INTERNATIONAL JOURNAL OF PSYCHOLOGY, 1999, *34* (5/6), 434–446

Redintegration and Response Suppression in Serial Recall: A Dynamic Network Model

Stephan Lewandowsky
University of Western Australia, Nedlands, Australia

This article presents a dynamic network model of redintegration and response suppression. Redintegration is the process that disambiguates partially retrieved memorial information into an overt response, and response suppression renders retrieved items temporarily unavailable for further report. Most models of serial recall assume the presence of both processes, but few explain how they are performed. Exploration of the network revealed that it can predict recency in serial recall of lists of varying lengths, the pattern of the associated transposition, omission, intrusion, and repetition errors, and the temporal dynamics of retrieval. The network thus augments any existing model of serial recall that assumes a distributed vector representation but does not specify the processes underlying redintegration and response suppression.

Cet article présente un modèle de réseau dynamique de reconstruction et de suppression de réponse. La reconstruction est le processus qui convertit les informations partiellement récupérées en réponses et la suppression de réponse rend les items récupérés temporairement non disponibles au rappel. La majorité des modèles de rappel sériel présupposent l'existence de ces deux processus, mais peu d'entre eux en expliquent l'éxécution effectuée. L'exploration du réseau révèle qu'il peut prédire l'effet de récence pour le rappel sériel de listes de différentes longueurs, les différents patrons d'erreurs et la dynamique temporelle de récupération. Le réseau peut être incorporé aux modèles existants de rappel sériel en spécifiant les processus sous-jacents à la reconstruction et à la suppression de réponse.

There is widespread agreement among memory theorists that serial recall involves not one but two distinct stages of processing. The first stage, called the *associative* stage in this article, is conceptualized as providing access to a memory trace and retrieval of at least partial information about the desired item and its ordinal position. The second stage, called here *redintegration*, refers to the disambiguation of that partial memorial information into a unique item and selection of an overt response. Most theorists also agree that redintegration is followed by response suppression, which renders recalled items temporarily unavailable for further retrieval.

Current theories of serial order memory that encompass those two stages fall into three broad classes, differentiated by their representational assumptions and explanatory emphasis: first, there are purely descriptive models that seek to identify the relative contributions of the associative and redintegrative stages through estimation of parameters (e.g. Brown & Hulme, 1995; Schweickert, 1993). These models do not specify any underlying cognitive processes, but by comparing parameter estimates across conditions, they can isolate the effect of experimental variables on the two stages. For example, redintegration has been consistently identified as the sole locus of word frequency and word lexicality effects in serial recall (e.g. Brown & Hulme, 1995; Hulme et al., 1997). However, these models do not predict the shape of the serial position curve (although Hulme et al. predicted the magnitude of the word frequency effect across serial positions), and they cannot capture the distribution of errors and the delicate balance among transpositions, omissions, and intrusions that are observed in serial recall (see Brown & Hulme, 1995, p. 617).

The second class of theories includes localist models that are specified at a process level. These models assume that each item is stored in a different location in memory and is represented by a distinct identifiable entity (e.g. Burgess & Hitch, in press; Henson, 1998; Page & Norris, in press a, b). Localist theories have much in common: all provide a process account of the associative and redintegration stages, all assume the presence of response suppression, and all handle a broad range of data, including the exact shapes of the serial position curve and underlying error gradients. These theories are,

Requests for reprints should be addressed to Stephan Lewandowsky, Department of Psychology, University of Western Australia, Nedlands, WA 6907, Australia (E-mail: lewan@psy.uwa.edu.au) (personal home page at URL: http://www.psy.uwa.edu.au/user/lewan/).

I wish to thank James S. Nairne for providing the data of his 1992 study in electronic form. Thanks are also due to Rik Henson, Ben Murdock, and Klaus Oberauer for their comments on an earlier draft of this paper. Simon Farrell and Mike Mundy assisted with manuscript preparation and their help is gratefully acknowledged.

This article was written while the author held a position as a Visiting Professor at the University of Potsdam, Germany, financed by funds from the Deutsche Forschungsgemeinschaft (German Research Foundation; Grant INK 12/A1, Project C).

however, limited to immediate serial recall of typically short lists.

Finally, distributed memory models represent items as vectors of features that, in contrast to localist models, are assumed to share a single storage location in memory (Brown, Preece, & Hulme, in press; Lewandowsky & Murdock, 1989; Murdock, 1982, 1983, 1993). Like their localist counterparts, these models specify the associative stage in quantitative detail and with conceptual precision. Like localist models, they assume that recalled items are suppressed. Distributed memory models also successfully account for benchmark data on serial recall; for example, the explanatory power of the recent model by Brown et al. (in press) rivals that of any localist theory. However, distributed theories have so far employed descriptive surrogates for redintegration rather than complete process implementations[1].

This article seeks to redress that potential shortcoming by presenting a stand-alone dynamic network model of redintegration that makes minimal assumptions about the associative stage. By remaining uncommitted to an associative stage, the network can potentially augment any model of memory in which a vector containing partial memorial information must be redintegrated (e.g. Brown et al., in press; Lewandowsky & Murdock, 1989; Murdock, 1993). In related precedents, stand-alone models that implement parts of a sequence of cognitive processes have been successfully used to constrain recognition models (e.g. Hockley & Murdock, 1987).

The article first reviews the evidence and theoretical considerations in support of a separate redintegration stage. I then examine the empirical and theoretical grounds for response suppression and outline how it can account for recency. Next, the redintegration network is applied to several benchmark findings. The model is shown to handle serial position effects across a range of list lengths and retention intervals. It also predicts the underlying pattern of transposition, intrusion, omission, and repetition errors. Finally, the model is shown to provide a natural account of retrieval latencies in serial recall. To conform to space constraints, the present article is restricted to single-trial effects in experiments involving visual presentation of phonologically nonconfusable items. Lewandowsky and Farrell (in press) applied the model to a variety of other paradigms, including the effects of articulation rate and word frequency.

THE NEED FOR REDINTEGRATION

The need for a separate redintegration process is obvious and most pressing in distributed models of memory, such as TODAM (e.g. Lewandowsky & Murdock, 1989), TODAM2 (Murdock, 1993), or OSCAR (Brown et al., in press). Central to these models is the assumption that the to-be-studied vectors are superimposed, by some mathematical process, in a common memory trace. This superimposed storage implies that the memorial information retrieved by the associative stage is, typically, a fuzzy approximation of the target that requires redintegration.

In TODAM, the early version of the theory (Lewandowsky & Murdock, 1989; Murdock, 1983) represented seriation by pairwise associations between adjacent items on the list. Items are associated by a process known as convolution, whereby the two constituent vectors are combined into another vector containing their symmetric association (for details, see Murdock, 1982). The convolved vector, and the constituent items, are then added to a common memory vector. At retrieval, cueing with one member of an association retrieves an approximation of the other item. Thus, in serial recall, the pairwise associations are consecutively probed to retrieve one item after the other. This chaining view of seriation has been repeatedly criticized (e.g. Henson, Norris, Page, & Baddeley, 1996; Mewhort, Popham, & James, 1994; Nairne & Neath, 1994), and in a major revision of the theory (TODAM2: Murdock, 1993), simple chaining was abandoned in favour of several other ways of representing serial order information (Murdock, 1995).

In OSCAR, by contrast, list items are associated not with each other but to a dynamically advancing timing signal provided by a set of oscillators. Unlike TODAM, associations are formed by computation of the outer product of the constituent vectors, which is then added to a common memory matrix. After study, the matrix contains the associations between each list item and the state of the timing signal at the time of its presentation. At retrieval, the dynamic timing signal is rewound to its initial state, and items are retrieved from the memory matrix in response to successive cueing with the advancing timing signal.

In both models, the output obtained in response to a cue is never exactly identical to any of the studied items or any other possible response. The model's performance therefore cannot be scored without some further mechanism to select a response[2]. The early TODAM (Lewandowsky & Murdock, 1989) and the current version of OSCAR (Brown et al., in press) sequentially compare the model's output to a set of available response candidates. The candidate that provides the best match is selected. One conceptual drawback of this scheme is that it assumes the existence of a lexicon of items, in which each response candidate has an identifiable and separate representation. This negates one of the basic premises of distributed memory models, that all information is stored in a common trace. In addition, the particular version of sequential comparison in the early TODAM (Lewandowsky & Murdock, 1989) has been the subject

[1] Chappell and Humphreys (1994) presented a model of memory that included a redintegration component. However, their emphasis was on recognition memory and they only reported a single simulation of recall involving paired associates. It is unclear whether their model could be extended to handle serial recall.

[2] This comment does not apply to the analytic version of TODAM that computes predicted recall probabilities without simulating individual study and test trials.

of critical attention (Mewhort et al., 1994; Nairne & Neath, 1994).

The need for redintegration following noisy selection of a response candidate arises also with localist models. For example, the Primacy Model (Page & Norris, in press a), which relies on differentially activated item nodes, relies on a second response stage to model the detrimental effects of phonological similarity. One critical attribute of phonological similarity is that, when lists contain a mix of similar and dissimilar items, recall of the dissimilar items is not impaired relative to a list containing only dissimilar items. Instead, the detrimental effect of similarity is limited to transpositions among the phonologically confusable items (e.g. Henson et al., 1996). Page and Norris (in press b, p. 13) reported that explorations of several single-stage models to handle those data were unsuccessful, concluding that ". . . the data alone appear to force us to accept a two-stage model. . . ." The same conclusion was reached for similar reasons by Henson (1998) during development of his Start-End Model (SEM). As in the Primacy Model, SEM postulates a second output competition that follows initial retrieval of a response candidate and that affects phonologically similar items[3].

Notwithstanding their explanatory diversity, most models thus agree on the need for a separate redintegration process. The models additionally agree that a recalled item is in most cases (at least temporarily) suppressed.

THE NEED FOR RESPONSE SUPPRESSION

Response suppression is required for a variety of empirical and theoretical reasons, predominant among them the frequent occurrence of pairwise transposition errors during recall. To illustrate with an example provided by Houghton and Hartley (1996), suppose a list of letters, such as T R A P, is actually recalled as T A R P, as is common in immediate serial recall. For this error to occur, A must be produced instead of R on the second retrieval. The further fact that A is then *not* produced in its correct third position suggests that its first recall was followed by suppression. If A had not been suppressed, errors such as T A A R P or T A A P should be quite common. In fact, however, these erroneous repetitions are rarely observed in serial recall (e.g. Henson et al., 1996; Vousden & Brown, 1998). Moreover, when erroneous repetitions do occur, they are most often separated by three or four intervening retrievals. Immediate repetitions are exceedingly rare (Vousden & Brown, 1998). This is compatible with the contention that an item, once recalled, is suppressed, and that the suppression gradually wears off while other items are recalled.

Turning to theoretical considerations, there is a pressing need for suppression in localist models that use a competitive cueing mechanism (e.g. Burgess & Hitch, in press; Houghton & Hartley, 1996; Page & Norris, in press a, b). In those models, the strongest or most active item is reported at each recall attempt. On the further assumption that activity is maximal for the first item and then decreases across list position, this mechanism will inevitably commence recall with the first item. However, to proceed through the list to progressively less active items, each retrieval must be followed by response suppression. A similar consideration applies to TODAM (Lewandowsky & Murdock, 1989) with its symmetric pairwise associations. Without suppression of previously recalled items, cueing with a list item would elicit not only the desired next item in the sequence but also the previous one. A final pragmatic reason for inclusion of response suppression is that it can provide a means of explaining recency.

VARIETIES OF RECENCY

Process explanations of recency in serial recall[4] tend to fall into two broad classes, referred to here as edge effects and response suppression. Edge effects refer to the inability of the terminal item to be involved in a transposition in more than one way because there are no adjacent items beyond the end of the list. The reduced frequency of transpositions necessarily results in recency. This mechanism principally contributes to recency in the Primacy Model (Page & Norris, in press a) and OSCAR (Brown et al., 1998). By themselves, edge effect explanations have two limitations. First, because items can exhibit recency only to the extent that the list boundary curtails their ability to transpose, the finding that relatively few transpositions involve nonadjacent items implies that edge effects cannot predict much recency beyond the terminal item. This runs counter to the observation that, even with visual presentation, recency often extends to the penultimate and even antepenultimate item (Madigan, 1971; Watkins & Watkins, 1977). Second, edge effects predict that recency should occur regardless of list length, which also appears to run counter to the data (e.g. Farrell, 1997).

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

[3] Technically, Henson (1998, Appendix B) postulates five stages in his model. However, these stages can be subsumed under two umbrella processes corresponding to associative retrieval (Henson's Stages 1 and 2) and redintegration (Stages 3 through 5).

[4] It is critically important to differentiate between two classes of recency. On the one hand, the recency observed in free recall, backward recall, and probed recall is associated with items that are recalled *first*. In these situations it is conceivable that recency reflects a contribution from some type of "short-term memory" or "rehearsal buffer". In forward serial recall, on the other hand, the terminal list items are necessarily recalled *last*. Moreover, ignoring omissions, the lag (i.e. the combined number of study and recall events) between study and report of an item is constant across all serial positions. This rules out any contribution of short-term memory to recency in forward recall, which renders it particularly theoretically interesting. This article focuses exclusively on forward recall.

alternatives remain, thus facilitating choice of the correct item. Henson (1998) acknowledged the contribution of response suppression to recency in his SEM model. Brown (personal communication, 18 August, 1998) suggested that suppression may likewise contribute to recency in OSCAR, although the contribution is difficult to differentiate from edge effects: Because suppression in OSCAR is all or none, removal of suppression would inevitably lead to an unreasonable increase in erroneous repetitions. Finally, Lewandowsky and Murdock (1989) relied entirely on response suppression to produce recency (see Mewhort et al., 1994; Nairne & Neath, 1994, for a cogent critique of that mechanism). 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.

A DYNAMIC CONNECTIONIST MODEL OF REDINTEGRATION AND RESPONSE SUPPRESSION

The discussion thus far permits the following conclusions. First, redintegration plays a major theoretical role in serial recall, with a process implementation being particularly indispensable for distributed memory models. Second, most models of serial recall assume that redintegration is followed by response suppression, primarily to accommodate the observed pattern of transpositions. Third, response suppression has been linked to the occurrence of recency, and may present a more powerful explanation than edge effects. The remainder of this article therefore presents a dynamic nonlinear auto-associative network that implements redintegration and response suppression (see Lewandowsky & Li, 1994, for an early sketch of this model).

Overview

The redintegration model was instantiated as a dynamic auto-associative network known as the "brain-state-in-a-box" (BSB) model (Anderson, Silverstein, Ritz, & Jones, 1977). The BSB has the key property that when cued with an arbitrary starting vector, the obtained output is iteratively fed back into the network until a stable state, known as an attractor, is reached. Attractors comprise all previously studied items (under the conditions of orthogonality assumed here) plus a number of additional "spurious" attractors.

To model redintegration, possible response candidates are first encoded in the BSB. In the present simulations, response candidates comprised the list items for a given trial. At the time of recall, it is assumed that the associative stage provides a partial response vector (call that $\mathbf{f}'$) that serves as starting vector for the BSB. If $\mathbf{f}'$ is sufficiently similar to the correct response, and it falls within the basin of attraction that surrounds the correct item, an exact copy of the target ($\mathbf{f}$) will be recovered. If $\mathbf{f}'$ 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. In all cases, once an attractor is reached, there is no longer any ambiguity about the identity of the network's response.

Because all attractors consist of symmetric binary vectors (i.e. sequences of -1, $+1$, . . .), they can be thought of as vertices in hyperspace. Moreover, because all activations are restricted to the range -1 through $+1$, the state of the network at any given time is confined to lie within the box formed by the attractors at the vertices—hence the name "brain-state-in-a-box." The presence of the box in conjunction with the iterative update dynamics guarantees that an attractor is eventually reached from any (nonpathological) starting point.

The BSB provides a natural way to implement response suppression through a mechanism known as anti-learning. Anti-learning refers to a selective *reduction* of the strength (through a negative learning rate) of an attractor. In previous applications, anti-learning has been used to model the multi-stable perception of the Necker cube (Anderson, 1991) or the change in perceived meaning of ambiguous words (Kawamoto, 1993). In the present context, whenever an attractor is reached and redintegration is complete, its strength is reduced through anti-learning.

Unlike response suppression in existing distributed memory models, the extent of anti-learning is determined by a parameter, thus allowing for *partial* suppression of recalled items. Partial suppression, in turn, can subsequently be reversed (e.g. Anderson, 1991). Although not relevant to the present set of single-trial simulations, this "release from suppression" is assumed to follow completion of recall and occurs through amplification of the weight matrix. Thus, release from suppression affects the entire response set rather than individual items. The implications of this mechanism were explored elsewhere (Farrell, 1997).

Associative Component: Assumptions

The redintegration model is known to function in conjunction with the early TODAM (Lewandowsky & Li, 1994). Here, emphasis was on extending the generality of the BSB by designing a stand-alone model that made only minimal assumptions about the associative stage. In line with virtually all models of serial recall, it was assumed that the quality of the information available in the associative stage decreased across serial position. This decreasing function has been variously described as a gradient of activation (Page & Norris, in press a), a decline in attention (Brown et al., in press), or a decreasing effectiveness of rehearsal (Lewandowsky & Murdock, 1989). For the present simulations, this assumption was embodied by:

$$s_j = c \times j^{-\lambda} \qquad (1)$$

where s represented the similarity between the output provided by the associative stage and the correct response. The constants c and λ were two free parameters representing, respectively, a starting value and the rate of decline of the similarity (s) across serial positions (j). Equation 1 represents only one possible form of the decreasing effectiveness assumption; see Lewandowsky and Li (1994) for an alternative. The similarity value, s, was used to create a random starting vector ($\mathbf{f}'$), by Monte Carlo means, with that specified similarity to the correct response ($\mathbf{f}$). The starting vector was then redintegrated using the following BSB dynamics.

BSB Dynamics

Study

During study, the weight matrix for the BSB, $\mathbf{A}$, was formed by superimposition of auto-associations of all list items using standard Hebbian learning:

$$\mathbf{A}_j = \mathbf{A}_{j-1} + w_j \ \mathbf{f}_j \mathbf{f}_j^{\mathrm{T}}, \tag{2}$$

where $\mathbf{f}_j \mathbf{f}_j^{\mathrm{T}}$ represents the outer product of the jth item with itself. For parsimony, the w_js were related to the similarity values in Equation (1) by the function $s_j = f^{\,s}\ w_j$, where $f^{\,s}$ is a parameter that remained fixed at .15 for most simulations.

$\mathbf{A}$ was of constant dimensionality 128 and initialized to zero at the outset. For each simulated list, study items were randomly sampled without replacement from a vocabulary formed by the complete set of Walsh vectors. Walsh vectors are symmetric binary vectors (i.e. −1, +1, . . .) that are mutually orthogonal (e.g. Golubov, Efimov, & Skvortsov, 1987). The size of the vocabulary was therefore equal to the dimensionality (128) of the weight matrix.

Recall and redintegration

At retrieval, for each serial position j, the starting vector $\mathbf{f}'_j$ was derived from the correct item $\mathbf{f}_j$ according to the similarity specified by Equation (1), such that the dot product between $\mathbf{f}'_j$ and $\mathbf{f}_j$ was equal to s_j. The length of $\mathbf{f}'_j$ was set to .001.[5] Creation of this vector modeled retrieval from the associative stage. The starting vector $\mathbf{f}'_j$ then migrated towards an attractor using standard BSB dynamics, with the "state" vector $\mathbf{x}$ at any time t given by:

$$\mathbf{x}(t) = g\,(\beta\, \mathbf{x}(t\text{-}1) + \varepsilon\, \mathbf{A}\, \mathbf{x}(t\text{-}1) + \delta\, \mathbf{f}'_j), \tag{3}$$

where $\mathbf{x}(t\text{-}1)$ is the preceding state at time t-1, $\mathbf{f}'_j$ is $\mathbf{x}$ at time t=0, and β, ε, and δ are fixed parameters. The function g truncates all activations to the range −1 to 1, thus constraining the network state to lie within the box formed by the attractor vertices. Given a non-zero initial input, the system is guaranteed to converge to some attractor, located in a vertex (i.e. all elements of $\mathbf{x}$ either −1 or +1), in a finite number of steps.

In the simulations, a response was scored as correct when the state vector reached the correct item. Additionally, the number of steps taken to reach an attractor provided a natural account of predicted response latency (Anderson, 1991; Ratcliff, Van Zandt, & McKoon, 1999).

Response suppression

To implement response suppression, once the system had redintegrated $\mathbf{f}'_j$ to some vector $\mathbf{x}$, that attractor was removed from $\mathbf{A}$ using:

$$\mathbf{A}_j = \mathbf{A}_{j-1} - \eta\, w_j \ \mathbf{x}\, \mathbf{x}^{\mathrm{T}}, \tag{4}$$

where η was a parameter and the w_js were as in Equation (2), except that j here indexed output position not (input) serial position.

Parameters

Fixed parameters comprised ε, β, and δ, set to .2, .9, and 1.0, respectively. There were three *free* parameters: c, the starting value for the similarity, λ, the rate of decline of similarity across serial positions, and η, the anti-learning fraction. Additionally, $f^{\,s}$, the scaling constant relating encoding weights and associative retrieval similarity, was increased for two of the simulations in order to capture the particularly high level of accuracy shown by subjects. For the demonstrations reported here, all parameter values (summarized in Table 1) were obtained by manual adjustment. All results were based on 1000 replications, each using a different random sample of list items drawn from the vocabulary of Walsh vectors.

Summary

To clarify the operation of the network, it is helpful to contrast the two time scales at which events occur. First, each presentation of an item for study corresponds to an update of the weight matrix $\mathbf{A}$ according to Equation (1) (although not modeled here, a parallel update would occur in the associative stage). Similarly, each completed redintegration is followed by a modification of $\mathbf{A}$ according to Equation (4) (without a necessary parallel in the associative stage). Both events are indexed by the positional index j, and both correspond to study or retrieval events in models such as OSCAR or TODAM. Within this coarse time scale, a second type of iteration takes place while a given item j is being redintegrated. The matrix $\mathbf{A}$ remains static across iterations within this finer time scale. At each time t at this finer scale, the state vector $\mathbf{x}$ is updated by renewed probing of $\mathbf{A}$ until an attractor is reached. This dynamic redintegration does not have a parallel in existing models of serial recall.

[5] In most studies, list items are drawn from the same class (e.g. digits, letters) and thus share defining features. To represent this residual similarity among list items, $\mathbf{f}'_j$ was derived from a linear combination of all list items, with unit weight for the correct item ($\mathbf{f}_j$) and a weight of .2 for the others. The simulation results remained qualitatively unchanged if $\mathbf{f}'_j$ was derived from the correct item only.

TABLE 1
Summary of Parameter Values Used across All Simulations

			Parameter Values			
Simulation	*List Length*	*Condition*	*c*	λ	η	*Comments*
1.1	6		.6	.4	.999	$f^s = .4$
1.2	9		.8	.8	.999	
1.3	24		1.0	.5	.999	
1.4	6		1.0	.6	.900	
1.5	6		1.0	.6	.750	
1.6	6		1.0	.6	.500	
2.1	6		1.0	.6	.999	
2.2	5	Immediate	1.0	.3	.900	
		4 hours	.5	.3	.900	
		24 hours	.3	.3	.900	
3.1	4 to 9		1.0	.6	.999	$f^s = .4$

SIMULATIONS

Simulation 1: Recency and List Length

Lewandowsky and Li (1994) showed that the BSB, when combined with the early TODAM, can produce recency through response suppression. The first simulation sought to extend this finding to lists of varying lengths. The results (Simulations 1.1–1.3 for list lengths 6, 9, and 24) are shown in the left panel of Fig. 1.

The figure shows representative data in its right panel. The data for list lengths six and nine were taken from Henson et al. (1996) and Hitch, Burgess, Towse, and Culpin (1996), respectively. Data for the longest list were taken from Farrell (1997). The redintegration model handled the recency that was observed with short and intermediate lists and, simultaneously, the observed elimination of recency with longer lists. This effect emerged without manipulation of parameters because anti-learning was identical in all cases. The elimination of recency resulted from the fact that with longer lists, there was a greater likelihood of intrusion errors occurring before terminal list items had to be recalled. Suppression of intrusions does not affect the strength of nonrecalled items; by implication, for longer lists, there is a relatively

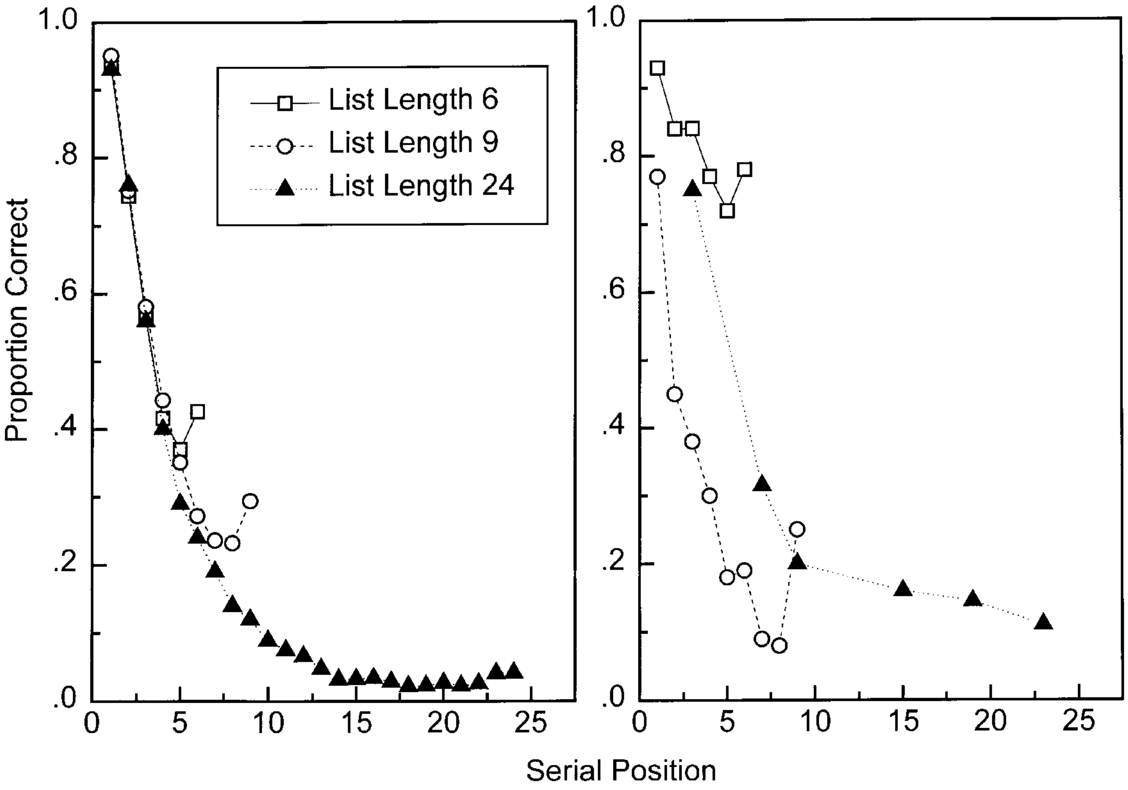

FIG. 1. Serial position curves for various list lengths. The left panel shows predictions of the redintegration model and the right panel shows the corresponding data.

greater number of strong attractors that compete for responses even towards the end of recall.

To confirm that response suppression was responsible for the occurrence of recency, another set of simulations using six-item lists was conducted that manipulated the value of the anti-learning parameter (η). To accentuate the effects of η, the other parameters were chosen to ensure steep primacy. The results (Simulations 1.4–1.6) are shown in Fig. 2 for values of η of .9, .75, and .5.

It is clear that the extent of recency in the redintegration model is tied to the magnitude of response suppression. Comparison of the serial position curves for η=.5 and η=.9 reveals that when response suppression is greatest, it exerts its effects on the last three to four list positions, despite the fact that proper recency (defined as an increase in performance from item n to item $n+1$) may not emerge until the last or penultimate position. By implication, models that explain recency primarily through edge effects, but which also include response suppression, may not in fact yield recency for the primary reasons cited (e.g. Brown et al., in press; Page & Norris, in press a).

Simulation 2: Errors in Serial Recall

There has been much recent emphasis on errors, owing to their presumed diagnostic value in differentiating between rival models of serial recall (Henson et al., 1996), in particular between chaining models and their alternatives. The second simulation analyzed the error pattern for six-item lists using the parameter values shown in Table 1 (Simulation 2.1). Four types of errors were examined: transpositions (first report of a list item in an incorrect position), intrusions (reporting an item not on the list), omissions (not reporting anything at a particular position), and erroneous repetitions (reporting a unique list item for the second time).

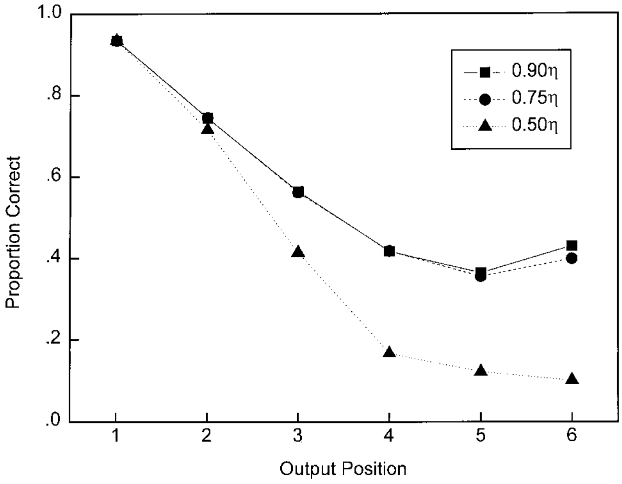

FIG. 2. The effect of reducing anti-learning (captured by the parameter η) on recency. All parameters other than η remain unchanged.

Transpositions

Figure 3 shows the predicted transposition gradients. Each plotted parameter refers to an output position and shows the proportion of items reported from each serial position. The peaks of each function represent correctly recalled items and, when connected, exhibit the typical serial position curve.

The model captured the pervasive finding that output positions tend to cluster around the items' serial positions. That is, items may erroneously migrate to adjacent positions but they are unlikely to be reported far from their true list position. Henson et al. (1996) called this the *locality constraint*.

The redintegration model satisfied the locality constraint without any item-to-item or item-to-position associations, but as a direct result of the differential strength of encoding at study. In the BSB, the size of the basins of attraction surrounding the list items is a function of the weighting received at study [see Equations (1) and (2)]. Thus, the first item has the largest basin of attraction and the last item the smallest. If the starting vector ($\mathbf{f}'_j$) is sufficiently different from the target ($\mathbf{f}_j$) to fall outside its basin of attraction, then the likelihood of $\mathbf{f}'_j$ reaching any of the remaining attractors is a sole function of their strength (ignoring for now the role of inter-item similarity). For example, if on the first retrieval the starting vector falls outside the correct basin of attraction, the second and third items are the most likely candidates for redintegration, thus giving rise to the observed positional gradient involving later list items. If, on the other hand, redintegration of the first item was successful, it is suppressed and thus no longer competes (much) for erroneous responses during subsequent retrievals. The positional gradient involving earlier items results from incomplete (or incorrect) prior suppressions.

A direct consequence of this mechanism is the "fill-in" phenomenon (Henson et al., 1996; Page & Norris, in press a). 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 items 2 and 3, the first item is most likely recalled next. Page and Norris reported a reanalysis of data by Henson et al. 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). In Simulation 2.1, the same pattern was observed without parameter manipulation: fill-ins outnumbered relative position reports by a factor of 3.99.

Intrusions and Omissions

It is well established that the number of intrusions and omissions increases across output position (e.g. Henson, 1998, fig. 6; Page & Norris, in press a, fig. 5). The redintegration model naturally handles intrusions through the inevitable presence of "spurious" attractors that do not represent studied items. Modeling of omissions is less straightforward because any starting vector will reach

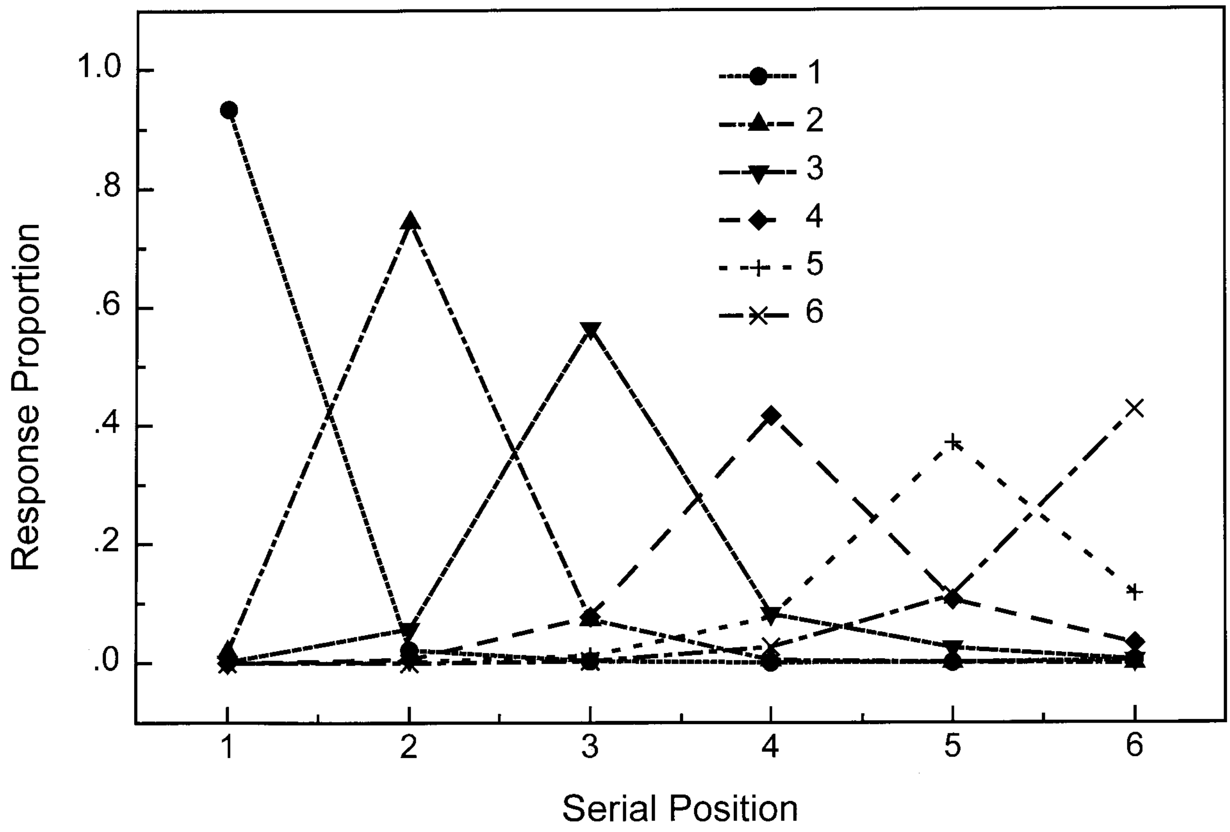

FIG. 3. Transposition gradients predicted by the redintegration model. Each parameter represents a unique output position and shows the proportion of responses from each original serial position.

an attractor in a finite number of steps. Given the implausibility of arbitrarily designating some attractors to represent nonresponses, omissions were instead classified as any response that exceeded a maximum number of iterations (set here to 10) before reaching a vertex. Figure 4 shows the predicted number of intrusions and omissions as a function of output position (note the different scales for the two types of errors).

The model captured the basic empirical pattern of omissions and intrusions: The steep increase in the number of omissions across output position corresponded to the functional relationship observed in the data (e.g. Henson et al., 1996). Likewise, the predicted negatively accelerated function for intrusions resembled the data reported by Page and Norris (in press a, fig. 5)[6]. Finally, predicted intrusions outnumbered omissions by a factor of about 10: This is identical to the ratio predicted by Henson's (1998, fig. 12) SEM model for an immediate test.

[6] Their figure reported data from an unpublished experiment in which omissions, intrusions, and repetition errors were summed together. The net contribution of intrusions is therefore not discernible from their data. However, Lewandowsky, Neely, Amos, and Ver Wys (1998) reported similar output position functions for intrusions in experiments in which subjects were not permitted to omit responses.

Repetition Errors

Erroneous repetitions of an item occur very infrequently—indeed, their low incidence was cited earlier in support of response suppression. For example, Henson (1996) reported that erroneous repetitions constituted 2% of all responses, and Vousden and Brown (1998) cited a figure of 5%. Nonetheless, repetition errors have a distinct distribution, with most repetitions involving early list items that are reported a second time late

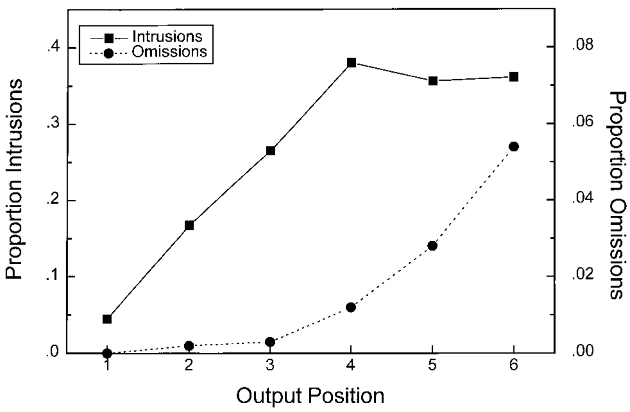

FIG. 4. Proportion of intrusions (scale on the left ordinate axis) and omissions (scale on the right) predicted by the redintegration model as a function of output position.

in recall. In consequence, repetition errors are typically separated by three or four output positions (e.g. Henson et al., 1996, observed an average of 3.34 positions apart).

This pattern was captured by the redintegration model. With the set of parameter values for Simulation 2.1, repetition errors constituted 0.7% of all responses and were separated by 3.57 output positions on average. As would be expected, a reduction of response suppression while keeping all other parameters constant was accompanied by an increased number of repetition errors, with η values of .8, .75, and .7 giving rise to 1.4%, 4.6%, and 11.1% erroneous repetitions, respectively. Remarkably, at a finer grain of analysis for one of those cases (Simulation 1.5; η=.75), the last output position included more reports of the first item (3.9%) than of the second and third (0.4% and 0.5%, respectively). This violation of the locality constraint occurs when a moderate number of repetition errors is present (Henson et al., 1996).

Overall, as with transpositions, intrusions, and omissions, the redintegration model predicted the correct pattern of repetition errors without parameter manipulation. Interestingly, the correctly predicted separation of repetitions occurred without a reduction of response suppression across output positions. This account goes beyond that offered by the Primacy Model, which lumps omissions, intrusions, and repetitions together as ''item'' errors (Page & Norris, in press a, b).

Retention Intervals

Many of the more recent models of serial order have been explicitly restricted to immediate recall from short-term memory (Burgess & Hitch, in press; Henson, 1998; Page & Norris, in press a, b). In some cases, this limitation was mandated by intrinsic problems, for example the properties of the context signal in the model by Burgess and Hitch (Brown et al., in press, p. 23). This limited scope presents a problem in the light of data showing the qualitative uniformity of serial recall performance across retention intervals (Nairne, 1992).

The redintegration model was applied to the data by Nairne (1992) by manipulating the starting value (c) of the primacy gradient [see Equation 1]. The change in c reflected the uncontroversial assumption that the strength of memorial information decreases over time. Figure 5 shows the results (Simulation 2.2) together with the data (Nairne, 1992). The redintegration model captured the observed flattening of transposition gradients over time without any change to response suppression.

Simulation 3: Response Latency in Serial Recall

In some areas of research, such as recognition memory, response latency has played a significant role in guiding and constraining theories for several decades. Accordingly, recognition memory is characterized by a large database of latency measurements (e.g. Murdock & Anderson, 1975). This contrasts sharply with serial recall in which latency analyses are rare. One exception is a recent paper by Dosher and Ma (1998) that explored the relation between total retrieval time and memory span.

The redintegration model was applied to the methodology of Dosher and Ma (1998) without additional assumptions or mechanisms. Total retrieval times were computed by summing the number of cycles [t in Equation (3)] required for redintegration across all output positions. To conform to Dosher and Ma's methodology, no omissions were permitted. The results (Simulation 3.1) are shown in Fig. 6 together with data from one experimental condition (word lists recalled by keypress).

The left panel of Fig. 6 shows that the model captured the general shape of the memory span function, although it under-predicts performance at intermediate list lengths and over-predicts performance for the longest lists. This may be a consequence of the unusually high performance levels observed by Dosher and Ma; memory span functions for words are typically much steeper (e.g. Crannell & Parrish, 1957). The high accuracy in the data again necessitated an increase in f^s.

The right panel shows total retrieval time as a function of list length. To express predictions in real time, redintegration cycles were divided by eight. The slightly quadratic component in the data is significant (Dosher & Ma, 1998, p. 322), and it was at least qualitatively captured by the model: The predicted retrieval time *per item* increased across list lengths from 691msec (length 4) to 858msec (length 9).

One implication of the latter result is that all changes in latency are assumed to arise during redintegration, and that processing at the associative stage must take a constant amount of time. (Because if list length additionally affected the associative stage, the predictions shown in Fig. 6 would necessarily diverge from the data.) This constant-duration assumption is compatible with models such as TODAM or OSCAR, in which retrieval is based on a single operation (convolution or matrix post-multiplication) whose duration is unaffected by experimental manipulations.

GENERAL DISCUSSION

Potential Criticisms

At least two major criticisms can be leveled against the redintegration model. First, the model may be seen to beg the entire question of seriation because it relegates the—arguably—most crucial component of memory retrieval to an hypothetical associative stage that is characterized by assumptions that are conveniently compatible with the redintegration architecture. Despite that, the redintegration model required three or four free parameters. Second, although the scope of the redintegration model extends beyond the current simulations (cf. Lewandowsky & Farrell, in press), the field already has

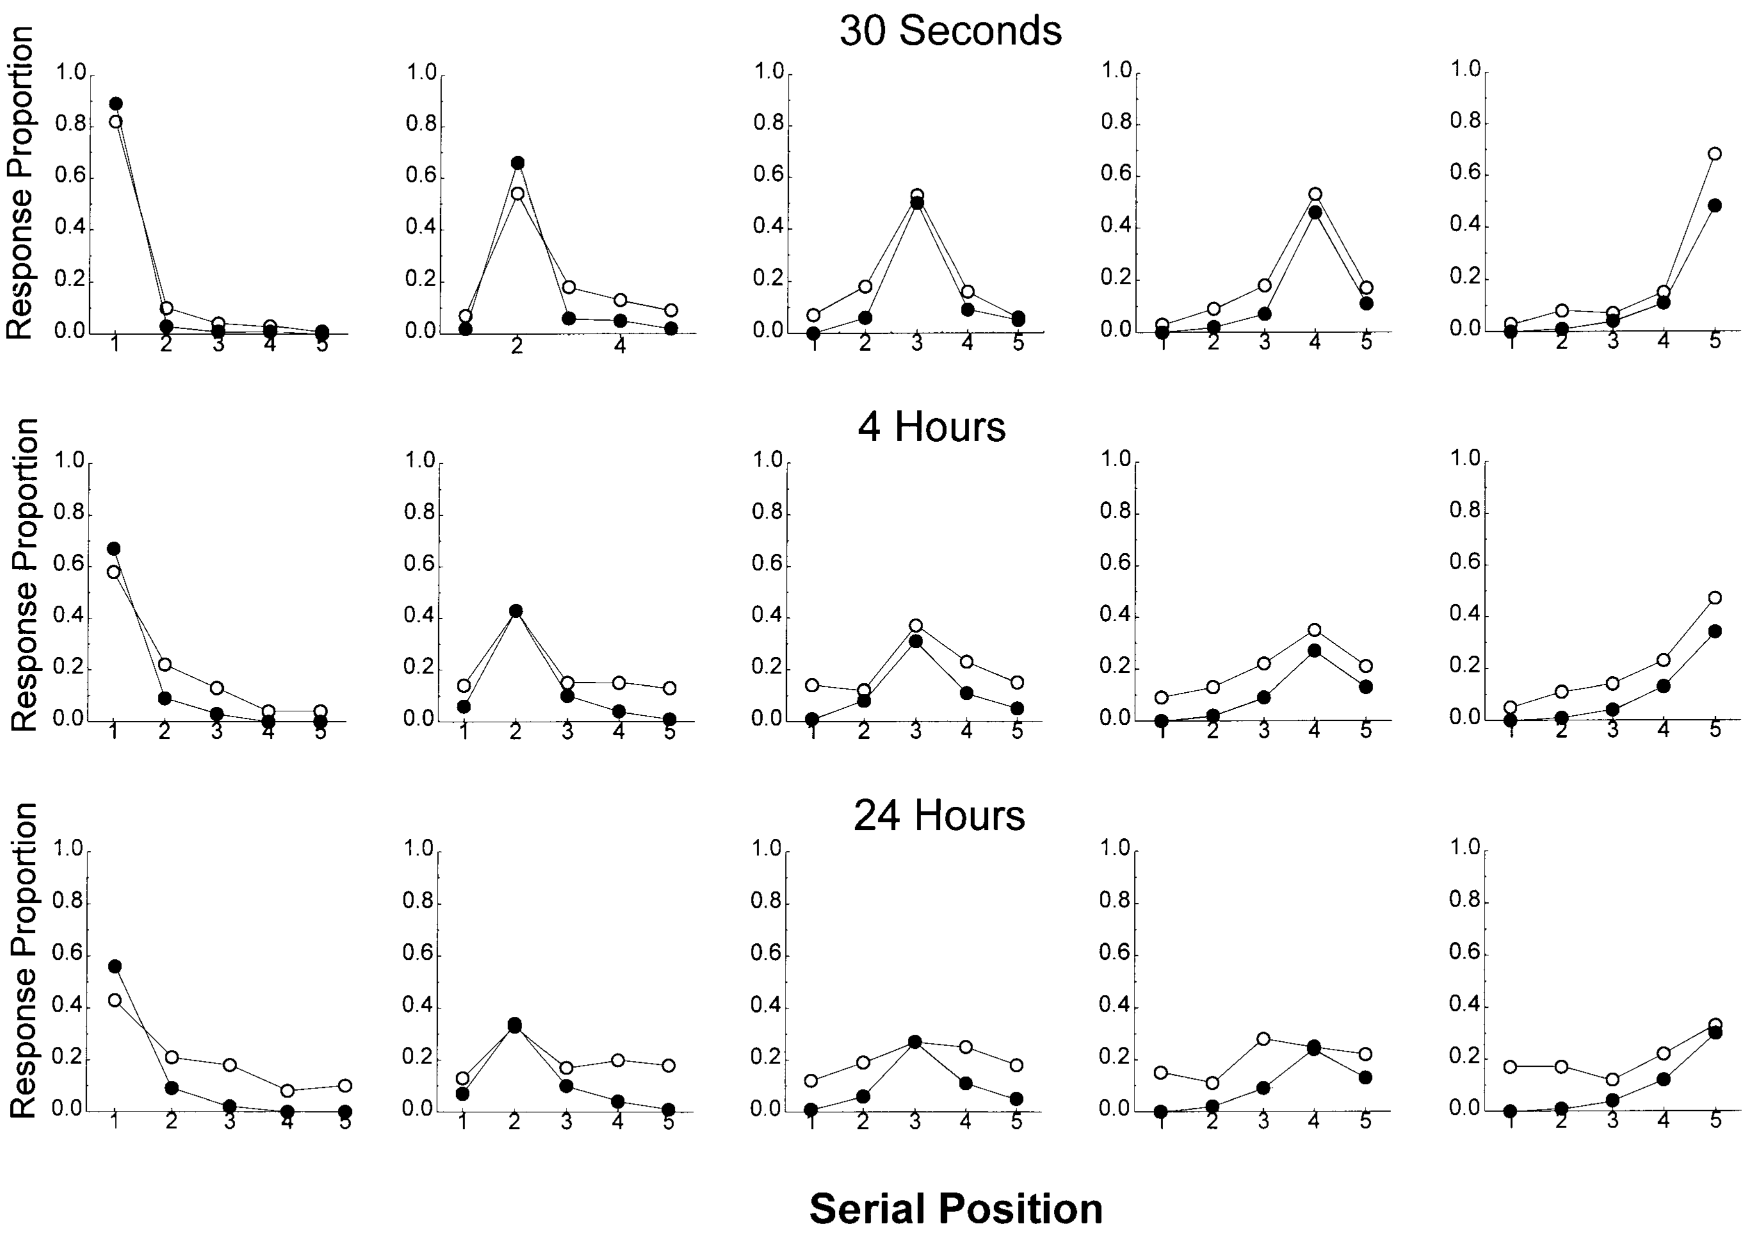

FIG. 5. The effects of retention interval on transposition gradients. The observed data (Nairne, 1992) are represented by open circles and the predictions of the redintegration model by filled circles.

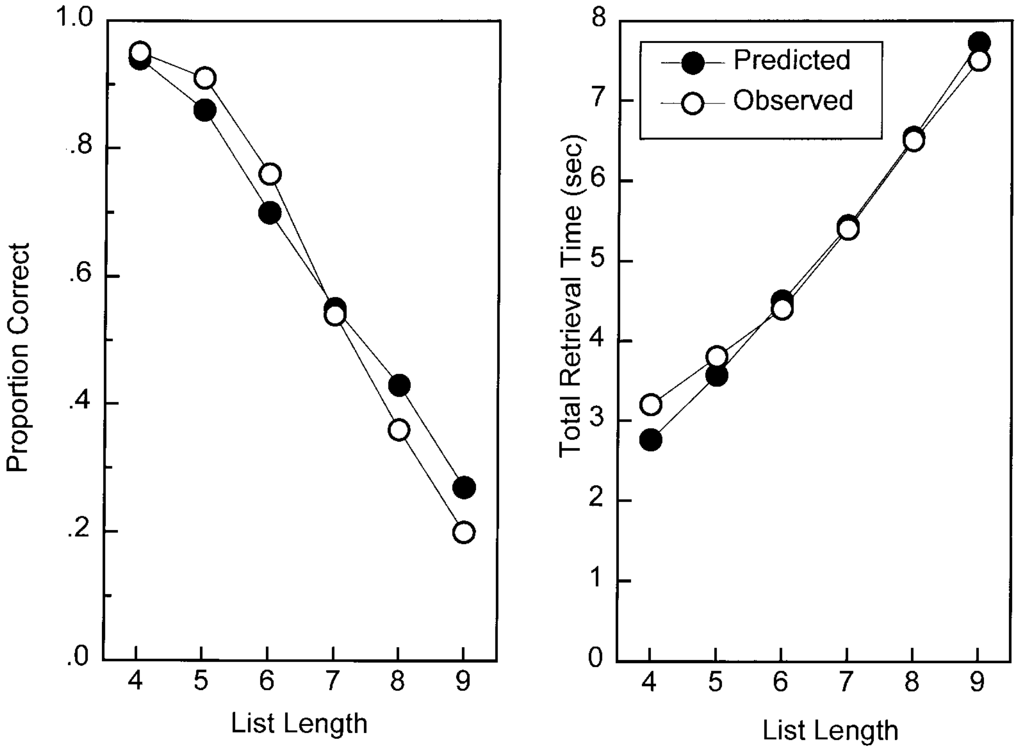

FIG. 6. The left panel shows the memory span functions observed by Dosher and Ma (1998) together with the corresponding predictions of the redintegration model. The right panel shows observed and predicted total time to recall as a function of list length.

access to a number of quite powerful process models. It might therefore be argued that there is little need for an additional contender.

The first issue can be settled by analysis of two candidate models of the associative stage, OSCAR and TODAM, both of which are demonstrably compatible with the redintegration model. TODAM has already been shown to function with the model in an integrated manner (Lewandowsky & Li, 1994). OSCAR has not been integrated with the redintegration model, but it is known to provide output of the general form specified by Equation (1) (see Brown et al., in press, fig. 11c). Two of the current BSB parameters, c and λ, would be replaced by those intrinsic to TODAM or OSCAR if an associative stage were implemented. The redintegration model therefore has only a single intrinsic free parameter (η), which in turn would replace at least one parameter intrinsic to OSCAR or TODAM, suggesting that the BSB provides a parsimonious account of several benchmark findings.

The second issue, concerning the need for a new (and only partial) model of serial recall, can be addressed in several ways. Most important is the fact that the redintegration model provides a natural—and so far successful—account of retrieval dynamics. Given the known theoretical diagnosticity of latency measures, it is essential to have available a model of retrieval dynamics: At present, the redintegration model is the sole existing candidate. Moreover, the redintegration model may be more general than some current localist theories. For example, Henson's (1998) SEM and the model by Burgess and Hitch (in press) are explicitly restricted to *immediate* serial recall, with little apparent opportunity for extension to longer retention intervals. Finally, the redintegration model can handle the observed pattern of errors at a greater level of detail (i.e. by differentiating between intrusions, omissions, and repetition errors) than the Primacy Model (Page & Norris, in press a).

Taken together, there is a clear need for a process model of redintegration and response suppression that is of general applicability, that can handle many list lengths and retention intervals, and that can model different classes of errors at a detailed level.

Constraints and Current Limitations

The redintegration model has at least one in-principle constraint. Many core predictions of the model are irrevocably tied to the decreasing strength with which successive list items are encoded. Although this assumption is virtually ubiquitous (e.g. Brown et al., in press; Henson, 1998; Houghton & Hartley, 1996; Lewandowsky & Li, 1994; Lewandowsky & Murdock, 1989; Page & Norris, in press a, b), there is little independent justification for it. Brown et al. (in press, p. 43) sought to provide such justification by appealing to the "intuition that each successive item . . . is progressively less 'surprising' or attention-demanding than the previous one". That intuition, in turn, was said to be consonant with the demands on an adaptively rational organism. At present, this reasoning must suffice to accept the decreasing-strength constraint as a necessary assumption of the redintegration model.

The redintegration model presently also has limited scope. It has not been applied to lists with repeated items; it does not account for any similarity effects, phonological or semantic; and it does not account for any of the variables surrounding phonology (e.g. modality effects and the effects of articulatory suppression). It is unclear how many of those effects will remain beyond the scope of the model, and thus need to be addressed by the associative stage, versus how many can be explained with further development work.

Relationship to Other Models

On the one hand, the redintegration model provides a process implementation of the descriptive approaches to redintegration pursued by Brown and Hulme (1995), Hulme et al. (1997), and Schweickert (1993). Those models estimated the relative contributions of memory retrieval and redintegration in a variety of situations, concluding that the effects of lexicality (word vs. nonword lists) and word type (e.g. word length and word frequency) primarily involved redintegration. By implication, the redintegration model should accommodate lexicality and word type effects: This was shown to be the case by Lewandowsky and Farrell (in press).

On the other hand, the redintegration model could also be modified to subsume a fairly simple associative stage, in which case it would resemble the Primacy Model (Page & Norris, in press a, b). There already are a number of similarities between the Primacy Model and the redintegration model: both postulate that item strengths decrease with serial position, both assume that item strength contributes to the likelihood of recall, both postulate response suppression as an integral part of retrieval, and neither contains any item-to-item or item-to-position associations. The principal difference between the models is that in the Primacy Model localist representations compete for output of the strongest item, whereas in the distributed redintegration model a retrieval cue is required to elicit, typically, the strongest item. By implication, if the redintegration model could recall items in their order of strength without external cues, it would represent a distributed implementation of the Primacy Model with the added capability of modeling retrieval dynamics. It turns out that, in principle, the redintegration model can "cue itself" by using a completely random vector as the starting state. Anderson (1995) illustrated the process by which a random cue moves to the attractor representing the strongest item (for a related idea, see Murdock, 1983, p. 320). This possibility would deserve exploration if one sought to incorporate the associative stage into the redintegration model.

CONCLUSION

The redintegration model was shown to account for several benchmark findings in memory for serial order: recency across a wide range of list lengths; transposition gradients across a wide range of retention intervals; the pattern of intrusions, omissions, and repetition errors in immediate recall; and the temporal dynamics of retrieval. It follows that these phenomena need no longer be addressed by a model of the associative stage. Indeed, this article showed that the single necessary contribution of such a model is to retrieve information with decreasing accuracy across serial positions.

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: Lawrence Erlbaum Associates Inc.

Anderson, J.A. (1995). *An introduction to neural networks.* Cambridge, 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.

Brown, G.D.A., & Hulme, C. (1995). Modeling item length effects in memory span: No rehearsal needed? *Journal of Memory and Language, 34*, 594–621.

Brown, G.D.A., Preece, T., & Hulme, C. (in press). Oscillator-based memory for serial order. *Psychological Review.*

Burgess, N., & Hitch, G. (in press). Memory for serial order: A network model of the phonological loop and its timing. *Psychological Review.*

Chappell, M., & Humphreys, M.S. (1994). An auto-associative neural network for sparse representations: Analysis and application to models of recognition and cued recall. *Psychological Review, 101*, 103–128.

Crannell, C.W., & Parish, J.M. (1957). A comparison of immediate memory span for digits, letters, and words. *Journal of Psychology, 44,* 319–327.

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, and Cognition, 24,* 316–335.

Farrell, S. (1997). *Serial and associative recall: Dissociating the effects of repeated testing.* Unpublished honours thesis, University of Western Australia.

Golubov, B., Efimov, A., & Skvortsov, V. (1987). *Walsh series and transforms: Theory and applications.* Dordrecht, The Netherlands: Kluwer Academic Press.

Henson, R.N.A. (1998). Short-term memory for serial order: The Start-End Model. *Cognitive Psychology, 36,* 73–137.

Henson, R.N.A., Norris, D.G., Page, M.P.A., & Baddeley, A.D. (1996). Unchained memory: Error patterns rule out chaining models of immediate serial recall. *Quarterly Journal of Experimental Psychology, 49A,* 80–115.

Hitch, G., Burgess, N., Towse, J., & Culpin, V. (1996). Grouping and immediate recall: A working memory analysis. *Quarterly Journal of Experimental Psychology, 49A,* 116–139.

Hockley, W.E., & Murdock, B.B. Jr. (1987). A decision model for accuracy and response latency in recognition memory. *Psychological Review, 94*, 341–358.

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, and Cognition, 23,* 1217–1232.

Kawamoto, A.H. (1993). Nonlinear dynamics in the resolution of lexical ambiguity: A parallel distributed processing account. *Journal of Memory and Language, 32,* 474–516.

Lewandowsky, S., & Farrell, S. (in press). A redintegration account of word-length, lexicality, and articulatory suppression effects in short-term serial memory. *Psychological Research*.

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.

Lewandowsky, S., Neely, J. H., Amos, A., & Ver Wys, C. (1998). *Failures to see repetition blindness in RSVP.* Manuscript in preparation.

Madigan, S.A. (1971). Modality and recall order interactions in short-term memory for serial order. *Journal of Experimental Psychology, 87,* 294–296.

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. (1982). A theory for the storage and retrieval of item and associative information. *Psychological Review, 89*, 609–626.

Murdock, B.B. Jr. (1983). A distributed memory model for serial-order information. *Psychological Review, 90*, 316–338.

Murdock, B.B. Jr. (1993). TODAM2: A model for the storage and retrieval of item, associative, and serial-order information. *Psychological Review, 100,* 183–203.

Murdock, B.B. Jr. (1995). Developing TODAM: Three models for serial-order information. *Memory & Cognition*, *23,* 631–645.

Murdock, B.B. Jr., & Anderson, R.E. (1975). Encoding, storage, and retrieval of item information. In R.L. Solso (Ed.), *Information processing and cognition: The Loyola Symposium* (pp. 145–194). Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

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.

Page, M.P.A., & Norris, D. (in press a). The Primacy Model: A new model of immediate serial recall. *Psychological Review.*

Page, M.P.A., & Norris, D. (in press b). 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.* Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Ratcliff, R., Van Zandt, T., & McKoon, G. (1999). Connectionist and diffusion models of reaction time. *Psychological Review*, *106*, 261–300.

Schweickert, R. (1993). A multinomial processing tree model for degradation and redintegration in immediate recall. *Memory & Cognition, 21,* 168–175.

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 and Memory, 3,* 712–718.