

An endogenous distributed model of ordering in serial recall

SIMON FARRELL and STEPHAN LEWANDOWSKY
University of Western Australia, Crawley, Australia

We introduce a distributed model of memory for serial order, called SOB, that produces ordered serial recall by relying on encoding and retrieval processes that are endogenous to the model. SOB explains the basic shape of the serial position curve, the pattern of errors during recall (including the balance between transpositions, omissions, intrusions, and erroneous repetitions), the effects of list length on the distribution of errors, the overall level of recall and response latency, and the effects of natural language frequency on recall performance. In addition, contrary to several recent suggestions, SOB demonstrates that distributed representations can support unambiguous recall, selective response suppression, and novelty-sensitive encoding.

There has been considerable progress in our understanding of short-term memory for serial order. Some 10 years ago, theories could only describe broad patterns of results, such as the serial position curve, without consideration of the underlying pattern of errors (e.g., Lewandowsky & Murdock, 1989). Current models, by contrast, can account for intricate details of the data, such as the delicate balance between omission, intrusion, and repetition errors during serial recall (e.g., Brown, Preece, & Hulme, 2000; Burgess & Hitch, 1999; Henson, 1998b; Lewandowsky, 1999; Page & Norris, 1998b).

Although discussion of these models has typically focused on underscoring their differences, some theoretical convergence has occurred, and two major assumptions are now shared by several models. First, the strength with which items are encoded is assumed to decrease across serial positions along a *primacy gradient* (e.g., Brown et al., 2000; Henson, 1998b; Lewandowsky, 1999; Lewandowsky & Murdock, 1989; Page & Norris, 1998b). Second, items that have been recalled are thought to be suppressed and are, thus, at least temporarily unavailable for report (e.g., Brown et al., 2000; Burgess & Hitch, 1999; Henson, 1998b; Lewandowsky, 1999; Lewandowsky & Murdock, 1989; Page & Norris, 1998b).

Emphasis in most models has been on the exact representation of order information, and the relative advantages and disadvantages of different representational schemes have been discussed at length (e.g., Henson, 1999). This emphasis has entailed the expense that the processes underlying serial recall have not been fully specified; in particular, the mechanisms underlying the two major assump-

tions above have often been implemented by parameters, rather than by specified processes (e.g., Brown et al., 2000; Henson, 1998b; Lewandowsky, 1999; Lewandowsky & Murdock, 1989). In response, this article presents a distributed network architecture that implements those assumptions within a common mechanism. This architecture, called SOB, uses the energy of its activation patterns to compute a primacy gradient and to determine the extent of response suppression. We call this energy-gating process *endogenous* because it considers only information shared between incoming items and the memory store; all information is, thus, intrinsic to the model. Because we are interested in the core ability of this endogenous process to explain serial recall, we explore a limited set of benchmark phenomena with a minimal set of assumptions, in preference to demonstrating the breadth and quantitative precision of the model.

This article first analyzes the predominant encoding and retrieval assumptions of existing models in some detail. The analysis reveals several shortcomings, which are taken as a stimulus for further theory development. We then present the SOB model, which addresses these hitherto ignored issues. We show that its endogenous processes can account for the shape of the serial position curve, the pattern of transpositions observed during recall, the relative proportion of different types of errors, the effects of list length, and the effects of natural language frequency. We then show by parameter sensitivity analysis that these predictions result from the intrinsic properties of the architecture. A final section compares SOB to other existing theories of memory and draws links to related models in other arenas, such as language processing and semantic priming.

TWO COMMON MECHANISMS FOR ENCODING AND RETRIEVAL

Two principal encoding and retrieval assumptions prevail among models of memory. At encoding, a primacy

This research was supported by a Large Grant from the Australian Research Council to the second author. We thank Murray Maybery for his comments on this manuscript. Correspondence concerning this article should be addressed to S. Lewandowsky, Department of Psychology, University of Western Australia, Crawley, W.A. 6009, Australia (e-mail: lewan@psy.uwa.edu.au).

gradient ensures that successive items are stored in memory with progressively decreasing strength. At retrieval, response suppression renders recalled items temporarily unavailable.

Primacy Gradient

Most models of serial recall assume that the quality of information available for retrieval of an item decreases across serial positions. This is typically achieved through one or two weighting parameters that reduce encoding strengths across successive items (e.g., Brown et al., 2000; Henson, 1998b¹; Houghton & Hartley, 1996; Lewandowsky, 1999; Lewandowsky & Li, 1994; Lewandowsky & Murdock, 1989; Page & Norris, 1998a, 1988b).

In most models, the primacy gradient is motivated by the need to predict a recall advantage for early list items and is thought to reflect a strategic decline in attention or a rehearsal tradeoff across serial position, rather than a principled assumption about memory processes (e.g., Brown et al., 2000; Lewandowsky, 1999; Lewandowsky & Murdock, 1989). Brown et al. justified the primacy gradient by appealing to the “intuition that each successive item . . . is progressively less ‘surprising’ or attention-demanding than the previous one” (p. 151). That intuition, in turn, Brown et al. considered to be consonant with the demands on an adaptively rational organism. A strength of this view is its potential generality, because it can accommodate other situations, outside the primacy context, in which to-be-encoded events may differ in the extent to which they are surprising or attention demanding (cf. Dennis & Kruschke, 1998). However, no formal or computational implementation has been offered to explain or justify this suggestion.

Some models, known collectively as competitive cuing models, go even further and rely exclusively on a primacy gradient to represent serial order (e.g., Houghton, 1990; Page & Norris, 1998a, 1998b). In those models, it is always the strongest item in memory that is retrieved next. On the additional assumption that each recalled item is suppressed and, thus, unavailable for further report, a primacy gradient will necessarily yield forward serial recall.

In their primacy model, Page and Norris (1998a, 1998b) motivated the existence of the primacy gradient in two ways. On the one hand, they suggested that it “might be thought of as resulting from association of each list item with some representation of the start-of-list context, with the strength of association decreasing with list position” (Page & Norris, 1998b, p. 763). On the other hand, in a connectionist implementation of their model (Page & Norris, 1998a; see also Appendix A of Page & Norris, 1998b), the primacy gradient arose from the decaying activation of a modulating node that was additionally weighted by the number of items already activated. Specifically, the modulating node was maximally activated at list onset and then decayed exponentially over time. That activation, in turn, was multiplied by a linearly decreasing function of the number of items

already encoded, thus yielding the final activation level of the next to-be-encoded item.

Although this fully specifies a mechanism for a primacy gradient, several problems and limitations can be identified. First, the mechanism relies on knowledge of the number of already-encoded items, but the way in which that knowledge is represented or obtained is not specified within the network architecture (see Figure 1 in Page & Norris, 1998a). Second, the mechanism applies only to generation of a primacy gradient; it has no potential for generality outside the context for which it was created. That is, it is difficult to see how the postulated decay of the modulating node could be sensitive to anything other than the passage of time, or how the weighting by the number of encoded items could be sensitive to anything other than encoding of additional items. It follows that the mechanism cannot be extended to handle known attentional effects in memory—for example, the fact that people favor a cue that is highly distinctive but relatively rare over one that is common but less distinctive (Dennis & Kruschke, 1998). Finally, the mechanism relies on localist representations; that is, it assumes that the modulating node is a discrete entity in a unique location in memory that is selectively subject to decay. Arguments in support of localist representations have been provided by Page (2000); however, in light of opposing views (e.g., Farrell & Lewandowsky, 2000), the exclusive linkage of the Page and Norris (1998a) mechanism to localist representations deserves mention.

Overall, we consider two limitations of current primacy mechanisms to be most relevant. The intuition offered by Brown et al. (2000), that less “surprising” items are encoded less strongly, is attractive because of its potential generality, but it suffers from a lack of formal specification. This difficulty is shared by the models by Henson (1998b), Lewandowsky (1999), and Lewandowsky and Murdock (1989). Conversely, the fully specified primacy mechanism proposed by Page and Norris (1998a) offers no possibility of extension to other situations in which encoding strengths may differ between items.

In response to these limitations, we present a model below in which item storage is sensitive to the current state of memory. This permits endogenous computation of a primacy gradient, without the limitations of the mechanism offered by Page and Norris (1998a). The endogenous process is of considerable generality and thus provides a computational implementation of the broad intuition offered by Brown et al. (2000).

Response Suppression

Response suppression refers to the assumption that each recalled item is temporarily suppressed and unavailable for further report. In the primacy model (e.g., Page & Norris, 1998a, 1998b) and other competitive cuing models, in which the strongest or most active item is reported at each recall attempt, suppression is essential for maintaining serial reproduction. Without suppression, only the first

(and strongest) item would be produced at each successive retrieval attempt.

Other models rely on suppression to account for the frequent occurrence of pairwise transposition errors (recall *ACBD* instead of *ABCD*) and the relative infrequency of erroneous repetitions (e.g., Brown et al., 2000; Henson, 1998b). In the few cases in which erroneous repetitions (e.g., *ABCA* instead of *ABCD*) do occur, they are typically separated by three or four intervening retrievals, with immediate repetitions being exceedingly rare (Vousden & Brown, 1998). This is compatible with the notion that an item, once recalled, is suppressed and that the suppression gradually wears off while other items are recalled.

Response suppression is also implied by the pervasive finding that people are reluctant to repeat themselves during recall, even when a list contains repetitions (e.g., Henson, 1998a). This deficit for report of a repeated item occurs even in conditions in which people are told to expect repetitions in advance and in which people can detect repetitions, if instructed to do so, more than 85% of the time (Henson, 1998a). Finally, the repetition deficit persists when guesses—which might be biased against reporting repetition—are eliminated from consideration (Henson, 1998a). This further supports the notion of mandatory and automatic response suppression.

Finally, response suppression has been used to model recency (e.g., Lewandowsky, 1999; Lewandowsky & Li, 1994; Lewandowsky & Murdock, 1989). Process explanations of recency in serial recall² tend to fall into two broad classes, referred to here as *edge effects* and *response suppression*. Edge effects occur because the terminal item, which has no neighbors beyond the end of the list, cannot be involved in a transposition in more than one way. The reduced frequency of transpositions necessarily results in recency. This mechanism principally contributes to recency in the primacy model (Page & Norris, 1998b), OSCAR (Brown et al., 2000), and the model by Burgess and Hitch (1999). One limitation of edge effect explanations is that recency is tied to the truncation of possible transpositions by the list boundary. By implication, given that few transpositions involve nonadjacent items, edge effects cannot predict any but the smallest amount of recency beyond the terminal item. This runs counter to the observation that recency may be moderately large for the penultimate or antepenultimate item even with visual presentation (Madigan, 1971; Watkins & Watkins, 1977). Another limitation of edge effect explanations is their inherent symmetry. Without additional assumptions, edge effects are equal for both list boundaries, and this necessarily predicts a symmetry of primacy and recency that is absent in the data (cf. Burgess & Hitch, 1999).

Explanations based on response suppression, on the other hand, exploit the fact that as more and more items are recalled—and hence, suppressed—fewer response alternatives remain, thus facilitating choice of the correct item. Henson (1998b) acknowledged the contribution of response suppression to recency in his SEM model.

Lewandowsky and Murdock (1989) relied entirely on response suppression to produce recency. Unlike edge effects, response suppression can handle recency that extends further into the list, because the size of the response set continually decreases across serial position.

It is noteworthy in this context that Page (2000) has claimed that distributed representations cannot accommodate selective suppression of individual items. This claim has already been called into question by the models presented by Lewandowsky (1999), Lewandowsky and Farrell (2000), and Lewandowsky and Li (1994). Here, we take the distributed approach further by presenting a model that includes item-specific suppression in addition to novelty-sensitive encoding. The model is derived from the brain-state-in-a-box (BSB) model of J. A. Anderson, Silverstein, Ritz, and Jones (1977). We therefore call the model SOB, for “serial-order-in-a-box.”

SOB: A DYNAMIC DISTRIBUTED MODEL OF ORDERING IN SERIAL RECALL

Overview

SOB inherited its basic architecture and aspects of the nonlinear dynamics that govern each study and retrieval event from the BSB model. Briefly, the model assumes that each item in memory is represented by a list (or *vector*) of features. Items are encoded in memory by adding their representations to a common weight matrix that is fully interconnected. That is, the core architecture of the model consists of connections between each feature of an item and all its other features that are superimposed onto items already presented.

In addition, and most critically, retrieval from the model involves nonlinear iterative dynamics. Memory is probed by presenting cue vectors to the weight matrix. In contrast to many other networks, the first response of the model is not taken to be its final answer. Instead, the output is fed back into the weight matrix across multiple iterations until a stable state, known as an *attractor*, is reached. In the case of a correct recall, that final attractor state will be identical to the target item, and in the case of incorrect recall, the final attractor will be different from the target.

We now provide an overview of the architecture and dynamics in turn, before presenting the details of the simulations.

Architecture. SOB uses a linear autoassociative network with a Hebbian learning rule (e.g., O’Toole, Diefenbacher, Valentin, & Abdi, 1994). At study, each item is first associated with itself by forming the outer product of its vector representation and its transpose. The outer product is then added to an excitatory weight matrix that stores all studied items. Successive list items are encoded with progressively less strength, as determined by the *energy* (described later) associated with each item.

Items are retrieved by cuing the weight matrix with a randomly constructed starting vector that initiates the nonlinear dynamics. The representation of the recalled item

is then attenuated to model response suppression. This is achieved by subtracting the autoassociation of the recalled item from a separate matrix of *inhibitory* weights. Excitatory and inhibitory weights are added together to determine the nonlinear retrieval dynamics.

Nonlinear dynamics. SOB has the key property that each retrieval event involves nonlinear iterative dynamics. At each iteration, the output obtained in response to cuing of the weights is fed back into the matrix for renewed cuing until a stable state, known as an attractor, is reached. Under the conditions of orthogonality assumed here, those attractors include all learned items plus a number of additional *spurious* attractors.

At each iteration of this dynamic process, the state of the network is characterized by a vector whose elements are constrained to lie between the values $+1$ and -1 . By implication, all possible states of the network lie within a hyperspatial box (hence “brain-state-in-a-box”) whose vertices consist of binary vectors (e.g., $+1, -1, \dots$). For reasons described below, all attractors are necessarily vertices of the hyperspatial box.

Interaction between architecture and dynamics. To model serial recall, list items are first encoded in the autoassociative weight matrix. At recall, if the randomly constructed starting vector falls within the basin of attraction surrounding the correct response (e.g., the first list item during the first recall attempt), SOB responds with the appropriate item. If the starting vector falls into a different basin of attraction, it will reach another attractor representing either a different list item or, in the case of a spurious attractor, an extra-list intrusion.

Despite the random nature of the starting vector, SOB can model accurate recall because successive list items are encoded with progressively decreasing strength and because recalled items are suppressed. In consequence, at any stage during recall, the correct response is most likely to have associated with it the largest basin of attraction, and the random starting vector is most likely to fall within that basin.

Network Architecture and Weight Update

The simulations used vectors of 256 units to represent list items and responses. Units were fully interconnected by two separate weight matrices, \mathbf{W} and \mathbf{A} (each of dimensionality 256×256), which represented excitatory memory connections and inhibitory weights, respectively. Thus, each unit was connected to every other unit with a weight from \mathbf{W} and one from \mathbf{A} . These weights were set to $\mathbf{0}$ at the outset, and \mathbf{A} was also reset to $\mathbf{0}$ before each list was recalled. The items serving as a vocabulary on each study–test trial were randomly sampled without replacement from the possible space of 256 Walsh vectors. Walsh vectors are binary vectors (i.e., $-1, +1, \dots$) that are mutually orthogonal (e.g., Golubov, Efimov, & Skvortsov, 1987). Mutual orthogonality implies that the similarity between any two vectors in the set is exactly nil (i.e., pairwise dot product or cosine of zero).

Pretraining. To reflect preexperimental knowledge that subjects bring to bear on serial order tests, SOB was pre-trained on a randomly chosen subset of 50 of the Walsh vectors. Pretraining consisted of Hebbian learning of excitatory weights with a constant encoding strength η_p (set to .001 throughout):

$$\mathbf{W}_k = \mathbf{W}_{k-1} + \eta_p \mathbf{f}_k \mathbf{f}_k^T, \quad (1)$$

where \mathbf{f}_k is the vector representing item k , and $\mathbf{f}_k \mathbf{f}_k^T$ its autoassociation formed by taking the outer product of the item and its transpose. Each of the 50 pretraining items was presented and encoded 20 times.

Study of list items. After pretraining, a subset of the pretrained items was selected for presentation as list items. Similar to pretraining, the autoassociation of each list item was added to \mathbf{W} , using Hebbian learning during a single presentation (see Equation 1). Unlike pretraining, the encoding strength for list items (η_k) was computed anew for each item, using the energy dynamics described below. The inhibitory matrix, \mathbf{A} , was not involved during study.

Recall. At recall, a cue vector, \mathbf{f}' , was formed by randomly setting its elements to $+1$ or -1 and then normalizing its length to .0001. The initial state of the network (represented by the vector \mathbf{x}) was then obtained for time $t = 1$ by cuing the sum of both weight matrices:

$$\mathbf{x}(t) = (\mathbf{A} + \mathbf{W})\mathbf{f}'. \quad (2)$$

Equation 2 represents conventional linear network retrieval processes. The state vector \mathbf{x} then migrated toward an attractor, using the iterative nonlinear dynamics given by Equation 3:

$$\mathbf{x}(t+1) = G[\beta \mathbf{x}(t) + \varepsilon(\mathbf{A} + \mathbf{W})\mathbf{x}(t)], \quad (3)$$

where $\mathbf{x}(t+1)$ is the new state of the network at time $t+1$, and β and ε are invariant weighting constants. The second term on the right-hand side of the equation represents feedback to the units in the state vector through both sets of weights (inhibitory weights in \mathbf{A} and excitatory memory weights in \mathbf{W}).

The function G is a piecewise linear function that clips all activations in \mathbf{x} to the range of -1 to 1 , thus confining the network state to lie within the box formed by the vertices. At the level of individual elements of \mathbf{x} (x_i), this is formalized by

$$G(\mathbf{x}) = \begin{cases} +1 & \text{if } x_i > +1 \\ x_i & \text{if } -1 \leq x_i \leq +1 \\ -1 & \text{if } x_i < -1 \end{cases} \quad \text{for all } x_i \in \mathbf{x} \quad (4)$$

It is a feature of these dynamics that, given a nonzero input, the state vector \mathbf{x} is guaranteed to converge on an attractor (all units saturating at $+1$ or -1) in a finite number of steps (Golden, 1986).³

In intuitive terms, this process corresponds to a blindfolded person seeking to find the corner of a room. If the person continues to walk in a straight line, a wall will be

encountered sooner or later, whereupon the forward motion is deflected into a sideways slide along the wall until the person becomes trapped in the corner of the room and all forward motion ceases.

In the simulations, if no attractor was reached within a maximum number of iterations (uniformly set to 12), recall was discontinued, and an omission was scored. Otherwise, if the network reached an attractor, it was compared with the correct item to allow classification of the response as a correct recall, a transposition, or an extralist intrusion (convergence on a nonlist-item attractor). If the network converged on the reflection of a list item (i.e., the opposite vertex diagonally across the box; $-\mathbf{f}_j$ instead of \mathbf{f}_j), it was considered to have converged on that item for the purpose of scoring (cf. J. A. Anderson, 1995, p. 512). This scoring method is appropriate because the outer product of a vector ($\mathbf{f}_k \mathbf{f}_k^T$; see Equation 1) is not affected by directional reversal ($-\mathbf{f}_k - \mathbf{f}_k^T = (-1)^2 \mathbf{f}_k \mathbf{f}_k^T = \mathbf{f}_k \mathbf{f}_k^T$). By implication, $-\mathbf{f}_k$ and \mathbf{f}_k are identical with respect to a given weight matrix \mathbf{W} .

In addition, the iterative nature of the SOB allowed predictions of recall latency to be derived, through consideration of the number of steps required for the state to converge on an attractor (J. A. Anderson, 1991; Ratcliff, Van Zandt, & McKoon, 1999). Unlike in many other models, the latency predictions are intrinsic to SOB and tightly coupled with its accuracy predictions.

Response suppression. The attractor associated with a recalled item was attenuated by adding the autoassociation of the recalled item to the inhibitory matrix \mathbf{A} with a negative encoding strength. This effectively reduced the basins of attraction in \mathbf{W} , because \mathbf{A} was used during the iterative recall dynamics in summation with the memory weights in \mathbf{W} (see Equations 2 and 3). Hence, a negatively weighted autoassociation in \mathbf{A} offset its positively weighted counterpart in \mathbf{W} . The presence of two sets of weights is functionally isomorphic to using a single set with different learning rates (i.e., positive for learning vs. negative for suppression) and follows relevant precedent (e.g., J. A. Anderson, 1995; Begin & Proulx, 1996).

Energy in the SOB

In SOB, each state of the network can be characterized by its energy with respect to the weight matrix. Energy was used to determine encoding strength (η_k), as well as the extent of response suppression. Energy reflects the consistency of the network's current state with respect to the weights and is generally given by

$$E' = -(\varepsilon/2) \sum_i \sum_j w_{ij} x_i x_j, \quad i \neq j, \quad (5)$$

where ε is as in Equation 3, x_i and x_j are the activation values of units i and j of the state vector \mathbf{x} , respectively, and w_{ij} the corresponding connection weight in \mathbf{W} .

To understand why E' reflects the consistency of a state of activation with respect to its weights, consider the possible relationships among values of w_{ij} , x_i , and x_j . Sup-

pose that x_i and x_j are both positive or both negative: Because they are activated in unison, consistency with the weights would be greatest if w_{ij} were positive. A negative value of w_{ij} , by contrast, would be inconsistent with the positive correlation between x_i and x_j . Conversely, suppose that one of x_i and x_j was negative and the other was positive: This negative correlation would be consistent with a negative w_{ij} , but not with a positive weight. Examining the cumulative implications of those cases in Equation 5 confirms that energy has the greatest negative value if consistency, as just defined, is high. Energy has been used extensively as a tool for analysis of the behavior of networks, and it plays a central role in the dynamics of several models (e.g., Hinton & Sejnowski, 1986; Hopfield, 1982, 1984; Kawamoto, 1993).

There are two differences between Equation 5 and related applications of energy (e.g., Haykin, 1994; Kawamoto, 1993). First, the diagonal of the weight matrix (where $i = j$) is disregarded. Since the SOB is autoassociative, the self-connecting weights along the diagonal are always positive and, therefore, not informative for the purpose of assessing consistency between the weight matrix and the state vector. Second, for computational convenience, the term $-(\varepsilon/2)$ was set to $-1/2$ in the present simulations (cf. Kawamoto, 1993). The notation E' (rather than E) was chosen to reflect those two minor differences.

Another important implication of Equation 5 is that energy is free to vary in the absence of any weight changes when the vector \mathbf{x} changes across iterations. It follows that the migration of \mathbf{x} toward a vertex of the box can be reexpressed as a movement along an energy landscape, formed by a constant set of weights and a variable state vector \mathbf{x} . Accordingly, it has been shown that reaching a vertex of the box corresponds to a movement toward a minimum on the energy landscape (Golden, 1986). Figure 1 provides a representation of the isomorphism between the box and the energy metaphors for SOB's nonlinear dynamics.

The surface at the top of the figure shows energy for all possible states of a two-dimensional vector \mathbf{x} with respect to a particular matrix \mathbf{W}_x . The possible values of the two elements in \mathbf{x} define the plane below the saddle-shaped surface, and the third axis represents E (not E' , because the tiny matrix mandated inclusion of the diagonal). The matrix \mathbf{W}_x was formed by adding the autoassociations of the two column vectors $\{-1, 1\}$ and $\{1, 1\}$, with encoding weights .04 and .02, respectively, to an empty ($\mathbf{0}$) matrix (see Equation 1). Inspection of the saddle-shaped surface reveals the two attractors: a particularly large one located at the "front" corner $\{-1, +1\}$ and a smaller one in the "back" corner on the right $\{1, 1\}$. Those two corners are the two encoded vectors, and the depth and size of the associated basins of attraction reflect the associated encoding strengths. (The symmetry of the energy surface underscores the identity between each attractor and its reflection, which was discussed earlier.)

The bottom part of Figure 1 shows the (two-dimensional) box with the two attractors represented by the corresponding vertices. The dots that are arrayed from the central

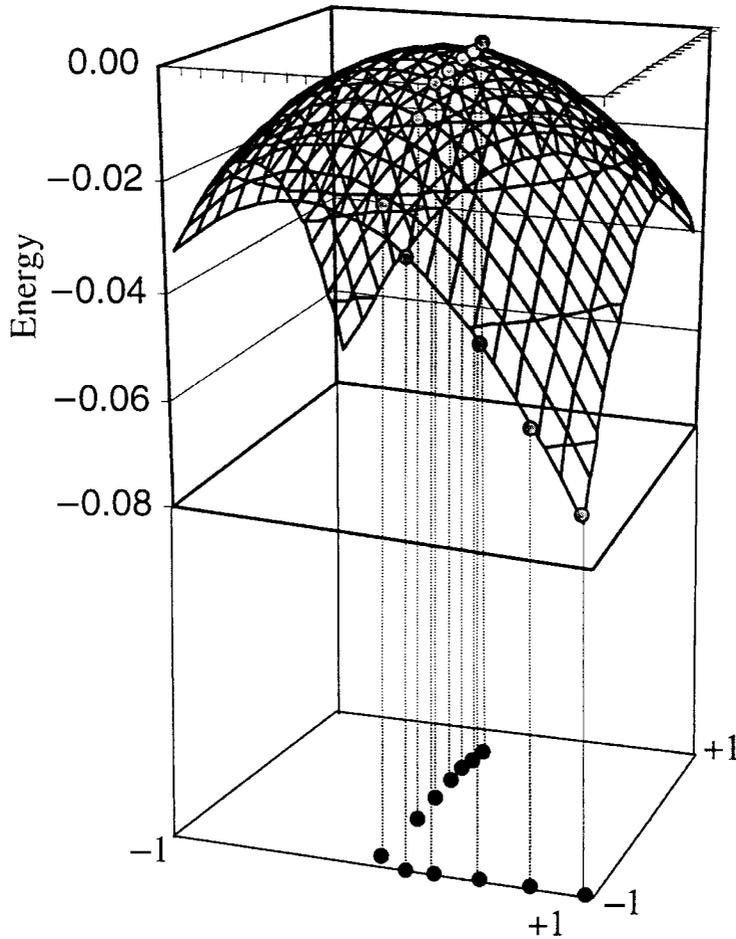


Figure 1. Illustration of the relationship between energy and network dynamics. The top half of the figure uses the energy surface metaphor to show all possible states of a simple network. Attractors, with minimal associated energy, are located at the corners of the state space. The bottom half of the figure uses the *brain-state-in-a-box* metaphor to illustrate the convergence of a random starting point to an attractor. Each point represents the state of the network at each iteration. The connecting lines between panels indicate the correspondence between each state within the box and its associated energy. See the text for further details.

region to the front right corner of the box describe the path taken by the state vector \mathbf{x} upon cuing with the starting location $\{-1, -2\}$, using the dynamics of Equation 3 (with β and ε set to unity). In this particular case, the attractor was reached after 53 iterations (not all iterations are represented by dots). The isomorphism between the box and the energy metaphors is emphasized by the dotted vertical lines that connect selected states in both representations.

Visual presentation of the energy surface can also illuminate the effects of response suppression: Figure 2 shows the same energy surface for the same weight matrix \mathbf{W}_x after the attractor for $\{-1, +1\}$ has been suppressed by adding it to \mathbf{W}_x with *negative* encoding strength $-.3$. (Note that, in our simulations, response suppression is implemented in a separate matrix \mathbf{A} ; this difference does not

affect interpretation of Figure 2.) The figure shows that following suppression, the other attractor, for the item $\{1, 1\}$, is now the stronger of the two.

In summary, the iterative dynamics of SOB can be alternatively understood as a nonlinear path toward a corner or a descent along an energy surface. Energy, in turn, reflects the consistency of a state vector with respect to the network's weights and thus reveals an item's relationship to what is already encoded in memory. Analysis of energy thus provides a natural way of obtaining a primacy gradient.

Energy and the Primacy Gradient

With orthogonal vectors, additional learning decreases the magnitude of the energy (i.e., E' becomes less negative), as defined by Equation 5. This is because each of the off-diagonal elements in a Hebbian autoassociative

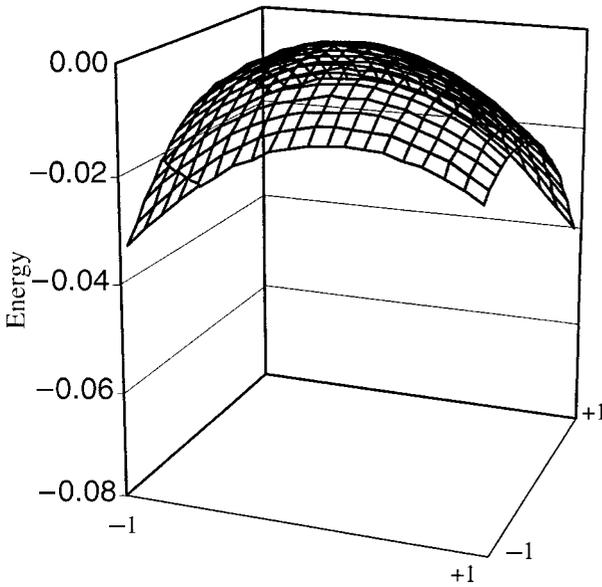


Figure 2. The same energy surface as that in Figure 1 after the initially stronger item has been suppressed. See the text for further details.

matrix \mathbf{W} , which are the ones considered for computation of E' , captures the correlation of activations for a given connection across all encoded stimuli (McClelland & Rumelhart, 1988, p. 85). For orthogonal vectors, the expected value of all those correlations is zero in the long run. Hence, study of additional items reduces the absolute magnitudes of the off-diagonal elements in \mathbf{W} (J. A. Anderson, 1995, p. 507), thus enabling the network to derive its own primacy gradient by considering each incoming item in turn.

This was achieved by determining the encoding strength for a particular item k as follows:

$$\eta_k = -E'_k / \phi_e, \tag{6}$$

where E'_k is the energy of item k (setting $\mathbf{x} = \mathbf{f}_k$ in Equation 5) with respect to the memory matrix \mathbf{W} and ϕ_e is a scaling constant that remained invariant in all the simulations reported in this article.

The general shape of the gradient is illustrated in Figure 3, which shows the computed values of the encoding strength (η_k) for one replication of Simulation 1. As we show after presenting all the simulations, the general shape of the function is unaffected by the value of ϕ_e .

The energy-gated encoding mechanism formally specifies the suggestion by Brown et al. (2000), discussed earlier, that a rationally adaptive organism would pay increasingly less attention to successive items. The adaptive utility of this mechanism can be clarified by considering what would happen if the network did not consider energy. In that case, all items would be stored with the same encoding strength, and by implication, all attractors would be of the same size. In consequence, owing to the unresolvable competition among studied items, the network ei-

ther would converge on a nonstudied attractor or would fail to converge within the allotted number of cycles. This was confirmed by an exploratory simulation using an equal encoding strength for all items, which resulted in a combined intrusion and omission rate of nearly 80%. The energy-gating mechanism circumvents this general memory problem while also giving rise to a primacy gradient and, by implication, the ability to recall items in order.

The use of a self-governing encoding mechanism is not without precedent. For example, Metcalfe (1993) renormalized the learning term in her CHARM model to regulate the variance in memory trace activations. Although Metcalfe's approach was motivated by different concerns, it shares with SOB the basic principle of making encoding contingent upon the current contents of memory. A similar approach was taken in the *closed-loop* version of TODAM (e.g., McDowd & Murdock, 1986), in which the weighting given to new information was reduced by the extent to which it had already been learned, thus enabling TODAM to show improved performance across multiple learning trials (e.g., Lewandowsky & Murdock, 1989). SOB extends these precedents by gating the encoding of incoming information through its endogenous novelty detection process.

Energy and Response Suppression

Just as energy was used to determine the extent of learning, so too energy was used to determine the extent of suppression of each recalled item. Specifically, the

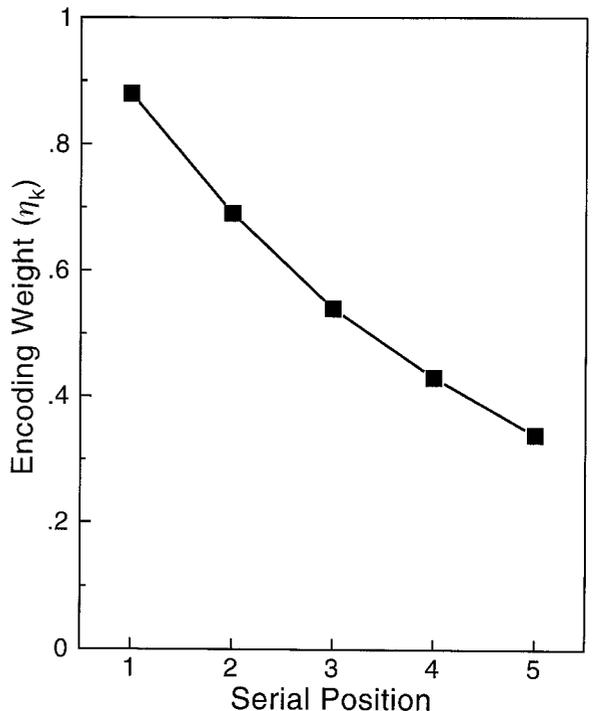


Figure 3. Representative energy-gated encoding function for a five-item list, showing the change in encoding strength for each item across input positions.

energy of the k th recalled item was assessed with respect to \mathbf{W} (see Equation 6), and the autoassociation of the recalled item was then added to the inhibitory matrix \mathbf{A} , using Hebbian learning (see Equation 1), with encoding strength γ_k determined by the following equation:

$$\gamma_k = -E_k / (E_1 \phi_s), \quad (7)$$

where E_k represents the energy of item k , ϕ_s is a scaling constant, and E_1 is the energy of the first recalled item. Equation 7 mirrors the form of the earlier Equation 6—except that the encoding strength γ_k takes on negative values to implement suppression—and thus synchronized the extent of suppression given to a recalled item with the strength with which it was encoded. Without such synchronization, recalled items might receive more suppression than their original encoding strength, which could create eigenvectors with negative eigenvalues in $\mathbf{W} + \mathbf{A}$, which in turn would threaten the stability of the nonlinear dynamics (J. A. Anderson, 1995).

The fact that the learning rule specifies the process by which response suppression is accomplished clearly distinguishes SOB from other models of serial recall. Those other models variously rely on the removal of items from a competitor pool (Brown et al., 2000; Lewandowsky & Murdock, 1989) or assume some reduction in activation of recalled items, without, however, specifying the agent underlying that outcome (Burgess & Hitch, 1999; Henson, 1998b; Page & Norris, 1998a, 1998b). In further distinction to these models, suppression in SOB occurs by using the same process, over the same units and weights, as the one used in learning. A final feature of this response suppression in conjunction with distributed representations is that the use of continuous inhibitory weights permits *partial* suppression of recalled items. Partial suppression, in turn, can subsequently be reversed (e.g., J. A. Anderson, 1991). Although not relevant to the present set of single-trial simulations, this *release from suppression* can be assumed to follow completion of recall and occurs through resetting the inhibitory weight matrix \mathbf{A} to $\mathbf{0}$.

Lewandowsky (1999) and Lewandowsky and Farrell (2000) explored a precursor of SOB, which only modeled the reintegration (i.e., disambiguation) of a memorial response that was assumed to have been retrieved by an unspecified associative memory stage. Although that model also used antilearning for suppression, it differs from SOB in several critical ways. First, SOB places additional emphasis on energy to gate response suppression. Second, SOB includes a complete encoding mechanism that obviates the need for an unspecified associative memory stage. Finally, encoding is also energy gated, which creates a conceptual link between encoding and retrieval.

SIMULATIONS

Parameters

The present article focuses on demonstrating the inherent power of the architectural principles of SOB. Hence, no free parameters were estimated from the data.

The constants of the iterative updating function (Equation 3), β and ε , were set to .2 and .7, respectively. The two scaling constants used to compute energy (Equations 6 and 7), ϕ_e and ϕ_s , were set to 600 and 1.4, respectively, for all simulations, unless noted otherwise below.

General Method

All simulation results were aggregated across 200 replications, each using a different random sample of items drawn from the vocabulary of Walsh vectors. Each replication consisted of pretraining, followed by study and recall of list items.

Recall of each item entailed cuing with a random vector (\mathbf{f}'), followed by energy-gated suppression of the recalled item, regardless of what was recalled. Thus, omissions were also suppressed, but because their nonconvergence resulted in lower energy (because at least some $|\mathbf{x}_i| < 1$; see Equation 6), the extent of their suppression was necessarily low.

Choice of Benchmark Phenomena

There is no universally agreed set of phenomena against which theories of memory for serial order are to be evaluated. We therefore selected to-be-simulated phenomena according to several constraints. First, we sought to maximize diagnosticity by applying SOB to the detailed pattern of errors in serial recall, whose diagnosticity in differentiating between theories is well established (e.g., Henson, 1998b; Henson, Norris, Page, & Baddeley, 1996). Second, to ensure some generality, we simulated the standard serial position curve for different list lengths and simulated the effects of natural word frequency. Third, we excluded phenomena that prior theoretical work had identified as requiring more than a single stage of processing. This choice reflects our focus on the endogenous mechanism in SOB, in preference to developing extensions that might broaden its scope, but at the expense of added complexity and additional parameters.

This decision eliminated the effects of phonological similarity from consideration, because all existing attempts at explanation with a single-stage mechanism have been unsuccessful. There is now widespread agreement that phonological similarity involves a second, item-based stage of processing that is not directly involved in order memory (Henson, 1998b; Page & Norris, 1998a). Likewise, we chose not to apply the model to grouping or chunking effects (e.g., Hitch, Burgess, Towse, & Culpin, 1996; Ryan, 1969). Although those effects—enhanced overall recall and “mini” serial position curves within each chunk—are readily produced by inserting a temporal gap between groups of study items, there is agreement that they are beyond the scope of the architecture explored here. Specifically, grouping effects are typically thought to arise from a hierarchical context signal (e.g., Brown et al., 2000; Burgess & Hitch, 1999) that cannot be represented in SOB without additional assumptions. Finally, for related reasons, we chose not to model backward recall, which is con-

sidered beyond the scope even of some more complex theories (e.g., OSCAR; Brown et al., 2000).

In summary, we sought to apply the simplest possible version of SOB to the largest possible number of diagnostic findings. This does not preclude future development of an enhanced version of the theory that can handle additional findings, such as similarity and grouping effects or backward recall, in the same manner as other theories do—namely, by incorporating constructs such as a context signal or a phonological confusion stage, each governed by an additional set of assumptions and parameters.

Simulation 1: Serial Position Curve and Transposition Gradient

The first simulation sought to demonstrate that SOB is capable of accounting for the serial position curve. The resulting simulated serial position curve for five-item lists is shown in Figure 4. These data are representative of patterns of serial recall for short lists (e.g., Estes, 1991; Henson et al., 1996).

It is clear from the figure that SOB predicts the bow-shaped serial position curve that has been observed in countless experiments. The simulation thus confirmed the utility of response suppression as an explanatory mechanism for recency in forward serial recall (cf. Lewandowsky, 1999; Lewandowsky & Li, 1994). One implication of this result is that models that explain recency primarily through edge effects but also include response suppression (e.g., Brown et al., 2000; Page &

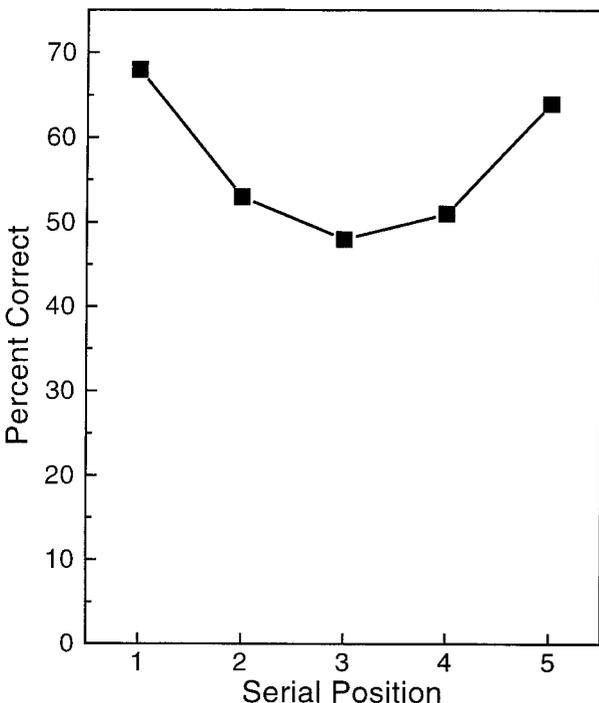


Figure 4. Serial position curve predicted by SOB for a five-item list (see Simulation 1 for details).

Norris, 1998b) may not, in fact, yield recency for the primary reasons cited. Instead, recency in those models may arise from suppression as well as edge effects.

The level of accuracy exhibited by the model deserves brief commentary because human performance on a five-item list is likely to be far better.⁴ A known property of the BSB architecture is that accuracy of recall with a random starting vector is a direct function of the size of the box relative to the length of the starting vector (J. A. Anderson, 1995, p. 512). In the extreme case of an infinitely large box—for example, when the clipping function G is removed from Equation 3—recall in SOB is isomorphic to the power method for computing eigenvectors, which is known to be arbitrarily accurate (e.g., J. A. Anderson, 1995, p. 499). Several additional runs of Simulation 1 confirmed that accuracy increases when the size of the box increases relative to the length of the cuing vector (\mathbf{f}'). For computational convenience, these simulations used a unit-sized box but reduced the length of the starting vectors. As the length of \mathbf{f}' decreased from its standard value of .0001 to 10^{-6} , 10^{-9} , and 10^{-11} , accuracy on the first recalled item increased from .68 to .73, .83, and .89, respectively. The shape of the serial position curve remained unaltered, with distinct primacy and recency being present regardless of overall level of accuracy. These additional simulations confirmed that SOB is capable of matching human recall performance; however, because we were particularly interested in error patterns, we purposely kept accuracy low by using a slightly longer length of the cuing vector in all the remaining simulations.

The transposition error gradients that accompany the predicted serial position curve in Figure 4 are shown in Figure 5, together with comparison data from an experiment by Nairne (1992). Each plotted parameter refers to an output position and shows the proportion of items reported from each input position. Thus, SOB predicts that the first recalled item is the first list item in about 70% of all cases, the second list item in 20%, and so on. The peaks of each plotted parameter represent items that were recalled in the correct output position and, when connected, form the serial position curve shown in Figure 4.

The model captured the pervasive finding, demonstrated in Nairne's (1992) data, that output positions tend to cluster around an item's serial position. Although items may erroneously migrate to adjacent positions during report, they are unlikely to be recalled far from their actual position. Henson et al. (1996) called this the *locality constraint*.

SOB satisfied the locality constraint without any associations between an item and a list marker (cf. Henson, 1998b) or a timing signal (cf. Brown et al., 2000), but as a direct result of the differential strength of encoding of list items. In SOB, the size of the basin of attraction surrounding each list item is a function of its encoding strength, which in turn is determined by the energy associated with a presented item (see Equations 1 and 5 and Figure 1). Thus, the first item has the largest basin of attraction, and the last item the smallest. At retrieval, given

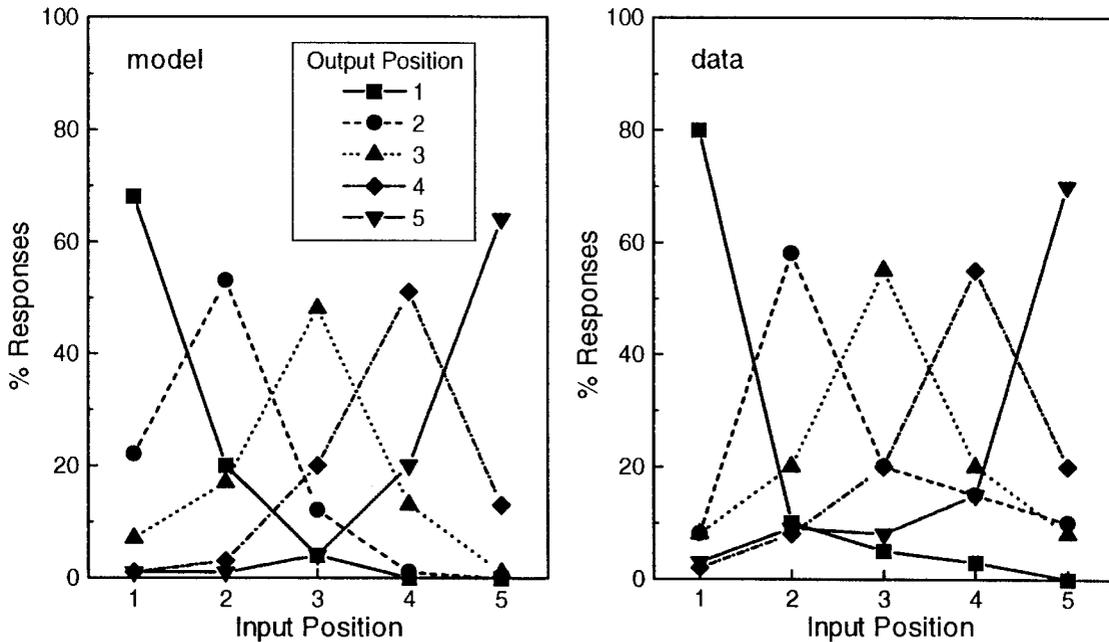


Figure 5. Transposition gradients for five-item lists, indicating the proportion of recall of each input item at every output position. The left panel shows the predictions of SOB, whereas the right panel illustrates representative data from Nairne (1992).

that the cue (f') is randomly constructed, the likelihood of the state vector x reaching any of the attractors is a sole function of their strength. Thus, if on the first retrieval the random starting vector, by chance, falls outside the correct basin of attraction, the second and third items are the most likely candidates for report. This automatically gives rise to the predicted positional gradient for the first output position (see Figure 5). If, on the other hand, the first item was successfully reported, it has been suppressed and can thus no longer compete (much) for report during subsequent retrievals. Hence, for those subsequent retrievals, the strongest items are those that were studied in the corresponding list position.

One consequence of this mechanism is a phenomenon known as *fill-in* (Henson et al., 1996; Page & Norris, 1998b). Fill-in occurs when, say, the second list item is erroneously recalled first. Rather than being followed by report of the third item, which would preserve the order between the second and third items, the first item is most likely to be recalled next because, not having been suppressed, its basin of attraction continues to be strongest. Page and Norris (1998b) reported a reanalysis of data from Henson et al. (1996), showing that when a list such as "123456" is recalled incorrectly, output sequences of the type "21xxxx" (fill-in) are three times as frequent as "23xxxx" (relative position maintained). In the present simulation, the same pattern was observed, with fill-ins outnumbering relative position reports by a factor of five.

Although this ratio of fill-ins to relative position transpositions is qualitatively similar to that found in the data,

closer inspection of the simulation results suggests that the ratio is overly large—that is, that the model suffers from *over-fill-in*. This is witnessed in Figure 5 as adjacent-position transpositions being at least as frequent as those in the data, whereas those between nonadjacent items are less frequent than is found empirically. This is particularly pronounced for the second output position, whose asymmetry is in the opposite direction to that found in the data. A similar pattern of over-fill-in appears to be found in Page and Norris (1998b, their Figure 4), although it is less apparent because the high absolute level of recall in their simulations constrained the number of transpositions.

Simulation 2: Varieties of Errors

Among theorists of serial order, there has been much emphasis on the errors that arise during recall, owing to their presumed diagnostic value in differentiating between rival models (e.g., Henson et al., 1996). The second simulation explored the interacting pattern of transpositions, intrusions (reporting an item not on the list), omissions (not reporting anything at a particular position), and repetition errors (erroneously reporting an item twice).

Proportions of different classes of error. Comparative data for this simulation are from Henson (1996, his Figure 3-5) and are shown in the left panel of Figure 6. Although transpositions are not often presented as a function of output position, the total number of transposition errors generally exhibits the pronounced inverse U-shape shown in the figure (e.g., Henson, 1996; Henson et al., 1996). Likewise, intrusions often show a slight inverse

U-shape, reflecting a decrease in intrusions toward the end of recall. Finally, omissions typically increase across output position in a monotonic manner (Henson, 1996).

SOB handled intrusions quite naturally through the inevitable presence of *spurious* attractors that did not represent studied items. Any convergence on a nonstudied attractor was thus readily scored as an extra-list intrusion. Any response that did not converge on an attractor—spurious or otherwise—within the limit of 12 iterations stated earlier was counted as an omission. The right panel of Figure 6 shows the predicted proportion of transpositions, intrusions, and omissions as a function of output position.

The model captured the basic empirical pattern of errors: Omissions increased monotonically with output position, intrusions exhibited a slight inverse U-shaped function, and transpositions showed a pronounced inverse U-shape. Taken together with the results of the first simulation, SOB clearly captured the pattern of errors that are observed in the serial recall of short lists. With the possible exception of exaggerated fill-in, the predictions were in close quantitative correspondence to the data.

Repetition errors. The final class of errors involves erroneous repetitions of an item (e.g., report the list A B C as A B B C). Erroneous repetitions occur very infrequently, with estimates of their incidence ranging from 2% of all responses (Henson, 1996) to 5% (Vousden & Brown, 1998). Nonetheless, there is a distinct distribution of repetition errors, with most repetitions involving early list items that are reported a second time late in recall. Henson (1996) reported that erroneous repetitions were, on average, 3.34 positions apart.

This pattern was captured by SOB. In Simulation 2, repetition errors constituted 0.1% of all responses and were separated by four output positions on average. That margin increased to 0.5% for lists with a length of six. The increase of the proportion with list length mirrors the data, although the total number of erroneous repetitions is underpredicted overall. At a finer grain of analysis for six-item lists, for the last output position SOB included more reports of the first item (2%) than of the second and third (1% for both). This violation of the locality constraint occurs when there are a moderate number of repetition errors (Henson et al., 1996). Thus, although SOB underpredicted the overall incidence of erroneous repetitions, it correctly predicted the increase with list length and the associated highly specific violation of the locality constraint. It should be noted that, unlike SEM (Henson, 1998b), which is the only other serial recall model that has been explicitly applied to repetitions, SOB does not require release from suppression to allow repetitions to occur. This is because response suppression in SOB is continuous, rather than all or none.

Overall, the detailed account offered by SOB goes beyond that of the primacy model, which lumps omissions, intrusions, and repetitions together as *item* errors (Page & Norris, 1998a, 1998b). The SOB account thus rivals that provided by SEM (Henson, 1998b), except that SOB computes its primacy gradient endogenously and does not require an additional recency gradient.

Simulation 3: List Length Effects

Serial position curves. In the next simulation, the effects of list length were examined. The basic procedure remained unchanged, except that list length varied between

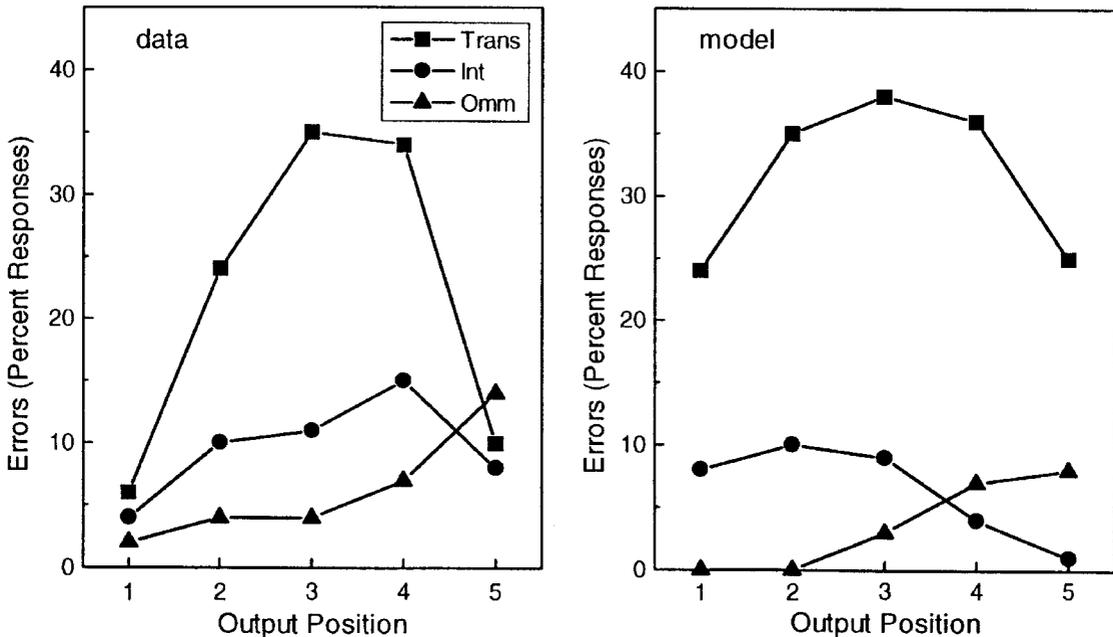


Figure 6. Patterns of different types of error across output position, as a percentage of all responses. The left panel gives representative data from Henson (1996); the right panel illustrates the predictions of SOB.

three and seven. Figure 7 shows the resulting serial position curves in two ways. The left panel uses a strict scoring criterion, which only considers correct responses up to the first erroneous retrieval, regardless of the type of error. The right panel uses a less stringent criterion that considers a response correct if it occurs in the appropriate output position, regardless of whether an error has been made earlier.

For the conditional scoring (left panel), the model predicted a limited separation of the serial position curves. This effect, known as *fanning*, refers to a decrease in performance at a given serial position as list length increases. Traditionally, on the basis of the results of Drewnowski and Murdock (1980), a large extent of fanning has been considered an empirical benchmark for computational models (e.g., Lewandowsky & Murdock, 1989). However, Murdock (2001) recently reevaluated the methodology of Drewnowski and Murdock. Prompted by the discovery of possible artifacts, Duncan and Murdock (2000) repeated the study under improved conditions. The maximal extent of fanning in their study, observed when subjects were aware of list length prior to presentation, was much smaller than that observed by Drewnowski and Murdock and was quantitatively similar to the predictions by SOB shown in Figure 7.

For the lenient scoring (right panel), used in all other simulations reported here, SOB likewise predicted limited fanning, but in conjunction with substantive recency whose extent remained constant across list length. For the

range of lengths considered here, this prediction is consonant with the data (e.g., Henson et al., 1996; Hitch et al., 1996). Note, however, that Lewandowsky (1999) reported an elimination of recency with very long lists (i.e., 40 items) that are beyond the purview of SOB's focus.

Latency data. The dynamics of retrieval in SOB (see Equation 3) provide a natural way of generating latency predictions for every type of response. With the exception of the symbolic model by J. R. Anderson and Matessa (1997), no other existing model predicts latency at that level of detail. It would therefore be particularly informative to compare the predictions of SOB with behavioral results.

Notwithstanding the acknowledged utility and diagnosticity of response latencies (e.g., Ratcliff et al., 1999), very few detailed examinations have been reported in the serial memory literature. Considering first the overall effects of list length, Doshier and Ma (1998) showed that total retrieval time for all items is a quadratic function of list length, although the shape of this curve varies with the materials used and the recall method employed. The quadratic component of the curve can be very small and may escape immediate visual detection. Representative data from the study by Doshier and Ma are shown in Figure 8, along with the output times predicted by SOB for various list lengths. The figure shows that SOB closely captured the relationship between latency and list length without any parameter manipulation.

Turning to a more fine-grained analysis of latencies, we are aware of only four studies that have reported latency

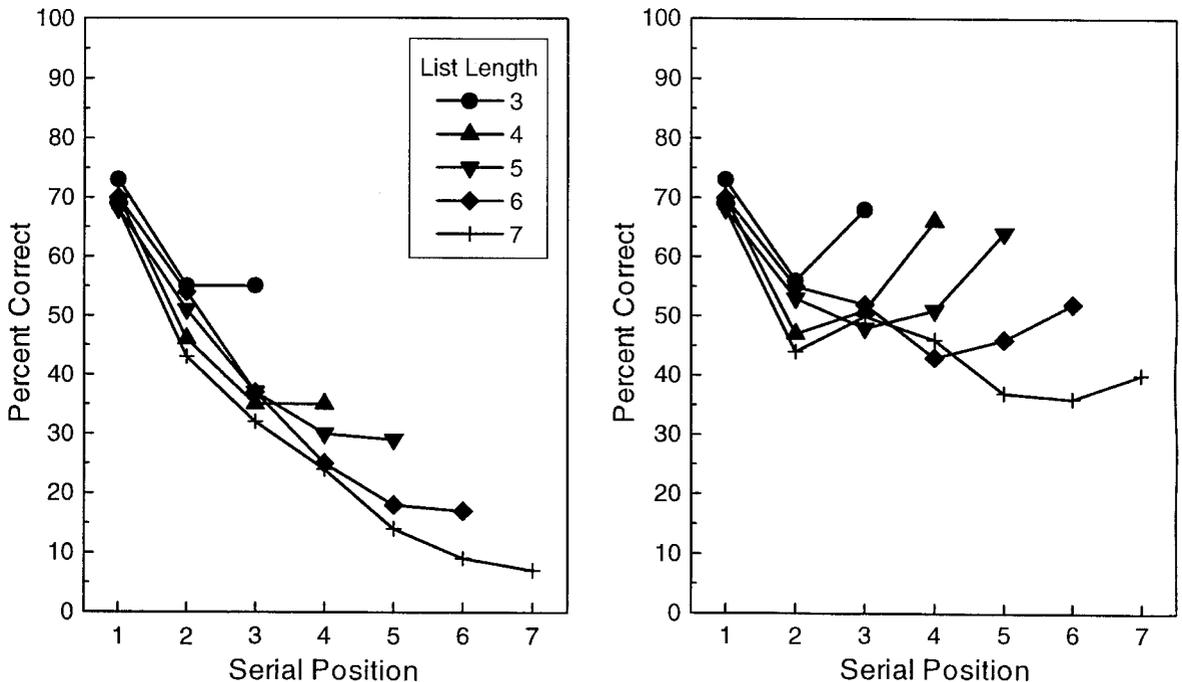


Figure 7. Predicted serial position curves for various list lengths, according to different scoring criteria. The left panel shows predicted list length effects when recall is scored up to the first error (see, e.g., Drewnowski & Murdock, 1980). The right panel shows standard serial position curves for various list lengths.

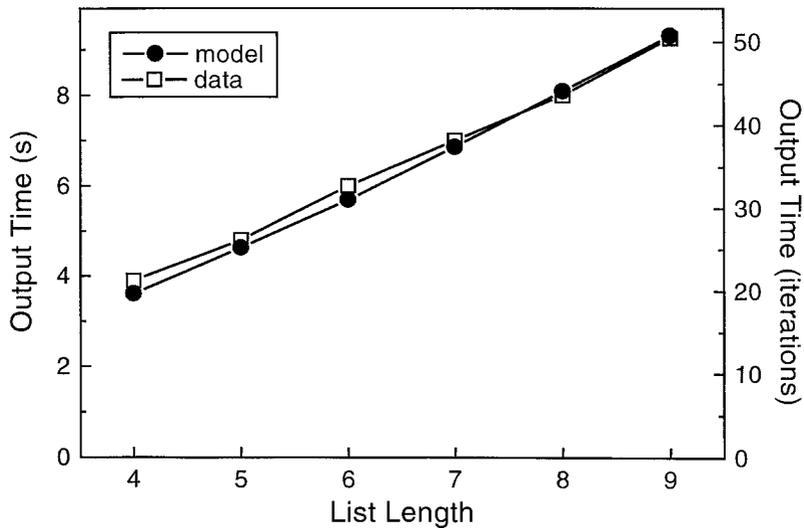


Figure 8. Predicted total recall duration for various list lengths, with data for word lists from the keypress condition of Doshier and Ma (1998). Both predictions and data were aggregated across lists that contained no omissions.

at the level of individual serial positions (J. R. Anderson & Matessa, 1997; Cowan et al., 1998; Doshier, 1999; Maybery, Parmentier, & Jones, in press). Because SOB models only retrieval processes, but not pronunciation operations, and because pronunciation durations are sometimes found to covary with pauses in between retrievals (e.g., Cowan et al., 1998, p. 158), we focus on methodologies in which subjects recalled items by pressing keys on a keyboard or another output device. We take keypress latencies to be particularly sensitive measures of retrieval dynamics that are independent of pronunciation and articulatory set-up processes. Similarly, because SOB does not represent grouping during list presentation, we only consider methodologies that involved ungrouped presentation. This reduces the pool of relevant studies to those by Doshier and by Maybery et al. Both studies yield remarkably similar results, with the data from Maybery et al. shown in the top panel of Figure 9. The data represent the cumulative response latencies for correct responses across serial positions for list lengths three through six.

Cumulative latencies are shown because the extremely long latency observed for the first item, which presumably includes general set-up processes prior to initiating recall, are reflected in the intercept. This makes it easier to align the predictions of SOB, which does not model preretrieval set-up processes. The predictions with the standard set of parameter values are shown in the bottom left panel of Figure 9. It is clear that although SOB captures the overall list length effect, it does not capture the fanning of cumulative latency functions that is observed in the data. This lack of fanning, which parallels the accuracy predictions, occurs because of the steep primacy gradient in the model, which renders early-list attractors very strong regardless of how many other, relatively weak items are encoded subse-

quently. Hence, accuracy and latency of recall of early list items is not affected (much) by list length.

Support for this analysis is provided in the bottom right panel of Figure 9, which shows predicted latencies with the encoding parameter (ϕ_e) set to 1,200, double its standard value. Under those circumstances, fanning of latency functions occurs because the energy-gated primacy gradient is less steep, and early attractors, being weaker, are thus more affected by the number of subsequently encoded items (for a sample of a corresponding accuracy serial position curve, see Figure 11, below).

In summary, Simulation 3 demonstrated that SOB captures the basic effects of list length on serial recall, whether performance is assessed by a strict or lenient recall criterion or by a cumulative latency measure. SOB also accommodated the shape of the accuracy serial position curves under both scoring modes, but it did not simultaneously capture the fanning with list length that is observed empirically for latency serial position curves.

Simulation 4: Effects of Word Frequency

It is well known that short-term serial memory performance is a function of the nature of the to-be-remembered material. For example, natural language frequency affects recall performance, with high-frequency words being recalled better than their low-frequency counterparts (e.g., Hulme et al., 1997). The magnitude of that high frequency advantage is known to increase across serial positions, with the effect often being minuscule for the first three serial positions (e.g., Hulme et al., 1997, Figure 5).

In Simulation 4, natural language frequency was modeled by manipulating the extent of pretraining. Although the standard encoding strength was used ($\eta_p = .001$), 25 randomly selected vocabulary items were presented 20

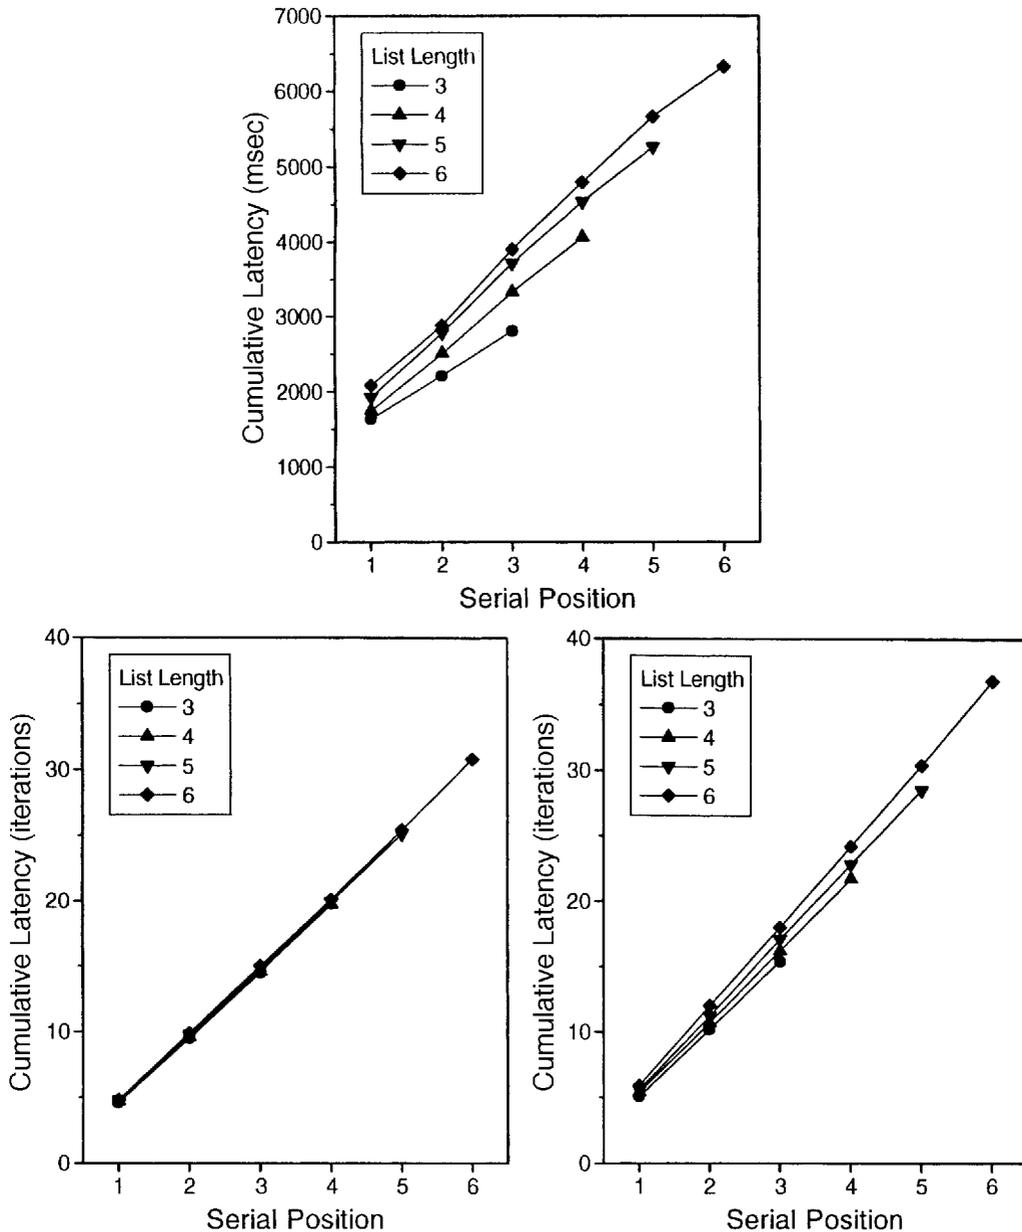


Figure 9. Obtained and predicted cumulative latency serial position curves. The top panel shows the results of Maybery, Parmentier, and Jones (in press). The bottom left panel shows the predictions of SOB with standard parameter values. The bottom right panel shows the predictions of SOB when the encoding constant (ϕ_e) was doubled from 600 to 1,200.

times to model high natural language frequency, and the remaining 25 items were presented 15 times to represent low-frequency words. Five words were then randomly selected from the appropriate pool to model high- or low-frequency study lists.

The results of the simulation are shown in Figure 10. The left panel shows the predicted serial position curves for high- and low-frequency words. The right panel shows the corresponding number of different classes of errors made by the network.

It is clear from the figure that SOB captured the effects of word frequency reported by Hulme et al. (1997): High-frequency words were recalled better than low-frequency words, and this effect increased past the first three serial positions (left panel). The increasing advantage for high-frequency words was directly tied to the decrease of encoding strength across serial position (see Equation 6). Because items with lower encoding strengths are more dependent on additional strength from pretraining for successful recall, the more extensively pretrained high-

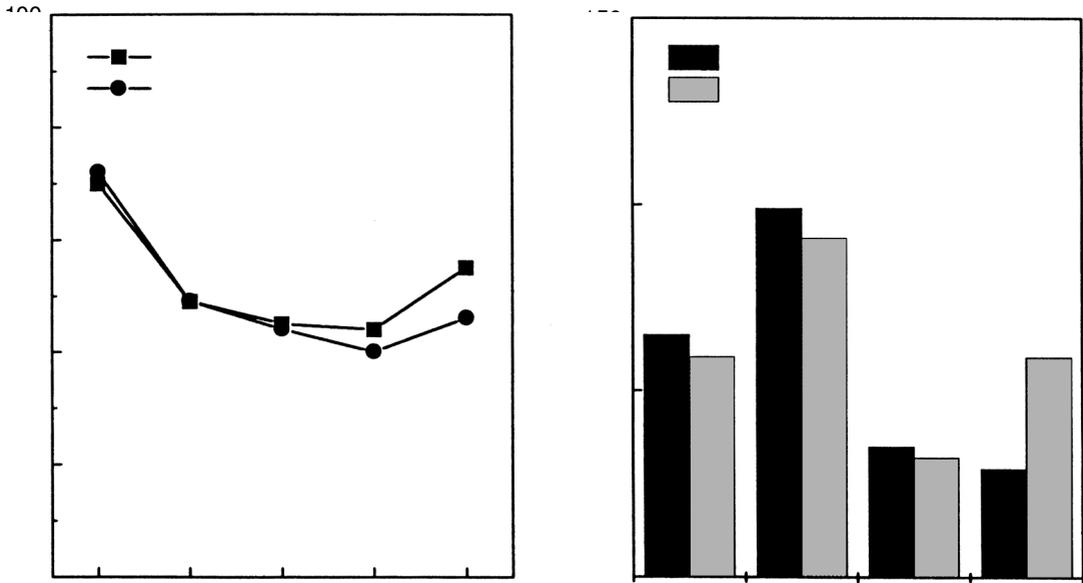


Figure 10. Predicted effects of word frequency on serial recall. The left panel illustrates the effects of frequency on the serial position curves. The right panel shows the difference in various error types for high- and low-frequency words.

frequency words enjoyed a selective advantage late in the list.

The left panel of Figure 10 reveals another very subtle effect of word frequency among the first three serial positions: High-frequency words are recalled with fractionally less accuracy than are low-frequency words at the first serial position, whereas the reverse is true for the third serial position, with that high-frequency advantage increasing over subsequent positions. Although the effect is very small and, hence, may not merit much exploration, it should be noted that a similar minute reversal of the word frequency effect was observed by Hulme et al. in some conditions (1997, their Figure 4).

The predictions regarding different types of responses (right panel of Figure 10) were also consistent with the data: Hulme et al. (1997) reported that word frequency primarily affected omission rates, rather than any other class of error. This effect is clearly present in the figure. Again, SOB’s predictions were found to be consistent with several intricate features of a well-established effect.

It should be noted that the findings of Hulme et al. (1997) contrast with those of Watkins and Watkins (1977), who found a *decreasing* effect of word frequency across serial position. The reason for this disparity can be inferred from the results of Watkins and Watkins, whose recency effect (with auditory presentation) was nearly as large as primacy. This suggests that Watkins and Watkins’s instructions to recall the words in order, in writing, were not effective in preventing subjects from recalling the last few items on the list first (but reporting them in the last few

positions on the output sheet). This contrasts with the procedure of Hulme et al., who forced serial recall by requiring subjects to recall lists verbally. The recall regime simulated by SOB is directly analogous to the task requirements of Hulme et al.

The primacy model has also been applied to word frequency effects, although in an arguably less principled manner. In the primacy model (Page & Norris, 1998b), word frequency was modeled by manipulating an output threshold: Although this turned out to capture the pattern in the data, the manipulation did not directly reflect the differing nature of the items themselves.

Simulation 5: Robustness of the Model

We have shown the general applicability of SOB and the endogenous processes it relies on to govern encoding and retrieval. We now examine the extent to which the model’s predictions are tied to particular parameter values, an approach known as parameter sensitivity analysis (e.g., Li, Lewandowsky, & DeBrunner, 1996). Although no parameters were estimated for the predictions reported in Simulations 1–4, it is possible that the range of possible values for the principal parameters, the encoding (ϕ_e) and suppression (ϕ_s) constants, is small and that values outside that narrow range would engender nonsensical predictions.

Figure 11 shows the predicted serial position curves across a range of values for ϕ_e (left panel) and ϕ_s (right panel). For the encoding parameter, the serial position curve exhibits the required primacy and recency for all values within the range of 500–1,000. For the retrieval parameter,

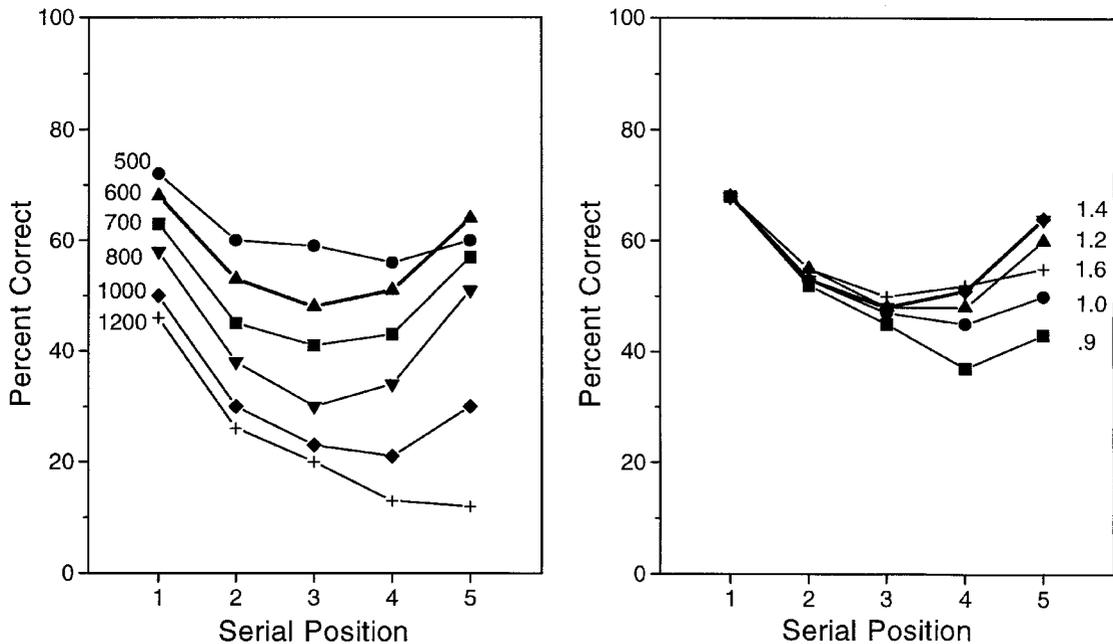


Figure 11. Predicted serial position curves across a wide range of values for ϕ_e (left panel) and ϕ_s (right panel). The bold lines represent the values used in all the simulations reported in this article ($\phi_e = 600$ and $\phi_s = 1.4$).

the serial position curve retains the appropriate shape across the range of 0.9–1.6.

To put these results into perspective, note that the predictions remain qualitatively unchanged across a parameter variation of almost 100% (i.e., from $\phi_e = 500$ –1,000 and $\phi_s = 0.9$ –1.6). In the context of the large changes in prediction that some models exhibit in response to much smaller variations in parameters, sometimes in the range of 1–10% (e.g., Li et al., 1996), the robustness of SOB must be noted.

Two conclusions can be drawn on the basis of these results. First, the predictions of SOB must result primarily from the properties of its architecture—in particular, the endogenous gating of encoding and retrieval. Had the predictions been attributable to the suitable setting of parameters, they would not have withstood a variation of 100%. Second, by way of tradeoff, that robustness may limit SOB's ability to adapt its behavior and, thus, maximize its quantitative fit to the data through parameter estimation. Given our focus on the inherent properties of endogenous encoding and retrieval, the former strength outweighs the latter potential difficulty.

GENERAL DISCUSSION

Summary of Simulations

In the first simulation, SOB accounted for the general shape of the serial position curve and the underlying pattern of transposition errors. SOB also correctly predicted that fill-ins outnumber relative position reports, albeit by

a greater margin than is commonly observed. In the second simulation, the pattern of errors was examined in more detail, and it was confirmed that SOB predicted the qualitatively correct tradeoff between transpositions, intrusions, omissions, and erroneous repetitions across output position. The third simulation demonstrated that SOB predicted the correct amount of fanning among serial position curves for different list lengths. SOB also captured the quadratic relationship between total retrieval time and list length. The fourth simulation showed that SOB can accommodate the effects of natural language frequency by predicting an increasing superiority of high-frequency words across serial position. Critically, the fifth and final simulation showed that these predictions were relatively insensitive to manipulation of parameters across a wide range.

Notwithstanding these apparent strengths, the model comes with its own set of limitations. We first discuss those limitations, before comparing SOB with other existing theories.

Model Limitations

Lack of associations. In most models of serial recall, order is represented by associating the to-be-remembered information with some other entity that is then used to cue retrieval. This other entity can be a set of oscillators (e.g., Brown et al., 2000), list markers (e.g., Henson, 1998b), other items on the list (e.g., Lewandowsky & Murdock, 1989), or an abstract representation of context (e.g., Burgess & Hitch, 1992, 1999). We have deliberately avoided

the use of any such associative cuing mechanism. Although associative cuing would undoubtedly extend the abilities of the model, it would contribute nothing to our conclusion that SOB provides an endogenous process account of the two crucial assumptions that are shared by several models whose associative cuing assumptions are quite diverse.

That said, results from other tasks reveal the limitations of our choice to omit associative cuing from the model. For example, in probe tasks a cue (e.g., “A-?”) is presented to elicit recall of the item that followed the cue at study (e.g., “B”). In these experiments (e.g., Murdock, 1968), people can clearly recall a specified item without overtly retrieving the entire list up to that position. Moreover, when response availability is experimentally controlled, people are equally adept at recalling the item that preceded a cue on the list and recalling the item that followed it (Rizuto & Kahana, 2001). There presently is no mechanism in SOB that would permit modeling of these tasks.

Related to this, SOB is unable to account for backward recall, since it would require encoding strength to decrease, rather than increase, across input position. Because the primacy gradient is computed endogenously, this reversal could not be achieved even if a theoretical justification for it could be adduced. This stands in contrast to some other models, which, by parameter adjustment, could produce a recency gradient as readily as they now produce a primacy gradient.

However, there is some evidence that forward and backward serial recall rely on entirely different retrieval processes (Li & Lewandowsky, 1993, 1995). It follows that, by recognizing that other processes are responsible, SOB may be able to account for backward recall notwithstanding its unidirectional primacy gradient. One possible solution was foreshadowed by Page and Norris (1998b), who showed that their primacy model could produce backward report by successive covert forward recalls of decreasing lengths that “peeled off” the terminal items on each occasion. Murdock (1995) proposed a similar mechanism to implement backward recall in a version of TODAM known as the chunking model.

Lack of a similarity mechanism. All the simulations reported here used orthogonal vectors, and any relaxation of that orthogonality constraint could adversely affect performance of the model. By implication, SOB cannot be applied to any similarity effects, including the pervasive effects of phonological similarity (e.g., Henson et al., 1996). In light of the widespread agreement that a single-stage mechanism, such as that embodied by SOB, cannot account for similarity (Henson, 1998b; Page & Norris, 1998a), the lack of a suitable mechanism in SOB need not prevent an even comparison to single-stage similarity-free versions of other models.

Lack of chunking. Another limitation of SOB is that it cannot accommodate the presence of chunking. At first glance, this runs counter to the strong evidence that even short lists may be encoded in separate chunks under some

circumstances. Specifically, chunks clearly occur with repetition (e.g., Martin & Noreen, 1974) and grouped list presentation (Frick, 1989; Henson, 1996; Hitch et al., 1996; Ryan, 1969). Accordingly, several contemporary models—for example, OSCAR (Brown et al., 2000) and SEM (Henson, 1998b)—include theoretical provisions for chunking. In the case of SEM, the provision consists of an additional set of markers that identify the group to which an item belongs. At study, items are associated with the group markers, in addition to the conventional list markers.

However, under the experimental conditions simulated in this paper, people may, in fact, not form chunks during encoding. Murdock (2001) presented a descriptive model whose parameters, when estimated from the data, could reveal the presence of chunking. When applied to the data of Duncan and Murdock (2000), which are representative of the conditions under consideration here, the model strongly implied that chunking did not occur. Thus, much as in the case of phonological similarity, the lack of a chunking mechanism in SOB need not prevent an unbiased comparison with other models, under the conditions simulated here.

Summary of limitations. As it stands, SOB cannot explain probe task performance and backward recall, and it also cannot represent chunks of items. Moreover, SOB cannot accommodate the effects of phonological similarity.

It is likely that solutions could be found for these limitations, in the same way that other models account for those results by postulating additional mechanisms. Those possible solutions were not explored here because our emphasis was on the endogenous capabilities of the architecture. We now place SOB into the context of existing theorizing by comparing it with versions of other models that likewise do not include provisions for chunking and phonological similarity.

Comparison With Other Models

TODAM (Lewandowsky & Murdock, 1989). We discuss TODAM primarily to provide historical context. TODAM accounted for a number of phenomena, but as compared with SOB and other contemporary models, it did not provide a detailed account of error patterns and other intricacies of the data, and it required several free parameters even for its coarse account. These included the assumption that encoding strength decreases across serial positions and that the number of competing responses decreases with the number of items recalled (for a critique, see Mewhort, Popham, & James, 1994; Nairne & Neath, 1994).

SOB provides an endogenous process instantiation of several of the problems raised in connection with TODAM, and it provides a more detailed account of the data.

OSCAR (Brown et al., 2000). Like TODAM, OSCAR explains many effects on the basis of a well-specified cuing mechanism. Conceptually, OSCAR shares some of its parameters with TODAM—most notably, the assumption that encoding strength decreases across input posi-

tions. Both models rely on this primacy gradient to yield plausible serial position curves. OSCAR also implements response suppression in much the same way as TODAM did.

In contrast to TODAM, OSCAR provides a reasonable account of the distribution of errors in serial recall, with the exception of intricate details, such as fill-in and repetition errors. OSCAR also handles grouping effects that are beyond the purview of both TODAM and SOB.

SEM (Henson, 1998b). This recent model accounts for more detail in the data than does either OSCAR or TODAM. However, SEM has two notable limitations. First, it only characterizes the outcome of the processes underlying serial recall by setting up associative confusion gradients, but it arguably fails to specify the nature of the processes underlying performance. SOB, by contrast, specifies the mechanisms and processes that are assumed to underlie serial recall. Second, SEM relies on the questionable assumption that people know at all times during encoding how close they are to the end of the list. This assumption is required in order to compute the strength of the end marker during item encoding. SOB, by contrast, does not require knowledge about list length or any other external information.

Primacy model (Page & Norris, 1998a, 1998b). Notwithstanding the obvious similarities between the primacy model and SOB, the models also differ in crucial ways. Most important among those is the fact that in the primacy model, localist representations compete for output of the strongest item, whereas in the distributed SOB, all items are simultaneously, and to varying extents, elicited by the random retrieval cue. This is of particular interest in light of the claim by Page (2000) that distributed representations do not lend themselves to selective suppression of items. Clearly, SOB shows otherwise.

Moreover, the distributed representations in SOB provide a natural account of extra-list intrusions through the presence of spurious attractors. By contrast, it is unclear how the localist representations in the primacy model would explain extra-list intrusions. Indeed, to date, the primacy model has considered extra-list intrusions together with omissions and repetition errors in the broad class of *item errors*.

Finally, the primacy gradient in SOB is computed by the same process that is also used to govern response suppression during retrieval. On balance, therefore, SOB goes beyond the primacy model in at least those two respects, without abandoning the basic competitive cuing approach.

Theoretical Implications

Energy-gated storage and retrieval. In SOB, the primacy gradient is critical not only because it provides a mechanism for seriation, but also because, without a gradient, SOB could recall little in response to a random cue. Unlike the primacy model, which would recall an item at random if the primacy gradient were flat, SOB primarily

produces intrusions and omissions if items are encoded with constant strength. Thus, the energy-gated mechanism that generates the primacy gradient in SOB is crucial to the model's functioning and must operate in all situations to which it is applied.

In psychological terms, the on-line analysis of the contents of memory that underlies energy gating is best understood as an attentional process. One likely manifestation of an attentional process in memory is the *von Restorff effect* (e.g., Hunt, 1995), which refers to the recall advantage of a particularly distinct list item (e.g., the word *hippopotamus* on a list of vegetables). The energy-based learning mechanism in SOB might accommodate this effect by adjusting encoding strengths on the basis of each item's novelty, as compared with what has already been acquired. The greater novelty associated with the distinct item would enhance its encoding strength, as compared with other list items.

Another attentional effect in memory was recently reported by Dennis and Kruschke (1998), and SOB may again provide a natural account of that effect. In the study by Dennis and Kruschke, subjects first studied word triples that consisted of an imperfect predictor, a perfect predictor, and a target. Following standard anticipation-learning methodology, the predictors were shown together, and the subject's task was to "predict" (i.e., recall) the target before it appeared. In a second learning phase, people again learned the same word triples, along with a second set of stimuli that shared the same imperfect predictor but were combined with a different perfect predictor and target. At test, when the subjects were presented with both perfect predictors, they tended to respond with the target associated with the later presented predictor, even though the triples involving the first predictor had by then been presented twice as often. This could be explained by the energy-gated mechanism in SOB, if it is assumed that the total attention given to a triple is held constant. Since the imperfect predictor will have less novelty attached to it when the second set is presented, the network would, on the assumption of constant attention between triples, learn the new perfect predictor with a greater endogenously computed encoding strength.

This potential generality of the energy-based learning mechanism provides a fertile avenue for further exploration.

Absence of cues. Our results suggest a reevaluation of at least one prominent theoretical construct—namely, the need for an encoded cue to initiate retrieval that is embodied in many models. For example, in the early TODAM, the first list item was associated to a start signal that was used to prompt recall (e.g., Lewandowsky & Murdock, 1989). In OSCAR, each item is associated with an autonomous timing signal that is rewound at retrieval (e.g., Brown et al., 2000). A similar mechanism is present in the models by Burgess and Hitch (1992, 1999). SOB, by contrast, relies on a random cue that is generated prior to each retrieval and that bears no relationship or resemblance

to anything encoded in memory. A precedent for this approach was provided by Murdock (1983), who also used a cue generated at retrieval—in that case, a delta vector (i.e., central element set to unity, the remaining elements set to zero)—to reliably elicit the strongest list item without any encoded cues.

Although the cueless approach has clear limitations—for example, it is difficult to see how SOB could accommodate grouping effects without encoding of group markers—the present simulations showed that a considerable number of phenomena can be explained without cues.

Iterative dynamics. Another contribution of SOB is that it naturally accounts for the dynamics of serial recall. Indeed, in contrast to most existing models (see J. R. Anderson & Matessa, 1997, for an exception), SOB cannot predict recall accuracy without also making predictions for retrieval times. These necessary predictions tightly constrain future tests of the model and suggest a fertile avenue for future empirical research.

Although other serial order models could be extended to make latency predictions—for example, in OSCAR, latency predictions could be based on the similarity between the correct response and the model's output—this would presumably require several additional parameters and assumptions. Similar comments would seem to apply to all the other models of serial recall.

Contribution of long-term memory representations. It is not uncommon for contemporary models of serial recall to disregard or minimize the role of long-term memory representations (e.g., SEM; Henson, 1998b). This renders it difficult to accommodate certain classes of semantic effects on short-term memory recall (e.g., Haarmann & Usher, 2001), which are the likely result of long-term memory representations. Although we did not systematically explore the role of long-term memory in SOB, it must be noted that the model includes a rudimentary long-term memory component. The pretrained corpus of items in SOB forms a vocabulary of long-term memory representations that are instantiated by (small) attractors. When some of those items are presented for memorization as part of a list, their pretrained attractors are significantly strengthened. Importantly, the same attractor landscape handles both long-term *semantic* and short-term *episodic* representations. In this regard, SOB is related to the computational model of Haarmann and Usher (2001), which views short-term memory “as activated LTM [long-term memory] representations in the prefrontal cortex” (Haarmann & Usher, 2001, p. 569). Although the Haarmann and Usher model does not address the issue of serial ordering, it uses recurrent dynamics that partially resemble those in SOB.

Related approaches in semantic memory. Masson (1995) presented a distributed memory model of semantic priming that accounted for several findings in the literature. For example, the model accommodated the finding that the standard semantic priming effect (the word DOCTOR is processed faster if preceded by *nurse* than by *bread*) is eliminated if an unrelated word briefly appears

between presentation of the prime and the target. The model also handled the boundary conditions of that effect, which are a function of the extent to which the intervening word is processed relative to the prime.

Critically, Masson's (1995) model was based on a Hopfield (e.g., 1982) network that shares many properties with the present SOB architecture. For example, Masson's model used the Hebbian learning rule of SOB, and identification of a word was modeled as a descent on an energy landscape similar to that of SOB. Masson's model also included a pretrained vocabulary similar to the pretraining in SOB.

The original Masson (1995) model has been subject to extension, modification, and debate (e.g., Borowsky & Masson, 1996; Dalrymple-Alford & Marmurek, 1999a, 1999b; Masson, 1999), but there is little doubt that recurrent network dynamics constitute a powerful approach to semantic effects such as priming. This is further supported by the work of Kawamoto (e.g., 1993) on ambiguity resolution. Using the BSB architecture underlying SOB, Kawamoto (1993) that the model could handle various aspects of ambiguity resolution (i.e., identifying BANK as a place to deposit money, rather than as the edge of a river), including its time course and the role of context.

Taken together with the model of Haarmann and Usher (2001), those related applications in the semantic memory arena clearly show that a dynamic iterative approach to memory representation can accommodate a much broader range of findings than the serial recall phenomena reviewed here.

CONCLUSIONS

We introduced an endogenous mechanism for forward serial ordering, embodied in the SOB architecture, that was shown to account for a collection of benchmark phenomena. At study, items are encoded with decreasing strength across serial positions, and at retrieval, each recalled item is suppressed. SOB derives the strengths with which items are encoded and suppressed by analysis of the current contents of memory.

The fact that SOB can accommodate numerous phenomena at a fairly detailed level of analysis challenges other models of memory for serial order that have not provided a process implementation of the primacy gradient and response suppression on which they rely to model serial order.

REFERENCES

- ANDERSON, J. A. (1991). Why, having so many neurons, do we have so few thoughts? In W. E. Hockley & S. Lewandowsky (Eds.), *Relating theory and data: Essays on human memory in honor of Bennet B. Murdock* (pp. 477-507). Hillsdale, NJ: Erlbaum.
- ANDERSON, J. A. (1995). *An introduction to neural networks*. Bradford, MA: MIT Press.
- ANDERSON, J. A., SILVERSTEIN, J. W., RITZ, S. A., & JONES, R. S. (1977). Distinctive features, categorical perception, and probability learning: Some applications of a neural model. *Psychological Review*, *84*, 413-451.
- ANDERSON, J. R., & MATESSA, M. (1997). A production system model of serial memory. *Psychological Review*, *104*, 728-748.

- BEGIN, J., & PROULX, R. (1996). Categorization in unsupervised neural networks: The Eidos model. *IEEE Transactions on Neural Networks*, **7**, 147-154.
- BOROWSKY, R., & MASSON, M. E. J. (1996). Semantic ambiguity effects in word identification. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **22**, 63-85.
- 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 & 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.
- COWAN, N., WOOD, N. L., WOOD, P. K., KELLER, T. A., NUGENT, L. D., & KELLER, C. V. (1998). Two separate verbal processing rates contributing to short-term memory span. *Journal of Experimental Psychology: General*, **127**, 141-160.
- DALRYMPLE-ALFORD, E. C., & MARMUREK, H. H. C. (1999a). More on semantic priming in a fully recurrent network: A response to Masson (1999). *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **25**, 795-803.
- DALRYMPLE-ALFORD, E. C., & MARMUREK, H. H. C. (1999b). Semantic priming in fully recurrent network models of lexical knowledge. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **25**, 776-794.
- DENNIS, S., & KRUSCHKE, J. K. (1998). Shifting attention in cued recall. *Australian Journal of Psychology*, **50**, 131-138.
- DOSHER, B. A. (1999). Item interference and time delays in working memory: Immediate serial recall. *International Journal of Psychology*, **34**, 276-284.
- DOSHER, B. A., & MA, J.-J. (1998). Output loss or rehearsal loop? Output-time versus pronunciation-time limits in immediate recall for forgetting-matched materials. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **24**, 316-335.
- DREWNOWSKI, A., & MURDOCK, B. B. (1980). The role of auditory features in memory span for words. *Journal of Experimental Psychology: Human Learning & Memory*, **6**, 319-332.
- DUNCAN, M., & MURDOCK, B. B., JR. (2000). Recognition and recall with precuing and postcuing. *Journal of Memory & Language*, **42**, 301-313.
- ESTES, W. K. (1991). On types of item coding and sources of recall in short-term memory. In W. E. Hockley & S. Lewandowsky (Eds.), *Relating theory and data: Essays on human memory in honor of Bennet B. Murdock* (pp. 155-173). Hillsdale, NJ: Erlbaum.
- FARRELL, S., & LEWANDOWSKY, S. (2000). The case against distributed representations: Lack of evidence. *Behavioral & Brain Sciences*, **23**, 476-477.
- FRICK, R. W. (1989). Explanations of grouping in immediate ordered recall. *Memory & Cognition*, **17**, 551-562.
- GOLDEN, R. M. (1986). The "brain-state-in-a-box" neural model is a gradient descent algorithm. *Journal of Mathematical Psychology*, **30**, 73-80.
- GOLUBOV, B., EFIMOV, A., & SKVORTSOV, V. (1987). *Walsh series and transforms: Theory & applications*. Dordrecht, Netherlands: Kluwer.
- HAARMANN, H., & USHER, M. (2001). Maintenance of semantic information in capacity-limited item short-term memory. *Psychonomic Bulletin & Review*, **8**, 568-578.
- HAYKIN, S. (1994). *Neural networks: A comprehensive foundation*. New York: Macmillan.
- HENSON, R. N. A. (1996). *Short-term memory for serial order*. Unpublished doctoral dissertation, University of Cambridge.
- HENSON, R. N. A. (1998a). Item repetition in short-term memory: Ranschburg repeated. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **24**, 1162-1181.
- HENSON, R. N. A. (1998b). Short-term memory for serial order: The start-end model. *Cognitive Psychology*, **36**, 73-137.
- HENSON, R. N. A. (1999). Positional information in short-term memory: Relative or absolute? *Memory & Cognition*, **27**, 915-927.
- HENSON, R. N. A., NORRIS, D. G., PAGE, M. P. A., & BADDELEY, A. D. (1996). Unchained memory: Error patterns rule out chaining models of immediate serial recall. *Quarterly Journal of Experimental Psychology*, **49A**, 80-115.
- HINTON, G. E., & SEJNOWSKI, T. J. (1986). Learning and relearning in Boltzman machines. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. 1, pp. 282-317). Cambridge, MA: MIT Press.
- HITCH, G. J., BURGESS, N., TOWSE, J. N., & CULPIN, V. (1996). Temporal grouping effects in immediate recall: A working memory analysis. *Quarterly Journal of Experimental Psychology*, **49A**, 116-139.
- HOPFIELD, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, **79**, 2554-2558.
- HOPFIELD, J. J. (1984). Neurons with graded response have collective computational properties like those of two-state neurons. *Proceedings of the National Academy of Sciences*, **79**, 3088-3092.
- 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* (pp. 287-319). London: Academic Press.
- HOUGHTON, G., & HARTLEY, T. (1996). Parallel models of serial behaviour: Lashley revisited. *Psyche*, **2**(25). [Symposium on implicit learning and memory]. <http://psyche.cs.monash.edu.au>
- HULME, C., ROODENRYS, S., SCHWEICKERT, R., BROWN, G. D. A., MARTIN, S., & STUART, G. (1997). Word-frequency effects on short-term memory tasks: Evidence for a redintegration process in immediate serial recall. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **23**, 1217-1232.
- HUNT, R. R. (1995). The subtlety of distinctiveness: What von Restorff really did. *Psychonomic Bulletin & Review*, **2**, 105-112.
- KAWAMOTO, A. H. (1993). Nonlinear dynamics in the resolution of lexical ambiguity: A parallel distributed processing account. *Journal of Memory & Language*, **32**, 474-516.
- LEWANDOWSKY, S. (1999). Redintegration and response suppression in serial recall: A dynamic network model. *International Journal of Psychology*, **34**, 434-446.
- LEWANDOWSKY, S., & FARRELL, S. (2000). A redintegration account of word-length, lexicality, and articulatory suppression effects in short-term serial memory. *Psychological Research*, **63**, 163-173.
- LEWANDOWSKY, S., & LI, S.-C. (1994). Memory for serial order revisited. *Psychological Review*, **101**, 539-543.
- LEWANDOWSKY, S., & MURDOCK, B. B., JR. (1989). Memory for serial order. *Psychological Review*, **96**, 25-57.
- LI, S.-C., & LEWANDOWSKY, S. (1993). Intralist distractors and recall direction: Constraints on models of memory for serial order. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **19**, 895-908.
- LI, S.-C., & LEWANDOWSKY, S. (1995). Forward and backward recall: Different retrieval processes. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **21**, 837-847.
- LI, S.-C., LEWANDOWSKY, S., & DEBRUNNER, V. E. (1996). Using parameter sensitivity and interdependence to predict model scope and falsifiability. *Journal of Experimental Psychology: General*, **125**, 360-369.
- MADIGAN, S. A. (1971). Modality and recall order interactions in short-term memory for serial order. *Journal of Experimental Psychology*, **87**, 294-296.
- MARTIN, E., & NOREEN, D. L. (1974). Serial learning: Identification of subjective subsequences. *Cognitive Psychology*, **6**, 421-435.
- MASSON, M. E. J. (1995). A distributed memory model of semantic priming. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **21**, 3-23.
- MASSON, M. E. J. (1999). Semantic priming in a recurrent network: Comments on Dalrymple-Alford and Marmurek (1999). *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **25**, 776-794.
- MAYBERY, M. T., PARMENTIER, F. B. R., & JONES, D. M. (in press). Grouping of list items reflected in the timing of recall: Implications for models of serial verbal memory. *Journal of Memory & Language*.
- MCCLELLAND, J. L., & RUMELHART, D. E. (1988). *Explorations in parallel distributed processing*. Cambridge, MA: MIT Press.
- MCDOWD, J. M., & MURDOCK, B. B., JR. (1986). Mathematical models of memory and the problem of stimulus variation: A comparison of MINERVA2 and TODAM. *Acta Psychologica*, **62**, 177-188.
- METCALFE, J. (1993). Novelty monitoring, metacognition, and control in a composite holographic associative recall model: Implications for Korsakoff amnesia. *Psychological Review*, **100**, 3-22.
- MEWHORT, D. J. K., POPHAM, D., & JAMES, G. (1994). On serial recall: A

- critique of chaining in the theory of distributed associative memory. *Psychological Review*, **101**, 534-538.
- MURDOCK, B. B., JR. (1968). Serial-order effects in short-term memory. *Journal of Experimental Psychology Monographs*, **76** (1, Pt. 2).
- MURDOCK, B. B., JR. (1983). A distributed memory model for serial-order information. *Psychological Review*, **90**, 316-338.
- MURDOCK, B. B., JR. (1995). Developing TODAM: Three models for serial-order information. *Memory & Cognition*, **23**, 631-645.
- MURDOCK, B. B., JR. (2001). Analysis of the serial position curve. In H. L. Roediger III, J. S. Nairne, I. Neath, & A. M. Surprenant (Eds.), *The nature of remembering: Essays in honor of Robert G. Crowder* (pp. 151-169). Washington, DC: American Psychological Association Press.
- NAIRNE, J. S. (1992). The loss of positional certainty in long-term memory. *Psychological Science*, **3**, 199-202.
- NAIRNE, J. S., & NEATH, I. (1994). A critique of the retrieval/deblurring assumptions of the theory of distributed associative memory. *Psychological Review*, **101**, 528-533.
- O'TOOLE, A. J., DEFFENBACHER, K. A., VALENTIN, D., & ABDI, H. (1994). Structural aspects of face recognition and the other-race effect. *Memory & Cognition*, **22**, 208-224.
- PAGE, M. P. A. (2000). A localist manifesto. *Behavioral & Brain Sciences*, **23**, 443-467.
- PAGE, M. P. A., & NORRIS, D. (1998a). Modeling immediate serial recall with a localist implementation of the primacy model. In J. Grainger & A. M. Jacobs (Eds.), *Localist connectionist approaches to human cognition* (pp. 227-255). Hillsdale, NJ: Erlbaum.
- PAGE, M. P. A., & NORRIS, D. (1998b). The primacy model: A new model of immediate serial recall. *Psychological Review*, **105**, 761-781.
- RATCLIFF, R., VAN ZANDT, T., & MCKOON, G. (1999). Connectionist and diffusion models of reaction time. *Psychological Review*, **106**, 261-300.
- RIZZUTO, D. S., & KAHANA, M. J. (2001). An autoassociative neural network model of paired-associate learning. *Neural Computation*, **13**, 2075-2092.
- RYAN, J. (1969). Grouping and short-term memory: Different means and patterns of grouping. *Quarterly Journal of Experimental Psychology*, **21**, 137-147.
- VOUSDEN, J. I., & BROWN, G. D. A. (1998). To repeat or not to repeat: The time course of response suppression in sequential behaviour. In J. A. Bullinaria, D. W. Glasspool, & G. Houghton (Eds.), *Proceedings of the fourth neural computation and psychology workshop: Connectionist representations* (pp. 301-315). London: Springer-Verlag.
- WATKINS, O. C., & WATKINS, M. J. (1977). Serial recall and the modality effect: Effects of word frequency. *Journal of Experimental Psychology: Human Learning & Memory*, **3**, 712-718.

NOTES

1. Henson's (1998b) model is noteworthy because it postulates a recency gradient in addition to a primacy gradient, each represented by a distinct context marker. However, because the former typically far outweighs the latter (Henson, 1998b, p. 87), we include the model here among others that postulate a primacy gradient.
2. It is critical to differentiate between two classes of recency. On the one hand, in free recall, backward recall, and probed recall, recency is associated with items that are recalled *first*. In these situations, recency may therefore reflect a contribution from a separate *short-term memory* or *rehearsal buffer*. On the other hand, for the forward serial recall situations considered here, the terminal list items are necessarily recalled *last*. In addition, ignoring omissions, the lag (i.e., the combined number of study and recall events) between study and recall of an item is constant across all serial positions. This rules out any contribution of short-term memory to recency.
3. In some implementations, the starting vector (\mathbf{f}') is continually added into the dynamics specified in Equation 3 (e.g., J. A. Anderson, 1991). This was omitted here because the random cues used in SOB contain no relevant information. In the extremely unlikely event of a random cue's being orthogonal to all items presented in the matrices, the present dynamics thus would not converge on an attractor, because $\mathbf{x}(t)$ would be $\mathbf{0}$ throughout.
4. We thank Jeffrey Bowers for pointing this out to us.

(Manuscript received December 2, 1999;
revision accepted for publication April 25, 2001.)