

Sequential Retrieval and Inhibition of Parallel (Re)Activated Representations: A Neurocomputational Comparison of Competitive Queuing and Resampling Models

Eddy J. Davelaar

Department of Psychology, University of Maryland

Sequential behavior is observed in various domains of cognitive psychology, including free recall paradigms. In this article, within a neurocomputational framework, resampling (RS) mechanisms are compared with competitive queuing (CQ) mechanisms. While both types of implementations select the most active representation, the subsequent inhibition is at the level of selection for RS models and at the level of (re)activation for CQ models. It is shown that despite the overwhelming success of CQ models in serial recall (with regard to types of sequencing error) RS models outperform CQ models with regard to inter-response times in a free recall task. Additional analyses show that decay of response suppression reduces the difference between the models. The RS model is sensitive to the size of the search set and accounts for memory selection performance in patients with Alzheimer's dementia or Huntington's disease. Finally, a non-mnemonic clustering behavior is observed, which is related to the dynamical process of the selection mechanism.

Keywords competitive queuing · memory retrieval · recall latencies · resampling · Alzheimer's dementia · Huntington's disease

1 Introduction

In many complex systems that exhibit sequential ordering of actions, a selection mechanism converts simultaneously activated representations into a sequence of actions. The simplest form this selection mechanism can take is one by which one action is selected and then executed, after which another action is activated (from zero), selected and executed, and so on. However, research within psychology has demonstrated that this

type of selection mechanism is inconsistent with detailed error analyses in serial recall (Henson, Norris, Page, & Baddeley, 1996) and typing (see Salthouse, 1986), as it would, for example, be unable to account for the smooth transitions between successive actions. Instead, evidence points towards a two-stage mechanism, by which in the first stage all task-appropriate representations are activated in parallel (instead of sequentially), and in the second stage representations are selected one at a time to produce its corresponding

Correspondence to: Eddy J. Davelaar, School of Psychology, Birkbeck, University of London, WC1E 7HX, London, UK.
E-mail: e.davelaar@bbk.ac.uk; eddy.davelaar@gmail.com

Copyright © 2007 International Society for Adaptive Behavior (2007), Vol 15(1): 51–71.

DOI: 10.1177/1059712306076250

Figures 4, 5, 9, A1 appear in color online: <http://adb.sagepub.com>

action. Although most of the initial knowledge was gathered from motor behavior, with typing as one familiar example (see the review by Salthouse, 1986 and references therein), researchers have suggested that this two-stage mechanism may also be central to other forms of sequential behavior, from walking through a room or grasping a mug to producing speech, preparing coffee, or even memorizing words. Given this centrality of serial ordering in our daily lives, it may come as no surprise that computational modelers have developed models through which simultaneously activated representations produce serially ordered actions. Two of these models are highlighted in this article and relate to sequential retrieval from memory.

The reason for focusing on memory retrieval is twofold. First, a long tradition of mathematical sampling models has led to a wide range of memory models that are currently used to account for memory phenomena. During the 1990s, a new generation of models emerged that were embedded in a connectionist framework and used a particular mechanism called “competitive queuing”, which is able to transform an activation gradient over to-be-reported memoranda into a sequential order of reports. As is argued in this article, the two mechanisms make different predictions with regard to the timing of actions. This difference may be important beyond the memory literature in other areas in which sequential selection of parallel-activated actions is a common implementational choice. A modeler would need to know which mechanism to use to account for a particular type of data and empirical researchers (psychologists, biologists) would need to know what the behavioral signatures of these mechanisms are. For example, it is useful to know that if a certain mechanism underlies higher-order behavior (e.g. parallel activation and sequential selection), specific implementational details may lead to different behavior at a more fine-grained level (in this article, the profile of the timing between successive actions). Secondly, in recent years, evidence has been found that favors some components and disfavors other components of competitive queuing. Given the wide applicability of a two-stage activation–selection mechanism, it could be beneficial to re-examine the components with computational analytical tools that are currently available. This article aims at directly comparing two well-known mechanisms that have never before been compared in this way. The memory literature contains much of the data needed for this comparison; this arti-

cle uses these data in an attempt to validate these models.

In this article, the term “action selection” is used loosely and relates to the selection of representations that when executed have a behavioral consequence. Within this definition, retrieving a word from memory will eventually lead to the spoken or typed version of that word. Similarly, retrieving a plan of actions (e.g. making coffee) will eventually lead to the selection of a series of goal-directed movements. This article therefore deals only with situations in which a higher-order action plan (retrieving memoranda) activates lower-order plans (producing a word) that eventually lead to a behavioral response that can be measured in real time. The article does not deal with situations that require unpacking of a motor sequence, such as walking or throwing a ball.

Section 2 outlines the competitive queuing and sampling mechanisms, which is then followed by a brief introduction of the retrieval latencies that differentiate these mechanisms. A generic computational model is presented in which both mechanisms are implemented. This model is then used to demonstrate the impact of the different mechanisms on retrieval latencies, and is used to show its ability to qualitatively account for psychological data through a number of simulations. The article closes with implications of the computational evaluation for models of memory, action selection, and other domains.

2 Competitive Queuing Versus Resampling

This section presents the two mechanisms that will be compared in a generic model. Before describing these mechanisms, it should be noted that models of recall memory fall into either one of two almost separate literatures. On the one hand are the models that deal with serial recall, while on the other hand are models that deal with free recall. Each literature has its own debates, model implementations, empirical methods and focus on what is relevant in the data. This article focuses on free recall or sequential selection in any order for three reasons. First, neurocomputational models of free recall are being developed (Davelaar, Goshen-Gottstein, Ashkenazi, Haarmann, & Usher, 2005; Davelaar, Haarmann, Goshen-Gottstein, & Usher, 2006), but still rely on the non-dynamic retrieval process, as used in current mathematical models. The question is how to

extend this new neurocomputational work by borrowing dynamical mechanisms of sequential selection from the literature on serial recall. Secondly, detailed analyses of retrieval latencies exist for free recall (see Section 3), and these form a benchmark against which to evaluate the models. Although data for serial recall exist (Farrell & Lewandowsky, 2004; Kahana & Jacobs, 2000), similar rigorous analyses have not yet been conducted. Thirdly, as free recall is by definition free from any constraints on output order, the task is highly suitable to investigate mechanisms of sequential selection. Just as the properties of inherent circadian rhythm in agents (man, animal, plants) can be investigated by taking away external cues of time, so can the sequential selection mechanism of agents be investigated by taking away external constraints of sequencing (i.e. the requirement to produce actions or report words in a pre-specified order). Given that sequential selection is central to behavior in general, the analyses in the following support the view that measuring latencies between successive behaviors can provide insight into the type of underlying selection mechanism that led to the sequence of behaviors.

2.1 Competitive Queuing

Grossberg (1978a, 1978b) developed a model in which sequential output can be obtained from parallel activation in a dynamical model with feedforward excitatory and feedback inhibitory connections. The basic architecture is shown in Figure 1. The profile of activations in the activation layer is fed into the selection layer. This profile may originate from an activation-based short-term buffer, or from long-term memory with or without a static signal. There are three critical components. First, the activation profile shows the degree of activations for all yet-to-be executed plans, with the desired sequential order being from most active to least active. Second, the activations compete, that is, the representations in the activation layer all inhibit each other. A non-specific arousal signal arrives at the selection layer, which brings the activations above a response threshold. Third, as soon as a plan is selected (e.g. for moving the fingers to the desired key for typing, or flexing the muscles for articulation of a word), the representation in the selection layer inhibits the representation it received input from. This prevents perseverative behavior and allows the next-strongest representation to become selected for output.

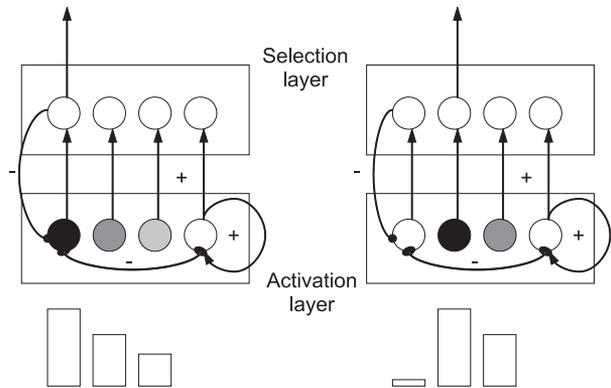


Figure 1 Architecture of the model used in Grossberg (1978a, 1978b). Each unit in the activation layer has self-recurrent excitatory connection and inhibits all other units in the same layer. The feedforward one-to-one connections from the activation layer to the selection layer are excitatory, whereas the feedback one-to-one connections are inhibitory. At the lower end is a snapshot of the activation gradient in the activation layer. On the left, the leftmost unit is the most active and leads to a response. After that, the feedback inhibition causes that unit to be lowered in activation. Because of a decrease in the amount of inhibition in the activation layer, the unit with the next-highest activation level can increase in activation and will lead to the next response. This cycle continues until all units who received activation have produced an output.

This two-layer output module produces all actions corresponding to the activated representations in the correct serial order with smooth transitions from one selected action to the next. In addition, by adding noise to any of the three critical components, it can account for errors in sequencing seen in human behavior, such as seen in typing (e.g. Salthouse, 1986) and memory for serial order.

In the 1990s, starting with Houghton (1990), this architecture received much attention in the memory literature (Brown, Preece, & Hulme, 2000; Burgess & Hitch, 1999; Henson, 1998; Page & Norris, 1998; for a review, see Page & Henson, 2001) and related areas – planning (Cooper & Shallice, 2000), spelling (Glasspool & Houghton, 2005) and speech production (Hartley & Houghton, 1996) – leading to a range of models employing the select-then-inhibit dynamics (for a review, see Glasspool, 2005). The models are collectively referred to as competitive-queuing (CQ) models, as the activations are placed in a competitive queue before being selected.¹

It should be noted, however, that these models of serial ordered recall are different from the original model of Grossberg (1978b) in at least two critical aspects. First, in the Grossberg model, the inhibition of the representation in the activation layer is complete (i.e. total response suppression), whereas the CQ models were aimed at capturing errors, such as repetitions, and thus assumed a need for a gradually decaying inhibition (for a non-decaying inhibition account, see Farrell & Lewandowsky, 2002). If the set of activated representations in the activation layer is called the search set, then the Grossberg model could be referred to as a sampling-without-replacement model and the CQ models could be referred to as sampling-with-delayed-replacement models. Secondly, whereas the Grossberg model focuses on producing sequential behavior given a single activation gradient, the aforementioned CQ models (except the primacy model by Page & Norris, 1998, which follows Grossberg in this regard) have addressed various forms of changing the source of activation during the course of retrieval. A contextual representation is incorporated in those models that changes along a certain dimension (e.g. temporal, absolute, or relative position in the list) and becomes linked with the representations of the to-be-remembered items. During retrieval, this context signal is replayed, leading to serial recall. Not surprisingly, this context signal may contribute greatly to variations in the timing of actions. However, these models have yet to be extended to account for retrieval latencies.²

Neurophysiological studies support the dynamics assumed in these types of models. For example, Averbeck, Chafee, Crowe, & Georgopoulos (2002) trained macaque monkeys to draw geometric figures (triangle, square, trapezoid, inverted triangle). After training, the recorded neural firing patterns of neurons in the prefrontal cortex during drawing suggested that representations of all segments of a figure were coactivated before the initiation of the first action. During drawing, the neurons corresponding to the upcoming action increased and then decreased in firing rate before and after the action, respectively, in the same manner as predicted by the Grossberg model and CQ models that use a single context representation. However, the results did not show that the activation of the selected-and-then-inhibited representations gradually increased, as is critical in CQ models to account for repetition errors. Although it is certainly possible that the motor task of drawing figures by monkeys is not comparable to

human memory performance, recent research on human memory has questioned the existence of a gradual decaying inhibition and employed other forms of response suppression (Duncan & Lewandowsky, 2005; Farrell & Lewandowsky, 2002, 2004). Nevertheless, the simulation study in this paper (specifically in the Appendix) considers both the sampling-without-replacement (CQ⁰) and the sampling-with-gradual-replacement (CQ^Δ) versions.

2.2 Resampling Models

Although the CQ models gained much success, a different approach to producing sequential retrieval has been employed in global memory models, such as Search of Associative Memory (SAM; Raaijmakers & Shiffrin, 1980, 1981), MINERVA 2 (Hintzman, 1984), and Theory of Distributed Memory (TODAM; Murdock, 1982). Here, SAM is used to exemplify the retrieval process, which is broken down into three critical components. First, a search set is defined based on available retrieval cues, which in most models is the list context, but could also be categories (e.g. Gronlund & Shiffrin, 1986) or chunks (e.g. Anderson & Matessa, 1997; Anderson, Bothell, Lebiere, & Matessa, 1998). Second, a single trace is selected (sampled) from the search set and used to recover the full memory representation. This could correspond with recovering the phonemic representation in order to utter the word or the execution of the motor plan for typing the word. Third, after successful recovery, the trace is allowed to compete for resampling (RS) if the same cue is used to probe the memory system, but will not produce an output. Because of the possibility of RS (but not recovery) of an item, SAM could be seen as a RS or selection-with-replacement model. In order to compare the CQ model with the RS model, it is assumed that the first two components of both models are equivalent, and that the critical difference lies in the post-response suppression mechanism. The assumption in SAM that a resampled representation never leads to an output can be relaxed and the assumption of gradual decay of inhibition could be used instead. Such a model has never been used in the SAM framework, but will be used in the simulation study for completeness.

Functional imaging data support the view that the brain has separate components for activation and selection (Badre, Poldrack, Pare-Blagoev, Insler, & Wagner, 2005). For example, in a study by Badre et al., participants performed a number of tasks that have

previously been used in investigations on memory retrieval and have been argued to rely to different degrees on the ability to select the correct response among distractors. The authors were interested in whether selection and activation mechanisms are subserved by the same or by different neural substrates. Factor analysis on the behavioral results produced two factors that were labeled as selection and non-selection components. These behaviorally defined factors were associated with variance in different brain regions, with the anterior ventrolateral pre-frontal cortex (aVLPFC) and the inferior-temporal cortex (IT) related to activation of information and the mid-ventrolateral pre-frontal cortex (mVLPFC) related to selection of task-appropriate representations. This two-stage model of retrieval supports the global memory models, but does not necessarily falsify the CQ mechanism. Nevertheless, the retrieval dynamics are different for the two mechanisms, as discussed in Section 3.

3 Retrieval Dynamics

This section summarizes the data patterns that are used to compare and validate the above two models. Recall latencies provide information on the underlying memory processes and distinguish different patient populations.

3.1 Recall Latencies

In a series of analytical studies, Rohrer and Wixted (Rohrer, 1996, 2002; Rohrer & Wixted, 1994; Wixted & Rohrer, 1993) investigated the temporal dynamics in free recall of words. This work, which extends the limited number of previous empirical investigations (e.g. Murdock & Okada, 1970; Patterson, Meltzer & Mandler, 1971; for a review, see Wixted & Rohrer, 1994), provided critical insight regarding the sampling and recovery of items. In a typical experiment, a participant memorizes a number of words and after a distractor task reports as many words as possible in any order. During the retrieval phase, the time taken to retrieve an item is measured as the main dependent variable. The basic findings are as follows. The time needed to retrieve an item increases with the number of items that were memorized (set size effect). The inter-response time (IRT), the elapsed time between two consecutive responses, increases during the retrieval phase and is a function of the number of words yet to be recalled (see

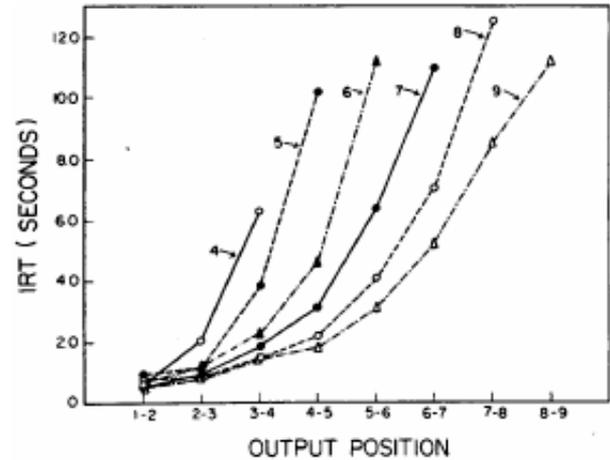


Figure 2 Mean IRT as a function of ordinal position in output. The list had 20 items and because participants successfully report a variable number of items, the data are partitioned according to the total number of words recalled (4–9). Note that independent of the total number of items reported from a 20-word list, the IRTs increase with the number of words reported so far. (From Murdock, B. B., & Okada, R. (1970, p. 265). *Interresponse times in single-trial free recall*. *Journal of Experimental Psychology*, 86, 263–267. Published by APA. Reprinted with permission.)

Figure 2). These findings have been interpreted to support the RS mechanism of selection, where the recall latency reflects the size of the search set. The larger the search set, the longer it takes to select an item from the set.

3.2 Utility of Retrieval Latencies

The analysis of recall latencies has provided an empirical tool for measuring the size of the memory search set, dissociating the effect of episodic and semantic cues (Rohrer, 2002), and measuring the loss of semantic memory in patients with Alzheimer's dementia (AD). After analyzing the recall latencies (minus the first recall latency), Rohrer, Wixted, Salmon, & Butters (1995) argued that patients with AD have a structural memory deficit. The loss of neural tissue in AD leads to a smaller size of the memory set, which in turn is observed as shorter recall latencies for patients with AD compared to controls. This is in contrast to the longer recall latencies observed in patients with Huntington's disease (HD; Rohrer, Salmon, Wixted, & Paulsen, 1999). Whereas patients with AD and HD have lower total recall compared to controls, the retrieval latencies revealed marked differences.

3.3 Model Predictions

As mentioned above, CQ and RS models have a stage in which all target representations are activated by a cue. It takes time for the first item to be selected and produce an output. Given that no differences exist between these two models at this initial stage, both models will produce the same first retrieval latency. However, after the first item is retrieved, in CQ models the representation of this item is inhibited (i.e. deleted from the queue). As every output is followed by inhibition of the retrieved item representation, there is an ever-decreasing competition during the retrieval phase. This model therefore predicts that the time needed for selecting a new item decreases: IRTs become shorter. In RS models, the item representation is still activated by the cue and can therefore be sampled, even though it will not be output. The more items that have been output, the lower the probability (and therefore the longer it will take) that a new not-yet retrieved item will be sampled. This model predicts that the time needed to report the next item depends on the activation gradient, and therefore reflects the size of the memory set. This model therefore predicts that the time needed for selecting a new item increases: IRTs become longer.

4 Generic Model

In order to evaluate the mechanisms, a generic model is presented in which RS and CQ mechanisms can be explored. The main dependent measures are the IRTs and the distributions of recall latencies. As Simulation 1 makes clear that the CQ mechanism does not capture the IRTs, Simulations 2, 3, and 4 only address the RS model to illustrate its applicability.

4.1 Model Architecture

The model consists of four layers (see Figure 3). Each layer contains 20 localistic representations that are connected to corresponding representations in other layers. Each unit corresponds to a large number of neurons that together participate in the neural code for that particular representation. For every unit in each layer, its current activation depends on the activation value on the previous time step, the self-recurrent excitation, the inhibition felt from every other unit in

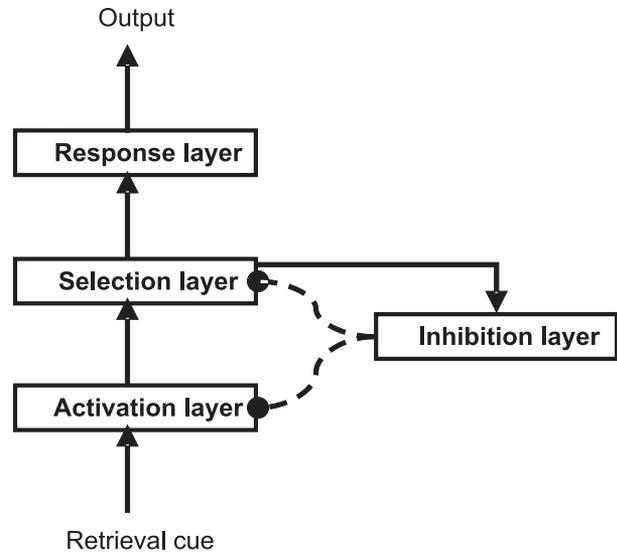


Figure 3 Model architecture. The model consists of four interconnected layers. The arrows ending in pointed heads and circled heads are excitatory and inhibitory connections, respectively. A number of representations are activated at the activation layer, which feeds activation to the selection layer. When a unit in the selection layer reaches a selection threshold, the input to the inhibition and response layers receive a sharp pulse. Depending on the employed mechanism, the output of the inhibition layer inhibits the selection (in the RS model) or the activation (in the CQ model) layers.

the same layer, the external input and some random noise. The activations of all units in the model are updated at each time step according to the following differential equation (see also Davelaar et al., 2005, 2006; Usher & McClelland, 2001):

$$x_i(t+1) = \lambda x_i(t) + (1 - \lambda) \{ \alpha F[x_i(t)] + I_i(t) - \beta \sum_j F[x_j(t)] + \xi \}. \quad (1)$$

Here, x represents the internal activation, $\lambda = 0.98$, α is the self-recurrent excitatory connection, $F(x) = x / (1 + x)$, for $x > 0$ represents the output activation function. $I_i(t)$ represents the external input at time step t to unit i . Units in each layer compete for activation, which is governed by the within-layer inhibition parameter β . Each unit receives inhibition from every other unit in the same layer. The activation of each unit (only in the memory layer) is supplemented with zero-mean Gaussian noise ξ , with standard deviation

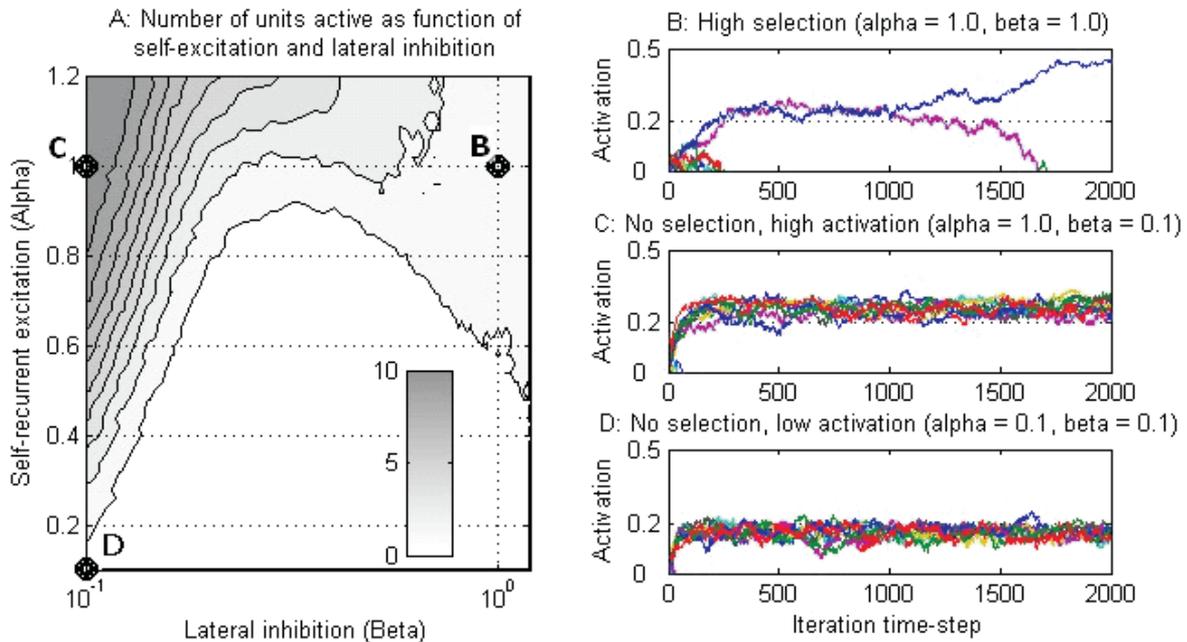


Figure 4 Profiles of the influence of the self-recurrency and lateral inhibition parameters on the activation of 10 activated units. (A) Two-dimensional contour map of the number of activated units (at 2,000 iteration time steps). Note that the values for the lateral inhibition are on a logarithmic scale. The three dots are parameter combinations that were used to illustrate activation trajectories. (B) Activation trajectories for high selection. Note that, at first, two (out of 10) units are most active, followed by one unit increasing in activation; a selection has taken place. (C) Activation trajectories with low inhibition, but high activation leading to no selection, but all units being above a 0.2 threshold. (D) Activation trajectories with low inhibition and low activation, leading to no selection and all units being below a 0.2 threshold.

σ . The parameter values vary between layers, but are chosen to minimize epiphenomenological dynamics that may obscure the comparison between the CQ and RS mechanisms. For example, in order to compare the models, each model should be able to produce an output. If the connection between the activation and selection layers is too weak, the CQ model, but not the RS model, produces an output; if the response suppression is too weak, the RS model, but not the CQ model, produces an output. Each layer can be interpreted as representing a set of brain areas that have been implicated in memory activation, memory selection, and action selection.

Figure 4 shows the influence of the parameters α and β on the activation of the representations, which eventually govern the overall system behavior. Figure 4A shows, for a variety of values for α and β , the average number of representations (out of 10, $I_i = 0.33$ for all) that are still activated above a fixed threshold (0.2, as used in previous work) after 2,000 iterations

over 100 simulation runs ($\sigma = 0.1$). Figures 4B–D show, for three points from Figure 4A, the corresponding activation trajectories. As can be seen, increasing the self-recurrent excitation will increase the number of representations that are still active above the threshold at the end of the simulation. However, this is only true for low levels of inhibition, as an intermediate level of inhibition dampens the overall activation, putting an upper bound to the total number of activated representations. With high inhibition, the system enters a winner-take-all selection mode.

4.1.1 Activation Layer The activation layer contains representations of the words that are to be recalled. The memory representations may already be in an active state (they are in the short-term buffer; Davelaar et al., 2005) or are (re)activated by a cue, which could be a category name (semantic cue, as in a verbal fluency task) or a contextual reference (episodic cue, as in a list

memory task). In this article, only cued activations are considered. The cue activates the first 10 items in the memory layer range from 0.28 to 0.37 (0.01 difference between units). The activated representations compete weakly with each other. In the simulations, zero-mean Gaussian noise is added to the activations of the memory units ($\alpha = 0$, $\beta = 0.1$, $\sigma = 0.1$).

4.1.2 Selection Layer The selection layer (modeled after Usher & McClelland, 2001) contains representations that receive weighted input, $W_{ms}F(x)$, from those in the activation layer, where W_{ms} represents the connection weight between the activation and the selection layers. The activated representations compete strongly ($\beta = 1.0$) with each other, and have moderate ($\alpha = 1.0$) self-connections. Both the activation and selection layers could be implemented in pre-frontal areas, such as the VLPFC (Badre et al., 2005).

4.1.3 Output Layer The output layer contains representations that receive weighted input, $W_{so}H[F(x),0.4]$, from those in the selection layer, with $W_{so} = 2.0$ and $H[F(x),0.4] = 1$, when $F(x) > 0.4$, 0 otherwise. The units are inert with no dynamics other than activation decay after a unit in the selection layer has provided a pulse to the output layer. This allows for obtaining an exact point in time in which a response is initiated. The connection between the selection and the output layers represents the “direct pathway” or the “selection pathway” (Gurney, Prescott, & Redgrave, 2001) in the basal ganglia.

4.1.4 Inhibition Layer The inhibition layer contains representations that receive input, $H[F(x),0.4]$, from those in the selection layer. To maximize the influence of the inhibition layer on the retrieval dynamics and thereby make the whole system sensitive to differences related to the RS and CQ mechanisms only, units in the inhibition layer have strong self-excitation and do not inhibit each other ($\alpha = 2.0$, $\beta = 0$). This implements a form of output buffer in which all selected items are maintained throughout the retrieval phase and prevents perseverations and repetitions. The neural substrate of the inhibition layer is assumed to include subcortical areas of the basal ganglia that are in the “indirect” or “control” pathway. In addition, the self-

recurrent connection may be interpreted as including cortical projections to the subthalamic nucleus. Recent work by Frank (2006) and Frank, Sherman, and Scheres (in press) suggests that a critical function of the subthalamic nucleus is to prevent (too early) responding.

4.2 Simulation 1: RS Versus CQ Mechanisms

In the first simulation, the weighted output of the inhibition layer, $W_{inh}F(x)$, is sent to the selection layer or the memory layer to implement the RS and CQ models, respectively. Figure 5 shows a noiseless simulation of both implementations. As can be seen, of the 10 activated representations in the memory layer, only eight produce a response in the RS model, whereas all produce a response in the CQ model. Figure 6(A) presents a comparison of the first recall latency and the subsequent IRTs for both noiseless simulations. The results are striking. The RS model shows increasing IRTs throughout the retrieval phase, whereas the CQ model shows a gradual decrease in IRTs. The RS model produces increased IRTs because, with each response, the responsible unit in the activation layer continues to send activation to the corresponding unit in the selection layer. As this unit is inhibited, the next-highest unit will win the competition, but as it receives less activation, it takes longer to reach the response threshold. This continues throughout the recall phase, leading to ever-increasing IRTs. The CQ model, however, reduces the number of activated units in the activation layer, thereby reduces the competition in the selection layer. With less competition, the units in the selection layer can more readily reach the response threshold, resulting in ever-decreasing IRTs.

The models were run with noise ($\sigma = 0.1$) in the activation layer. To obtain a wide range of total recall, 1,000 simulations were run with $W_{ms} = [1.4 : 2.0, \text{step } 0.1]$, which is justified under the assumption that the activations of the memory representations are modulated by attentional (Usher & Davelaar, 2002) or motivational factors. Despite this range, the CQ model produced all 10 responses in 99.2% of the simulations. Figures 6B and C show the IRTs as a function of the total recall and the recall interval for the RS and CQ models, respectively. These figures mimic the noiseless results in Figure 6A. Human behavioral data are in accordance with the results of the RS

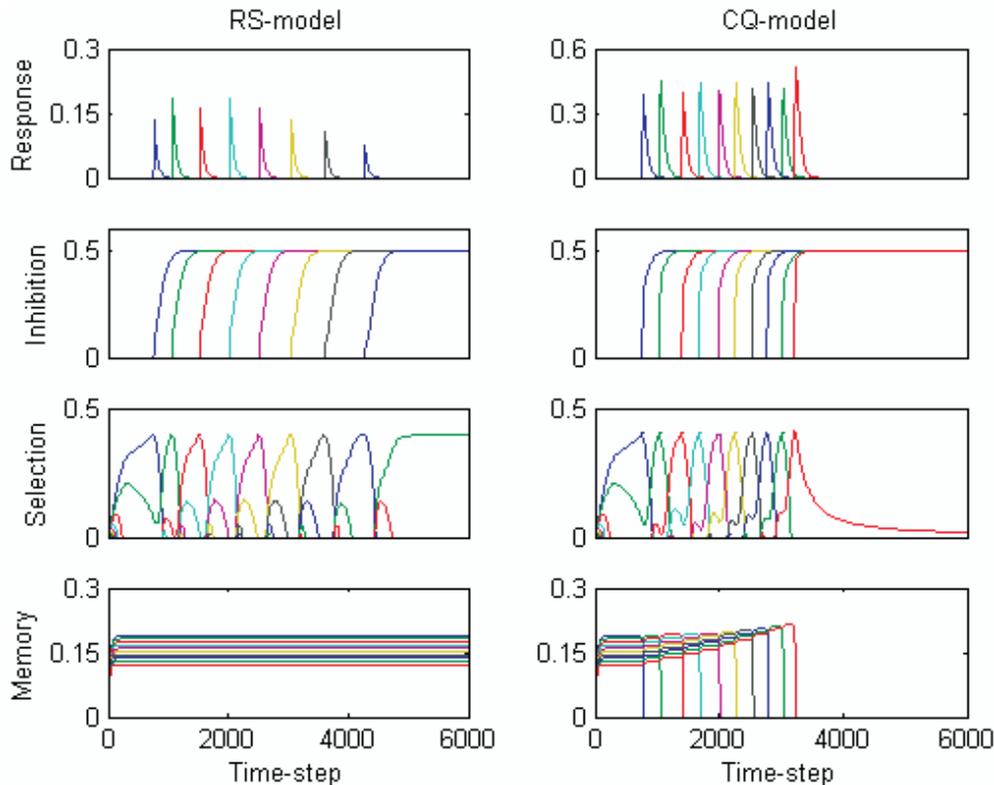


Figure 5 Activation trajectories of the RS and CQ models. Activation between layers goes from “memory” to “selection” to “output”, with “inhibition” suppressing the activation at either the “memory” (CQ model) or the “selection” (RS model) layer. Note that the RS model does not produce all responses (eight spikes in the output layer), whereas the CQ model does. Note also that the activation in the memory layer for the CQ model shows the signature of the CQ mechanism (after an output the corresponding trace is inhibited, which causes increase in activation for the remaining traces).

model (e.g. Murdock & Okada, 1970; as shown in Figure 2).

The results of Simulation 1 are expanded with additional analyses in the Appendix, in which the slope of the IRT profile is compared in four model implementations: RS and CQ models with (RS^A , CQ^A) and without (RS^0 , CQ^0) gradual decay of response suppression. The analyses reveal that the RS models produce the positive slope of the IRT function more often than the CQ models and that this difference between RS and CQ models is smaller if the response suppression decays gradually during the retrieval phase.

4.3 Simulation 2: Set size Effects

Given that the RS model captures the basic profile of IRTs, the question arises whether it also shows the critical sensitivity to set size. Rohrer and colleagues

(Rohrer, 1996, 2002; Rohrer and Wixted, 1994; Rohrer et al., 1995, 1999; Wixted & Rohrer, 1993) have shown that the distribution of recall latencies reflect the size of the memory set, which could be utilized to identify loci for memory deficits in patient populations. To this end, 1,000 simulations of the RS model were run with 5 or 15 memory units being activated (the five highest activated units were used in both set sizes, and therefore any effect on first recall latencies can only be attributed to set size). Inhibition in the selection layer was lowered to $\beta = 0.8$ and $W_{ms} = 4.0$. As can be seen in Figure 7, the RS model is sensitive to differences in set size (average recall of 0.918 and 0.307 for setsize 5 and 15, respectively). In particular, the model is slower when the memory set is larger. The reason for this is that with larger set size, more items activate units in the activation layer, and therefore in the selection layer. The overall increase in competing repre-

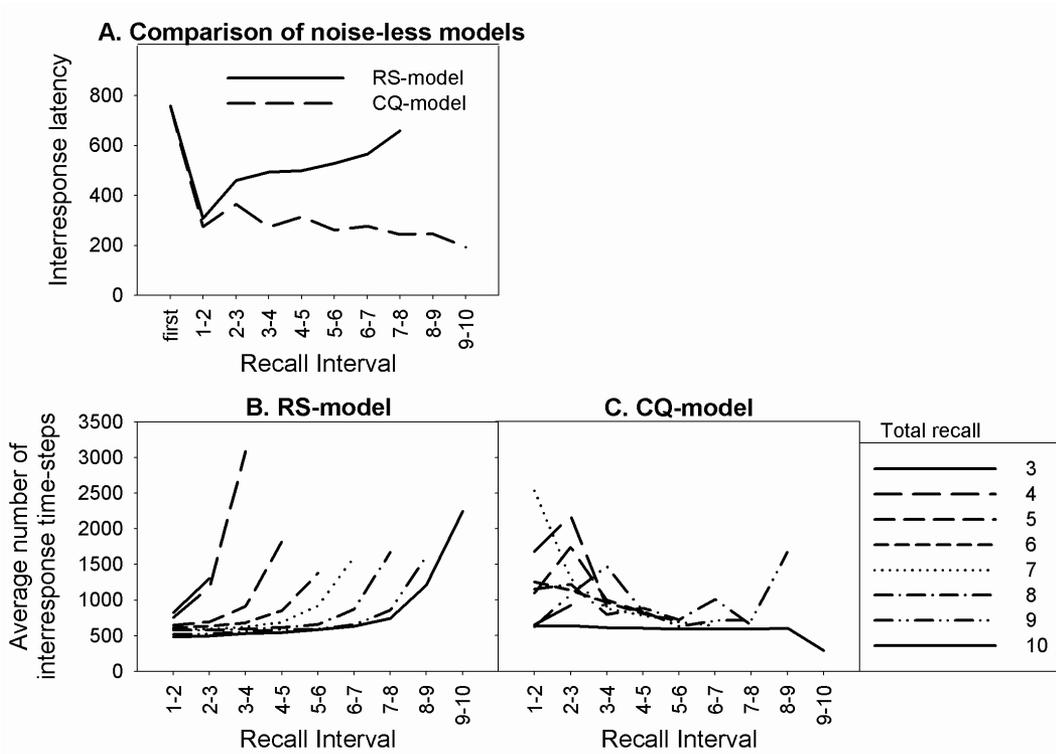


Figure 6 Results of Simulation 1. (A) First recall latencies and IRTs for the noiseless RS and CQ models. Note the increase in IRT for the RS model and the decrease in IRT for the CQ model. (B) IRTs as a function of total recall and recall interval for the RS model. (C) IRTs as a function of total recall and recall interval for the CQ model.

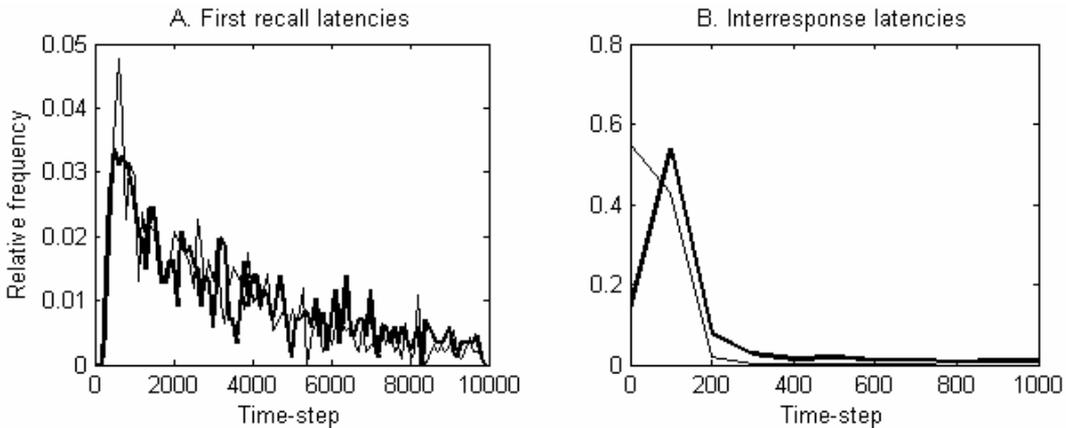


Figure 7 Results of Simulation 2. Comparison of small (thin line) and large (thick line) memory set sizes. (A) First recall latency distributions. The two distributions overlap considerably. (B) IRT distributions.

sentations leads to a slower rise in activation for each of the units in the selection layer. This continues throughout the recall phase, leading also to longer

IRTs. For the first recall latencies, the average latency was 3,164 time steps with size = 5, and 3,456 time steps with size = 15; for the IRT, 99 time steps with

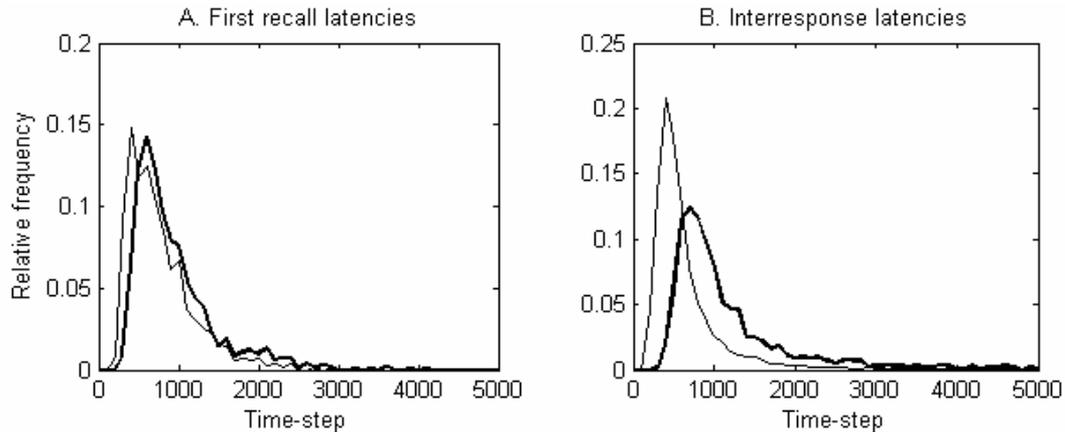


Figure 8 Results of Simulation 3. Comparison of normal selection threshold (thin line) and high selection threshold (representing Huntington's disease; thick line). (A) First recall latency distributions. The mean recall latency is slower for the HD simulation. (B) IRT distributions for normal and HD simulations.

size = 5, and 603 time steps with size = 15. The RS model not only captures the IRT profile, but is also sensitive to the size of the memory set, thereby providing computational validation of the theoretical analyses of Wixted and Rohrer (1994).

4.4 Simulation 3: Alzheimer's Dementia Versus Huntington's Disease

Patients with AD suffer from increased loss of neural tissue that represents long-term memory (see, for example, Fleischman & Gabrieli, 1999). This includes aspects of episodic and semantic memory. As in memory retrieval tasks, recall latencies reflect the size of the memory set, patients with AD are faster in retrieving items, but have a lower total recall. Simulation 2 has already shown the effect of set size and is not repeated here. HD is caused by a loss of striatal neurons, resulting in a decreased output to the thalamus (for a review, see Alexi et al., 2000). In the model, HD would affect the selection layer and is approximated by increasing the selection threshold, which slows down the retrieval process without affecting the memory set.

The RS model with selection $\beta = 1.0$ and $W_{ms} = 2.0$ was used (10 items; 1,000 runs) with the difference that the selection threshold was increased from 0.40 to 0.44. As can be seen in Figure 7, the HD simulation is indeed slower in retrieving items. For the first recall latencies, the average latency was 991 time steps (HD) compared with 825 time steps (baseline); for the

IRT, 1,262 time steps (HD) compared with 700 time steps (baseline). This contrasts with the results of patients with AD (set size effect).

4.5 Simulation 4: Non-Mnemonic Clustering

During model exploration, an interesting form of clustering was observed. When the weight between the activation and selection layers is high (or the inhibition in the selection layer is low), a number of representations may still be active when a representation is selected. As the selection layer is assumed to receive modulated input (for a model of attentional modulation of response selection, see Usher & Davelaar, 2002), the model links attentional and motivational factors to the speed of retrieval. However, this pattern could only be observed in the distributions if the bin size for the distributions was small enough.

An illustrative simulation of these parameters was conducted. Five items were activated in a simulation with two levels of activation–selection weight, $W_{ms} = [1.5, 2.0]$. Figure 9 presents activation trajectories of a single trial (with noise, $\sigma = 0.1$; the same random seed was used in both simulations), while Figures 10(A) and (B) present the recall latency distributions (from 1,000 simulation runs) with large (1,000 time steps) and small (100 time steps) bin size, respectively. It is immediately apparent that the model is able to produce a form of clustering that is independent of the mnemonic structure of the to-be-retrieved representations.

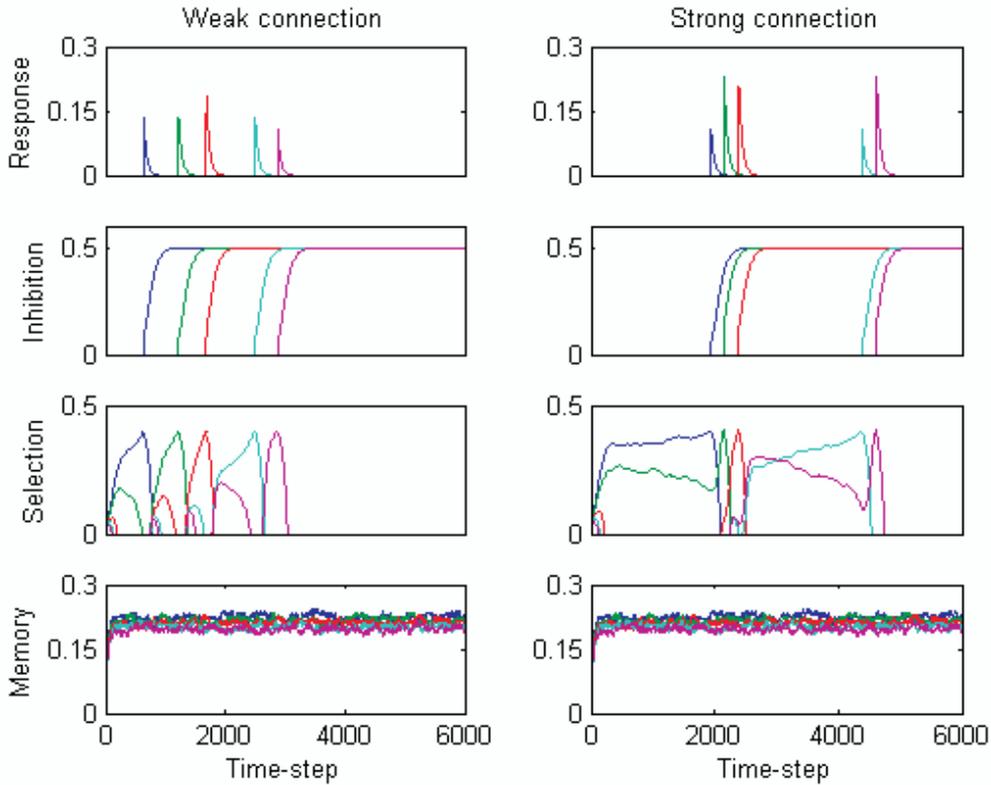


Figure 9 Activation trajectories of the RS model, with strong connection weight between memory and selection layer (left) compared to a baseline (right). Note the bursting behavior of the model despite equivalent mnemonic structure, which is a result of not having a full selection in the selection layer (two units are active simultaneously before the first response is made).

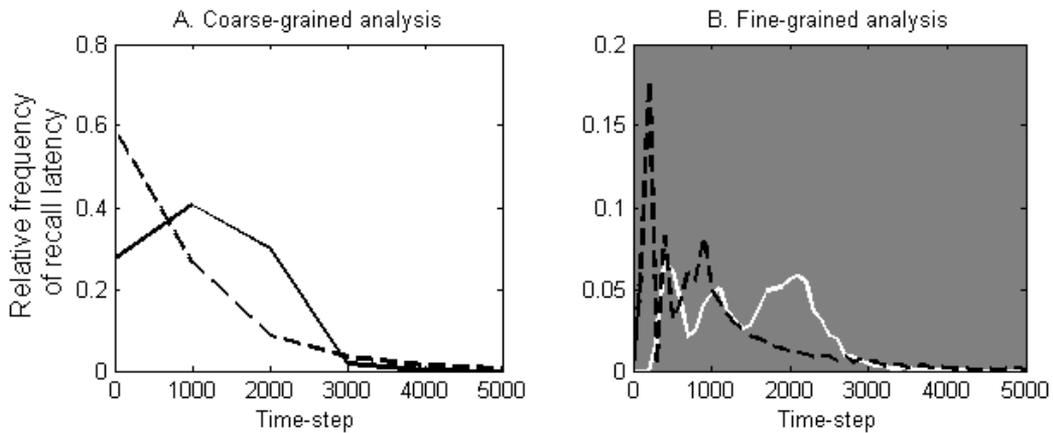


Figure 10 Results of Simulation 4. Recall latency distributions for weak (solid line) and strong (dotted line) W_{ms} : (A) coarse-grained bin size (1000 time steps); (B) fine-grained bin size (100 time steps). Note the faster latencies with increased strength and the appearance of clustering with small bin size.

In fact, there is no mnemonic structure. Instead, this type of clustering reflects the pressure that the repre-

sentations feel in the selection layer combined with the self-recurrency. The simulation also shows that in

order to observe this profile, an adequate temporal resolution is required. Therefore, it is recommended that in empirical studies, looking at retrieval latencies in free recall tasks, care is taken to obtain precise timing measurements. It remains an open question whether the clustering profile can be observed in neurobiological studies with animals.

5 Discussion

The aim of this article is to compare two commonly used mechanisms of sequential memory retrieval and to evaluate their ability to capture the IRTs in free recall. A generic model has been used in which the two mechanisms are implemented. Simulations have revealed that the RS mechanism provides a better overall qualitative match to the published data on retrieval latencies in a free recall paradigm than the CQ mechanism. A comparative analysis of four model implementations (RS^Δ , CQ^Δ , RS^0 , and CQ^0 ; see Section 4.2 and Appendix) reveals that the RS models produce the positive slope of the IRT function more often than the CQ models, and that this difference between RS and CQ models is smaller if the response suppression decays gradually during the retrieval phase. This latter mechanism makes the CQ^Δ model (CQ with gradual decay of suppression) more similar to the RS models (RS^0 and RS^Δ). The RS mechanism has subsequently been tested on its sensitivity to the size of the activated memory set and its success is promising for further research in the specific neural implementation of the mechanism. The use of recall distributions has been shown to provide a better understanding of the memory deficits in AD and HD (Section 4.4) and to resolve theoretical debates that have focused mainly on total recall as the dependent variable. The RS model captures the differences in recall latency distributions between patients with AD or HD. The important utility of IRT analyses in this domain justifies further developments and analyses of the computational architecture to quantitatively account for the neuropsychological data. The implications of the present model comparison follow next.

5.1 Implications for Dynamical Models of Memory

The CQ mechanism has had great impact on the memory literature, but the simulations show that for pro-

files of IRTs in free recall, the mechanism does not fare well compared with the RS mechanism, unless additional assumptions are in place. Even though the evaluation between the two types of models in Simulation 1 may imply that CQ models are inappropriate for free recall, this does not invalidate CQ models. In fact, given the neurophysiological support for CQ dynamics, it is conceivable that CQ dynamics may play a dominant role at the response level and less so at the memory level. Note that the successes of CQ models are found in tasks that require execution of a well-learned motor program, whether it is drawing geometric shapes (Averbeck et al., 2002), or pronouncing words (Hartley & Houghton, 1996). As such, the use of CQ dynamics may be more related to the unpacking of chunks, where the content of chunks could be words (as in serial recall), movements (as in typing or spelling; Glasspool & Houghton, 2005; Hartley & Houghton, 1996) or even other chunks (as in planning; Cooper & Shallice, 2000). Within the literature on serial recall, not much emphasis has been placed on retrieval latencies (but see Farrell & Lewandowsky, 2004; Kahana & Jacobs, 2000). Nevertheless, in a production system, Anderson and Matessa (1997) explicitly modeled the IRTs as the unpacking of chunks, with equal time needed for each item within a chunk and with additional time to move from one to the other chunk. In a follow-up paper (Anderson et al., 1998) some aspects of free recall were modeled, but not retrieval latencies.

An important assumption of some CQ models is that after it has made a response, the representation leading to an overt response is inhibited. This response suppression decays gradually over time, leading to reselection. Despite this critical feature (i.e. decay of response suppression), the neurophysiological data that support competitive queuing in general (Averbeck et al., 2002) do not show this pattern. Although repetitions in movements in making geometric shapes are unlikely, decay of response suppression seems not to be needed to account for repetition errors in serial recall (Botvinick & Plaut, 2006; Farrell & Lewandowsky, 2002) and is even directly contested empirically (Duncan & Lewandowsky, 2005).

Although the modeling work in this article has highlighted a critical difference in employment of CQ and RS mechanisms, the model itself requires extensions in order to account for a variety of other memory-related findings in sequential retrieval. These include semantic (Davelaar et al., 2006) or episodic clustering,

strategic retrieval (idiosyncratic cues: Gronlund & Shiffrin, 1986; output order: Dalezman, 1976) or semantic dementia (Forde & Humphreys, 2002). However, the evaluation presented here was of a qualitative nature and was aimed at providing a means to justify the choice of mechanism used, with analyses of retrieval latencies as the main dependent measure that constrains the modeling choice. Future work may merge the current generic model into other more specified models of recall memory to address the other findings.

Recent work on recall memory has incorporated new theoretical approaches to working memory (Davelaar et al., 2005, 2006). In this research, the content of working memory is defined as the activated part of long-term memory (e.g. Cowan, 2001) and the modeling work was focused on the encoding of information in episodic memory. In that work, episodic memory is the matrix of connection weights between a contextual system (related to the medial-temporal lobe) and the cortical long-term memory system. In these models of free recall, the retrieval process was approximated by a selection and recovery phase, akin to that of global memory models. The dynamical model of retrieval presented in this article provides the back-end to those previous models of encoding. It is noteworthy that the level of inhibition in the retrieval component produces some form of clustering. A few items are reported in rapid succession, but are not interconnected. The size of this cluster is directly related to the number of items that are active simultaneously in the selection layer (see Figure 9), which could potentially be used in empirical investigations on working memory capacity (reviewed in Cowan, 2001).

5.2 Implications for Dynamical Models of Sequential Selection

As mentioned in the introduction, studying free recall may be a more useful paradigm than serial recall to study the properties of sequential retrieval from memory, as serial recall places an external constraint on the inherent selection mechanism. However, under the assumption that this is indeed the case, the question then arises how the inherent selection mechanism is influenced by this external constraint of output order. In other words, assuming that an RS mechanism underlies all types of sequential behavior, how does the need for a particular correct serial order affect the operations of the selection? A partial answer to this is

revealed by a closer look at how the slopes of the RS models are affected by the parameters (see the Appendix). In free recall, a positive slope of the IRT function is observed, whereas a negative slope is observed for serial performance such as typing (Salthouse, 1986). To capture a negative slope (i.e. ever faster responding) in the RS^0 model, the higher self-recurrency would normally lead to more items becoming active, but the higher lateral inhibition prevents this (together with the lower weight from the activation layer to the selection layer). With every retrieved item (in the selection layer) being inhibited after reaching a threshold, the next-highest item in activation shoots up. As the overall competition within the selection layer decreases during the retrieval phase, the IRT between successive retrievals becomes shorter. The same dynamics occur in the RS^A model, where because of the gradual decay of response suppression, the overall selection mechanism needs to be liberal and with a low response threshold. Interestingly, the amount of self-recurrency is also lower.

In a recent computational study of response selection, Usher and Davelaar (2002) have proposed that the parameters for self-recurrency and lateral inhibition in a selection layer are modulated by norepinephrine (NE) in a positive manner (both parameters are positively correlated with the level of NE). Application of that study to the current discussion produces the hypothesis that the instruction to perform in a particular ordered sequence leads to a modulation in NE, which translates into a higher or lower self-recurrency and lateral inhibition in the selection layer, thereby producing an observable negative or positive slope in the IRT function of actions, respectively. Note that the covariation between self-recurrency and lateral inhibition as a function of the IRT slope is negative for the CQ models, critically distinguishing the CQ and RS models on their reliance on these two parameters. Granted that the analyses in the Appendix are not full proof and require further extensions incorporating noise and a consideration of other types of behavior (e.g. error profiles), the current hypothesis is not inconsistent with previous work on response selection (Usher & Davelaar, 2002) and memory for free recall (Davelaar et al., 2005). Linking the idea of neuromodulation to the different IRT profiles within a single system may be a useful vehicle for the investigation in the control of selection among competing behavioral alternatives. Recent work has linked NE and the locus coeruleus with

conflict monitoring and attentional control (Botvinnick, Braver, Carter, Barch, & Cohen, 2001) and with the balance between exploitation–exploration behavior in animals (e.g. Cohen, Aston-Jones, & Gilzenrat, 2004).

5.3 Relation to Other Psychological Domains

The evaluation of the memory retrieval dynamics could inform other domains within psychology, such as decision-making, neuropsychological assessment, and language/sequence learning/production. Recent work in the field of decision-making makes increasingly more use of memory theories (e.g. Dougherty, Gettys, & Ogden, 1999). For example, when a physician is generating a number of hypotheses about a possible diagnosis, the information about the symptoms is used to cue the memory system and search the activated part of memory to produce a number of likely diagnoses. This type of research could benefit from a deeper understanding of how the retrieval dynamics are affected by such variables as the number of possible hypotheses (search set), the motivation of the physician (focus of attention for selection), and the dual-task situation (affects the retrieval speed).

In the domain of neuropsychological assessment, a recent study (Scahill, Hodges, & Graham, 2005) has shown that current episodic memory tasks fail to differentially diagnose patients with AD from patients with semantic dementia (SD, a progressive neurodegeneration of the temporal lobe, characterized by loss of semantic memory). Of course, patients with SD can be distinguished from patients with AD using semantic tasks, but it is not certain yet whether episodic memory tasks may contribute in the differential diagnosis. Rohrer et al. (1999) argued that recall latencies contain information that separates patients with AD from patients with HD. The search for alternative (and additional) methods of differentiating between these two groups of patients can be facilitated by neurocomputational models of the type presented in this article, in which certain components have a clear neurological counterpart. Using simulations of lesions, they can create informed predictions to guide the development of more sensitive tests. This type of modeling extends beyond the measures related to sequential retrieval, and combined with other dynamical models of memory may prove to be of high value to clinicians.

Finally, Dominey (2005) reviewed his work on sequence learning and grammatical constructions.

Whereas the tasks focused more on the learning of sequences, the core component of the modeling comprises the learning of connections between internal states (retrieval cues) and the correct output response. The advantage of his model over standard recurrent networks is that the Dominey model iterates activation over multiple time steps and thereby is sensitive to the temporal structure of input sequences. In the same vein, the present model could be extended to produce actions (e.g. speech), and thus would require a mechanism to transform the activated abstract message into a sequence of function and content words.

5.4 Limitations and Extensions

The current model focuses on two proposed inhibition mechanisms involved in memory retrieval. To this end, the inhibition layer and the output layer have been oversimplified. This oversimplification was needed in order to reveal the consequences of either implementation. However, two pointers are given that may be of interest to those involved in modeling human behavior. First, adding within-layer inhibition within the inhibition layer leads the system to exhibit recycling or rehearsal behavior. This could be prevented by having the content of the inhibition layer function as retrieval cues (breaking the cycle), which echoes the use of retrieved cues in subsequent retrieval in global memory models, such as SAM. Secondly, the pulses in the output layer could be temporally extended through self-connections. These activation profiles could, in turn, activate subsequent selection mechanisms involved in sequencing actions at a lower level in the action hierarchy, such as speech production (or any other motor sequence), allowing examination of the relation between the information being searched and the speech rate, or speech confusions. These extensions go beyond comparing the two mechanisms and would have obscured the unique differences between them, which are relevant for both the psychological and non-psychological domains. Future work can take the RS model and investigate the implications of an output buffer to retrieval dynamics.

5.5 Conclusion

This article has compared two mechanisms by which simultaneously activated representations produce actions (here the recall of memoranda) in a sequential manner.

Although these two mechanisms have been used in a variety of models, they have not previously been compared directly with each other. By examining the dynamics of sequential selection through retrieval latencies, it has been shown that sampling-with-replacement captured human memory retrieval better than sampling-without-replacement. Both are further improved by including dynamics of response suppression that gradually fades. This evaluation provides the cognitive modeler with a tool with which to choose the components of a selection mechanism, and provides insights in the internal dynamics of a chosen implementation by addressing the patterns it produces given a certain parameter space. Finally, the analyses provide a handle to understand qualitative differences in sequential behavior through global modulation of critical parameters, which require further evaluation from empirical and computational studies in humans and animals.

Appendix

This appendix presents an analysis of four model implementations, the RS model and the CQ model, which differ in the locus of the post-response suppression and within each model; this suppression remains throughout the remainder of the retrieval phase or gradually decays. A qualitative analysis addresses the question of how central a particular pattern, in this case the slope of the IRT profile, is to a particular model and thereby provides a general answer to what type of implementation of sequential retrieval captures existing data on IRT profiles in a free recall paradigm best. Note that all four implementations exhibit sequential behavior and are compared to data from memory experiments. Different results may be obtained if the models are compared against data from motor planning. After a brief description of the computational method used, the results are presented and the Appendix closes with some evaluation of this method.

The analysis takes its inspiration from a recent paper by Pitt, Kim, Navarro, & Myung (2006). They presented a procedure called parameter space partitioning (PSP), which allows a computational modeler to analyze the parameter space of a given model. The procedure involves the modeler identifying a particular set of patterns and going through the multidimensional parameter space in search of those patterns. The full procedure was not used in this study and the reader is

referred to the paper by Pitt et al. for details on the procedure and how to use it. Briefly, the PSP procedure selects a set of parameters and evaluates the model's prediction (or pattern). The algorithm works such that it samples from each pattern at least once and uniformly. An interesting feature of this procedure is that it provides an estimate of the size of the parameter space that is occupied by a particular pattern. In other words, the procedure is able to address the question of how central a particular data pattern is to a given model, which is exactly what is needed in comparing CQ and RS models.

The inspiration from the PSP procedure was to go through the parameter space of the four models (CQ and RS models, with and without gradual decay of inhibition) and to address whether and how much the slope of the IRT function is influenced by (i) the structural differences between the models, (ii) the presence/absence of decaying inhibition, and (iii) the four chosen parameters. These parameters were chosen because initial explorations by hand showed that these parameters mattered most in the simulations in Section 4. These parameters were the connection weight between the activation layer and the selection layer, W_{ms} , the self-recurrency of the unit in the selection layer, α , the lateral inhibition between all pairs in the selection layer, β , and the threshold, θ , above which a response is made. In the models with decaying inhibition, the self-recurrency of the units in the inhibition layer was reduced to 1.2, which was shown to be a compromise between obtaining noticeable reactivation while preventing too many repetitions (perseverations). The models are called RS^0 and CQ^0 for those without decaying inhibition, and RS^Δ and CQ^Δ for those with decaying inhibition.

Procedure

Step 0. Explore all models by hand and define a parameter boundary and an initial set of parameters. The parameter vector used for the initial parameters was $P_0 = [W_{ms}; \alpha; \beta; \theta] = [2.0; 1.0; 1.0; 0.4]$. The boundaries were set to $W_{ms} = [1.2 : 2.2]$, $\alpha [0.8 : 2.0]$, $\beta [0.2 : 2.0]$, and $\theta [0.3 : 0.5]$.

Step 1. Patterns were defined by the slope of the IRT function and the number of items reported within a simulation run of 6,000 iterations. This yields 2 (slope > 0 , slope ≤ 0) \times 8 (three to 10 outputs) + (no output or

less than three outputs) = 17 patterns. The PSP program from Pitt et al. (2006) was used on the noiseless models with the initial set of parameters, P_0 , and default values. All models produced all 17 patterns. The PSP program produced values that indicate the volume of the region of all possible patterns occupied by a certain pattern (for details, see Pitt et al., 2006). Although the log(volume) favored the RS models, as measured with the number of times that the log(volume) is greater for the positive slope for a given total output ($RS^0 - 8/8$; $CQ^0 - 4/8$; $RS^\Delta - 8/8$; $CQ^\Delta - 3/8$), this result is not optimal as the models were run without noise (and the procedure is limited to non-stochastic simulation runs). This step also produced additional parameter sets (in matrix `xmcv` from the program), one for each of the 17 patterns, P_{psp1} .

Step 2. Parameter sets that were obtained under step 1 were used as initial parameters in a further search through parameter/pattern space. This guarantees that all regions will be sampled from. To speed up the simulations (each of the four simulations under Step 1 took 2.5 h of computing time) the number of patterns was reduced to three (slope > 0, slope ≤ 0, no output or less than three outputs). The program was also modified to store all unique combinations of parameters and the corresponding pattern.

As 17 PSP runs were carried out on each (noiseless) model; the averaged log(volume) was informative. The results are shown in Table A1, showing great consistency across the two steps. The RS models per-

form generally better than the CQ models, but this difference is greatly attenuated when a decaying of response suppression is implemented (for all comparisons, $p < 0.001$).

Volume estimates support the central thesis of this paper that when it comes to modeling IRTs, resampling is preferred over competitive queuing. However, as all models are able to produce all possible patterns, caution should be taken to prevent over-interpretation. As mentioned, all points that were being evaluated in the PSP algorithm were stored together with the results. This yielded more than 100,000 parameter sets for each model. Table A2 shows the number of parameter sets for each model against each pattern. Again, it can be observed that a positive slope is more central to the RS model than to the CQ model, and that with decay of response suppression implemented, this difference decreases.

Figure A1 shows the normalized distributions of relative frequencies for each model, for the valid patterns and for each of the four parameters. A uniform distribution would be a horizontal line with $y = 0.05$. This means that for all models the strength between the activation and selection layers does not affect the slope of the IRT function. Only the RS^0 model tends to be affected by this parameter. Note, however, that W_{ms} does seem to be relevant when it comes to clustering behavior in the output (see Simulation 4). The same holds for the parameter for the response threshold. The self-recurrency and the lateral inhibition in the selection layer greatly affect the slope of the IRT function,

Table A1 Result of the PSP algorithm under Steps 1 and 2. The percentages indicate the amount of the valid space (where the model produces slopes for the IRT function) and the percentage of that space that contains the pattern of positive slopes. Step 1 results are taken over all total output (eight data points). Step 2 results are averaged across 17 PSP runs with standard deviations in brackets.

		Model			
		Without decay		With decay	
PSP run		RS^0	CQ^0	RS^Δ	CQ^Δ
Step 1	Valid	54.9%	32.2%	17.2%	28.7%
	Slope > 0	85.0%	38.0%	83.9%	78.9%
Step 2	Valid	56.2% (3.6%)	37.8% (3.3%)	25.9% (2.6%)	34.2% (3.8%)
	Slope > 0	89.1% (5.8%)	42.5% (5.9%)	91.4% (7.4%)	77.3% (5.7%)

Table A2 Total number of parameter sets obtained in Step 2 for each model against each possible pattern.

	Model			
	Without decay		With decay	
	RS ⁰	CQ ⁰	RS ^Δ	CQ ^Δ
Invalid	34,357	43,409	49,980	43,878
Valid (slope ≤ 0)	22,498	36,861	16,377	17,966
Valid (slope > 0)	43,640	37,567	51,989	44,676

depending on the model. For RS⁰, positively sloped IRT functions have intermediate values of α and β , whereas for negatively sloped IRT functions the values of α and β are large. For CQ⁰, positively sloped IRT

functions have high α and low β (suggesting low selection), whereas for negatively sloped IRT functions, α is low and β is large (suggesting high selection). For RS^Δ, positively sloped IRT functions have

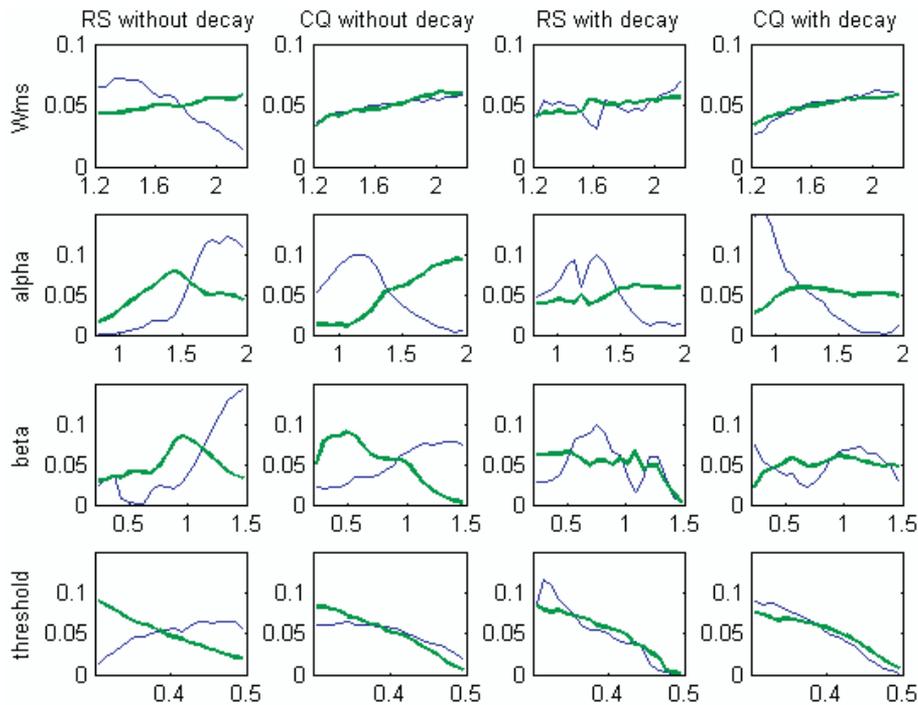


Figure A1 Normalized relative frequency distributions for each of the four parameters and for each of the four models. From top to bottom: W_{ms} , α , β , and threshold. The thick lines represent the distribution of the parameter values (shown on the x-axes) for which the corresponding model produced an IRT function with slope > 0. The thin lines represent the distribution of the parameter values for which the corresponding model produced an IRT function with slope ≤ 0. A uniform distribution would be $y = 0.05$. Note the strong influence of the self-recurrency (α) and lateral inhibition (β) parameters. Note also that models implementing decaying response suppression (two right columns) are less sensitive (flat line) to the values of the α and β parameters than the models that do not implement a decaying response suppression (two left columns).

undetermined values of α and β , whereas for negatively sloped IRT functions, the values of α and β are small. Finally, for CQ^A, positively sloped IRT functions have undetermined values of α and β , whereas negatively sloped IRT functions have low values of α and high values for β (suggesting high selection).

In summary, the values for α and β seem to suggest that the CQ models tend to settle into high or low modes of selection, producing negatively or positively sloped IRT functions, respectively, whereas RS models consistently keep the same mode of selection. These mode preferences are attenuated with decaying response suppression.

Limitations

Although this new procedure may prove to be a very useful addition to the modeler's toolbox, a number of simplifications have been applied here that should be noted. First, the models are stochastic models and the PSP assumes a stationary pattern space. To this end, the models were run without noise, which reduces the full model potential, but does allow for addressing the slope of the IRT function. Secondly, the number of patterns was greatly reduced compared with what the models can do. Besides having distributions of total recalls, the models are also able to show clustering behavior. This feature was not part of the pattern definitions, but was instead addressed in a full stochastic version in Simulation 4. Thirdly, default parameters for the PSP algorithm were used, and it can be argued that more optimal values are possible that allow better estimates of volume. Nevertheless, the volume results from Steps 1 and 2 were consistent and the algorithm produced a massive amount of data used in distributions of the parameters. Finally, only a subset of parameters was varied, limiting the size of parameter space. This should not be an issue, as the aim of this exercise was to compare models given the same four parameters that were found to be greatly affecting performance, some of which could be linked to brain damage.

These limitations aside, the procedure has given much more than an analysis by hand could have given. Future work may address the above problems more directly and give more principled answers to the question of how central the full (stochastic) pattern of IRT functions is to a given model.

Acknowledgments

The author would like to thank Michael Dougherty and Isaiah Harbison for stimulating discussions on memory retrieval, and Joanna Bryson and the reviewers for providing critical commentaries. Substantial parts of the simulation work were conducted while the author was at the University of California, San Diego.

Notes

- 1 The literature seems to be inconsistent in its definition of which post-Houghton (1990) models can be called CQ models. Here, the following definition is used: a CQ model is a model in which given the same context, the layer immediately preceding the selection (or sampling) layer (i) drives the selection, (ii) has multiple representations active, and (iii) has its activation profile altered as a consequence of the selection layer. This captures the notion of a queue that is altered with subsequent selection. Within this definition, models such as the Primacy model (Page & Norris, 1998) and the Start-End model (Henson, 1998), and also the Grossberg (1978b) model, are CQ models. Models of free recall, such as SAM (Raaijmakers & Shiffrin, 1980, 1981) do not alter the activation profile given the same context, and thus are not considered CQ models, but instead are here referred to as RS models.
- 2 The reader may also have noted that the timing signal itself could be implemented using a simple chaining mechanism. Whereas strong arguments have been raised against item-to-item chaining (Henson et al., 1996), all the models converged on using a temporal context signal that can be implemented using a simple chaining submodel (for an early CQ model of serial recall with random context change, see Burgess & Hitch, 1992). This issue has not been clarified in the CQ literature, and with the development of recurrent network models of serial recall that capture the error patterns (Botvinick & Plaut, 2006), a closer look at mechanisms of serial order seems warranted.

References

- Alexi, T., Borlongan, C. V., Faull, R. L. M., Williams, C. E., Clark, R. G., Gluckman, P. D., & Hughes, P. E. (2000). Neuroprotective strategies for basal ganglia degeneration: Parkinson's and Huntington's diseases. *Progress in Neurobiology*, *60*, 409–470.
- Anderson, J. R., Bothell, D., Lebiere, C., & Matessa, M. (1998). An integrated theory of list memory. *Journal of Memory and Language*, *38*, 341–380.
- Anderson, J. R., & Matessa, M. (1997). A production system theory of serial memory. *Psychological Review*, *104*, 728–748.

- Averbeck, B. B., Chafee, M. V., Crowe, D. A., & Georgopoulos, A. P. (2002). Parallel processing of serial movements in prefrontal cortex. *Proceedings of the National Academy of Sciences*, *99*, 13172–13177.
- Badre, D., Poldrack, R. A., Pare-Blagoev, E. J., Insler, R. Z., & Wagner, A. D. (2005). Dissociable controlled retrieval and generalized selection mechanisms in ventrolateral prefrontal cortex. *Neuron*, *47*, 907–918.
- Botvinick, M. M., Braver, T. S., Carter, C. S., Barch, D. M., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review*, *108*, 624–652.
- Botvinick, M. M., & Plaut, D. C. (2006). Short-term memory for serial order: a recurrent neural network model. *Psychological Review*, *113*, 201–233.
- 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. (1999). Memory for serial order: a network model of the phonological loop and its timing. *Psychological Review*, *106*, 551–581.
- Cohen, J. D., Aston-Jones, G., & Gilzenrat, M. S. (2004). A systems-level perspective on attention and cognitive control: guided activation, adaptive gating, conflict monitoring, and exploitation versus exploration. In M. I. Posner (Ed.), *Cognitive neuroscience of attention*. New York: Guilford Press.
- Cooper, R., & Shallice, T. (2000). Contention scheduling and the control of routine activities. *Cognitive Neuropsychology*, *17*, 297–338.
- Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, *24*, 87–185.
- Dalezman, J. J. (1976). Effects of output order in immediate, delayed, and final recall performance. *Journal of Experimental Psychology: Human Learning and Memory*, *2*, 597–608.
- Davelaar, E. J., Goshen-Gottstein, Y., Ashkenazi, A., Haarmann, H. J., & Usher, M. (2005). The demise of short-term memory revisited: empirical and computational investigations of recency effects. *Psychological Review*, *112*, 3–42.
- Davelaar, E. J., Haarmann, H. J., Goshen-Gottstein, Y., & Usher, M. (2006). Semantic similarity dissociates short- from long-term recency: Testing a neurocomputational model of list memory. *Memory and Cognition*, *34*, 323–334.
- Dominey, P. F. (2005). From sensorimotor sequence to grammatical construction: evidence from simulation and neurophysiology. *Adaptive Behavior*, *13*, 347–361.
- Dougherty, M. R. P., Gettys, C. F., & Ogden, E. E. (1999). MINERVA-DM: A memory processes model for judgments of likelihood. *Psychological Review*, *106*, 180–209.
- Duncan, M., & Lewandowsky, S. (2005). The time course of response suppression: no evidence for a gradual release from inhibition. *Memory*, *13*, 236–246.
- Farrell, S., & Lewandowsky, S. (2002). An endogenous distributed model of ordering in serial recall. *Psychonomic Bulletin and Review*, *9*, 59–79.
- Farrell, S., & Lewandowsky, S. (2004). Modelling transposition latencies: constraints for theories of serial order memory. *Journal of Memory and Language*, *51*, 115–135.
- Fleischman, D. A., & Gabrieli, J. (1999). Long-term memory in Alzheimer's disease. *Current Opinion in Neurobiology*, *9*, 240–244.
- Forde, E. M. E., & Humphreys, G. W. (2002). The role of semantic knowledge in short-term memory. *Neurocase*, *8*, 13–27.
- Frank, M. J. (2006). Hold your horses: A dynamic computational model for the subthalamic nucleus in decision making. *Neural Networks*, *19*, 1120–1136.
- Frank, M. J., Sherman, S. J., & Scheres, A. (in press). Understanding decision making deficits in neurological conditions: insights from models of natural action selection. *Philosophical Transactions of the Royal Society B: Biological Sciences*.
- Glasspool, D. W. (2005). Modelling serial order in behaviour: Evidence from performance slips. In G. Houghton (Ed.), *Connectionist model in cognitive psychology*. Hove: Psychology Press.
- Glasspool, D. W., & Houghton, G. (2005). Serial order and consonant-vowel structure in a graphemic output buffer model. *Brain and Language*, *94*, 304–330.
- Gronlund, S. D., & Shiffrin, R. M. (1986). Retrieval strategies in recall of natural categories and categorized lists. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *12*, 550–561.
- Grossberg, S. (1978a). Behavioral contrast in short-term memory: Serial binary memory models or parallel continuous memory models? *Journal of Mathematical Psychology*, *17*, 199–219.
- Grossberg, S. (1978b). A theory of human memory: Self-organization and performance of sensory-motor codes, maps, and plans. In R. Rosen & F. Snell (Eds.), *Progress in theoretical biology*, Vol. 5. New York: Academic Press.
- Gurney, K., Prescott, T. J., & Redgrave, P. (2001). A computational model of action selection in the basal ganglia. I. A new functional anatomy. *Biological Cybernetics*, *84*, 401–410.
- Hartley, T., & Houghton, G. (1996). A linguistically constrained model of short-term memory for words and non-words. *Journal of Memory and Language*, *35*, 1–31.
- Henson, R. N. A. (1998). Short-term memory for serial order: the start-end model. *Cognitive Psychology*, *36*, 73–137.

- Henson, R. N. A., Norris, D. G., Page, M. P. A., & Baddeley, A. D. (1996). Unchained memory: Error patterns rule out chaining model of immediate serial recall. *Quarterly Journal of Experimental Psychology*, *49*, 80–115.
- Hintzman, D. L. (1984). MINERVA 2: A simulation model of human memory. *Behavior, Research Methods, Instruments, and Computers*, *16*, 96–101.
- 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*. London: Academic Press.
- Kahana, M. J., & Jacobs, J. (2000). Inter-response times in serial recall: Effects of intraserial repetition. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *26*, 1188–1197.
- Murdock, B. B. (1982). A theory for the storage and retrieval of item and associative information. *Psychological Review*, *89*, 609–626.
- Murdock, B. B., & Okada, R. (1970). Inter-response times in single-trial free recall. *Journal of Experimental Psychology*, *86*, 263–267.
- Page, M. P. A., & Henson, R. N. A. (2001). Computational models of short-term memory: Modeling serial recall of verbal material. In J. Andrade (Ed.), *Working memory in perspective*. Hove, UK: Psychology Press.
- Page, M. P. A., & Norris, D. (1998). The primacy model: A new model of immediate serial recall. *Psychological Review*, *105*, 761–781.
- Patterson, K. E., Meltzer, R. H., & Mandler, G. (1971). Inter-response times in categorized free recall. *Journal of Verbal Learning and Verbal Behavior*, *10*, 417–426.
- Pitt, M. A., Kim, W., Navarro, D. J., & Myung, J. I. (2006). Global model analysis by parameter space partitioning. *Psychological Review*, *113*, 57–83.
- Raaijmakers, J. G. W., & Shiffrin, R. M. (1980). SAM: A theory of probabilistic search of associative memory. In G. Bower (Ed.), *The psychology of learning and motivation*, Vol. 14. New York: Academic Press.
- Raaijmakers, J. G. W., & Shiffrin, R. M. (1981). Search of associative memory. *Psychological Review*, *88*, 93–134.
- Rohrer, D. (1996). On the relative and absolute strength of a memory trace. *Memory and Cognition*, *24*, 188–202.
- Rohrer, D. (2002). The breadth of memory search. *Memory*, *10*, 291–301.
- Rohrer, D., Salmon, D. P., Wixted, J. T., & Paulsen, J. S. (1999). The disparate effects of Alzheimer's disease and Huntington's disease on semantic memory. *Neuropsychology*, *13*, 381–388.
- Rohrer, D., & Wixted, J. T. (1994). An analysis of latency and interresponse time in free recall. *Memory and Cognition*, *22*, 511–524.
- Rohrer, D., Wixted, J. T., Salmon, D. P., & Butters, N. (1995). Retrieval from semantic memory and its implications for Alzheimer's disease. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 1127–1139.
- Salthouse, T. A. (1986). Perceptual, cognitive, and motoric aspects of transcription typing. *Psychological Bulletin*, *99*, 303–319.
- Scahill, V. L., Hodges, J. R., & Graham, K. S. (2005). Can episodic memory tasks differentiate semantic dementia from Alzheimer's disease? *Neurocase*, *11*, 441–451.
- Usher, M., & Davelaar, E. J. (2002). Neuromodulation of decision and response selection. *Neural Networks*, *15*, 635–645.
- Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. *Psychological Review*, *108*, 550–592.
- Wixted, J. T., & Rohrer, D. (1993). Proactive interference and the dynamics of free recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *19*, 1024–1039.
- Wixted, J. T., & Rohrer, D. (1994). Analyzing the dynamics of free recall: an integrative review of the empirical literature. *Psychonomic Bulletin and Review*, *1*, 89–106.

About the Author



Eddy J. Davelaar holds M.Sc. degrees in biological health sciences and in psychology, both from Maastricht University in 1999. He holds a Ph.D. in psychology from Birkbeck, University of London. His main research interest is the control of memory and attention, using a variety of experimental and computational methods.