

Stephan Lewandowsky · Simon Farrell

A redintegration account of the effects of speech rate, lexicality, and word frequency in immediate serial recall

Received: 1 October 1998 / Accepted: 28 December 1998

Abstract Short-term serial recall performance is strongly affected by the nature of the items to be remembered. For example, memory span declines with decreasing speech rate (i.e., increasing pronunciation duration) of the items and, for a given speech rate, memory for non-words is poorer than for words. Similarly, words of high natural language frequency are recalled better than low-frequency words. Existing descriptive models have identified redintegration as underlying many of those effects. Redintegration refers to the process by which partially retrieved memorial information is converted into an overt response. This article presents a process model of redintegration based on a non-linear dynamic network, which is shown to handle the effects of speech rate, lexicality, and word frequency on memory span. Unlike previous descriptive efforts, the redintegration model also predicts the shape of the underlying serial position curves.

Introduction

The nature of the material to be remembered affects recall from immediate memory. For example, performance is a linear function of the speech rate of the material to be remembered, such that memory span declines as the time to pronounce the items increases (e.g., Baddeley, Thomson, & Buchanan, 1975). This speech-rate effect is easily accommodated by models that view memory span as equivalent to the amount of material that can be rehearsed – and thus prevented from decaying – within some constant amount of time (e.g., Baddeley, 1986; Schweickert & Boruff, 1986). Estimates for that constant are in the range of 1.5 to 2 s and re-

markably invariant (e.g., Baddeley et al., 1975; Schweickert & Boruff, 1986).

The rehearsal view also accommodates the observed elimination of the speech-rate effect when subjects recite irrelevant material during list presentation (e.g., by repeatedly counting aloud from 1 to 8; Baddeley et al., 1975). Under these “articulatory suppression” conditions, memory span is relatively poor but equal for items with short and long pronunciation durations, presumably because the rehearsal that could prevent trace decay is suppressed (Baddeley, 1986).

However, an unmodified rehearsal view has difficulty with several other benchmark effects relating to the type of material (for a detailed analysis, see Brown & Hulme, 1995). Most relevant here is the impact of lexicality and word frequency on the function relating memory span to speech rate. Although that function nearly always linearly increases when articulatory suppression is absent, its intercept varies considerably with the type of material. For example, words are generally remembered better overall than non-words, independent of pronunciation duration (e.g., Hulme, Maughan & Brown, 1991; Hulme, Roodenrys, Brown, & Mercer, 1995). Memory for non-words varies with the extent to which they resemble words: Non-words that sound like words (e.g., BRANE) give rise to a higher span than those that do not (e.g., SLINT; Besner & Davelaar, 1982). Likewise, memory for words varies with their familiarity: Words with high natural language frequency are remembered better than low-frequency words – again, independent of pronunciation duration (Hulme et al., 1997). Finally, the memorial disadvantage for non-words can be partially eliminated if their intended pronunciation (Hulme et al., 1995) or their intended pronunciation and meaning (Hulme et al., 1991) are known to subjects. Taken together, these results strongly challenge a pure rehearsal view, because they reveal the contributions of long-term semantic and phonological knowledge in memory-span tasks.

In consequence, recent theorizing has sought alternative accounts for the effects of speech rate, lexicality,

S. Lewandowsky (✉) · S. Farrell
Department of Psychology,
University of Western Australia, Nedlands, W.A. 6907, Australia;
e-mail: lewan@psy.uwa.edu.au;
URL: <http://www.psy.uwa.edu.au/user/lewan/>

and word familiarity. On the one hand, this effort has culminated in a detailed process model known as the feature model (e.g., Neath & Nairne, 1995; Neath, Surprenant, & LeCompte, 1998), which can explain the effects of speech rate without assuming trace decay. Instead, speech rate affects memory because longer items are thought to be more susceptible to errors during “assembly” of the trace at recall. On the other hand, Brown and Hulme (1995) and Schweickert (1993; Hulme et al., 1997) have provided purely descriptive models that also abandon the notion of rehearsal. The focus here is on these models because they complement our own approach (e.g., Lewandowsky, in press).

The descriptive models by Schweickert (1993) and Brown and Hulme (1995) divide memory retrieval into two broad stages of processing whose relative contributions are estimated by parameters. Comparison of parameter estimates across conditions can then isolate the effect of experimental variables on the two stages. The first stage, called the associative stage in this article, provides access to a memory trace and (at least partial) retrieval of the desired item and its ordinal position. The second stage, called here redintegration, disambiguates partial memorial information into an overt response. The distinction between associative retrieval and subsequent redintegration is shared by many theories: It is particularly important in distributed memory models (e.g., Brown, Preece, & Hulme, 1998; Lewandowsky & Murdock, 1989; Murdock, 1993) which require redintegration of the necessarily fuzzy information provided by a distributed memory.

When estimating model parameters from the data, in nearly all cases redintegration emerges as the primary or sole locus of lexicality and familiarity effects. For example, Hulme et al. (1997) empirically compared memory span for high- and low-frequency words of different speech rates. Although speech rate was strongly and positively related to span, as predicted by the ubiquitous rehearsal view (e.g., Baddeley, 1986), an additive effect of word frequency remained that was interpreted as a contribution from phonological long-term memory during redintegration. Hulme et al. (1997) rejected the possibility that high-frequency words were recalled directly from (episodic) long-term memory, because the items most likely to have been transferred to long-term memory – namely, early list items – showed the smallest word-frequency effect. Similarly, parameter estimates for Schweickert’s (1993) multinomial tree model ascribed the word-frequency effect on span as well as its increasing magnitude across serial positions entirely to redintegration.

Brown and Hulme (1995) modeled the standard speech-rate function by two opposing processes. On the one hand, longer items were assumed to suffer more forgetting than shorter ones, paralleling conventional trace decay (e.g., Baddeley, 1986). On the other hand, the success of redintegration was assumed to be greater for longer items, and more likely still for words than non-words. By adjusting the weightings of those two

processes, Brown and Hulme correctly predicted memory span to be greater for words than non-words (with familiarized non-words in between). In addition, the span-speech function was predicted to be steeper for non-words than for words. Importantly, all effects of lexicality and non-word familiarity were again confined to redintegration.

Overall, although the descriptive models rely on redintegration to account for lexicality and familiarity effects, they fail to describe the process beyond suggesting that “...it is likely to be a complex psychological operation...” (Brown & Hulme, 1995, p. 600). Moreover, as acknowledged by Brown and Hulme (p. 617), descriptive models do not handle any closely related findings such as the shape of the serial position function and the delicate balance among transpositions, omissions, and intrusion errors.

This article presents a stand-alone process model of redintegration, based on a non-linear dynamic neural network, that has previously been shown to account for various aspects of serial recall, including the serial position curve and the associated pattern of transpositions, omissions, and intrusion errors (Lewandowsky, in press). To maximize generality, the model does not implement an associative stage. Instead, the output from an hypothetical associative stage is simulated corresponding to the behavior of several fully-specified relevant models (e.g., OSCAR: Brown et al., 1998; TODAM: Lewandowsky & Murdock, 1989). To illustrate, consider the properties of OSCAR, which assumes that list items are associated with a dynamically advancing timing signal provided by a set of oscillators. At recall, the timing signal is rewound to its initial state, and items are retrieved from memory in response to successive cueing by the advancing oscillators. Because the encoding strength of items is assumed to decrease across serial positions (an assumption shared by TODAM and other models), the retrieved information bears increasingly less resemblance to the target item as recall proceeds. This general assumption of decreasing resemblance was used here to model the output from an associative stage.

We shall first show that the redintegration model handles the lexicality and familiarity effects previously ascribed to redintegration. In particular, we shall demonstrate that longer words are more easily redintegrated than short words, and that pre-training of the network with a corpus of words yields the standard lexicality effect. This implements one of the opposing processes postulated by Brown and Hulme (1995). We then show that when augmented with standard assumptions about the associative stage – the other opposing process of Brown and Hulme – the redintegration model predicts the function relating memory span to speech rate for words and non-words. Unlike descriptive models, the redintegration model simultaneously handles the underlying serial position curves. A final simulation extends the use of differential pre-training to word-frequency effects. The model handles the effect of word

frequency on memory span and the increasing effects of frequency across serial positions.

A dynamic network model of redintegration

Overview

The redintegration model was based on the “brain-state-in-a-box” (BSB) network (Anderson, Silverstein, Ritz, & Jones, 1977). The BSB is an auto-associative dynamic network that iteratively modifies its initial input until a stable state, known as an attractor, is reached. Under standard assumptions, any previously studied item serves as an attractor. Additional “spurious” attractors typically exist that do not represent studied items. The BSB achieves redintegration by using the partial memorial information provided by the associative stage as its starting state, which is then iteratively fed back into the network until an attractor – and hence, a disambiguated response – is reached.

At the outset, response candidates are encoded in the BSB. These include list items and, if pre-training is present, an additional vocabulary of lexical items. At recall, the associative stage is assumed to provided a partial response vector (call it \mathbf{f}') which serves as input to the BSB. If \mathbf{f}' is sufficiently similar to the correct response and thus falls within the basin of attraction that surrounds the correct item, 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 spurious attractor, an extra-list intrusion. In all cases, once an attractor is reached, there is no longer any ambiguity about the identity of network’s the response.

Associative stage: Assumptions

The BSB can augment any model of the associative stage in which a vector containing partial memorial information must be redintegrated (e.g., OSCAR: Brown et al., 1998; TODAM: Lewandowsky & Murdock, 1989; TODAM2: Murdock, 1993). In line with these models, it was assumed that the quality of the information available in the associative stage decreased across serial positions. This decreasing function has been variously described as a decline in attention (Brown et al., 1998) or a decreasing effectiveness of rehearsal (Lewandowsky & Murdock, 1989), and it was implemented here as:

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

where s represented the similarity between the output of 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). In lieu of fully specifying an associative stage, the simulations below

used a random starting vector (\mathbf{f}') – generated by Monte Carlo means – whose similarity to the correct response (\mathbf{f}) was equal to s . The starting vector was then redintegrated using standard BSB mechanics.

BSB dynamics

Study

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

$$\mathbf{A}_j = \mathbf{A}_{j-1} + w_j \mathbf{f}_j \mathbf{f}_j^T, \quad (2)$$

where $\mathbf{f}_j \mathbf{f}_j^T$ represents the outer product of the j th item with itself. For parsimony, the w_j s were related to the similarity values in Equation 1 by the function $s_j = f^s \times w_j$, where f^s was a parameter. Thus, the weighting of items learned by the BSB followed the same primacy gradient that was assumed for the associative stage (cf. Eq. 1).

Matrix \mathbf{A} was of dimensionality 333 and initialized to $\mathbf{0}$ at the outset. Word length – and, by implication, speech rate – was represented by the number of non-zero bits used in the vectors (\mathbf{f}_j) representing the study items. Short words were represented by 200 non-zero bits (randomly set to -1 or $+1$ with equal probability), medium words by 250 bits, and long words by 333 bits. All non-zero bits were contiguous and padded by zeros on the right. Using an arbitrary constant for pronunciation of 500 bits per second, word length translated into speech rates of 2.5, 2.0, and 1.5 per second, for short, medium, and long words, respectively.

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 cosine between \mathbf{f}'_j and \mathbf{f}_j was equal to s_j , and the length of \mathbf{f}'_j was equal to .0001.¹ The vector \mathbf{f}'_j migrated towards an attractor by iterative computation of the network’s state. The “state” vector \mathbf{x} of a BSB at time t is given by:

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

where $\mathbf{x}(t-1)$ is the 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 . Given a non-zero initial input, the system is guaranteed to converge to some attractor (i.e., all elements of \mathbf{x} either -1 or $+1$) in a finite number of steps. The attractor reached after convergence represents the output of the model. In

¹ To reflect the effects of 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 0.2 for all remaining lists items. The simulation results remain qualitatively unchanged if \mathbf{f}'_j is derived from the correct item only.

the current simulations, recall was considered correct only when a list item was recalled (i.e., its corresponding attractor was reached) in the correct serial position. Any other response (i.e., a transposition with another list item or an extra-list intrusion: omissions were not permitted to occur) was counted as an error. When assessing convergence or scoring a response, only those elements of \mathbf{x} were considered that were non-zero in the target item (cf. Kawamoto, 1993).

Response suppression

In addition, response suppression was implemented by attenuating the attractor reached during redintegration. Strong empirical reasons exist for the presence of response suppression – for example, the extremely low incidence of erroneous repeated reports of a list item (see Lewandowsky, in press, for more details). Thus, following a response, the redintegrated vector \mathbf{x} was partially removed from \mathbf{A} through anti-learning:

$$\mathbf{A}_j = \mathbf{A}_{j-1} - \eta w_j \mathbf{x} \mathbf{x}^T, \quad (4)$$

where η was a parameter (set to .9 in all simulations below) and the w_j s were as in Equation 1, except that j here indexed output position, not serial (input) position. In previous BSB applications, anti-learning has been used to model the multi-stable perception of the Necker cube (Anderson, 1991), the dynamic change in perceived meaning of ambiguous words (Kawamoto, 1993), and recency in serial recall and associated error patterns (Lewandowsky, in press; Lewandowsky & Li, 1994). Unlike response suppression in other models (e.g., OSCAR: Brown et al., 1998; TODAM: Lewandowsky & Murdock, 1989; TODAM2: Murdock, 1993), the extent of anti-learning in the BSB is determined by a continuous parameter, thus allowing for partial suppression. Moreover, when suppression is partial, it can subsequently be reversed (e.g., Anderson, 1991), allowing for the fact that suppression is only temporary.

Anti-learning can be problematic if it renders the BSB dynamics unstable. To ensure the stability of the dynamics, the weight matrix \mathbf{A} must contain only positive eigenvalues (Anderson, 1995, p. 502), a condition met by Equation 2 if all w_j s are positive. When \mathbf{A} contains negative eigenvalues, the sign of the feedback component in the updating function (i.e., the term $\mathbf{A}\mathbf{x}(t-1)$ in Eq. 3) changes at each iteration, thus interfering with the ability of the network to reach a vertex. Negative eigenvalues can be introduced by anti-learning (see Eq. 4) if it involves an item not previously studied by the BSB, as in the case where a spurious attractor – an attractor corresponding to a non-studied vector – is reached at recall. Empirically, this possibility did not present a problem in previous applications of anti-learning (e.g., Anderson, 1991; Kawamoto, 1993; Lewandowsky, in press; Lewandowsky & Li, 1994); however, in the present simulations, additional steps were taken to prevent the occurrence of negative eigenvalues by preventing

anti-learning of spurious attractors. This was achieved by analyzing the pre-convergence behavior of the BSB.

The movement towards an attractor can be expressed as a gradient descent along an energy surface (e.g., Golden, 1986; Kawamoto, 1993). The energy of the network state at time t is defined as:

$$E = -(1/2) \sum_i \sum_j \mathbf{A}_{ij} a_i a_j, \quad (5)$$

where a_i and a_j are the activations of units x_i and x_j (the elements of the state vector \mathbf{x}), respectively, summed here across the units that were non-zero as determined by word length.

The pattern of descent in energy space turns out to differ between list items and other responses, both in the steepness of the gradient and in the final energy of the converged state. In the current simulations, we first compared the mean energy of converged state vectors for list items and other responses. List items consistently yielded more extreme (i.e., lower) energy values than spurious attractors, allowing the definition of a criterion for energy (b):

$$b = -(c/\phi) \times n, \quad (6)$$

where c was as in Equation 1, ϕ was a free parameter, and n represented a scaling constant (the square of the number of non-zero bits for the given word length). The criterion b was used to gate response suppression, such that if a converged vector fell below b , it was suppressed through anti-learning, as specified in Equation 4. Conversely, if a converged vector fell above b , response suppression was omitted. Consequently, only intra-list responses tended to be suppressed. Owing to the overlap of the energy distributions of the two classes of converged vectors, spurious attractors could occasionally fall below b , in which case they were also suppressed. However, throughout the simulations, there was no evidence of the occurrence of negative eigenvalues.

Parameters

The constants in the iterative updating function were fixed parameters (e , b , and d ; set to .2, .9, and 1.0, respectively). There were four free parameters: c , the starting value for the similarity function; λ , the rate of decline of similarity across serial positions; f^s , the translation constant between similarity and encoding strength; and ϕ , the parameter for the energy criterion that gated response suppression. For the demonstrations reported here, parameter values were obtained by manual adjustment. All results were based on 500 replications, each using a different randomly created set of study vectors.

Simulations

The first simulation sought to establish the basic applicability of the redintegration model by examining its

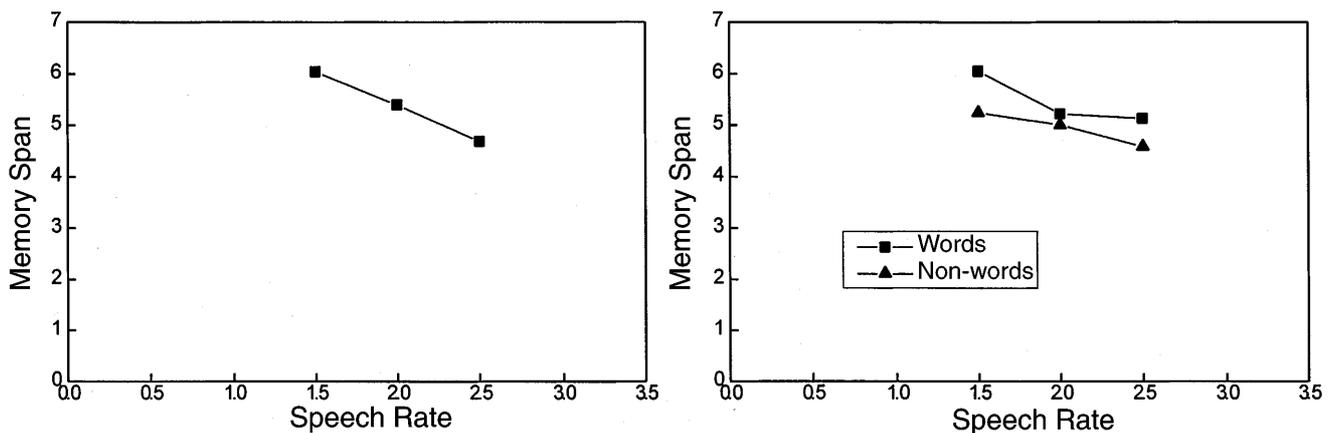
parameter-free behavior. Brown and Hulme (1995) suggested that long items should enjoy a redintegration advantage over short items, because loss of a single phoneme from, say, HIPPOPOTAMUS is less likely to be consequential than loss of a single phoneme from CAT. In addition, regardless of speech rate, words should be redintegrated more easily than non-words. Ideally, the redintegration model should produce these two effects without parameter manipulation.

Simulation 1: Effects of speech rate and lexicality on redintegration alone

To investigate the net effect of item length at redintegration, the first simulation used a constant set of parameter values ($c = .6$; $\lambda = .2$; $f^s = .3$, $\phi = 2.7$) across the three different speech rates. The results are shown in the left panel of Fig. 1. As in most behavioral experiments, simulated memory-span values were computed by linear interpolation of performance between the two list lengths bracketing 50% perfect ordered recall.

Each data point in the left panel represents an independent simulation run based on an initially empty (i.e., $A_0 = \mathbf{0}$) weight matrix, the only difference between data points being the number of non-zero bits in the study vectors to model different word lengths. It is clear from the figure that the BSB predicted a large performance advantage for long items (i.e., those with a low speech rate) over short items (high speech rate). However, in contrast to the descriptive model used by Brown and Hulme (1995), the BSB required no parameter manipulation to yield this effect. Instead, the advantage for long items was a natural consequence of the well-known fact that, in general, neural networks are better able to represent information as the number of elements in the vectors – and hence the number of connection weights – increases.

Fig. 1 Simulation 1: Predicted memory span values as a function of speech rate and lexicality under identical starting conditions. The observed advantage for long words, the reverse of the behavioral data, falls out of the BSB and is consonant with standard assumptions about the role of redintegration



Now consider the pattern shown in the right-hand panel of Fig. 1, which was obtained by pre-training the BSB to represent prior knowledge and the effects of lexicality. Pre-training involved study of items from a corpus of 150 randomly created “words” using a constant encoding weight w^L of .13 (see Eq. 2). To avoid cross-talk or interference, a separate set of 50 patterns was used for words of each length. Following pre-training, redintegration performance was simulated as before, using the same parameter values that gave rise to the pattern in the left panel (i.e., $c = .6$; $\lambda = .2$; $f^s = .3$), with the exception of the criterion parameter ϕ , which differed between words (2.15) and non-words (2.4). Word lists were obtained by sampling items from the appropriate pre-trained corpus, whereas non-word lists consisted of random novel vectors.

Several observations can be made about the effects of pre-training. First, it is clear that the advantage for words emerged without parameter manipulation. This lexicality effect reflected the fact that repeated presentation of an item increased the strength of its single associated attractor. All other things being equal, attractor strength maps into the probability of successful redintegration; consequently, (pre-trained) words were recalled better than (novel) non-words. Second, consider the difference in performance between the random novel items in the left panel and the identically created non-words in the right panel. It is clear that pre-training had a detrimental effect on redintegration of subsequently studied novel items. This occurred because pre-training created many additional attractors that, though smaller than those of studied items, competed with the correct target at redintegration. Words did not suffer from this competition because their attractors were among those being pre-trained.

The first simulation supported two conclusions. First, it showed that the BSB can provide a parameter-free account of the effects previously ascribed to redintegration by descriptive models (e.g., Brown & Hulme, 1995). Second, it supported the contention that the positive relation between memory span and speech rate must reflect an opposing effect of the associative stage that outweighs – and thus reverses – the role of redin-

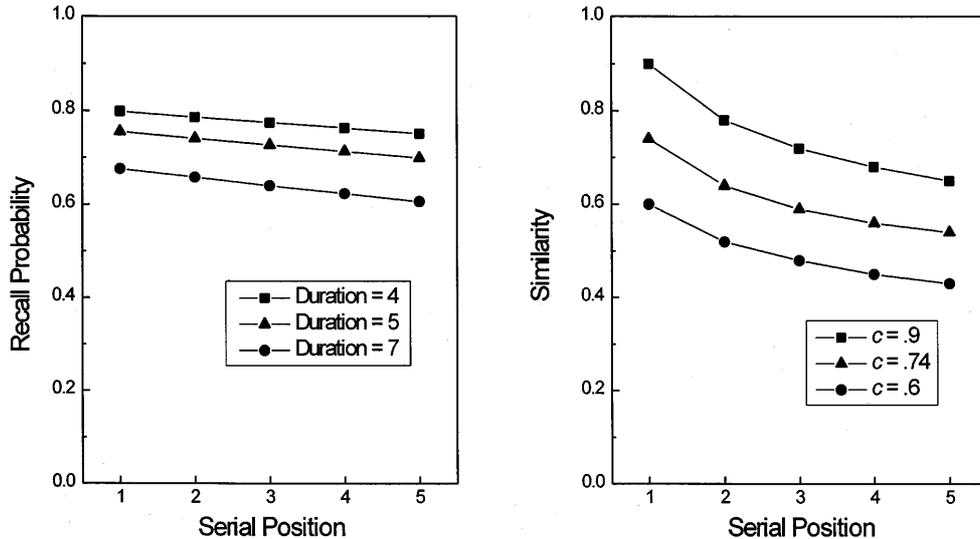


Fig. 2 The *left panel* shows the predicted serial position curves of the model by Brown and Hulme (1995) when redintegration is omitted. The *right panel* shows the presumed output of a putative associative stage that was used to generate starting vectors for the remaining simulations

tegration. This was further explored by the second simulation.

Simulation 2: Effects of speech rate and lexicality on recall

Brown and Hulme (1995) modeled the speech-rate effect by assuming that longer items were composed of a greater number of temporally defined “segments”. Once encoded, segments decay at a constant rate over time. With the further assumption that a word is recalled only if all constituent segments are recalled correctly, this representational assumption gives rise to impaired recall of longer items, notwithstanding their advantage at redintegration.

Although Brown and Hulme (1995) did not explicitly group components of their model into a single associative stage, the performance of all components preceding redintegration can be identified and analyzed separately. It is therefore conceivable that the output of the “associative” components of the Brown and Hulme model, when re-expressed as a vector or similarity measure, might provide the necessary starting conditions for the BSB to predict a positive relation between memory span and speech rate.

The left panel in Fig. 2 presents the predictions of the model by Brown and Hulme (1995: Demonstration 1) when redintegration was omitted. Predictions were generated from the set of parameter values (for input decay, output decay, initial trace strength, and forgetting probability) provided by Brown and Hulme (1995: Appendix A). Several comments can be made about the pattern. First, there was a clear effect of speech rate.

Second, the effect of speech rate was invariant across serial position. Third, the serial position curves are implausible because primacy is negligible and recency is absent.²

In light of these problems of the Brown and Hulme (1995) model, we followed a different approach and modeled the output from the associative stage by using Equation 1 with a constant λ (.2) but a different value of c for each speech rate. The number of parameters was equal to the model of Brown and Hulme. We are neutral concerning the mechanism (e.g., decay vs. interference) that gives rise to those different values of c at the associative stage. Several possibilities will be explored after all simulations have been presented.

The resulting serial position curves, representing the expected similarity (s_j) between \mathbf{f}'_j and \mathbf{f}_j , are shown in the right panel of Fig. 2 for values of c of .90, .74, and .60. Each data point summarizes an entire vector of retrieved memorial information that is of a specified similarity to the desired target. The serial position curves are compatible with the known behavior of process models of the associative state (e.g., Brown et al., 1998; Lewandowsky & Murdock, 1989). The absence of recency does not present a problem, because the redintegration model is known to predict recency via response suppression, even when its input is of monotonically decreasing similarity (Lewandowsky, in press; Lewandowsky & Li, 1994).

The results of using these similarity values to generate starting vectors for redintegration are shown in Fig. 3 (remaining parameter values: $f^s = .3$, $\phi = 2.3$). The figure shows the positive relation between memory span and speech rate that is typical of the behavioral data. This suggests that a detrimental effect of item length at the

²The overall serial position curves predicted by Brown and Hulme (1995), with the contribution from redintegration added, differ from those shown in the left panel of Fig. 2 only by a shift in intercept. Hence, the complete Brown and Hulme model predicts serial position curves with minimal primacy and with no recency.

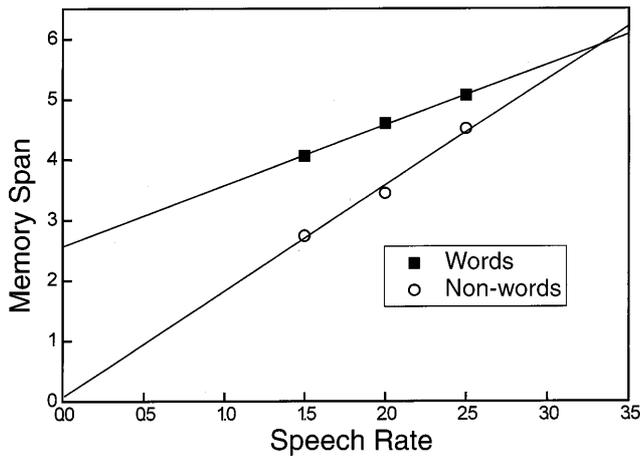


Fig. 3 Simulation 2: Predictions of the redintegration model when provided with starting vectors that conform to the assumed output of the associative stage (shown in the *right panel* of Fig. 2). The parameter c varies between different speech rates and between words and non-words, reflecting a further effect of lexicality on the associative stage

associative stage can override the redintegration advantage observed in Simulation 1.

As in Simulation 1, the network was pre-trained on a corpus of 150 words. Unlike Simulation 1, to enhance realism, there was a single corpus containing an equal number of words of each length. Initial simulation runs showed that, by itself, pre-training did not give rise to a sufficiently large lexicality effect, suggesting that lexicality may also affect encoding or retrieval in the associative stage. This appears particularly plausible because words are readily accessible from semantic memory, thus supporting immediate and presumably full encoding during list presentation. Non-words, by contrast, have to be constructed from scratch, in particular when their pronunciation is not known, which takes additional processing time and hence may lead to imperfect encoding. Assuming, furthermore, that this construction process occurs on a syllable-by-syllable (or phoneme-by-phoneme) basis, the non-word disadvantage also increases with item length. These two assumptions were captured in the final simulation run shown in Fig. 3, in which the value of c for non-words was the square of the corresponding value for words, yielding values of c of .81, .55, and .36 for short, medium, and long non-words, respectively. Because c affects both the associative stage (Eq. 1) and the strength of studied items in the redintegration model (via the translation parameter f^s), the manipulation embodied a consistent set of encoding assumptions. All other parameters remained unchanged.

The figure attests to the success of these additional encoding assumptions. Indeed, the predictions in Fig. 3 correspond exactly to the idealized pattern of results created from a review of the relevant literature by Brown and Hulme (1995). Because values of c for words and non-words were related by a constant (i.e., exponent 2), the predictions were based on a total of six free pa-

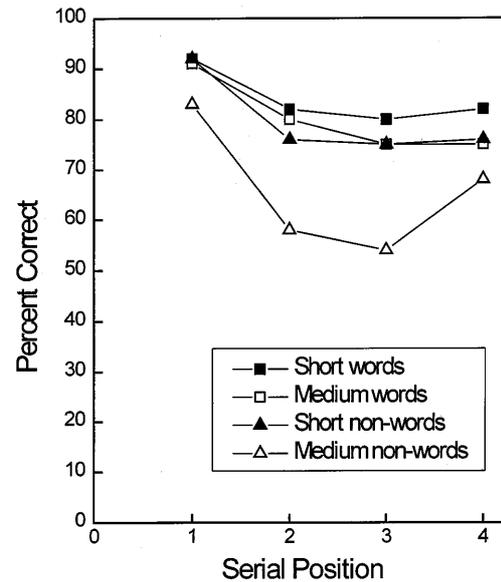


Fig. 4 Simulation 2: Predicted serial position curves for different word lengths (i.e., speech rates) for list-length 4

rameters, paralleling the descriptive account by Brown and Hulme. However, unlike that earlier account, the redintegration model simultaneously predicted the serial position curves underlying these memory-span values.

A representative set of predicted serial position curves for list-length 4 is shown in Fig. 4. They resemble those observed in behavioral experiments. Note in particular the presence of (slight) recency, which arises despite the monotonically decreasing similarity between starting vector and target across serial positions. As outlined elsewhere (Lewandowsky, in press), the occurrence of recency is one consequence of response suppression. In addition, although not shown here, the associated error transposition gradients also mirror the data. The redintegration model therefore provides far more explanatory power, with a comparable number of free parameters, than the descriptive model by Brown and Hulme (1995).

Simulation 3: Effects of word frequency

The final simulation extended the use of pre-training to account for the effects of word frequency. Like lexicality, word frequency affects memory span, such that high-frequency words are recalled better than low-frequency words (Hulme et al., 1997). For a given list length and speech rate, the magnitude of the frequency effect increases across serial positions.

As in the preceding simulation, the network was pre-trained with a corpus of 150 words composed of an equal number of words of each length. In addition, to model the frequency structure of natural language, 120 of the corpus words were presented once with an encoding weight (w^L) of .0017. These items represented low-frequency words. The remaining 30 items were

encoded with a weight of .17, thus representing high-frequency words whose frequency of occurrence was 100 times greater than their low-frequency counterparts. The different sizes of the two frequency pools approximated the distribution of frequencies in natural language.

Following pre-training, redintegration was modeled as before, using constant values for λ (.3), f^s (.4), plus three different values of c for short, medium, and long words (.62, .44, and .30, respectively). The criterion parameter ϕ differed with word frequency (1.75 vs. 2.45 for low and high frequency, respectively). High-frequency study lists were sampled from the high-frequency portion of the corpus, and low-frequency lists from the low-frequency portion.

The results are shown in Fig. 5. The top panel shows the effects of word frequency and speech rate on memory span, and the bottom panel shows a representative set of serial position curves for short words at list-length 5. Clearly, the redintegration model captured the effects of word frequency on memory span and on the serial position curve without parameter manipulation, simply by implementing a realistic pre-training regime. Importantly,

word frequency was not assumed to have an effect at the associative stage, confirming the prior theoretical analysis using a descriptive model by Hulme et al. (1997).

General discussion

Summary of simulations

The redintegration model explained the speech-rate effect by trading off a redintegration advantage for long words with their impoverished retrieval by the associative stage. The model explained the effects of lexicality (better memory for words than non-words) and word frequency (better memory for high- than low-frequency words) by making straightforward assumptions about the nature of pre-experimental knowledge. Parameters had to be manipulated only for unfamiliar non-words, which were assumed to be encoded with greater difficulty than pre-trained items.

In all cases, plausible serial position curves accompanied the predicted memory-span values. This compared favorably to descriptive models, which either could not address serial position effects (Hulme et al., 1997) or generated implausible predictions (Brown & Hulme, 1995).

Role of parameters

At most, seven free parameters were used, namely λ , f^s , two values of ϕ for different word types, plus three different values of c for different item lengths. This is comparable to the number required by earlier models with a more limited scope (e.g., Brown & Hulme, 1995). Note also that f^s did not vary between experimental conditions and only took on a different value for the final simulation. It is therefore better understood as a fixed parameter whose only effect is to set the absolute level of performance. Likewise, the parameter that determined the setting of the energy criterion (ϕ) is, in principle, replaceable by an on-line analysis of convergence behavior.

The primary role of the parameters was to model the output of the associative stage. By implication, one of the principal contributions of this article, beyond exploring a process model of redintegration, was to provide novel constraints on the design and operation of models of the associative stage. We shall turn to an examination of those constraints after analyzing the model's performance and after dealing with some potential criticisms.

Explanation of the performance of the model

Most of the effects arose from the nature of the pre-training corpus that represented pre-experimental

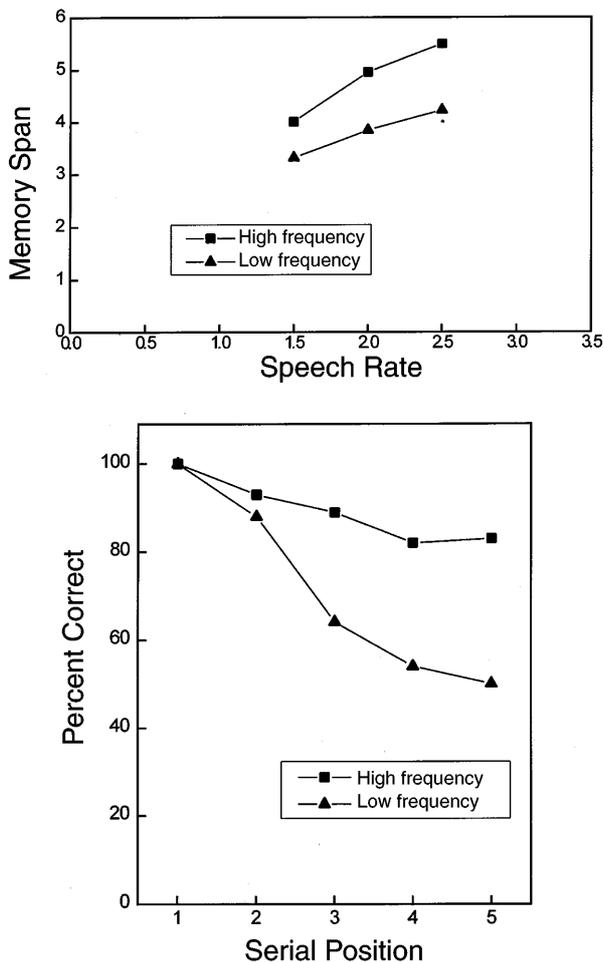


Fig. 5 Simulation 3: Predicted effects of word frequency and speech rate on memory span (top panel) and representative underlying serial position curves (short words at list-length 5, bottom panel)

knowledge. Consider first the absolute difference in performance due to lexicality and word frequency, which reflected interference from pre-trained items whose attractors competed for convergence at redintegration. The attractors of non-words, having received no training prior to list presentation, were readily overwhelmed by the pre-trained corpus. Words, by contrast, were protected from this interference because the strength of their attractors was combined across pre-training and study.

Interference also accounted for the difference between low- and high-frequency words. Although both were encoded equally at study, and although both were pre-trained, the more frequent words received stronger attractors at pre-training than their low-frequency counterparts.

The interaction between word length and lexicality – in particular, the steeper slope for non-words – can also be understood through pre-training. Regardless of lexical status, long items were encoded with less strength than short items. For words, encoding at study was sufficiently strong relative to pre-training to eliminate interference for all word lengths. However, for non-words, the minimal encoding at study rendered long items particularly vulnerable to pre-training interference.

Finally, pre-training also explains the increasing effect of word length, word frequency, and lexicality across serial positions. For all types of material, the decreasing encoding gradient (see Eq. 1) ensured that terminal list items were represented by less powerful attractors than initial items. In consequence, late list items were always more susceptible to competition from pre-trained attractors.

Limitations and potential criticism

Many of the core predictions of the redintegration model, in particular the shape of the serial position curve, arose from the decreasing strength with which successive list items were encoded (Eq. 1). Although this assumption is common throughout the literature (e.g., Brown et al., 1998; Houghton & Hartley, 1996; Lewandowsky & Li, 1994; Lewandowsky & Murdock, 1989; Page & Norris, in press a, b), it is based on sketchy independent justification. Brown et al. (1998) appealed to the “intuition that each successive item...is progressively less ‘surprising’ or attention-demanding than the previous one” (p. 43) – an intuition thought to be consonant with the environmental demands on an adaptively rational organism.

Another potential point of criticism would hold that the redintegration model begs the entire question of memory for serial order. A crucial component of memory retrieval is relegated to as hypothetical associative stage that is characterized only by assumptions that are conveniently compatible with the redintegration architecture. In response, it must be noted that two theories of the associative stage, OSCAR and TODAM, are de-

monstrably compatible with the redintegration model. TODAM functions with the redintegration model in an integrated manner (Lewandowsky & Li, 1994). OSCAR has not been computationally combined with the redintegration model, but its output is known to satisfy the general form specified by Equation 1 (see Brown et al., 1998, Fig. 11c). It follows that, although the associative stage was represented here by parameters, these parameters described the known behavior of existing models. It furthermore follows that some of the current parameters, in particular c and λ , would be replaced by those intrinsic to TODAM or OSCAR if an associative stage were implemented. The parsimony of the redintegration model is thus not an artificial consequence of the omission of an associative stage.

A related point of criticism would suggest that the redintegration model also begs the twin questions of rehearsal and speech rate – in particular because, by itself, the model predicts a reversal of the observed speech rate effect (see Simulation 1). In response, it must be noted that the associative stage was characterized by the same number of parameters such as those for the predictions of Brown and Hulme’s (1995) model without a redintegration component (see Fig. 2). Moreover, these parameters primarily captured the known behavior of models; it is only the additional assumption that speech rate affects the quality of the retrieved information that goes beyond the known properties of candidate models. Rather than seen as begging the questions of speech rate and rehearsal, this is best considered as providing a novel constraint for models of the associative stage.

Constraints for models of the associative stage

On the basis of the present simulations, models of the associative stage (e.g., TODAM or OSCAR) need to consider the following constraints if they are to handle the effects of speech rate, lexicality, and word frequency. First, word frequency need not be represented or considered by the associative models, because redintegration alone can produce these effects (Simulation 3). Second, lexicality need only be partially represented by associative models. The only requirement is that non-words suffer a disadvantage at encoding, presumably because their intended pronunciation may not be immediately available (Simulation 2). Finally, and perhaps most important, the associative stage must yield an effect of speech rate or word length that counters the redintegration advantage for longer words (i.e., different values of c : Simulation 2). Brown and Hulme (1995) suggested one possible way in which this might be modeled, namely, by decay of segmented representations. However, revisions to their theory are mandated by the implausibility of their predicted serial position curves (Fig. 2).

Segmented representations are also responsible for speech-rate effects in the feature model (Neath &

Nairne, 1995; Neath et al., 1998), although it is their erroneous assembly, rather than decay, that reduces memorability of long words. Although the feature model is a fully specified process model, the use of segmented vector representations makes it compatible with the redintegration model. In particular, some of the retrieval mechanisms of the feature model (sampling and recovery: Nairne, 1990) may be more elegantly accommodated within the BSB. To recall a list of n items, the feature model requires at least n^2 sequential vector comparisons between retrieval cues (in primary memory) and possible responses (in secondary memory). The redintegration model may provide a more efficient – though perhaps not isomorphic – mechanism by using each item in primary memory as the starting state for the BSB. The secondary memory in the feature model would thus be replaced by the redintegration model (and Eqs. 1–4 in Nairne, 1990, by the BSB dynamics).

Alternatively, the effect of speech rate may be a consequence of the limitations on rehearsal that are revealed when pronunciation time increases. Of course, this would constitute a return to some variant of the rehearsal model (e.g., Baddeley, 1986) introduced at the outset. However, contrary to the reasoning at the outset, the widely observed lexicality and word-frequency effects may not be incompatible with a rehearsal model that includes a redintegration stage along the lines explored here.

Finally, it has been suggested that forgetting during recall, rather than pronunciation time per sec, may be the underlying cause of the speech rate effect (e.g., Doshier & Ma, 1998). Doshier and Ma found that recall time was a slightly better predictor of memory span than speech rate. However, Doshier and Ma acknowledged that these effects were consistent with the notion, explored here and elsewhere (Brown & Hulme, 1995; Neath et al., 1998), that the poorer recall of long words, as well as their longer pronunciation duration, were both due to the larger number of elements in their segmented representation. Accordingly, Lewandowsky (in press) showed that the redintegration model can handle key aspects of the data by Doshier and Ma.

Conclusion

A redintegration model based on the BSB network was shown to account for many effects of stimulus material on immediate serial recall. Without manipulation of parameters, by implementing reasonable assumptions about prior experience, the model accounted for the effects of word frequency and, in part, the effects of lexicality. It follows that these lexical phenomena need no longer be addressed by a model of the associative stage. This article showed that the two necessary contributions of such a model are, first, to retrieve information with decreasing accuracy across serial positions and, second, to do so with greater accuracy for short items than long items, irrespective of their lexical status

or familiarity. Whether this results from decay, interference, or some other mechanism does not affect our conclusions.

Acknowledgements We wish to thank Mike Mundy for his assistance in manuscript preparation.

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*. 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.
- Baddeley, A.D. (1986). *Working memory*. Oxford: Oxford University Press.
- Baddeley, A.D., Thomson, N., & Buchanan, M. (1975). Word length and the structure of short-term memory. *Journal of Verbal Learning and Verbal Behavior*, *14*, 575–589.
- Besner, D., & Davelaar, E. (1982). Basic processes in reading: Two phonological codes. *Canadian Journal of Psychology*, *36*, 701–711.
- 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. (1998). *Oscillator-based memory for serial order*. Manuscript submitted for publication.
- Golden, R.M. (1986). The “brain-state-in-a-box” neural model is a gradient descent algorithm. *Journal of Mathematical Psychology*, *30*, 73–80.
- Doshier, 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.
- 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., Maughan, S., & Brown, G.D.A. (1991). Memory for familiar and unfamiliar words: Evidence for a long-term memory contribution to short-term memory span. *Journal of Memory and Language*, *30*, 685–701.
- Hulme, C., Roodenrys, S., Brown, G.D.A., & Mercer, R. (1995). The role of long-term memory mechanisms in memory span. *British Journal of Psychology*, *86*, 527–536.
- 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. (in press). Redintegration and response suppression in serial recall: A dynamic network model. *International Journal of Psychology*.
- 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.
- 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.

- Nairne, J.S. (1990). A feature model of immediate memory. *Memory & Cognition, 18*, 251–269.
- Neath, I., & Nairne, J.S. (1995). Word-length effects in immediate memory: Overwriting trace decay theory. *Psychonomic Bulletin & Review, 2*, 429–441.
- Neath, I., Surprenant, A.M., & LeCompte, D.C. (1998). Irrelevant speech eliminates the word length effect. *Memory & Cognition, 26*, 355–368.
- 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: Erlbaum.
- Schweickert, R. (1993). A multinomial processing tree model for degradation and redintegration in immediate recall. *Memory & Cognition, 21*, 168–175.
- Schweickert, R., & Boruff, B. (1986). Short-term memory capacity: Magic number or magic spell? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 12*, 419–425.