of the Model to Neuropsychological Phenomena

Earlier in the paper we considered a number of neurological disorders relevant to the control of intermediate domain action. Our discussion has thus far focused on just one of these, the action disorganisation syndrome investigated by Schwartz and colleagues. However, at least two of the other listed

disorders of intermediate domain action control—bradykinesia arising from Parkinson's disease and stereotypy arising from amphetamine psychosis—have been hypothesised to arise from inadequate operation of contention scheduling. We therefore appraise these hypotheses in the light of the model's behaviour following damage.

Parkinson's Disease

Robbins and Sahakian (1983) argue that activation of the striatal dopamine system corresponds to increased activation of schemas within a contention scheduling framework. As striatal dopamine is known to be deficient in patients with Parkinson's disease (Robbins, 1991), these arguments suggest that decreasing the parameter governing the ratio of self influence to lateral influence in the model's schema network (and thereby decreasing the activation of schemas) should result in behaviour similar to that shown by patients with Parkinson's disease.

One of the principal features of the behaviour of Parkinson's patients is bradykinesia: Willed initiation of action sequences is typically greatly slowed, but once an action sequence has been initiated it may proceed relatively normally. In some cases, once the action sequence has started further actions can be selected at roughly normal speed. Thus, Owen et al. (1992) found that mild and severe medicated Parkinson's patients were significantly slower to initiate action in the Tower of London task, but were not significantly slower in the time between the first movement and task completion.

To examine whether the model might show analogous behaviour in the coffee preparation domain the procedure for investigating the Internal:External, parameter was repeated with respect to the Self:Lateral, parameter⁴. Figure 10 shows the number of cycles to action onset and the number of cycles to task completion as the Self:Lateral, parameter varies from .00 to 1.00. When the parameter's value is below .27, schema activation profiles remain near their resting values and no actions are

⁴Although we assume that the Self:Lateral parameter is fixed across schema, object representation, and resource domains in normal functioning, there is no reason why this restriction should hold in the damaged system. Localised damage may affect one system only, resulting in the modification of a parameter's value in that domain, without affecting the parameter settings of other domains.

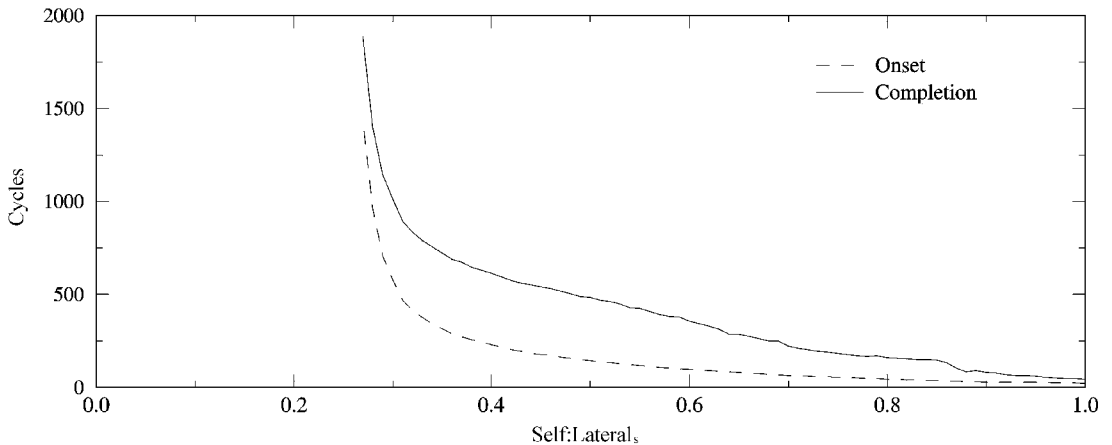


Fig. 10. Number of cycles to onset and to task completion as a function of the *Self:Lateral*_s parameter in the coffee preparation task.

selected. Successful behaviour is observed as the parameter ranges in value from .27 to .55. Within this range few errors are observed and the number of cycles to onset and to completion vary monotonically. At higher levels of the parameter (above .55), disorganisation of action appears. Within this region, the level of lateral influence is insufficient to ensure successful competition, and multiple competing schemas can simultaneously become highly active.

Of particular interest is behaviour in the “successful” range, .27 to .55. When the parameter’s value is near the high end of this range (e.g. .55) the mean number of cycles to onset is approximately 113. A further 390 cycles are required before the task is completed. However, when the value is near the low end of the range (e.g. .27) the mean number of cycles to onset rises 12-fold to 1439, but an equivalent increase is not seen in the mean time required for the remainder of the task, which rises only slightly to 445 cycles. Thus, action initiation is greatly slowed, but after initiation the slowing in action selection is far less pronounced. This behaviour is strongly suggestive of bradykinesia: It is difficult for the model to initiate action, but once started action occurs relatively normally. The behaviour therefore supports the hypothesis of Robbins and Sahakian (1983) concerning the relationship between striatal dopamine and schema excitation.

Bradykinesia arises in the model because at the beginning of the task there are several competing schemas with similar activation values. Under these conditions self influence plays an important role in separating competing schemas, and reduction in self influence leads to a general slowing in competition. However, once the initial competition has been resolved (and the first action has been performed), the dynamics of activation flow ensures that the average difference in activation of competing schemas is such that decreased levels of self influence have little effect. Any slowing of action selection after task onset is therefore slight.

The model presented previously was in fact modified to produce the graph in Fig. 10. The earlier model included an arbitrary ordering constraint, which ensured that coffee grinds were added to the water before milk and sugar. Figure 10 was obtained from a simplified model that excluded this ordering preference. The effect of ordering constraints on the model’s behaviour with reduced self influence is in fact highly significant. Without such constraints, bradykinesia arises. The effect disappears, however, when strong constraints (e.g. that coffee should be added to the boiling water before the other ingredients) are present. This is because ordering constraints effectively reduce the number of schemas competing at any one time. The model therefore predicts that bradykinesia will be significantly reduced when action is highly

constrained (and so when there is little competition between alternate actions).

Amphetamine Psychosis

A disorder that is in some ways opposite to bradykinesia arises with amphetamine psychosis. The effects of amphetamines on the control of action are well documented. In general, moderate doses lead to spatially diverse activity, but high doses lead to stereotyped movements that are spatially confined. Lyon and Robbins (1975) interpret these effects in terms of an increased rate of responding within a reduced number of response categories. At high dosages the rate of responding is very high, but responses are spatially confined because the number of response categories is very low. At moderate dosages, the number of response categories is sufficient to allow spatially diverse activity with a moderate increase in the rate of responding.

Several authors (notably Robbins, 1982, and Frith, 1992) have attempted to relate behaviours characteristic of amphetamine psychosis to the contention scheduling theory. Frith proposes that amphetamine affects inhibition within the schema network, such that (a) competing schemas are not sufficiently inhibited, leading to multiple active schemas, and (b) schema deselection does not lead to sufficient schema inhibition to temporarily stop the deselected schema from being immediately reselected.

The first of Frith's proposals relates directly to the Self:Lateral₁ parameter, and implies that amphetamine type behaviour should arise in the model when the parameter's value is greater than in the normally functioning system. This proposal receives further support from the neurophysiological effects of amphetamines, which include an increase in the concentration of dopamine and related neurotransmitters in the synaptic cleft (see, e.g. Robbins, 1982), and the relationship hypothesised earlier between the Self:Lateral₁ parameter and dopamine receptivity in Parkinson's disease.

Examination of the behaviour of the coffee preparation simulation when the Self:Lateral₁ parameter is high supports this position. Figure 11

shows the rate of behaviour (expressed in terms of number of actions per 100 cycles) as the value of the Self:Lateral₁ parameter varies from .00 to 1.00. The graph parallels Fig. 10. In the successful range (.27–.55) the rate of responding varies only slightly, increasing gradually as the value of Self:Lateral₁ parameter increases. This behaviour continues beyond the successful range, but the rate of responding increases dramatically when the value of Self:Lateral₁ parameter rises above .80, supporting the hypothesis that increased excitation within the schema network leads to increased rates of responding. Interestingly, this behaviour, unlike the Parkinsonian bradykinesia reported earlier, is not dependent upon the number of competing schemas. Although Fig. 11 was generated by the model without ordering constraints over subtasks, qualitatively equivalent behaviour is observed when ordering constraints are present. The model therefore predicts that the conditions that should reduce or prevent bradykinesia in Parkinson's patients (reducing the number of response options) will not affect response rates in amphetamine psychosis.

In its present form the model does not provide a clear answer to the question of whether the number of response categories decreases as the value of the Self:Lateral₁ parameter increases. The primary difficulty is that, even in normal functioning, few response categories are available in the coffee preparation task. However, a significant proportion of errors observed when the value of the Self:Lateral₁ is greater than .55 are of the perseverative type. This is suggestive of a tendency toward fewer response categories and stereotypy.

CONCLUSION

The problem of how to model the system controlling action selection and its integration with the rest of the cognitive system—in particular, the perceptual system—so as to represent the way that a complex multi-level activity can be successfully realised in a multi-object environment has not to our knowledge been previously tackled using a computational approach. Two key requirements in the domain are that there be multiple alternatives

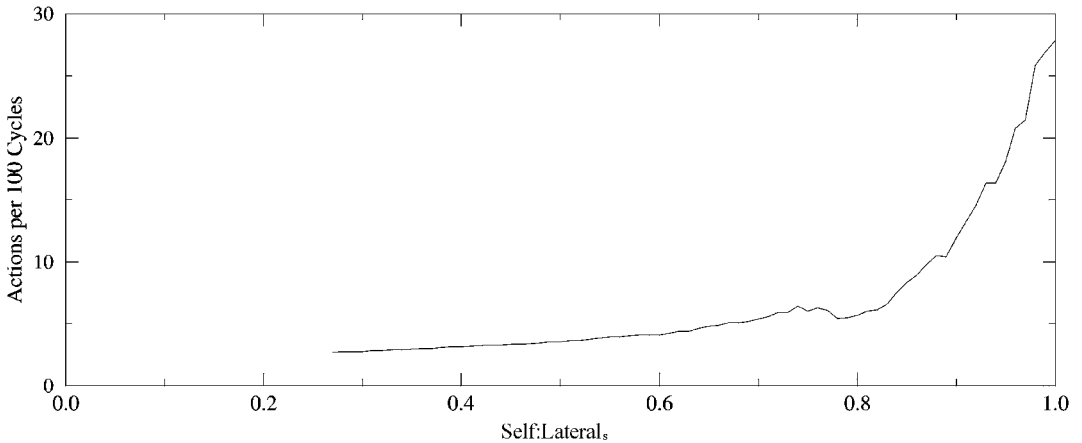


Fig. 11. Rate of responding (actions per 100 cycles) as a function of the *Self:Lateral_s* parameter in the coffee preparation task.

available to be selected in each of a number of subdomains (schemas, effectors, and arguments) and that these selections should interact appropriately. Moreover, the model needs to represent qualitative and quantitative changes in the world—that a sugar packet, for instance, can become empty.

The complexity of these requirements suggests that an initial model will inevitably be a considerable simplification of the processes it addresses. Indeed, there are a variety of areas where the current simulations are too rigid. For instance, there is no distinction between subgoals within a schema that are optional or that the environment may make unnecessary and ones that are always critical. In addition, there is no representation of the abstraction hierarchy orthogonal to the processing hierarchy through which “to add coffee from the coffee packet” and “to add sugar from the sugar packet” are both concrete realisations of “to add contents of condiment packet to target container.” To represent this will be of great value for modelling error correction. Further, there is no representation of objects coming into, being created in, or leaving the environment. We see these and similar computational requirements as being achievable without fundamental changes to the model. However, they remain to be carried out. Other developments such as the addition of error monitoring and correction functions will, however, on our approach, require implementing supervisory system functions.

Such functions are beyond the scope of contention scheduling.

Nevertheless, the simulations described in the present paper indicate that a model based on interactive activation principles and operating as a realisation of the contention scheduling component of the Norman and Shallice theory is a plausible candidate for more detailed development. The simulations show that the triple requirements of selecting sequentially appropriate action schemas, effectors, and arguments can be satisfactorily accomplished in a model task environment closely related to one characteristic of normal human routine action selection, and one which has indeed been employed with neurological patients. When parameters of the model are varied from the region in which task goals are satisfactorily achieved, the errors that the system makes are qualitatively similar to those produced by certain neurological patients. One may also conceive of the action lapses of normal participants as occurring when parameters of the system deviate from this region, due to dual task performance, stress, fatigue, and so on. The simulations further point to serious difficulties in how to score disorganised and errorful behaviour. In particular, two existing published scoring procedures, when applied to the same behaviour produced by the model, were found to yield very different characterisations of that behaviour.

Manuscript received 11 March 1998

Revised manuscript received 29 October 1998

Revised manuscript accepted 30 December 1998

REFERENCES

- Arbib, M.A. (1985). Schemas for the temporal organisation of behaviour. *Human Neurobiology*, 4, 63–72.
- Brazzelli, M., Colombo, N., Della Sala, S., & Spinnler, H. (1994). Sparing and impaired cognitive abilities after bilateral frontal damage. *Cortex*, 30(1), 27–51.
- Burgess, N., & Hitch, G. J. (1992). Toward a network model of the articulatory loop. *Journal of Memory and Language*, 31(4), 429–460.
- Charniak, E., & McDermott, D. (1985). *Introduction to Artificial Intelligence*. Reading, MA: Addison-Wesley.
- Cooper, R. (1998). Visual dominance and the control of action. In M.A. Gernsbacher, & S.J. Derry, (Eds.), *Proceedings of the 20th Annual Conference of the Cognitive Science Society* (pp. 250–255). Madison, WI: Cognitive Science Society.
- Cooper, R., Fox, J., Farrington, J., & Shallice, T. (1996). A systematic methodology for cognitive modelling. *Artificial Intelligence*, 85, 3–44.
- Cooper, R., & Shallice, T. (1997). Modelling the selection of routine action: Exploring the criticality of parameter values. In P. Langley, & M. Shafto (Eds.), *Proceedings of the 19th Annual Conference of the Cognitive Science Society* (pp. 131–136). Madison, WI: Cognitive Science Society.
- Cooper, R., Shallice, T., & Farrington, J. (1995). Symbolic and continuous processes in the automatic selection of actions. In J. Hallam (Ed.), *Hybrid problems, hybrid solutions* (pp. 27–37). Amsterdam: IOS Press.
- De Renzi, E., & Barbieri, C. (1992). The incidence of grasp reflex following hemispheric lesion and its relation to frontal damage. *Brain*, 25, 293–313.
- De Renzi, E., & Lucchelli, F. (1988). Ideational apraxia. *Brain*, 26, 1173–1188.
- Della Sala, S., Marchetti, C., & Spinnler, H. (1991). Right-sided anarchic (alien) hand: A longitudinal study. *Neuropsychologia*, 29(11), 1113–1127.
- Duncan, J. (1986). Disorganisation of behaviour after frontal lobe damage. *Cognitive Neuropsychology*, 3, 273–290.
- Elman, J.L. (1990). Finding structure in time. *Cognitive Science*, 14(2), 179–211.
- Elman, J.L. (1993). Learning and development in neural networks: The importance of starting small. *Cognition*, 48, 71–99.
- Frith, C.D. (1992). *The cognitive neuropsychology of schizophrenia*. Hove, UK: Psychology Press.
- Glasspool, D.W. (1998). *Modelling serial order in behaviour: Studies of spelling*. PhD thesis, Department of Psychology, University College London, UK.
- Goldberg, G., Mayer, N.H., & Togliola, J.U. (1981). Medial frontal cortex infarction and the alien hand sign. *Archives of Neurology*, 38, 683–686.
- Hartley, T., & Houghton, G. (1996). A linguistically constrained model of short-term memory for non-words. *Journal of Memory and Language*, 35, 1–31.
- Hashimoto, R., Yoshida, M., & Tanaka, Y. (1995). Utilization behaviour after right thalamic infarction. *European Neurology*, 35, 58–62.
- Head, H. (1926). *Aphasia and kindred disorders of speech*. Cambridge: Cambridge University Press.
- Hinton, G.E., & Shallice, T. (1991). Lesioning an attractor network: Investigations of acquired dyslexia. *Psychological Review*, 22(1), 74–95.
- Houghton, G. (1990). The problem of serial order: A neural network model of sequence learning and recall. In R. Dale, C. Mellish, & M. Zock, (Eds.), *Current research in natural language generation* (Ch. 11, pp. 287–319). London: Academic Press.
- Houghton, G., Glasspool, D. W., & Shallice, T. (1994). Spelling and serial recall: Insights from a competitive queueing model. In G.D.A. Brown, & N.C. Ellis, (Eds.), *Handbook of spelling: Theory, process and intervention* (pp. 365–404). Chichester, UK: John Wiley.
- Humphreys, G.W., & Forde, E.M.E. (1998). Disordered action schemas and action disorganisation syndrome. *Cognitive Neuropsychology*, 15, 771–811.
- Iberall, T., & Arbib, M.A. (1990). Schemas for the control of hand movements. In M.A. Goodale, (Ed.), *Vision and action: The control of grasping* (pp. 204–242). New York: Ablex Publishing Corporation.
- Jeannerod, M., Arbib, M.A., Rizzolatti, G., & Sakata, H. (1995). Grasping objects: The cortical mechanisms of visuomotor transformation. *Trends in Neurosciences*, 18(7), 314–320.
- Jordan, M.I. (1986). *Serial order: A parallel distributed processing approach*. Tech. rep. 8604, Institute for Cognitive Science, University of California, San Diego, CA.
- Katz, J., & Fodor, J. A. (1963). The structure of a semantic theory. *Language*, 39, 170–210.

- Lakatos, I. (1970). Falsification and the methodology of scientific research programmes. In I. Lakatos, & A. Musgrave, (Eds.), *Criticism and the growth of knowledge* (pp. 91–196). Cambridge: Cambridge University Press.
- Lhermitte, F. (1983). Utilisation behaviour and its relation to lesions of the frontal lobes. *Brain*, *106*, 237–255.
- Lyons, M., & Robbins, T.W. (1975). The action of central nervous system stimulant drugs: A general theory concerning amphetamine effects. In *Current Developments in Psychopharmacology* (pp. 80–163). New York: Spectrum Publications.
- Maes, P. (1989). How to do the right thing. *Connection Science*, *1*(3), 291–323.
- 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.
- Martin, N., Dell, G.S., Saffran, E.M., & Schwartz, M.F. (1994). Origins of paraphasias in deep dysphasia: Testing the consequences of a decay impairment to an interactive spreading activation model of lexical retrieval. *Brain and Language*, *47*, 609–660.
- 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.
- McLeod, P. (1977). A dual task response modality effect: Support for multiprocessor models of attention. *Quarterly Journal of Experimental Psychology*, *29*(4), 651–667.
- McLeod, P., & Dienes, Z. (1996). Do fielders know where to go to catch the ball or only how to get there? *Journal of Experimental Psychology: Human Perception and Performance*, *22*, 531–543.
- McLeod, P., & Posner, M.I. (1984). Privileged loops from percept to act. In H. Bouma, & D.G. Bouwhuis, (Eds.), *Attention and performance, X* (pp. 55–66). Hove, UK: Lawrence Erlbaum Associates Ltd.
- Norman, D.A. (1981). Categorisation of action slips. *Psychological Review*, *88*(1), 1–15.
- Norman, D.A., & Shallice, T. (1980). *Attention to action: Willed and automatic control of behaviour*. Chip report 99, University of California, San Diego, CA.
- Norman, D.A., & Shallice, T. (1986). Attention to action: Willed and automatic control of behaviour. In R. Davidson, G. Schwartz, & D. Shapiro, (Eds.), *Consciousness and self regulation: Advances in research and theory, Vol. 4* (pp. 1–18). New York: Plenum.
- Owen, A.M., James, M., Leigh, P.N., Summers, B.A., Marsden, C.D., Quinn, N.P., Lange, K.W., & Robbins, T.W. (1992). Fronto-striatal cognitive deficits at different stages of Parkinson's Disease. *Brain*, *115*, 1727–1751.
- Pashler, H. (1994). Dual-task interference and elementary mental mechanisms. In D.E. Meyer, & S. Kornblum, (Eds.), *Attention and performance, XIV* (pp. 245–264). Cambridge, MA: MIT Press.
- Poeck, K., & Lehmkuhl, G. (1980). Das Syndrom der ideatorischen Apraxie und seine Localisation. *Nervenartz*, *51*, 217–225.
- Reason, J.T. (1979). Actions not as planned: The price of automatization. In G. Underwood, & R. Stevens, (Eds.), *Aspects of consciousness* (Ch. 4, pp. 67–89). London: Academic Press.
- Reason, J.T. (1984). Lapses of attention in everyday life. In W. Parasuraman, & R. Davies, (Eds.), *Varieties of attention* (Ch. 14, pp. 515–549). Orlando, FL: Academic Press.
- Reason, J.T. (1990). *Human error*. Cambridge: Cambridge University Press.
- Robbins, T.W. (1982). Stereotypies: Addictions or fragmented actions? *Bulletin of The British Psychological Society*, *35*, 297–300.
- Robbins, T.W. (1991). Cognitive deficits in schizophrenia and parkinson's disease: Neural basis and the role of dopamine. In P. Willner, & J. Scheel-Krüger, (Eds.), *The mesolimbic dopamine system: From motivation to action* (pp. 497–528). New York: John Wiley & Sons.
- Robbins, T.W., & Sahakian, B.J. (1983). Behavioural effects of psychomotor stimulant drugs: Clinical and neuropsychological implications. In I. Creese, (Ed.), *Stimulants: Neurochemical, behavioural, and clinical perspectives* (pp. 301–338). New York: Raven Press.
- Rosenbaum, D.A., Vaughan, J., Barnes, H.J., Marchak, F., & Slotta, J. (1990). Constraints on action selection: Overhand versus underhand grips. In M. Jeannerod, (Ed.), *Attention and performance, XIII* (pp. 321–342). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Rumelhart, D.E., & Norman, D.A. (1982). Simulating a skilled typist: A study of skilled cognitive-motor performance. *Cognitive Science*, *6*, 1–36.
- Sandson, J., & Albert, M.L. (1984). Varieties of perseveration. *Neuropsychologia*, *6*, 715–732.
- Schank, R.C. (1982). *Dynamic Memory: A theory of reminding and learning in computers and people*. Cambridge: Cambridge University Press.

Schank, R.C., & Abelson, R. (1977). *Scripts, plans, goals and understanding*. Hove, UK: Lawrence Erlbaum Associates Ltd.

Schmidt, R.A. (1975). A schema theory of discrete motor skill learning. *Psychological Review*, 82(4), 225–260.

Schwartz, M.F., Montgomery, M.W., Buxbaum, L.J., Less, S.S., Carew, T.G., Coslett, H.B., Ferraro, M., Fitzpatrick-De Salme, E.J., Hart, T., & Mayer, N.H. (1998). Naturalistic action impairment in Closed Head Injury. *Neuropsychology*, 12(1), 13–28.

Schwartz, M.F., Montgomery, M.W., Fitzpatrick-De Salme, E.J., Ochipa, C., Coslett, H.B., & Mayer, N.H. (1995). Analysis of a disorder of everyday action. *Cognitive Neuropsychology*, 12(8), 863–892.

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. (1972). Dual functions of consciousness. *Psychological Review*, 79(5), 383–393.

Shallice, T. (1982). Specific impairments of planning. *Philosophical Transactions of the Royal Society of London*, B298, 199–209.

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

Shallice, T., & Burgess, P. (1996). The domain of supervisory processes and temporal organisation of behaviour. *Philosophical Transactions of the Royal Society of London*, B351, 1405–1412.

Shallice, T., Burgess, P.W., Schon, F., & Baxter, D.M. (1989). The origins of utilisation behaviour. *Brain*, 112, 1587–1598.

Shallice, T., Glasspool, D.W., & Houghton, G. (1995). Can neuropsychological evidence inform connectionist modelling? Analyses of spelling. *Language and Cognitive Processes*, 10, 195–225.

Tipper, S.P. (1985). The negative priming effect: Inhibitory priming by ignored objects. *Quarterly Journal of Experimental Psychology*, 37A, 571–590.

Tipper, S.P., Weaver, B., & Houghton, G. (1994). Behavioural goals determine inhibitory mechanisms of selective attention. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology*, 47A(4), 809–840.

Tipper, S.P., Weaver, B., Kirkpatrick, J., & Lewis, S. (1991). Inhibitory mechanisms of attention: Locus, stability, and relationship with distractor interference effects. *British Journal of Psychology*, 82(4), 507–520.

Viviani, P. (1990). Common factors in the control of free and constrained movements. In M. Jeannerod, (Ed.), *Attention and performance, XIII* (pp. 345–373). Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Wolpert, D.M., & Ghahramani, Z. (1997). Modular decomposition in visuomotor learning. *Nature*, 386, 392–395.

APPENDIX A

Details of Activation Flow

This appendix details the mathematics of activation flow within and between the various networks. Nodes in each network have up to four inputs: an internal influence, an external influence, a self influence, and a lateral influence. These are combined in a weighted sum, with the weights determined by parameters, to determine the net input to a node. An activation update function then determines how a node's net input affects its activation.

The Schema Network

The internal (i.e. top-down) influence I_s on a schema node is:

$$I_s = \begin{cases} 1 & \text{if the goal of } s \text{ is directly triggered by the} \\ & \text{supervisory system} \\ \frac{1}{N} \cdot A_{source} & \text{if the source schema of } s \text{ is selected and all} \\ & \text{preconditions of the goal of } s \text{ are achieved} \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

In equation (1) A_{source} is the activation of the source schema of s and N is the number of subgoals of that schema.

The external (i.e. environmental) influence E_s on a schema s is:

$$E_s = \text{squash} \left(\frac{\sum_{t: t \rightarrow s} \sum_{o \in t} A_o - A_{rest}}{|t|} \right) \quad (2)$$

In equation (2) $t \rightarrow s$ holds if the situation t triggers the schema s , $o \in t$ holds if o is an object or effector involved in the situation t , A_o is the activation of the object or effector o , A_{rest} is rest activation, and $|t|$ is the total number of objects and effectors involved in the situation t .

Equation (2) is obtained by the following procedure. The influence of each situation t on s is calculated. This is given by the average influence per object in t on s (This explains the central summation and division by $|t|$.) In the summation we employ a non-zero rest activation so that objects whose activation is below this rest activation can inhibit schemas. The average is used so that the influence of a triggering situation is independent of the number of objects in that situation. The situational influences are summed over all situations that could trigger a schema, and the result “squashed” so that it lies between -1 and $+1$. The squashing function used is:

$$squash(x) = \frac{e^x - 1}{e^x + 1} \tag{3}$$

The use of such a function prevents multiple situations from dominating the triggering of any schema: Two triggering situations have less than twice the effect of one triggering situation. It also ensures that triggering input is bounded.

The lateral influence LI_s on a schema s is:

$$LI_s = - \left(\sum_{c: c \text{ competes with } s} A_c - A_{rest} \right) \tag{4}$$

In equation (4) A_c is the activation of the competing schema c .

The self influence SI_s on a schema s is:

$$SI_s = \begin{cases} +A_s & \text{if } s \text{ is not selected or} \\ & \text{if } s \text{ has an unachieved subgoal} \\ -A_s & \text{otherwise} \end{cases} \tag{5}$$

In equation (5) A_s is the activation of the schema s .

The Object Representation Network

There is no internal influence on object representation nodes. The external influence (from schema nodes) on an object representation node o with respect to function s is:

$$E_o^f = squash \left(\sum_{s: o^f \text{ triggers } s} A_s - A_{rest} \right) \tag{6}$$

In equation (6) o^f triggers s holds if the object o serving function f triggers the schema s and A_s is the activation of schema s . Once again, the squashing function is used to ensure that the influence is bounded.

The lateral influence LI_o^f on an object representation o with respect to function f is:

$$LI_o^f = - \left(\sum_{c: c \text{ competes with } o} A_c^f - A_{rest} \right) \tag{7}$$

In equation (7) A_c^f is the activation of the competitor c for the function f .

The self influence SI_o^f on an object representation o with respect to function f is:

$$SI_o^f = A_o^f \tag{8}$$

In equation (8) A_o^f is the activation of the object representation o for the function f .

The Resource Network

There is no internal influence on resource nodes. Following equation (6), the external influence (from schema nodes) on a resource node r is:

$$E_r = squash \left(\sum_{s: r \text{ triggers } s} A_s - A_{rest} \right) \tag{9}$$

In equation (9) r triggers s holds if the schema s requires the resource r and A_s is the activation of schema s .

As discussed in the main text, competition within the resource domain has not been included. There is therefore no self influence or lateral influence on resource nodes.

Parameters

Activation flow within and between the schema, object representation, and resource networks is controlled by three principal “balance” parameters: Self:Lateral, Internal:External, and Competitive:Noncompetitive. The roles of these parameters are as follows:

- If SI and LI are, respectively, the self influence and lateral influence on a node, then the total competitive input to that node is:

$$C = \text{Self:Lateral} \times SI + (1 - \text{Self:Lateral}) \times LI \tag{10}$$

- If E and I are, respectively, the net external and internal influences on a node, then the total noncompetitive input to that node is:

$$NC = \text{Internal:External} \times I + (1 - \text{Internal:External}) \times E \tag{11}$$

- Given the above, the total input to a node is:

$$T = \text{Competitive:Noncompetitive} \times C + (1 - \text{Competitive:Noncompetitive}) \times NC \tag{12}$$

All activation sources yield excitation or inhibition in a restricted range. The use of proportions as indicated here ensures that the total influence remains in the same range, irrespective of the values of the balance parameters. This means that the balance parameters can be varied independently of those parameters that specify higher-level network characteristics (such as smoothness of activation profiles). This greatly simplifies the processes of determining appropriate parameter values and exploring the parameter space (cf. Cooper & Shallice, 1997).

Updating Activations

For each node the new activation is calculated from the current activation and the current input to that node (given by equations 1–12) on each processing cycle. The function underlying this calculation must have certain properties, but its precise form is not critical. The activation function employed in the simulations reported here is:

$$A_{t+1} = \bar{\vartheta} (P \cdot \bar{\vartheta}^{-1}(A_t) + IN_t) \tag{13}$$

or equivalently:

$$A_{t+1} = \bar{\vartheta} \left(\sum_{i=0}^{\infty} P^i \cdot IN_{t-i} \right) \tag{14}$$

In equations (13) and (14) P is persistence (a parameter between 0 and 1), A_t is the activation of the node on cycle t , IN_t is the net input to the node (i.e. T from equation 12 above, supplemented with normally distributed random noise), and $\bar{\vartheta}$ is the hyperbolic tangent function translated and scaled such that:

$$\bar{\vartheta}(+\infty) = 1$$

$$\bar{\vartheta}(0) = A_{rest}$$

$$\bar{\vartheta}(-\infty) = 0$$

This function is bounded by 0 and 1 and with zero net input tends to A_{rest} . The function has two critical features:

- It contains a persistence (i.e. momentum or decay) factor: The new activation of a node is calculated with respect to the node's current activation (or some fraction of it), rather than with respect to zero (or rest) activation. Persistence is necessary in order to maintain smoothly varying activation profiles. Activation functions without such a factor generally result in unstable activation profiles that oscillate between extreme values.

- With zero net input, it tends to a non-zero rest value. With negative input, nodes may be inhibited below this value. A non-zero rest activation was originally introduced to counter the effect of persistence at low levels of activation (which acted to prevent activations from rising once they fell to rest). As should be clear from the main text, rest activation has come to play a significant role in many aspects of the system.

With appropriate values for the parameters for A_{rest} and P (e.g. 0.10 and 0.80 respectively), equation (13) leads to a system in which activation profiles are smooth and bounded, and competition is robust.