A Computational Model of Prefrontal Control in Free Recall: Strategic Memory Use in the California Verbal Learning Task

Suzanna Becker and Jean Lim

Abstract

■ Several decades of research into the function of the frontal lobes in brain-damaged patients, and more recently in intact individuals using function brain imaging, has delineated the complex executive functions of the frontal cortex. And yet, the mechanisms by which the brain achieves these functions remain poorly understood. Here, we present a computational model of the role of the prefrontal cortex (PFC) in controlled memory use that may help to shed light on the mechanisms underlying one aspect of frontal control: the development and deployment of recall strategies. The model accounts for interactions between the PFC and medial temporal lobe in strategic memory use. The PFC self-organizes its own mnemonic codes using internally derived performance measures. These mnemonic codes serve as retrieval cues by biasing retrieval in the medial temporal lobe memory system. We present data from three simulation experiments that demonstrate strategic encoding and retrieval in the free recall of categorized lists of words. Experiment 1 compares the performance of the model with two control networks to evaluate the contribution of various components of the model. Experiment 2 compares the performance of normal and frontally lesioned models to data from several studies using frontally intact and frontally lesioned individuals, as well as normal, healthy individuals under conditions of divided attention. Experiment 3 compares the model's performance on the recall of blocked and unblocked categorized lists of words to data from Stuss et al. (1994) for individuals with control and frontal lobe lesions. Overall, our model captures a number of aspects of human performance on free recall tasks: an increase in total words recalled and in semantic clustering scores across trials, superiority on blocked lists of related items compared to unblocked lists of related items, and similar patterns of performance across trials in the normal and frontally lesioned models, with poorer overall performance of the lesioned models on all measures. The model also has a number of shortcomings, in light of which we suggest extensions to the model that would enable more sophisticated forms of strategic control. ■

INTRODUCTION

People have a remarkable ability to encode and retrieve information in a flexible manner. Understanding the neuronal mechanisms underlying strategic memory use remains a true challenge. Neural network models of memory have typically dealt with only the most basic operations involved in storage and recall. The goal of this work is to develop a computational model that will shed light on the neural mechanisms underlying strategic memory use in individuals with intact and lesioned frontal lobes. In the simulations reported here, we focus on a particular task, the free recall of lists of related words.

Evidence from patients with frontal damage indicates a crucial role for the prefrontal cortex (PFC) in the control of memory. The task that characterizes best the memory impairment of patients with frontal lobe damage is word list learning, in which free recall is impaired disproportionately relative to recognition performance (Stuss et al., 1994). Patients with frontal damage are impaired particularly on tasks requiring context-specific recall such as AB–AC list learning (Shimamura, Jurica, Mangels, Gershberg, & Knight, 1995), on tests of memory for temporal order (Shimamura, Janowsky, & Squire, 1990), and in tasks in which normal subjects would tend to make use of self-generated recall strategies. For example, when the task is to memorize a long list of words, people normally benefit from the categorical structure of the list, recalling more items from related lists than from unrelated lists (e.g., Bousfield, 1953). A key performance index in these tasks is categorical clustering—the tendency to recall items from the same taxonomic category consecutively at greater than chance levels. The California Verbal Learning Test (CVLT) (Delis, Kramer, Kaplan, & Ober, 1987) is a standardized version of this task, in which 16 words drawn from four categories are presented aurally in a mixed order for five repetitions with recall tested after each list presentation. On CVLT-like tests, whereas frontal patients do benefit from the categorical structure

McMaster University

© 2003 Massachusetts Institute of Technology

of word lists (Stuss et al., 1994; Jetter, Poser, Freeman, & Markowitsch, 1986), they tend to recall fewer categories in total, especially when tested after long delays, and in some cases they have lower semantic clustering scores (Hildebrandt, Brand, & Sachsenheimer, 1998; Jetter et al., 1986). The latter effect may depend on how obvious the categorical structure is (Moscovitch, 1994). Individuals with frontal lesions also recall fewer words in total, exhibit lower serial clustering scores, show less consistency in words recalled across trials, and make more intralist repetitions and intrusions from previously studied lists (Hildebrandt et al., 1998; Stuss et al., 1994). Healthy individuals performing a concurrent interfering task during both study and recall, a manipulation thought to disrupt frontal control functions, also recall fewer items and exhibit less categorical clustering (Moscovitch, 1994). These data support the view that the PFC is required to employ temporal context, organize information, and form retrieval strategies dynamically.

We propose that one of the functions of the PFC is to develop mnemonic codes rapidly according to task demands, driven by a reinforcement learning process. The reinforcement is derived from self-monitoring of performance relative to one's current motivational state and goals. These self-organized mnemonic codes can act as selective retrieval cues to the medial temporal lobe memory system.

The postulated role of the PFC in rapid self-organization of new internal codes is broadly consistent with Duncan's (2001) view of the PFC as a dynamic learning system, and Frith's (2000) view of the dorsolateral PFC as ''sculpting the response space.'' According to Frith, the dorsolateral PFC is responsible for the selection of appropriate actions when there are many possible alternative responses, rather than when an automatic response would be made. It is involved in forming arbitrary new categories of items, by enhancing some attributes while inhibiting others.

We have developed a connectionist model of the performance of intact and frontally lesioned individuals on free recall tasks that embodies the above assumptions. Our simulations show that a form of strategic control of memory emerges in a relatively simple neural network capable of rapid and highly flexible, reinforcement-driven learning. We have attempted to construct the simplest possible model that could exhibit strategic recall. Therefore, rather than building in fixed strategies, predefined retrieval cues, or component processes such as working memory, our model self-organizes its internal representations in the prefrontal module. They may persist over time, thereby providing a constraining context for retrieval.

Our model of controlled memory use combines a lexical/semantic memory (LSM) module, a medial temporal lobe memory (MTL) module, and a prefrontal cortical (PFC) module, as shown in Figure 1. We briefly describe the operation of our model here; details are provided in the Methods section. Activity flows bottom-up through the model during the processing of externally presented stimuli (Figure 1A) and top-down during memory retrieval and response generation (Figure 1B). When an externally presented word is processed by the model, the item activates a set of lexical and semantic features in the LSM module, which are associated with the current context into an episodic trace in the MTL module, and also encoded in the PFC module (see Figure 1). The context signal is a slowly varying random vector, representing the current internal state of an individual. The idea of a random or gradually evolving context vector has appeared in many previous memory models (e.g., Howard & Kahana, 2001; Brown, Preece, & Hulme, 2000; Murdock, 1997; Burgess & Hitch, 1992, 1996). During free recall, the previously recalled item activates the PFC, which cues the MTL module to recall an episodic trace, and activates a candidate word response in the LSM module. The model uses a ''generate-and-test'' strategy: The recency of the recalled item is evaluated by probing the MTL module with the recalled item combined with the current context. If the recency is too high, the item is assumed to be a repetition error, whereas if the recency is too low, an intrusion error is assumed. Otherwise, a response is generated and a new episodic trace of the recalled word in the current context is stored in the MTL module.

The key to our model's ability to perform strategic recall is in its learning mechanism. Whereas the LSM and MTL modules are trained via simple Hebbian learning, the PFC representation is learned rapidly according to task demands, based on self-monitoring of the model's memory-retrieval performance. The PFC module receives a reward signal when nonrepeated study list items are retrieved, and a punishment signal when nonlist or repeated items are retrieved. Occasionally, a repetition or intrusion error may go undetected by the model, resulting in a recall error. The model thereby learns to develop retrieval strategies dynamically in the course of both study and free recall of words.

Two processing constraints were added to the model to facilitate strategic learning in free recall. The first was motivated by primate PFC unit recording studies indicating sustained activations correlated with task-related working memory and response preparation (e.g., Constantinidis, Franowicz, & Goldman-Rakic, 2001). Thus, the PFC units in our model each had a ''fast bias'' that implemented a sustained response switch gated by reinforcement. A unit could thereby increase its bias over trials when repeatedly reinforced, and quickly reverse to a negative bias when a retrieval error is detected. Second, a response-suppression mechanism was included to prevent the same lexical unit being immediately recalled again. This was motivated by evidence for inhibition of return in visuospatial search tasks (Posner, Rafal, Choate, & Vaughan, 1985) and response suppression in serial recall tasks (e.g., Lewandowsky &

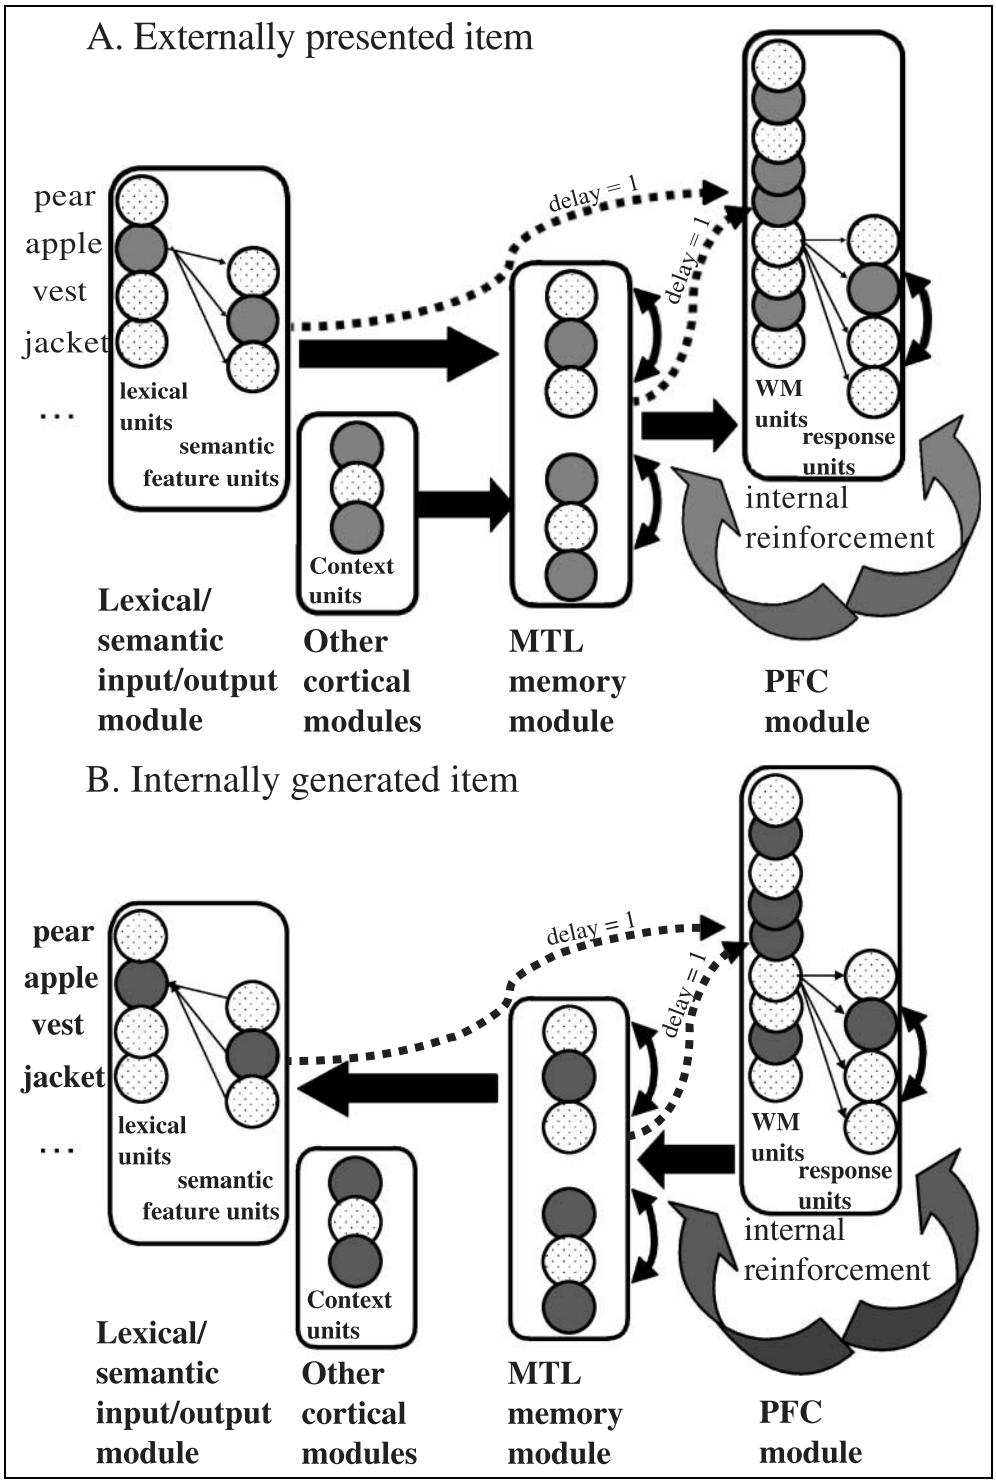

Figure 1. The architecture of the neural network model. The network used here had 800 LSM units, 300 context units, 1600 MTL units, and 10 PFC units (labeled "response units" in the figure). The PFC units received as input delayed copies of the LSM and MTL outputs (labeled "WM units" in the figure). Filled arrows indicate a reinforcement signal, which modulates the learning in all connections to, from, and within the PFC module. Arrows labeled "delay = 1" indicate pathways along which information from the previous time step is transmitted. (A) Bottom-up flow of activation during perception of an external stimulus. (B) The top-down flow of activation when a response is internally generated.

Duncan, 2002; Vousden & Brown, 1998). Response suppression has been incorporated in nearly all models of serial recall and serial response generation (e.g., Farrell & Lewandowsky, 2002; Brown et al., 2000; Vousden, Brown, & Harley, 2000; Burgess & Hitch, 1999; Henson, 1998; Page & Norris, 1998; Houghton & Hartley, 1996; Lewandowsky & Murdock, 1989). The fast bias and response-suppression mechanisms allowed the model to develop temporally persistent mnemonic cues in the PFC layer, while permitting a range of different responses to the same mnemonic cue, without response perseveration.

In three simulation experiments the model was tested on CVLT-like tasks with lists of 16 words for five study-recall trials. In Experiment 1, the model's correct recall and error scores are compared to that of two control networks, one lacking the response suppression, and the other lacking the fast bias for the PFC units. Experiment 2 examines the effect of frontal lesions on the

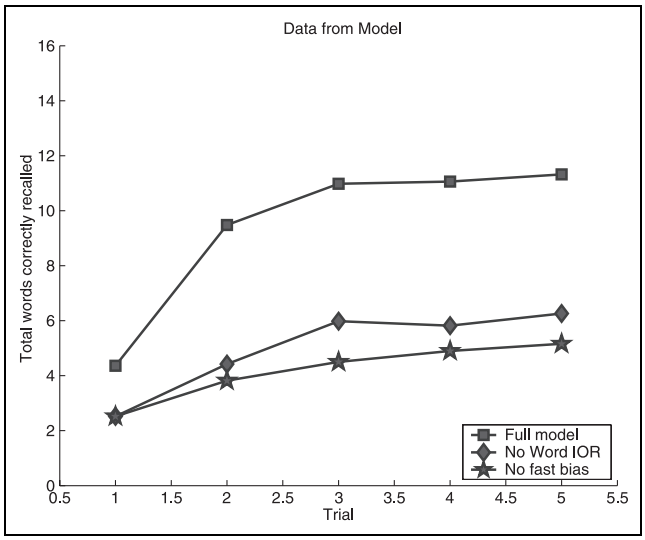

Figure 2. Number of total words correctly recalled by the full model and two control networks across five recall trials in a CVLT-like task, averaged over 50 runs. One control network, labeled "no word IOR," lacked response suppression/inhibition of return, and the other, labeled "no fast bias," lacked the rapidly learned, reinforcement-driven bias.

model; correct recall and semantic clustering scores are compared to human data from three experiments: Moscovitch's (1994) CVLT data for healthy subjects in control and divided attention conditions, and data from two studies of frontally lesioned individuals (Hildebrandt et al., 1998; Stuss et al., 1994). Experiment 3 extends these results by examining the model's performance on blocked and unblocked categorized lists of words, as well as lists of unrelated words. The results are compared to those of Stuss et al. (1994) for humans with control and frontal lobe lesions.

EXPERIMENT 1: RESULTS AND DISCUSSION

As shown in Figure 2, for all networks, recall increased over trials, with the greatest amount of learning in the first two trials. The full model learned to recall on average 11 to 12 of the 16 words by the fifth trial. However, the networks lacking either the fast bias in the PFC layer or the response suppression recalled considerably fewer words and had shallower learning curves. A one-way analysis of variance (ANOVA) revealed a significant main effect of the network architecture on total words recalled correctly, $F(2,147) = 352.4$, $p < .001$. All post hoc pairwise comparisons between correct scores for the three network types using Bonferroni-corrected t tests were significant ($p < .001$).

Figure 3 shows average repetition and intrusion error rates summed across the five trials. Without response suppression, the model makes many more repetition errors—both detected (hits) and undetected (misses). Intrusion error rates (bottom, Figure 3), relative to the total number of words recalled correctly (top, Figure 3), were similar across networks.

Sample responses generated by the full network in a typical recall session after five study sessions are shown in Table 1. In this case, 12 words were recalled, and 7 of

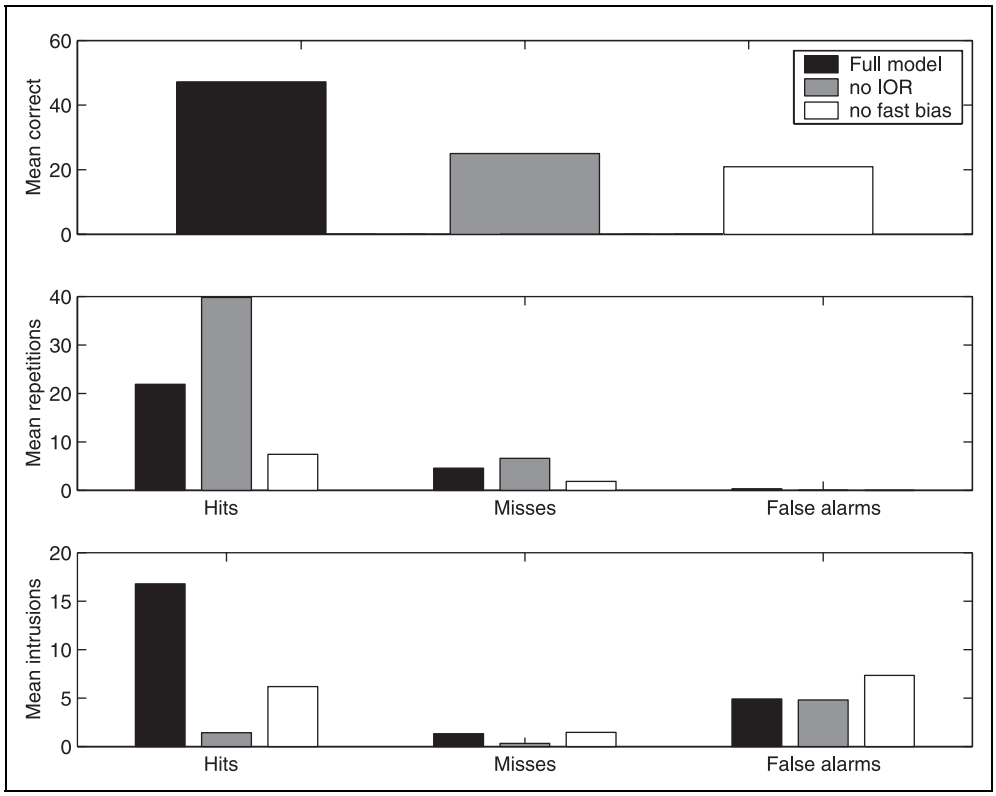

Figure 3. Number of total words correctly recalled, repetition errors, and intrusion errors made by the full model and two control models, one lacking response suppression/inhibition of return (labeled "no IOR") and one lacking the rapidly learned, reinforcement-driven bias (labeled "no fast bias") averaged over 50 runs.

Table 1. Responses Produced by the Model during a Typical Recall Session

Word Recalled	*Most Active PFC Unit*	*Recency in MTL*	*Intrusion Error*	*Repetition Error*	*Repeated Category*
Jacket	5	36.5			
Sweater	5	31.2			*
Slacks	5	41.8			*
Apricots	10	34.5			
Tangerines	10	30.9			*
Grapes	9	27.4			*
Plums	9	24.7			*
Pliers	2	30.9			
Wrench	2	28.2			*
Chisel	2	31.2			*
Chives	7	22.6			
Grapes	1	33.6		*	
XXXX	3	0.9	*		

those constituted categorical repetitions, yielding a clustering score of 7 for that session. Note that sustained responses by an individual PFC unit are associated with categorical clustering. The recency of each recalled word, a function of the match between the recalled word and the stored weights in the MTL module, tends to decrease with each successively recalled word, as the time since the word was studied increases. In this case the network generates one undetected repetition error, ''grapes,'' whose recency is not sufficiently higher than the last few words recalled so that it would be flagged as an error. The very last word recalled in a session is always a detected error, and terminates the trial when the model cannot generate any valid items. In this example, the final word recalled is an intrusion error, detected by its extremely low recency.

Neither control network showed this pattern in the PFC layer. The control network lacking the fast bias in the PFC layer sometimes would alternate amongst activating two or more different PFC units; it failed to develop the strategy of a sustained response in the PFC layer over several recalls that would permit clustered recall. On other occasions, this same network activated the same PFC unit for the duration of the recall session, failing to develop the strategy of switching PFC units after detecting an error. The control network lacking response suppression made many detected repetition errors (see Figure 3); each time such an error was detected, the network recovered by switching to a different PFC mnemonic unit. There was therefore no consistent pattern in the sequence of PFC units activated.

To summarize, the full model was able to learn a clustered recall strategy for the CVLT task. Both the fast bias in the PFC layer and the response suppression in the word unit layer played crucial roles in the development of this strategy, the former allowing the model to develop sustained categorical retrieval cues, and the latter preventing response perseveration in the face of a sustained retrieval cue.

EXPERIMENT 2: RESULTS AND DISCUSSION

Both the normal and frontally lesioned networks' recall and categorical clustering scores increased over trials (Figures 4 and 5). Lesioning one third of the incoming, outgoing, and internal PFC layer connections degraded the recall performance (Figure 4), as confirmed by a one-way ANOVA of the effect of the lesion on the total

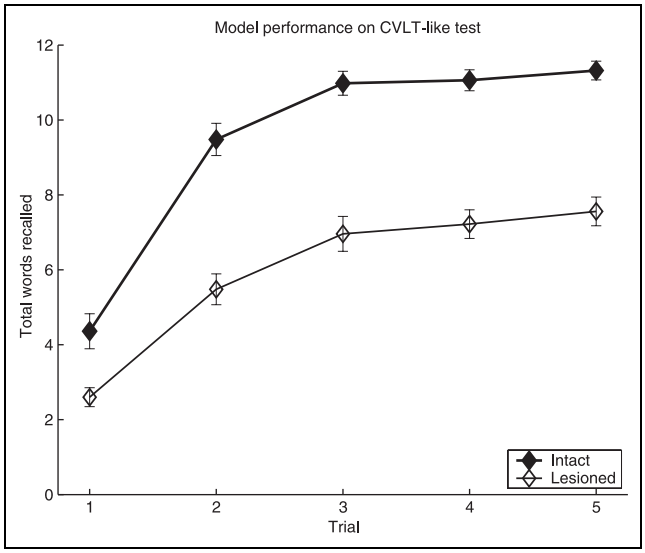

Figure 4. Total words recalled correctly in a test of free recall on a CVLT-like task, averaged over 50 runs, with standard error bars, for intact models and models with lesions to the PFC module.

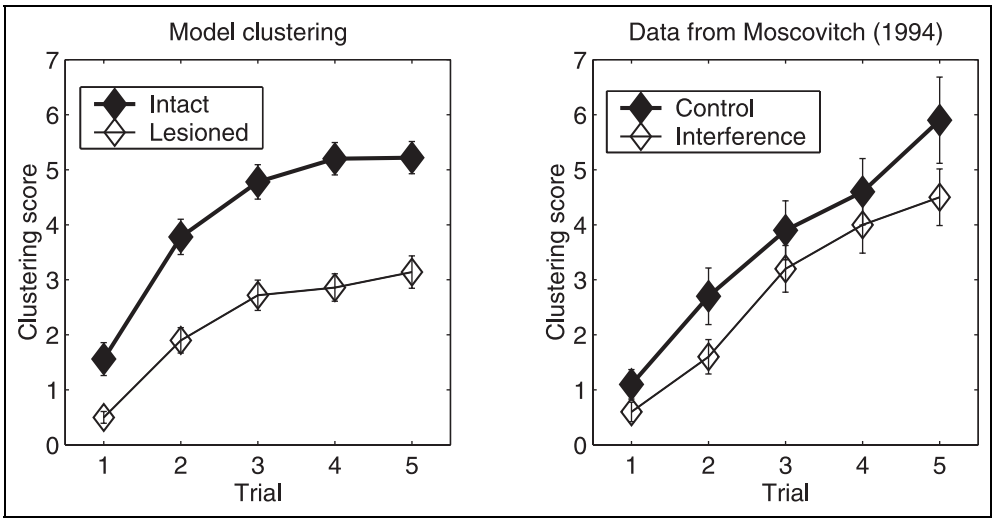

Figure 5. Categorical clustering scores versus trials for recall, with standard error bars. Left: intact versus frontally lesioned models on a CVLT-like task, averaged across 50 simulations. Right: data from human subjects on the standard version of the CVLT in two conditions, one with no interference, and one with dual task interference during both study and recall.

words recalled by each simulated subject, $F(1,98) = 127.6$, $p < .001$.

The simulated frontal lesion also affected the amount of categorical clustering (Figure 5) as revealed in a one-way ANOVA, $F(1,98) = 59.6$, $p < .001$. The categorical clustering score shown here is simply the total number of words recalled from the same category as the preceding word. This measure does not correct for the fact that recalling fewer words lowers the maximum possible clustering score. The maximum possible clustering score on a given trial is 3/4 of the total words recalled, because there are four words per category in the study list, and four words successively recalled from the same category would yield an uncorrected score of 3. Dividing the clustering score by the maximum possible clustering value given the total words recalled yields a corrected clustering score with a maximum possible value of 1. The corrected clustering scores on the first and last trial were .36 and .46 for the intact network, and .19 and .41 for the lesioned network. Thus, the normalized clustering score increases over trials for both networks, and is somewhat lower for the lesioned network.

Although the overall effect of the frontal lesion was to lower performance, the lesioned model could not be said to have a learning impairment per se; both recall and clustering scores increased over trials at roughly the same rate as for the intact model. The reduced recall and clustering in the frontally lesioned group, with relatively intact learning slopes for both measures, is suggestive of an impairment in the strategic use of memory.

The above results are broadly consistent, at least qualitatively, with the performance on CVLT-like tests of neurologically intact and frontally challenged humans. Quantitative comparisons are complicated by the fact that several factors vary across studies, including word lists and hence word similarity, number of study/test trials, and performance measures. We therefore compare our models' recall performance to the results of three different studies (see Table 2). Stuss et al. (1994) administered a CVLT-like test, but with only four study/test trials rather than the standard five, to patients with left, right, and bilateral frontal as well as nonfrontal control lesions; we have shown here data from their control lesion and bilateral frontal lesion groups only. Hildebrandt et al. (1998) used a German version of the CVLT and groups with left frontal lesions, left medial-temporal lobe lesions, and right-hemisphere control lesions. Moscovitch (1994) administered the standardized CVLT to neurologically intact individuals under conditions of divided attention, which is thought to disrupt executive control. Moscovitch's data reported in Table 2 were estimated from Figure 1 (Moscovitch, 1994); in this case,

Table 2. Free Recall Scores and Standard Deviations on CVLT-like Tasks for Computational Model (Experiment 3) and Human Data

	Control Group			*Frontal Group*		
Data Set	*Trial 1*	*Trial 2*	*Total Trials*	*Trial 1*	*Trial 2*	*Total Trials*
Model	4.4 (3.3)	9.5 (3.0)	47.2 (6.5)	2.6 (1.8)	5.5 (2.9)	29.8 (8.7)
Stuss et al. (1994)			38.1 (9.5)			25.9 (9.7)
Hildebrant et al. (1998)	5.8 (1.5)	12.7 (1.7)	51.3 (6.75)	3.8 (1.3)	8.6 (2.5)	31.9 (6.9)
Moscovitch (1994)	6	10		4	7	

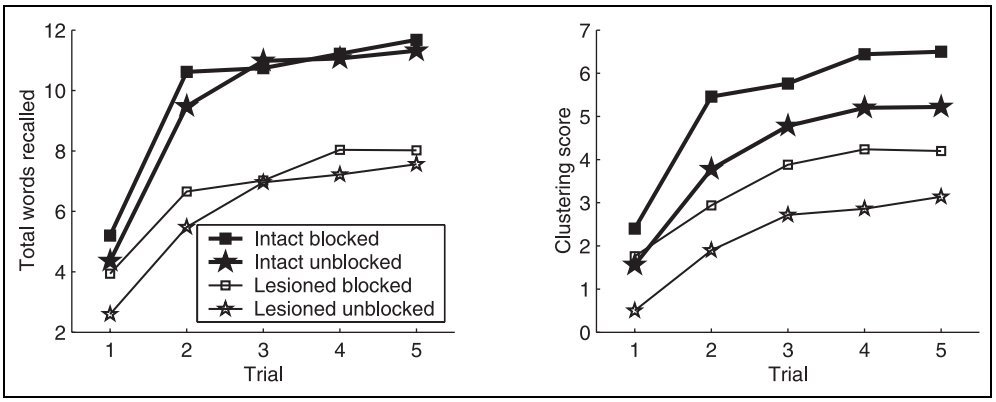

Figure 6. Total words recalled averaged over 50 simulations with standard error bars in a test of free recall for categorized blocked versus unblocked word lists.

"frontal" refers to Moscovitch's group with divided attention during both study and recall.

Frontally challenged individuals show a similar pattern of performance to that of the model with PFC damage: relatively intact learning across trials, but lower overall levels of recall. The model's pattern of categorical clustering across trials is very similar to Moscovitch's (1994) human data (see Figure 5), both qualitatively and quantitatively.

Some differences between the model and human data should be noted. The model's performance is somewhat higher on the first trial, and the learning curve is not quite as steep as the human data reported by Hildebrandt et al. (1998) and Moscovitch (1994). Various parameters of the training data and model, such as the learning rates for the various layers of the network, could be optimized to achieve much more accurate data fits. However, our goal in building the model was not to engage in a data-fitting exercise, but rather, to capture key aspects of human performance qualitatively, with as few built-in assumptions as possible, and subsequently apply the model to other tasks.

EXPERIMENT 3: RESULTS AND DISCUSSION

Experiment 3 extends the findings in Experiment 2 by comparing the model's performance on blocked and unblocked lists of categorized words, and lists of unrelated words. These results were compared to human data on the same sort of lists reported by Stuss et al. (1994). A two-way ANOVA of the effects of list type and lesion on total correct scores revealed significant main effects of lesion, $F(1,294) = 252.6$, $p < .001$, and of list type, $F(2,294) = 5.5$, $p < .004$, and a trend toward a list by lesion interaction, $F(2,294) = 2.9$, $p < .057$. For all list types, the intact models' learning slopes tended to be slightly steeper, and the overall error rates were lower than those of the lesioned models (Figure 6). Post hoc pairwise comparisons between correct scores for the three list types using Bonferroni-corrected t tests revealed significant differences between the recall of blocked and unblocked lists, $p < .05$, and between blocked and unrelated lists, $p < .01$, but not between unblocked and unrelated lists. This is very similar to the pattern of results Stuss et al. reported in comparing free recall of frontal-lobe-lesioned individuals to those with control lesions: poorer overall performance of the frontal groups in all conditions, and regardless of lesion site, superior recall of categorized lists when the words are blocked by category, as compared to unblocked or unrelated word lists, but no difference between the unrelated and unblocked conditions (see Table 3). Figure 6 shows the correct recall and categorical clustering scores of the normal and lesioned models on blocked and unblocked lists, plotted over trials. Although the differences in number of words recalled is relatively small between the blocked and unblocked conditions, the differences in clustering are quite large.

When one considers the magnitude of the blocking effect on words recalled, which was very large in both of

Table 3. Total Words Recalled, Summed over Four Trials, and Standard Deviations, by Computational Model (Experiment 3) and Human Subjects

	Model		*Stuss et al., 1994 Data*	
Condition	*Control*	*Lesioned*	*Control*	*Lesioned*
Blocked categorized list	37.8 (6.2)	25.7 (8.3)	52.8 (9.6)	34.2 (13.9)
Unblocked categorized list	35.9 (5.9)	22.3 (6.8)	38.1 (9.5)	25.9 (9.7)
Unrelated list	33.3 (5.1)	23.9 (7.9)	35.3 (7.8)	25.5 (10.7)

Stuss et al.'s (1994) patient groups and relatively small in the simulations, the quantitative fit between the model and human data is rather poor. One explanation for this may be that when people are given blocked categorized lists, the categorical structure is much more obvious and may encourage them to rely upon previously learned strategies, whereas our model has to discover its strategy by trial and error.

GENERAL DISCUSSION

Our model captures two key elements of strategic control not present in previous models, namely, the self-organization of mnemonic codes, and the ability to employ these mnemonic codes as memory cues in free recall tasks. When the PFC module is lesioned, the model performs similarly to humans with frontal lesions on CVLT-like tasks: There is no pronounced learning impairment, but fewer words are recalled and there is less categorical clustering, suggesting a deficit in strategic memory use.

Although the model can simulate some aspects of strategic memory use, it lacks many of the hallmark features of executive control that would be required to enable more complex strategies. For example, sequential strategies, while learnable in theory, are difficult for our model to learn within the time scale of a single experiment. It is likely that we learn this and other strategies gradually over the course of a lifetime and make use of them in new situations, rather than learning strategies anew every time as the model does. The ability to order responses sequentially might also be aided by unique architectural features of the PFC. Some computational models have incorporated architectural constraints that would facilitate the sequential organization of responses (e.g., Dehaene & Changeux, 1997; Bapi & Levine, 1994). Further research is needed to address other aspects of executive control, such as goal establishment, self-monitoring, task decomposition, and multitasking.

A number of other computational models of PFC involvement in memory have been proposed, though none has addressed the strategic use of memory in free recall and similar tasks. Some models have addressed specific component processes of frontal executive functions, such as maintaining state information in working memory (Dreher, Guigon, & Burnod, 2002; Rougier & O'Reilly, 2002; Durstewitz, Seamans, & Sejnowski, 2000; Moody, Wise, Pellegrino, & Zipser, 1997) and response sequencing (Dehaene & Changeux 1997; Bapi & Levine, 1994). Several models have accounted for both context maintenance and response selection/response inhibition, e.g., in the Wisconsin Card Sorting task (Berdia & Metz, 1998; Levine & Prueitt, 1989), Stroop interference task (Cohen, Dunbar, & McClelland, 1990), verbal response selection tasks (Gullapalli & Gelfand, 1995), delayed response tasks (Guignon, Dorizzi, B., Burnod, Y., & Schultz, 1995; Dehaene & Changeux, 1989), and continuous performance tasks (Rougier & O'Reilly, 2002; Braver, Barch, & Cohen, 1999; Cohen, Braver, & O'Reilly, 1996). In most of these models, PFC units are given prespecified meanings, e.g., ''sort cards by color'' or ''name the ink color.'' In contrast, our model adaptively discovers the best way to represent context, by self-organizing its internal representations of context according to task demands. Only two other PFC models we are aware of, by Rougier & O'Reilly (2002) and Guignon et al. (1995), support self-organized internal representations. While each of these models has several features in common with the PFC component of our model, they focus on modeling the maintenance in working memory of task-related variables. We address a rather different problem here: the self-organization of entirely new mnemonic codes, which may involve chunking or categorization of stimulus features, in the service of controlled memory storage and retrieval, and PFC interactions with the medial temporal lobes.

An important assumption in our model is that strategic learning occurs not only during study, but also during free recall: The model's recall strategy is dynamically organized and shaped by feedback from self-monitoring processes. Our model therefore predicts that strategic recall would be compromised if one were given multiple study episodes without the opportunity to recall after each study trial. This prediction is currently under investigation. Moscovitch (1994) found that subjects who performed an interfering task during both study and recall were impaired on the CVLT, whereas those who performed the interfering task at either study or recall alone were not. Our model predicts that either form of interference alone should affect performance. However, it is possible that Moscovitch's finger-tapping task was not sufficiently distracting to prevent strategic learning altogether.

Another important assumption in our model is that self-monitoring in free recall is reliant upon an intact MTL memory system. Individuals with MTL lesions should therefore have difficulty in employing a semantic clustering strategy in free recall, as it would result in large numbers of repetition errors. Interestingly, Hildebrandt et al. (1998) found that individuals with MTL lesions (left posterior cerebral artery incident), like those with frontal lesions, are very impaired on number of words recalled, but in addition, the former have a much greater tendency to use a serial clustering strategy in free recall as well as showing highly inflated recency effects.

The model described here could easily be extended to accommodate other verbal learning tasks such as AB–AC paired associate learning. The PFC module should learn to represent the AB or AC list context during study, and activate and maintain this context as a retrieval cue during recall.

METHODS

Architecture

The network used in all simulations reported here had 800 LSM units (100 lexical—the vocabulary size—and 500 semantic), 300 context units, 1600 MTL units (500 semantic and 300 context inputs, and an equal number of outputs), and 10 PFC units. These layer sizes were varied in preliminary simulations. Substantially fewer semantic features resulted in poorer recall scores, whereas fewer context features impaired the detection of intrusion and repetition errors. The size of the PFC layer was not critical; networks with 5 to 20 PFC units performed comparably. Each layer in the model was fully connected with the next layer, except that the input units to the MTL input layer were connected 1:1. The LSM and MTL inputs to the PFC units were delayed by one time step, so that during recall, the PFC module received context from the previously recalled word, shown as ''WM units'' in Figure 1. During study, the PFC units received nondelayed MTL input. The lesioned models had 33% of incoming and outgoing PFC connections randomly deleted. Varying amounts of damage in preliminary simulations caused graded deficits, rather than a steep drop at a critical level of damage.

Activation Rules

All units had binary (0, 1) states. The semantic, context, and MTL layers were constrained to have sparse activation levels (25% of units active) via a K-winner-take-all (kWTA) activation function. In the LSM word layer and PFC layer, only a single unit was active at one time. Before the experiments, the context units were intialized to sparse random binary states. After a word was recalled or studied, the context vector was gradually randomly updated by swapping the state of each unit with that of another unit with probability equal to .3 to maintain a fixed level of activation.

During the study phase, activity propagated bottom-up through the network (see Figure 1). An externally presented study item determined the states of the LSM units. The MTL input and output units' states were determined by the semantic and context features. In the PFC module, a single unit was selected probabilistically to become active, according to the ''softmax'' function (Bridle, 1990):

$$P_i = \frac{\exp(G \times \text{net}_i)}{\sum_j \exp(G \times \text{net}_j)}$$

where P_i is the probability of selecting the ith unit, G is a gain parameter (here $G = 100$), and $\text{net}_i = \sum_j W_{i,j} y_j + b_i^{(\text{slow})} + b_i^{(\text{fast})}$ is the weighted summed input to that PFC unit, where y_j is a presynaptic input activation, $W_{i,j}$ is the weight on that connection, and $b_i^{(\text{slow})}$ and $b_i^{(\text{fast})}$ are bias weights with small and very large learning rates, respectively. The slow bias' learning rate was the same for the other weights, W, permitting some PFC units to increase their response probability across trials, whereas the fast bias was reset to zero between recall trials. Each PFC unit was free to learn task-appropriate weights on all of these connections including the bias weights, so that it could in principle learn to exert self-excitation after being rewarded, and quickly switch to self-inhibition once the expected reward failed to materialize.

During recall there was a top-down flow of activity (see Figure 1). The PFC units' states were calculated as a function of the MTL and LSM module outputs from the previous trial. On the very first recall attempt, this input to the PFC represented the last word presented during the study phase. The input to the MTL was a weighted sum of the activations from the PFC layer, converted to a binary code using a kWTA activation function. Before propagating the activations through the MTL module, the off state of the input layer units was set to (−sparseness)/(1 − sparseness) rather than to zero; this reduces the learning equations to simple Hebbian rules (Willshaw & Dayan, 1990). The activation of the MTL output layer was a kWTA function of the net input from the MTL input layer. The activations of the LSM module's semantic feature units were simply a 1:1 copy of the corresponding MTL output units' activations. A single word response was then generated probabilistically in the LSM module's lexical layer, according to the softmax function (with Gain = 200) of the weighted inputs from the LSM semantic units.

To implement the response-suppression mechanism in the word layer, the last four words recalled did not enter into the softmax competition, and were thereby prevented from becoming active. More realistically, this could be simulated via self-inhibitory connections, or via long-range plastic connections from the PFC layer acting on local inhibitory interneurons. Recent neuropsychological evidence supports a role for the orbitofrontal cortex in mediating reward-related inhibition of return (Hodgson et al., 2002).

Because activations in the PFC and word layers were noisy, the network was allowed up to three retries on a given trial to generate a valid word response. After each unsuccessful attempt, the fast bias weights in the PFC layer were adjusted through reinforcement learning using a negative reinforcement signal (see below). If a successful recall attempt was made, a response was generated, and all of the weights in the network were adjusted with a positive reinforcement signal.

The model evaluated the recency of a recalled item by probing the MTL module with the item just recalled, combined with the current context. Recency is calculated as the harmony, or negative of the energy (Hopfield, 1982) in the MTL module:

$$\text{recency} = \sum_{i<j} W_{i,j}^{(MTL)} y_i^{(MTL)} y_j^{(MTL)}.$$

This was inspired by an idea of Joordens and Becker (1997) to use harmony of a retrieved item as a measure of familiarity. Recency is high when the probe closely matches the stored memory patterns in the weights, $W_{i,j}^{(\mathrm{MTL})}$. If the same item is repeatedly recalled, recency will be extremely high due to the close match to the current context. Thus, a repetition error was detected when the recency greatly exceeded that of the last few recalled words (recency > averecency + 6) where the average recency is updated after each recall by (1/3 recency) − (2/3 averecency). Similarly, an intrusion error was detected if recency was too low (recency < 0.5 × averencency).

Vocabulary

The network was trained on a vocabulary of 100 words, consisting of 16 CVLT ''list words'' drawn from four categories with four words in each, and 84 extra-list words. For each word, a set of semantic features was generated randomly with 25% sparseness, such that words within a category shared 40% of their features. Of the 84 extra-list words, 20 words were drawn from the same categories as the list words (but were not studied during the CVLT experiment), 8 words formed two new semantic categories containing 4 words each, and the remainder were unrelated words.

Learning Rules

All weights were initialized to zero, except for PFC weights, which were random between 0 and 1. The fast bias PFC weights were reset to zero at the end of each recall trial, and were fixed at zero during each study trial.

In a preexperimental phase, the LSM and MTL weights were pretrained in a single pass through the entire vocabulary of 100 words; for each vocabulary item, the states of the LSM units were set to the corresponding word features, and the lexical-to-semantic and reciprocal connection weights were then incremented by the following amount:

$$\Delta W_{i,j} = .005 y_i y_j$$

where y_i and y_j are the pre- and postsynaptic activities. The MTL module formed an episodic trace of each pretrained vocabulary item by autoassociating the vector of semantic features and current context; the within-MTL weights $W_{i,j}^{(\mathrm{MTL})}$ were updated via a Hebbian rule with weight decay:

$$W_{i,j}^{(MTL)} = \mathrm{Decay} \times W_{i,j}^{(MTL)} + \frac{1}{N} \mathrm{MTL}_i^{(input)} \mathrm{MTL}_j^{(output)}$$

where Decay = .96 and N is the MTL input layer size. In this preexperimental phase, each word was associated with a completely different random context in the MTL module.

During the experimental phase, the network studied only the 16 CVLT list words, and the context was changed gradually. The lexical-to-semantic and MTL connection weights were updated via the same Hebbian learning rules used in the preexperimental phase. The connection weights for PFC units were trained via Q-learning (Watkins, 1989), a version of reinforcement learning in which an agent learns to predict the ''value'' of each possible action in each situation. The value is defined as the expected sum of future rewards, temporally discounted by a factor γ (here $\gamma = 0.3$):

$$\mathrm{Value}(t) = r_t + \gamma r_{t+1} + \gamma^2 r_{t+2} + \ldots + \gamma^n r_{t+n} + \ldots$$

We interpret each PFC unit's net input $\mathrm{net}_i = \mathrm{Value}(i)$ as its estimated value of taking that ''action'' (or retrieval cue, in this case). Thus, the softmax function applied to this value probabilistically selects the PFC unit that predicts the greatest reward. The actual reward r received during free recall depends on the model's self-evaluation of the retrieved item: A non-repeated list word results in $r = +1.0$, and the selected PFC unit adjusts all of its weights. An intrusion or repetition error results in $r = -1.0$, and only the fast bias weight for the selected PFC unit is adjusted. During the study phase the reward is held fixed at 1, since the model was not permitted to generate a top-down response during this time. The Q-learning rule consists of a Hebbian term multiplied by a reward prediction error:

$$W_{i,j} = W_{i,j} + \mathrm{lrate}_{i,j} \mathrm{PFC}_i y_j \mathrm{error}$$

where PFC_i is the activation of the ith PFC unit (1 if that unit was selected to respond and 0 otherwise), y_j is the activation of the jth input to that unit ($y_j = 1$ for bias weights), and $\mathrm{lrate}_{i,j}$ is the learning rate for that connection. The reward prediction error on the ith trial (Watkins, 1989) is:

$$\mathrm{error}(i) = r(i+1) + \gamma \mathrm{Value}(i+1) - \mathrm{Value}(i)$$

where $\mathrm{Value}(i)$ is the net input to the selected PFC unit on the ith trial. For all weights on connections from the MTL to PFC units, the learning rate was 0.005. The top-down weights from the PFC to the MTL were constrained to be equal to the corresponding bottom-up MTL-to-PFC weights. For the PFC bias weights, the same Q-learning rule was used with a large learning rate of 5 when $r > 0$, and 50 when $r < 0$.

Q-learning has the following important consequences: (1) If more reward was received than expected, the weights to the selected unit for active input lines are increased, making that response to that input more likely in the future. (2) If less reward was received than expected, the same weights are decreased making that response less likely.

CVLT Procedure

Essentially the same procedure was used for Experiments 1, 2, and 3. Pretraining on the 100-word vocabulary was followed by five interleaved study and recall trials using the 16 simulated CVLT list words. At the start of each study and recall trial, the context vector was updated for five cycles, whereas after each word was recalled or studied, the context vector was updated for only one cycle, to generate a greater shift in context between than within trials. In Experiments 1 and 2, the studied words were always presented in an unblocked fixed order such that no two successive words came from the same category. In Experiment 3, in the blocked condition, study words were grouped by category, and in the unrelated condition, all study words were drawn from different categories. Fifty different simulated ''subjects'' were run in each condition, each with different random initial weights.

In the standard CVLT procedure, a trial ends when the subject cannot think of any more words. Similarly, in our simulations, a trial typically ended when the network could not generate any valid words. Although the network was given a maximum of 20 chances to recall words, it rarely if ever recalled more than 16 items, because it was able to detect and suppress intrusion and repetition errors with a high degree of accuracy.

The performance measures used were the total number of words recalled correctly, the number of intrusion and repetition errors, and the categorical clustering score. The latter was incremented each time a word was recalled correctly from the same category as that of the previously recalled word.

Acknowledgments

We thank Larry Jacoby for many inspiring discussions on the control of memory, which motivated the development of this model. This work was supported by a research grant from the Natural Sciences and Engineering Research Council of Canada to S. Becker.

Reprint requests should be sent to S. Becker, Department of Psychology, McMaster University, 1280 Main Street West, Hamilton, Ontario, Canada L8S 4K1, or via e-mail: becker@mcmaster.ca.

REFERENCES

Bapi, R., & Levine, D. (1994). Modeling the role of frontal lobes in sequential task performance: Basic structure and primacy effects. *Neural Networks, 7,* 1167–1180.

Berdia, S., & Metz, J. (1998). An artificial neural network stimulating performance of normal subjects and schizophrenics on the Wisconsin Card Sorting Test. *Artificial Intelligence in Medicine, 13,* 123–138.

Bousfield, W. (1953). The occurrence of clustering in the recall of randomly arranged associates. *Journal of General Psychology, 49,* 229–240.

Braver, T. S., Barch, D. M., & Cohen, J. D. (1999). Cognition and control in schizophrenia: A computational model of dopamine and prefrontal function. *Biological Psychiatry, 46,* 312–328.

Bridle, J. (1990). Probabilistic interpretation of feedforward classification network outputs, with relationships to statistical pattern recognition. In F. Fougelman Soulie & J. Herault (Eds.), *Neuro-computing: Algorithms, architectures and applications* (pp. 227–236). New York: Springer-Verlag.

Brown, G. D. A., Preece, T., & Hulme, C. (2000). Oscillator-based memory for serial order. *Psychological Review, 107,* 127–181.

Burgess, N., & Hitch, G. J. (1992). Towards a network model of the articulatory loop. *Journal of Memory and Language, 31,* 429–460.

Burgess, N., & Hitch, G. J. (1996). A connectionist model of STM for serial order. In S. E. Gathercole (Ed.), *Models of short-term memory* (pp. 51–72). Hove, UK: Psychology Press.

Burgess, N., & Hitch, G. J. (1999). Memory for serial order: A network model of the phonological loop and its timing. *Psychological Review, 106,* 551–581.

Cohen, J., Dunbar, K., & McClelland, J. L. (1990). On the control of automatic processes: A parallel distributed processing model of the Stroop effect. *Psychological Review, 97,* 332–361.

Cohen, J. D., Braver, T. S., & O'Reilly, R. C. (1996). A computational approach to prefrontal cortex, cognitive control, and schizophrenia: Recent developments and current challenges. *Philosophical Transactions of the Royal Society of London Series B (Biological Sciences), 351,* 1515–1527.

Constantinidis, C., Franowicz, M. N., & Goldman-Rakic, P. S. (2001). Coding specificity in cortical microcircuits: A multiple-electrode analysis of primate prefrontal cortex. *Journal of Neuroscience, 21,* 3646–3655.

Dehaene, S., & Changeux, J. (1989). A simple model of prefrontal cortex function in delayed-response tasks. *Journal of Cognitive Neuroscience, 1,* 244–260.

Dehaene, S., & Changeux, J. (1997). A hierarchical neuronal network for planning behavior. *Proceedings of the National Academy of Sciences, U.S.A., 94,* 13293–13298.

Delis, D. C., Kramer, J. H., Kaplan, E., & Ober, B. A. (1987). *The California Verbal Learning Test.* San Antonio, TX: Psychological Corporation.

Dreher, J.-C., Guigon, E., & Burnod, Y. (2002). A model of prefrontal cortex dopaminergic modulation during the delayed alternation task. *Journal of Cognitive Neuroscience, 14,* 853–865.

Duncan, J. (2001). An adaptive coding model of neural function in prefrontal cortex. *Nature Reviews Neuroscience, 2,* 820–829.

Durstewitz, D., Seamans, J. K., & Sejnowski, T. J. (2000). Dopamine-mediated stabilization of delay-period activity in a network model of prefrontal cortex. *The Journal of Neurophysiology, 83,* 1733–1750.

Farrell, S., & Lewandowsky, S. (2002). An endogenous distributed model of ordering in serial recall. *Psychonomic Bulletin and Review, 9,* 59–71.

Frith, C. (2000). The role of dorsolateral prefrontal cortex in the selection of action as revealed by functional imaging. In S. Monsell & J. Driver (Eds.), *Control of cognitive processes: Attention and performance XVIII.* (pp. 549–565). Cambridge: MIT Press.

Guignon, E., Dorizzi, B., Burnod, Y., & Schultz, W. (1995). Neural correlates of learning in the prefrontal cortex of the monkey: A predictive model. *Cerebral Cortex, 5,* 135–147.

Gullapalli, V., & Gelfand, J. (1995). A model of the dynamics of prefrontal cortico-thalamo-basal ganglionic loops in verbal

response selection tasks. In J. Grafman, K. Holyoak, & F. Boller (Eds.), *Structure and function of the human prefrontal cortex* (pp. 375–380). New York: New York Academy of Sciences.
Henson, R. N. A. (1998). Short-term memory for serial order: The start–end model. *Cognitive Psychology, 36,* 73–137.
Hildebrandt, H., Brand, A., & Sachsenheimer, W. (1998). Profiles of patients with left prefrontal and left temporal lobe lesions after cerebrovascular infarctions on California Verbal Learning Test-like indices. *Journal of Clinical and Experimental Neuropsychology, 20,* 673–684.
Hodgson, T. L., Mort, D., Chamberlain, M. M., Hutton, S. B., O'Neill, K. S., & Kennard, C. (2002). Orbitofrontal cortex mediates inhibition of return. *Neuropsychologica, 1431,* 1–11.
Hopfield, J. (1982) Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences, U.S.A., 79,* 2554–2558.
Houghton, G., & Hartley, T. (1996). Parallel models of serial behavior: Lashley revisited. *Psyche, 2*.
Howard, M. W., & Kahana, M. J. (2001). A distributed representation of temporal context. *Journal of Mathematical Psychology, 46,* 269–299.
Jetter, W., Poser, U., Freeman, R., & Markowitsch, H. (1986). A verbal long term memory deficit in frontal lobe damaged patients. *Cortex,* 229–242.
Levine, D., & Prueitt, P. (1989). Modeling some effects of frontal lobe damage—novelty and perseveration. *Neural Networks, 2,* 103–116.
Lewandowsky, S., & Duncan, M. W. J. (2002). The time-course of response suppression: No evidence for a gradual release from inhibition. *Proceedings of the Psychonomics Society, 7,* 69.
Lewandowsky, S., & Murdock, B. B. (1989). Memory for serial order. *Psychological Review, 96,* 25–57.
Moody, S., Wise, S., Pellegrino, G., & Zipser, D. (1997). A model that accounts for activity in primate frontal cortex during a delayed matching to sample task. *Journal or Neuroscience, 18,* 399–410.
Moscovitch, M. (1994). Cognitive resources and dual-task interference effects at retrieval in normal people: The role of the frontal lobes and medial temporal cortex. *Neuropsychology, 8,* 524–534.
Murdock, B. B. (1997). Context and mediators in a theory of distributed associative memory. *Psychological Review,* 839–862.
Page, M. P. A., & Norris, D. (1998). The primacy model: A new model of immediate serial recall. *Psychological Review, 105,* 761–781.
Posner, M. I., Rafal, R. D., Choate, L. S., & Vaughan, J. (1985). Inhibition of return: Neural basis and functions. *Cognitive Psychology, 2,* 211–238.
Rougier, N. P., & O'Reilly, R. C. (2002). Learning representations in a gated prefrontal cortex model of dynamic task switching. *Cognitive Science, 26,* 503–520.
Shimamura, A., Jurica, P., Mangels, J., Gershberg F., & Knight, R. (1995). Susceptibility to memory interference effects following frontal lobe damage: Findings from tests of paired-associate learning. *Journal of Cognitive Neuroscience, 7,* 144–152.
Shimamura, A. P., Janowsky, J. S., & Squire, L. (1990). Memory for the temporal order of events in patients with frontal lobe lesions and amnesic patients. *Neuropsychologia, 28,* 803–813.
Stuss, D., Alexander, M., Palumbo, C., Buckle, L., Sayer, L., & Pogue, J. (1994). Organizational strategies of patients with unilateral or bilateral frontal lobe injury in word list learning tasks. *Neuropsychology, 8,* 355–373.
Vousden, J. I., & Brown, G. D. A. (1998). To repeat or not to repeat: The time course of response suppression in serial behaviour. *Proceedings of the Fourth Neural Computation and Psychology Workshop: Connectionist Representations.* London: Springer-Verlag.
Vousden, J. I., Brown, G. D. A., & Harley, T. A. (2000). Serial control of phonology in speech production: A hierarchical model. *Cognitive Psychology, 41,* 101–175.
Watkins, C. (1989). *Learning from delayed rewards.* PhD thesis. King's College, Cambridge University, UK.
Willshaw, D., & Dayan, P. (1990). Optimal plasticity from matrix memories: What goes up must come down. *Neural Computation, 2,* 85–93.