Psychological Review
1994, Vol. 101, No. 3, 539–543

0033-295X/94/$3.00

Memory for Serial Order Revisited

Stephan Lewandowsky and Shu-Chen Li

The critiques by Mewhort, Popham, and James (1994) and Nairne and Neath (1994) identified at least 6 potentially serious problems with Lewandowsky and Murdock's (1989) Theory of Distributed Associative Memory (TODAM) model of memory for serial order. The authors show that the flaws attributed to the memory component of TODAM are less serious than claimed, whereas the problems attributed to the response selection stage necessitated a process implementation of the previously unspecified deblurring mechanism. The deblurring process, implemented by a dynamic autoassociative network, is shown to handle most of the problems identified by the critics without imperiling TODAM's ability to handle basic serial position data.

Mewhort, Popham, and James (1994) and Nairne and Neath (1994) revealed several limitations and potentially serious flaws in Lewandowsky and Murdock's (1989) Theory of Distributed Associative Memory (TODAM) model of memory for serial order. Figure 1 lists the criticisms put forward by Mewhort et al. (labeled **M**) and Nairne and Neath (labeled **N & N**) and identifies the component of TODAM to which they pertain.

The figure clarifies that most problems involved the R system, which is dedicated to response disambiguation and output, whereas only two difficulties were identified with the Q system, the central memory component of TODAM. In this rejoinder, we first show that the problems with the Q system are less serious than suggested by Mewhort et al. (1994), and we then propose how most of the remaining flaws in the R system can be circumvented by a process implementation. Paralleling the critics' main approach, we focus on simulations of TODAM in preference to the analytic version.

The Q System

Choice of Weighting Parameters

Mewhort et al. (1994) questioned Lewandowsky and Murdock's (1989) choice to link parameters γ and ω such that $\gamma = 1 - \omega$, suggesting that this would create situations in which item recognition (driven by γ) would be absent despite near-perfect associative recall performance ($\gamma = 0$ if $\omega_j = 1$).

Although Equation 5 in Lewandowsky and Murdock (1989) implies that $\gamma = 0$ whenever $j = 1$, it must be clarified that $j = 1$ indexes the context cue preceding the list, whereas the first list item corresponds to $j = 2$, thus yielding $\gamma > 0$ for all list items whenever ω_0 is less than unity. Indeed, ω_0 was less than one in all quantitative fits reported by Lewandowsky and Murdock, the only exception being the simulations and the qualitative fit reported in their Figure 30, where ω_0 was set to unity to reduce the number of free parameters. The broader implication of Mewhort et al.'s (1994) criticism, which points out that item information is encoded but not used by TODAM, is addressed by the response-deblurring process outlined later.

Cuing With the Facsimile Vector

One important conclusion of Lewandowsky and Murdock (1989) was that TODAM, unlike conventional associative models, could account for serial order memory through chaining. The chaining notion has been unpopular for quite some time; a primary reason being its obvious inability to allow for retrieval of the remaining list items once the chain has been broken by a recall failure. Lewandowsky and Murdock proposed that, in TODAM, retrieval can continue even after a break in the chain because an approximation of the preceding item (labeled f'_{j-1} by Lewandowsky & Murdock and *facsimile* by Mewhort et al., 1994) is always available as a cue. This ability to use partial information was identified as a principal strength of the distributed approach to memory embodied in TODAM (Lewandowsky & Murdock, 1989, p. 51).

It is therefore of particular concern that cuing with the facsimile vector, under the conditions initially cited in support of the idea, failed to raise performance above chance in Mewhort et al.'s (1994) analysis. We investigated the matter further by exploring the parameter space of the Mewhort et al. simulation. The results are summarized in Table 1.

The table also shows the values of the familiar TODAM parameters for all simulations, where N represents the number of extra-list competitors, N determines the dimensionality of the memory vector, $b - a$ corresponds to the width of the tolerance window centered around unity, α is the forgetting parameter, and λ is the rate constant of the exponential decline across serial positions of the associative attention component from starting value ω_0 (which was set to unity for all simulations). All simulations involved 1,000 replications. Probability of correct report following cuing with f'_{j-1} was compared with chance performance using a chi-square statistic.

Stephan Lewandowsky and Shu-Chen Li, Department of Psychology, University of Oklahoma.

We gratefully acknowledge comments on an earlier version of this article provided by Hervé Abdi, Douglas Mewhort, and Bennet B. Murdock. We also thank James A. Anderson for a helpful discussion of antilearning.

Correspondence concerning this article should be addressed to Stephan Lewandowsky, Department of Psychology, University of Oklahoma, Norman, Oklahoma 73019-0535. Electronic mail may be sent to lewan@uoknor.edu.

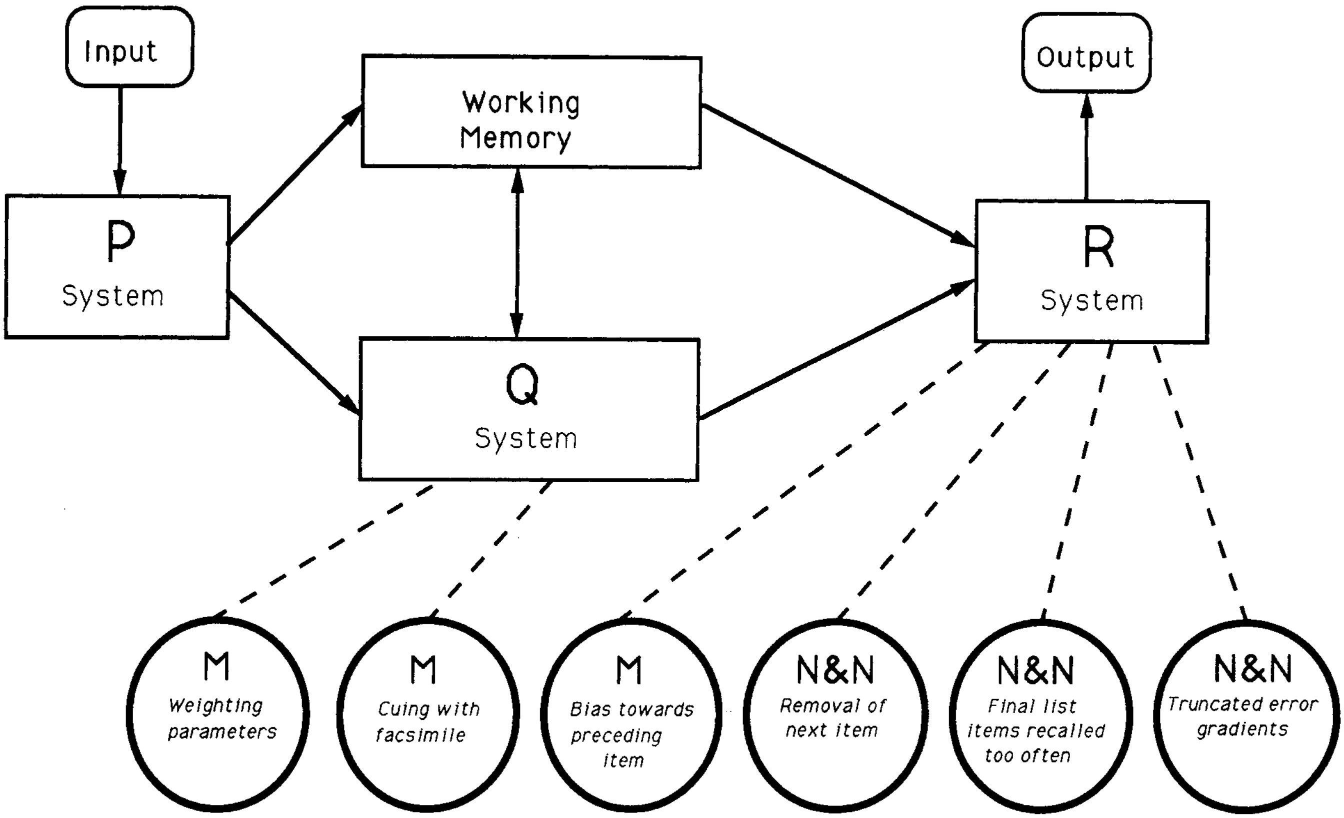

Figure 1. Overview of system architecture of the Theory of Distributed Associative Memory (TODAM) and of the problems identified by the critics Mewhort et al. (1994; circles labeled M) and Nairne and Neath (1994; circles labeled N & N). Problems are identified by subheadings or summary phrases used in the text.

The first row of Table 1 shows the outcome of a replication of Mewhort et al.'s (1994) simulation. The nonsignificant value of chi-square supports their conclusion that facsimile cuing is ineffective under those circumstances. Now consider the chi-square values reported in the remaining rows of the table, which correspond to methodologically identical simulations conducted with slightly different, but equally plausible, parameter settings. It is clear that in the majority of those additional simulations, the facsimile cue yielded above-chance performance. The extent of that effect can be illustrated with the results of Simulation 5: The facsimile cue yielded .23 correct (averaged across the last five serial positions), which compares to a chance level of .13. We therefore conclude that cuing with the facsimile vector serves to continue the retrieval chain in most situations,

Table 1
Parameter Values and Results of Simulations Examining Performance After Cuing With the Facsimile Vector

	Parameter values					Results		
Simulation	N	N	$b-a$	α	λ	χ^2	df	p
Mewhort et al. (1994)	0	197	6	0.90	0.70	3.08	6	.799
1	0	99	6	0.90	0.70	8.81	6	.180
2	0	99	6	1.00	0.70	66.92[a]	6	<.001
3	0	99	6	1.00	0.50	66.74[a]	6	<.001
4	5	1,699	6	0.90	0.15	575.86[a]	4	<.001
5	5	1,699	6	0.90	0.50	402.45[a]	4	<.001
6	5	1,699	2	0.98	0.65	582.69[a]	4	<.001
7	10	1,199	6	0.90	0.50	197.82[a]	5	<.001

[a] Significant values of chi-square.

much as initially claimed by Lewandowsky and Murdock (1989).

The R System

The remaining discussion should acknowledge that Lewandowsky and Murdock (1989) recognized the ad hoc nature of the competitors and their sequential removal (p. 34) but introduced that mechanism because a process implementation of deblurring was unavailable (p. 31). In support of their choice, Lewandowsky and Murdock claimed that the mechanism was consonant with data (p. 35) and yielded a large overall gain in explanatory power of the theory (p. 49). The critics' thorough explorations have rendered that choice indefensible and have necessitated a limited reformulation of the deblurring stage within a process model. The model conforms to the approach foreshadowed by Lewandowsky and Murdock (p. 31), and it rectifies the main problems identified by the critics while preserving the competitor-based approach to response selection that is central to TODAM.

The critics uncovered two principal problems with the R system that demand resolution: Mewhort et al.'s (1994) demonstration that retrieval is often biased toward the item preceding the cue, as opposed to the target following it, demands a solution more plausible than omitting comparisons between the facsimile vector and earlier list items. Similarly, Nairne and Neath's (1994) analysis convincingly calls for the use of a sampling-without-replacement scheme, in which the successful recall of an item, as opposed to its serial position, causes its removal from the set of competitors.[1]

Nonetheless, empirical grounds remain for postulating that items be removed from the set of competitors on recall. Perhaps most salient is the observation that subjects rarely recall items more than once. For example, in the Deese and Kresse (1952) studies cited by Nairne and Neath (1994), true intralist intrusions represented as little as 5% of the total number of errors (Deese & Kresse, 1952, Experiment 1). In addition, Mewhort and Armstrong (1993) demonstrated that repetition blindness—a reduction in memory performance for the second occurrence of a repeated item—reflects not an encoding deficit but a retrieval failure, which is readily modeled by the removal of a recalled item from the set of competitors.

Brain-State-in-a-Box

The deblurring model was adapted from the dynamic autoassociative network of Anderson, Silverstein, Ritz, and Jones (1977), which is known as the *brain-state-in-a-box* (BSB) model. The BSB model has the property that, when given a slightly distorted image of a previously learned item (i.e., f'), the network iteratively modifies its own output until an exact copy of the initial item (f) is recovered. If the input is distorted further, the network is still guaranteed to converge on some unambiguous state, although that state may then differ from the desired response.[2]

Previous BSB applications also point to a natural way in which to implement the removal of response competitors. For example, Anderson (1991) showed that the multistable perception of the Necker cube could be modeled by use of Hebbian "antilearning," which corresponds to a reduction of the strength (by use of a negative learning rate) of associative connections in the BSB network. In other applications, antilearning has served to limit the otherwise unbounded weight growth in standard Hebbian learning (e.g., Proulx & Begin, 1990), and it has been shown to provide the basis for a model of categorization (Proulx & Begin, 1993). In the present context, Hebbian antilearning provides a natural way to remove an item from the BSB network after recall.

Antilearning has several distinct advantages over the scheme used by Lewandowsky and Murdock (1989): First, the system itself, as opposed to some external mechanism, is responsible for removal of items. Second, the extent of antilearning can be guided by a parameter, thus allowing for partial suppression of competitors. Third, antilearning can be extended to TODAM's storage component, thus permitting parallel removal of an association once it has been used for retrieval. Finally, because antilearning was developed, in part, to satisfy neuro-physiological demands for brief inhibitory processes (e.g., Cox, Kakolewski, & Valenstein, 1969), it is often temporary (e.g., Anderson, 1991). Hence, it can model brief retrieval inhibitions, such as repetition blindness, with greater plausibility than a permanent removal of recalled items.

BSB Simulation

To demonstrate the basic feasibility of the BSB model as a deblurring tool, TODAM was combined with a BSB network in an exploratory simulation. During study, autoassociations of all items (five list items plus two context vectors) were stored in the BSB network, and pairwise associations between successive items were stored in TODAM's memory vector in the usual fashion. The familiar ω_j function (Equation 5 in Lewandowsky & Murdock, 1989) was used to determine encoding weights for both TODAM and the BSB. Because items were now represented in the BSB network, item information was omitted from TODAM. Further methodological details and parameter values are given in the Appendix.

Paralleling Nairne and Neath's (1994) simulation, recall was implemented according to Lewandowsky and Murdock's (1989) Equation 2a, except that the retrieved facsimile vector (f') was not compared with the competitors but was presented to the BSB network for deblurring. During deblurring, all as-yet unrecalled items, including those earlier in the list, were available in the associative matrix of the BSB model for potential output. The deblurred item, even if it did not correspond to

[1] It is clear from Nairne and Neath's (1994) article that the remaining problems, concerning the number of retrievals and the shape of error gradients, are either solved automatically, or are at least solvable in principle, under a sampling-without-replacement scheme.

[2] The name of the model derives from the nonlinearity of its iterative updating function, which confines the output to lie within a box formed by the vertex vectors (e.g., $\{-1, +1, -1\}$). In the genealogy of connectionist networks, the BSB model is a cousin of the Hopfield net, which Lewandowsky and Murdock (1989) cited as a potential future candidate for deblurring. The BSB model is also related to the linear associator, which, in turn, has been shown to be roughly isomorphic to TODAM (Lewandowsky, 1991).

the desired response, was removed from the BSB network using Hebbian antilearning, and the association used to cue retrieval was removed from TODAM's memory vector in identical fashion.

To illustrate, suppose f_3 has been correlated with the memory vector to yield f_4'. Suppose furthermore that f_4' has been deblurred correctly into the response f_4. Hebbian antilearning is then used to remove f_4 from the BSB matrix and to remove the corresponding association (the convolution $f_3 * f_4$) from TODAM's memory vector. The removal of the convolution by a parallel mechanism maximizes consistency between item and associative information and eliminates the bias toward preceding list items.

The results are shown in Figure 2: The top panel contains the serial position curve obtained from TODAM augmented by BSB deblurring, and the bottom panel shows the mean number of times each item was output during recall of the list. The two panels thus report data analogous to those shown in Figures 1 and 2, respectively, of Nairne and Neath (1994).

Resolved Problems and Remaining Limitations

The simulation resolved several major problems: (a) The bias toward preceding list items was eliminated without artificially constraining the set of competitors; (b) recency was obtained

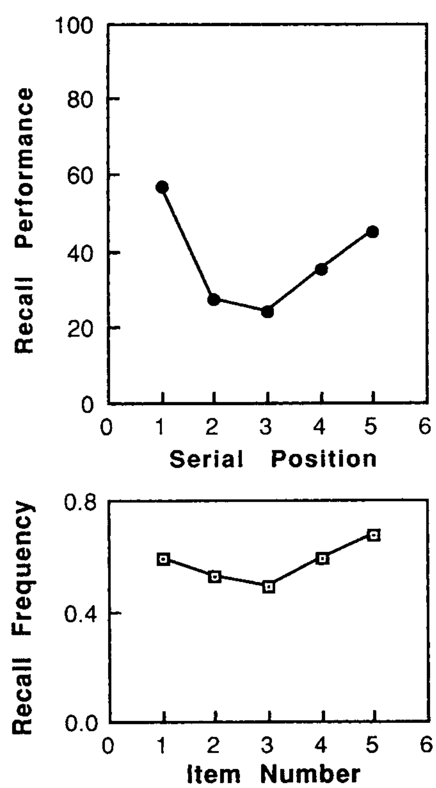

Figure 2. Serial position curve obtained by the Theory of Distributed Associative Memory (TODAM) with deblurring handled by a brain-state-in-a-box network (top panel) and the corresponding mean number of times each item was recalled (bottom panel).

despite implementation of a more realistic sampling-without-replacement scheme; and (c) recall frequencies remained roughly constant across serial positions, which was more consonant with data than the previous unrealistic increase.

In addition, the error gradients produced by the simulation (not shown in the figure) were no longer completely truncated, with earlier list items being reported occasionally. However, the problem remains partially unresolved because the simulated gradients continued to be highly asymmetrical, favoring later list items throughout, and were less sharply peaked than their behavioral counterparts. Similarly, the low overall level of performance underscores that this simulation must not be understood as a quantitative effort but as a demonstration tool. Nonetheless, we suggest that a deblurring process of this type may eventually be of considerable generality and might resolve similar problems with response selection and deblurring that are intrinsic to most connectionist models and distributed approaches to memory (Goebel & Lewandowsky, 1991).

Conclusions

Lewandowsky and Murdock's (1989) model provided a quantitative account for numerous aspects of memory for serial order. Subsequent explorations of the theory have yielded testable predictions (e.g., Li & Lewandowsky, 1993) as well as the serious problems identified by the present critics. In response, (a) we demonstrated that the flaws attributed to the storage component of TODAM were less serious than claimed, and (b) we acknowledged that the original Lewandowsky and Murdock response system was flawed, and we proposed a process implementation of deblurring that resolved several problems without appearing to imperil the strengths of TODAM.

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., Silverstein, J. W., Ritz, S. A., & Jones, R. S. (1977). Distinctive features, categorizational perception, and probability learning: Some applications of a neural model. *Psychological Review, 84,* 413–451.

Cox, V. C., Kakolewski, J. W., & Valenstein, E. S. (1969). Inhibition of eating and drinking following hypothalamic stimulation in the rat. *Journal of Comparative Physiology and Psychology, 68,* 530–535.

Deese, J., & Kresse, F. H. (1952). An experimental analysis of errors in rote serial learning. *Journal of Experimental Psychology, 44,* 199–202.

Goebel, R. P., & Lewandowsky, S. (1991). Retrieval measures in distributed memory models. In W. E. Hockley & S. Lewandowsky (Eds.), *Relating theory and data: Essays on human memory in honor of Bennet B. Murdock* (pp. 509–527). Hillsdale, NJ: Erlbaum.

Lewandowsky, S. (1991). Gradual unlearning and catastrophic interference: A comparison of distributed architectures. In W. E. Hockley & S. Lewandowsky (Eds.), *Relating theory and data: Essays on human memory in honor of Bennet B. Murdock* (pp. 445–477). Hillsdale, NJ: Erlbaum.

Lewandowsky, S., & Murdock, B. B., Jr. (1989). Memory for serial order. *Psychological Review, 96,* 25–57.

Li, S.-C., & Lewandowsky, S. (1993). Intralist distractors and recall direction: Constraints on models of memory for serial order. *Journal*

of Experimental Psychology: Learning, Memory, and Cognition, 19, 895–908.
Mewhort, D. J. K., & Armstrong, I. (1993, July). *Is repetition blindness encoding failure?* Paper presented at the Third Annual Meeting of the Canadian Society for Brain, Behaviour, and Cognitive Science, Toronto, Ontario, Canada.
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.
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.
Proulx, R., & Begin, J. (1990). A new learning algorithm for the BSB model. In M. Caudill (Ed.), *Proceedings of the International Joint Conference on Neural Networks* (Vol. 1, pp. 704–706). Hillsdale, NJ: Erlbaum.
Proulx, R., & Begin, J. (1993). *The use of Hebbian and anti-Hebbian learning for categorization in neural networks.* Manuscript submitted for publication.

Appendix

Overview

The simulation maximized the size of all vectors and the number of list items within the constraints of the available computing resources. This permitted use of item vectors of size 128 (hence N = 255 in TODAM and N = 16,384 in the BSB matrix) and a list of five items plus two context markers. To satisfy the constraints of Hebbian learning, exactly orthogonal bipolar binary {i.e., −1, +1, . . .} vectors were used. The same seven items were presented, in a different random order, in all 100 replications. A full description of the BSB network cannot be given here; see Anderson (1991) or Anderson et al. (1977) for a detailed mathematical treatment.

BSB Dynamics

During study, the associative weight matrix **A** was formed by adding autoassociations of items using standard Hebbian learning:

$$\mathbf{A}_j = \mathbf{A}_{j-1} + \omega_j \boldsymbol{f}_j \boldsymbol{f}_j^{\mathrm{T}},$$

where the ω_js were provided by the familiar TODAM weighting function using $\omega_0 = 1.0$ and $\lambda = .25$. **A** was set to **0** at the outset.

At test, the facsimile vector $\boldsymbol{f}'$ produced by TODAM was passed to the BSB network, which then deblurred the vector using the following dynamics. The output, or "state," vector **x** of a BSB network at any time t is given by

$$\mathbf{x}(t) = g(\beta\mathbf{x}(t-1) + \epsilon\mathbf{A}\,\mathbf{x}(t-1) + \delta\boldsymbol{f}'),$$

where $\mathbf{x}(t-1)$ is the preceding state at time $t - 1$, $\boldsymbol{f}'$ is the initially presented input, **A** is the matrix of studied associations, and β, ϵ, and δ are fixed parameters (set here to .85, 1, and .2, respectively). The function g is a simple squashing function that limits all activations to fall within the range −1 to 1. Given a nonzero initial input, the system is guaranteed to converge to some vertex (i.e., all elements of **x** either −1 or +1) in a finite number of steps. Once the system had deblurred $\boldsymbol{f}'$ to some vector **x**, it was removed from **A** using

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

where η was set to .95 throughout.

TODAM Implementation

This implementation differed from earlier versions of TODAM in the following ways: (a) To satisfy the constraints of the BSB model, binary {−1, +1, · · ·} input vectors had to be used; (b) because items were represented by the BSB network, item information (and hence the parameter γ) was omitted from the otherwise standard storage equation (which used α = 1.0 throughout); and (c) after $\boldsymbol{f}'$ was deblurred, the association that had retrieved the facsimile vector was removed from the memory vector in a fashion analogous to antilearning in the BSB network (η again set to .95):

$$\mathbf{M}_j = \mathbf{M}_{j-1} - \eta\,\omega_j(\boldsymbol{f}_{j-1} * \boldsymbol{f}_j).$$

Parameters

The number of fixed parameters remained comparable to previous applications of TODAM, with the tolerance limits a and b being replaced by the parameters ϵ, β, and δ for the BSB update function. The number of free parameters remained unchanged, with η (the antilearning fraction) taking the place of the earlier competitor parameter, N. Because α and ω_0 were set to unity, the results reflect the operation of only two free parameters (η and λ). The data change little if η is set to unity, thus allowing use of a single free parameter.

Received September 16, 1993
Accepted October 22, 1993 ■