Interference and the representation of order within associations

Mayank Rehani[1] and Jeremy B. Caplan[2]

[1]Department of Psychology, University of Alberta, Edmonton, AB, Canada
[2]Department of Psychology & Centre for Neuroscience, University of Alberta, Edmonton, AB, Canada

In the presence of interference, recall of pairs can critically depend on the diagnostic power of memory of the order of items within the pair. Models of pair memory make different assumptions about whether and how such order information is stored, from convolution-based models, which assume no explicit storage of order, to matrix models and several models that assume a pair is learned by concatenating the representations of the constituent items, which lead to perfect within-pair order memory (given retrieval of the pair). Here we investigate memory for associations and within-pair order by examining the relationship between forward and backward probes of pairs subject to order-dependent associative interference in a double-function list paradigm. Associative interference disrupted the high correlation between forward and backward recall accuracy that is typically observed in standard paired-associate learning, challenging matrix and concatenation-based models. However, participants could overcome some interference due to within-pair order ambiguity, challenging directionally ambiguous convolution-based models. Unexpectedly, the test–retest correlation was reduced for pairs under the influence of interference compared to control pairs. This finding is incompatible with all existing implementations of the model classes we consider. Any model must include an assumption that order encoding (but not retrieval) is unreliable, and the form of this additional mechanism may depend intimately on how a given model is designed. In sum, our findings suggest that within-pair order memory is neither poor nor perfect, pointing to a fallible mechanism for within-pair order learning in verbal association-memory tasks.

Keywords: Paired-associate learning; Serial-order memory; Associative symmetry; Double-function lists; Interference; Positional coding.

There are numerous models of memory for pairs designed to explain memory performance on cued recall tests. That is, having studied a pair of items, A−B, the participant can be asked to recall B given A as the probe (forward probe; A−?) or to recall A given B as the probe (backward probe; ?−B). Because a pair comprises only two items, the order of the constituent items within

the pair (A before B versus B before A) is superfluous to performance in standard cued recall of pairs. Nonetheless, models of the cued recall task make various different assumptions about whether and how within-pair order information is learned. We first consider these representational assumptions and how they stand up to existing empirical findings and then test them further against novel empirical data from an associative interference (double-function list) paradigm.

1. *Convolution-based models:* no within-pair order learning. Models that use the convolution operation to store pairwise associations (e.g., Borsellino & Poggio, 1971; Gabor, 1968; Longuet-Higgins, 1968; Metcalfe Eich, 1982; Murdock, 1982) cannot reproduce presentation order due to the mathematical properties of convolution and correlation (the operation used at retrieval). If Items A and B are represented by **a** and **b**, respectively, where boldface lower-case letters denote column vectors, the association is learned as **w = a ∗ b,** where ∗ denotes the convolution operation. For example, circular convolution (Plate, 1995) produces the vector

$$w_j = \sum_{k=0}^{n-1} a_k b_{j-k}$$

Subscripts are modulo-n. Retrieval is carried out with the correlation operation, denoted by # either in the forward direction (**a # b**): or in the backward direction **b # w**. Because convolution is commutative (**a ∗ b = b ∗ a**), the forward (A → B) and backward (A → B) associations are identical in this formalism, and within-pair positions are not obtainable at test.

2. *Matrix/tensor and concatenation-based models:* perfect within-pair order learning. In a number of cases, modellers have assumed that to learn a pair, one concatenates the item representations of the pair of items (Gillund & Shiffrin, 1984; Hintzman, 1984, 1986; Mensink & Raaijmakers, 1988, 1989; Rizzuto & Kahana, 2000, 2001). Thus, the association is learned as **a ⊕ b**, where ⊕ denotes the column-wise concatenation operation.

Learning the pair is akin to learning a single item that is twice as long as the constituent items. The assumption has been that the items are concatenated according to presentation order, an assumption that is challenged by our data. In this class of model, within-pair positions are perfectly retrievable, assuming the association is retrieved. The same type of reasoning applies to matrix models (Anderson, 1970) and tensor models (Humphreys, Bain, & Pike, 1989). These models use the outer product, $M = \mathbf{ab}^T$—outer product is the matrix, M, defined by all pairwise products of features values of the two vectors: the (i,j) the element of M, $M_{ij} = a_i b_j$—to store associations, where ^{T} denotes the transpose operation. To retrieve, one can multiply from the right with a probe-item vector. For a single-pair example, $M = \mathbf{ab}^T$, if we probe with the same **b** vector that was initially learned by multiplying the memory matrix from the right, $M\mathbf{b} = \mathbf{a}$, this simulates cued recall in the backward direction. Note that $M\mathbf{a} = 0$ if **a** and **b** are dissimilar (orthogonal). Likewise, one can multiply from the left with the transpose: $\mathbf{a}^T M = \mathbf{b}$ while $\mathbf{b}^T M = 0$; thus multiplying from the left can simulate forward cued recall. What this means is that, in a manner analogous to concatenation models, order is hardcoded into the association (and thus recoverable at test), in this case by the left- and right-handed items being encoded into different vector spaces. Because concatenation and matrix/tensor models treat within-pair order in a way that is formally similar, from here on we discuss all these models as a single class of model.

3. *Positional-coding models:* moderate level of within-pair order because order memory is assumed to be the basis of association memory. Positional-coding models were initially proposed to explain memory for serial lists. Whereas associative chaining models (Ebbinghaus, 1885/1913; Lewandowsky & Murdock, 1989) assume that a serial list is learned by "chaining" together nearest neighbour pairs of items (thus, an association-basis

for serial-order memory), positional-coding models assume that memory for a serial list is constructed by learning to associate each list item with a positional or ordinal or ordered-contextual code. Recently, it was proposed that positional-coding models could be used to learn lists of pairs by assigning constituent items to very close position codes (Caplan, 2005; Caplan, Glaholt, & McIntosh, 2006; Howard, Jing, Rao, Provyn, & Datey, 2009). Because this class of model is newer to association memory, a schematic representation of how a simple version of these models functions is provided in Figure 1. In this view, memory for pairs is derived from memory for order; a list of pairs is a special case of serial-list memory in which items are grouped into pairs along the position code, with constituent items of a pair being assigned very similar position codes ($s_W \ll s_B$ in Figure 1a; where s_W is the positional separation of items within pairs, and s_B is the positional separation of items between pairs). Positional-coding models assume some learning of within-pair order but could separately modulate association memory versus within-pair order memory.

It is possible to test participants for within-pair order in a paired-associate task using associative recognition. In associative recognition, participants are asked to judge whether a pair of items given as a probe was presented during study. In a typical associative recognition experiment, pairs are presented, which we denote as $A_1B_1 + A_2B_2 + A_3B_3$. . . . At test, "intact" pairs such as A_1B_1 are considered "old", and the chief "new" lure pairs are recombined from a left-handed item from one presented pair paired with a right-handed item from a different presented pair—for example, A_1B_2. This is considered a test of association memory because in order to rule out the recombined pairs, the participant must remember not only which items were presented, but which items were paired with one another. Like cued recall, the most common form of associative recognition does not test for within-pair order. We found only two published studies that report

accuracy data for associative recognition comparing intact with *reversed* (B_1A_1) pairs (Green & Tussing, 2001; Kounios et al., 2003). To rule out reversed probes, the participant must be able to retrieve within-pair order information. Both studies reported accuracy that was neither at ceiling nor at floor. The finding that participants have a moderate ability to reject reversed probes contradicts the prediction of current convolution models that accuracy should be at chance on this task. The implications for the remaining model classes are not clear, however, because the experiments that included reversed probes did not also include recombined probe pairs. The predictions of matrix, concatenation, and positional-coding models all regard the accuracy of memory for within-pair order *relative to* memory for associations (specifically, memory for order given that the association is retrieved), so these findings have limited diagnostic value. Mandler, Rabinowitz, and Simon (1981) found that when given the opportunity to free-recall a list of pairs, participants mainly recalled paired items together, and when they did so, order was always intact. This is consistent with near-perfect storage of within-pair order information, although it is possible that the recall procedure led to participants conservatively screening their retrievals and choosing not to free-recall pairs for which they did not have confident order information.

Also speaking to how within-pair order is learned is the question of whether associations are directional—namely, whether a single association links items A and B or two directional associations link the items, one enabling the participant to use A as a probe for B (forward association) and the other enabling the participant to use B as a probe for A (backward association). Two chief measures have been used to test for directionality effects in cued recall of pairs. The first is symmetry versus asymmetry in mean accuracy. Asch and Ebenholtz (1962) held that Gestalt theories of pair memory predicted symmetric mean performance in forward and backward directions. While the relevance to Gestalt learning has been challenged (Kahana, 2002), this robust finding has been replicated in many studies (see Kahana,
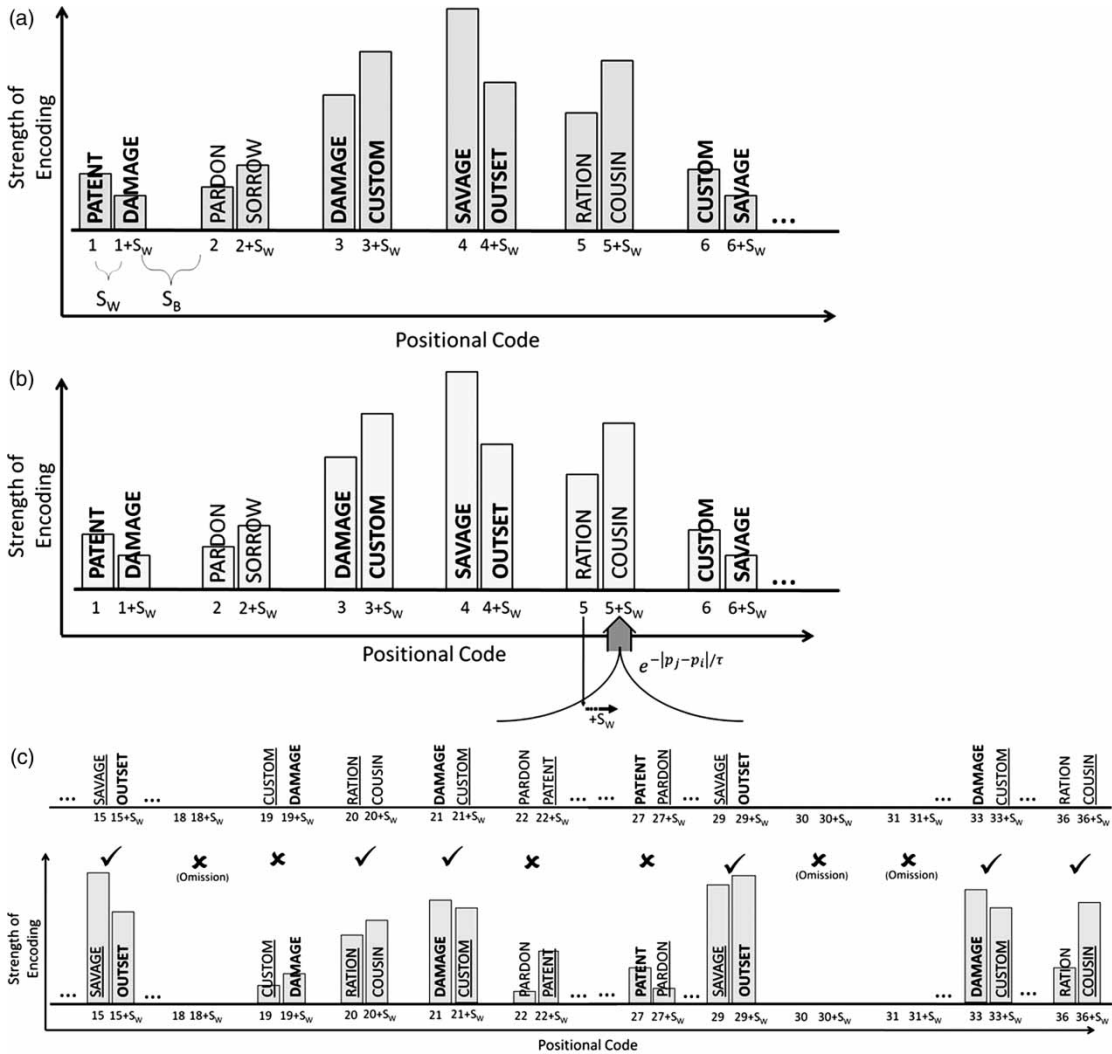
Figure 1. *Schematic representation of a simple strength–based positional–coding model based on the model that was applied to explain dissociations between memory for pairs and triples (Caplan, 2005; Caplan et al., 2006). Depicted are (a) a study trial, (b) a test trial, and (c) an illustration of output encoding during test. Items from double-function pairs are set in boldface font. It is a strength model, which learns by assigning a reliable, scalar positional code and noisy strength to each list item. In order to handle double-function pairs, in which constituent items are presented twice within the study phase, the prior model (Caplan et al., 2006) had to be modified to allow a single item to be assigned more than one positional code; we assume that the strengths of an item's association to each position are statistically mutually independent (i.e., encoding strengths drawn separately from a Gaussian distribution). To facilitate cued recall, paired items are assigned very similar positional codes whereas items from different pairs are assigned more distinct positional codes. This relative positional spacing is controlled by a single free parameter in the model previously described as the degree of isolation of studied pairs, $I = 1 - s_W/s_B$, where s_W is the positional separation of items within pairs, and s_B is the positional separation of items between pairs. Thus, the unitization effect (approximate associative symmetry) is not absolute, nor a property of the basic storage operation (as is true for convolution models), but emerges as a graded property from the relatively closer spacing of items within than between paired items. At test, the model probes with an item, retrieves the probe's positional codes, updates the positional codes by shifting them in the desired direction (forward or backward), and probes with this new position (weighted by encoding strengths) to retrieve candidate items for output. Probing with a position retrieves not only the item with the closest positional code, but also items that had been stored at nearby positions according to a positional similarity function, an exponential function parameterized by a positional decay constant, τ.*

2002, for a review). Jones and Pashler (2007) extended this finding to direct prediction and retrodiction of Markov-chain probabilistic sequences, suggesting that temporal symmetry (equality of forward- and backward-directed retrieval) may be quite general. The second measure of within-pair directionality effects was the *correlation* between forward and backward performance of a single pair over successive testing (i.e., testing each pair twice, in the same or different probe directions). Kahana introduced this measure as a more direct test of Gestalt-like pair memory. The predicted high correlation has been observed in verbal paired-associate learning by Caplan et al. (2006), Madan, Glaholt, and Caplan (2010), and Rizzuto and Kahana (2000), and in object-location pairs by Sommer, Rose, and Büchel (2007). Note that rather than conceptualizing A–B as a Gestalt, an alternative view would be to assume a single, bidirectional association or two separate but highly correlated associations. In any case, to produce a high correlation between forward (A–?) and backward (?–B) probes, forward probes must test the same underlying variability in encoding as backward probes. Mathematically the two measures—asymmetry in mean performance, and forward–backward correlation—are distinct measures. To understand why, consider hypothetical Participant 1 who was presented with the pair RATION–COUSIN. During test, Participant 1 responded correctly on the forward probe (RATION–?) but incorrectly on the backward probe (?–COUSIN); on another pair, REVOLT–VIRTUE, Participant 1 responded incorrectly on the forward probe (REVOLT–?) but correctly on the backward probe (?–VIRTUE). For Participant 1, mean accuracy in both directions is the same (50%); however, the correlation between forward and backward

cued recall is not high and positive (it is in fact negative; success in one direction is predictive of failure in the opposite direction). In contrast, Participant 2 receiving the same study and cued recall test responded correctly on both the forward and backward probes of RATION–COUSIN and incorrectly on both the forward and backward probes of REVOLT–VIRTUE; Participant 2 would also achieve 50% mean accuracy, but, in contrast to Participant 1, with a high positive correlation between forward and backward probes. Furthermore, the mean-accuracy and correlation measures are separably modifiable by experimental manipulations when measured empirically in human data (Madan et al., 2010). For these reasons, we consider the forward–backward probe correlation in addition to mean performance measures in this paper. The present work starts with constrained versions of the three aforementioned classes of association-memory models that are required to fit the two empirical properties of cued recall behaviour: symmetry in mean cued recall accuracy and a high forward–backward correlation in successive tests.

1. *Convolution-based models* need no modification to account for these two properties of cued recall behaviour since the convolution operation already embodies bidirectionality, or "associative symmetry" (Asch & Ebenholtz, 1962)—namely, no distinction between forward and backward associations.

2. *Matrix and concatenation-based models* must treat the left-handed ("A") items the same as the right-handed ("B") items to achieve symmetry in mean cued-recall accuracy. This means assuming that encoding levels are equal on average, which would produce symmetry in the mean accuracy measure. It

---

(Figure 1, continued) *The similarity function reflects uncertainty in position information similar to what others have found (e.g., Estes error gradients, Lee & Estes, 1977). Critically, we do not do anything differently for double-function than for single-function pairs, except insofar as items from double-function pairs are assigned two positional codes during study. Consequently, any dissociations between memory for single- and double-function pairs must follow from the structure of the study set rather than from differential study or test processes. Output encoding, which occurs whenever the model makes a response, functions identically to study, except that the pairs stored are based on the model's recall (which could be erroneous). Note that although this illustration is given for a positional-coding model, the same kind of mechanism could straightforwardly be applied with similar outcome to the other model classes.*

also means yoking average encoding-variability (and retrieval-variability, if present) levels between the first and last halves of the concatenated vector representation of the association. A high correlation between forward and backward recall directions may already be a property of some concatenation models such as MINERVA 2 (Hintzman, 1984, 1986) and models based on Search of Associative Memory (Gillund & Shiffrin, 1984; Mensink & Raaijmakers, 1988, 1989), whereas concatenation models based on the matrix model require a nearly symmetric encoding weight matrix (Rizzuto & Kahana, 2000, 2001).

3. *Positional-coding models* do not necessarily produce a high forward–backward correlation. Because each item is assumed to be stored independently (i.e., with independently drawn encoding strengths), forward and backward cued recall would be completely uncorrelated unless additional assumptions are added such that forward and backward probes are susceptible to the same levels of variability in encoding (for example, by requiring that each item-to-position strength be identical to its corresponding position-to-item strength), and within-pair positional separation (s_W in Figure 1a) is far smaller than between-pair positional separation (s_B in Figure 1a), termed the "isolation principle" (Caplan, 2005). Note that in simulations, the value of isolation, defined as $1 - s_W/s_B$, has not needed to take on an extreme value, which might have made this assumption less plausible. For instance, a ratio of .88 was sufficient to fit the high forward–backward correlation measured in empirical data (Caplan et al., 2006).

Kahana (2002) demonstrated that increasing the number of competing associations could reduce cued recall accuracy and implied that this effect on mean performance might not disrupt the high correlation between forward and backward probes. However, as we show shortly, the prediction for the correlation measure is model-dependent.

Here we investigate the double-function list paradigm (Primoff, 1938) because in this

paradigm, order is not explicitly demanded, but within-pair order memory is necessary to overcome associative interference. In this paradigm, we manipulate the number of episodic associations learned in the same study set. Some studied pairs have no other episodic associations—these are termed "single-function" pairs, denoted A–B, C–D, E–F, G–H. Other pairs each have one competing episodic association that involves each of the constituent items from the same study set—these are termed "double-function" pairs, denoted I–J, J–K, K–L, L–I. Because single-function pairs are typically symmetric in the mean-accuracy measure, and associative interference in this design is balanced, all models that incorporate symmetric association strengths should predict symmetric mean accuracy for double-function pairs. Prior double-function studies have also reported an overall reduced accuracy for double-function compared to single-function pairs (Howard et al., 2009; Primoff, 1938; Slamecka, 1976; Young, 1961), a property that is consistent with convolution-based models due to directional ambiguity as well as models with imperfect order memory such as positional-coding models (Caplan, 2005, Caplan et al., 2006). Models with perfect within-pair order memory, such as matrix and concatenation-based models, would not necessarily predict reduced accuracy, but could be modified (in an arguably uninteresting way) with a simple assumption that participants are challenged by the ambiguity in double-function pairs and study them less efficiently.

No study has measured the correlation between forward and backward cued recall of double-function pairs, the main goal of the experiments we report here. Kahana (2002) suggested that the effect of competing associations would not reduce the usual high forward–backward correlation observed for verbal associations, but Caplan (2005) suggested it would do precisely that. The prediction in fact depends on how order is represented within the association. We consider each of the aforementioned three classes of models in turn, in terms of the predictions they would make for mean accuracy and forward–backward correlation measures for double-function pairs.

1. *Convolution-based models.* Because the associations include no directional information, a probe item vector would be applied without regard to direction. Assume the model has learned $\mathbf{w} = \ldots + \gamma_{ij}\mathbf{i} * \mathbf{j} + \gamma_{jk}\mathbf{j} * \mathbf{k} + \gamma_{kl}\mathbf{k} * \mathbf{i} + \ldots$, where the γ values are assumed to be encoding strengths drawn independently for each pair. A probe of the pair J–K in the forward direction would be implemented as $\mathbf{j} \# \mathbf{w}$ and, as argued for serial lists previously (Caplan, 2005), would be susceptible to substantial levels of competition from the I–J pair, $\mathbf{i} * \mathbf{j}$, proportional to γ_{ij}. A probe in the backward direction would be implemented as $\mathbf{k} \# \mathbf{w}$ and would be susceptible competition from the K–L pair, $\mathbf{k} * \mathbf{l}$, proportional to γ_{kl}. Although both forward and backward cued recall operations are susceptible to common variability in learning (γ_{jk}), the level of competition depends on different sources of variability in learning, γ_{ij} versus γ_{kl}, respectively. Thus, a convolution-based model would predict that independent sources of competition should *reduce* the correlation between forward and backward cued recall probes of double-function pairs compared to single-function pairs.

2. *Matrix and concatenation-based models.* In current implementations of matrix and concatenation-based models, order within a pair is perfectly stored. For concatenation models, probing J–K in the forward direction would be implemented, for example, with $\mathbf{j} \oplus \boldsymbol{\eta}$ as the probe vector, where $\boldsymbol{\eta}$ is a "noise" vector used as a place-holder, filled with random values. A backward probe would be $\boldsymbol{\eta} \oplus \mathbf{k}$. In current concatenation models, assuming near-orthogonality of item vectors, $\mathbf{j} \oplus \boldsymbol{\eta}$ would tend to be extremely similar to the target pair, $\mathbf{j} \oplus \mathbf{k}$, but dissimilar to the other pair that shares the $\mathbf{j}$ (probe) item, $\mathbf{i} \oplus \mathbf{j}$; likewise, for the backward probe, $\boldsymbol{\eta} \oplus \mathbf{k}$ would tend to be extremely dissimilar to the potentially conflicting pair, $\mathbf{k} \oplus \mathbf{l}$. The consequence is that forward and backward probes will be correlated due to both probe directions being sensitive to a common source of variability:

γ_{jk}, the encoding strength of the concatenated target pair $\mathbf{j} \oplus \mathbf{k}$, and only negligible independent source of variability due to incidental similarity between probe items and other list items. Because the same logic applies to double-function as to single-function pairs, the prediction based on concatenation-based models is that the forward–backward correlation should *not* be reduced for double-function pairs. For matrix models, the same type of argument applies, considering that a forward probe with $\mathbf{j}$ by multiplying from the left would retrieve $\mathbf{k}$ but not $\mathbf{i}$. In order to retain the near-perfect correlation between forward and backward probes of single-function pairs, one has to assume either that the model can probe the same outer-product term (e.g., $\mathbf{jk}^{\mathrm{T}}$) from both the left (forward probe) and the right (backward probe), or that probing is carried out only by multiplying from one side but both forward and backward associations are stored in separate memories but with (nearly) perfectly correlated encoding strengths. This also means that the forward and backward probe will be driven by the same chief source of variability. Current concatenation-based and matrix models thus lead to the alternative prediction that the forward–backward correlation will be *equal* for double-function and single-function pairs.

3. *Positional-coding models.* As mentioned above, Caplan (2005) proposed that the high forward–backward correlation observed for cued recall of pairs (as contrasted with cued recall of serial lists) could be explained in a positional-coding model if within-pair positional codes were relatively much closer together than between-pairs positional codes (the "isolation principle"). Caplan (2005) further argued that the repetition of items in a double-function paradigm would result in positional ambiguity for double-function probe items, resulting in an independent source of encoding variability entering into forward versus backward cued recall. Thus, similar to convolution-based models, but in contrast to matrix and concatenation-based models, a

simple extension of the positional-coding model to accommodate repeated items leads to the prediction of *reduced* forward−backward correlation for double-function relative to single-function pairs.

We present three double-function experiments with successive testing to assess the relationship between forward and backward associations and report mean-performance measures and correlation measures that test assumptions of convolution, matrix/concatenation, and positional-coding based models. Experiment 1 was an experiment conducted in an undergraduate class, which produced the initial results. Because there was no randomization of list stimuli, we feared that the results might be due to idiosyncrasies in the lists. Experiments 2 and 3 differed in presentation rate; we thought that faster presentation rates should produce more of a challenge from associative interference than would slower presentation rates. Thus, we wanted to know whether the pattern of findings would replicate. All experiments yielded the same pattern of correlations. Thus, the three experiments suggest broad boundary conditions for the correlations findings: They occur whether mean accuracy is asymmetric (Experiment 1) or symmetric (Experiments 2 and 3) and for slow (Experiments 1 and 2) as well as fast (Experiment 3) presentation rates. To foreshadow the results, no model class provides a complete account of the resulting pattern of behaviour. In the Discussion we consider ways in which each model class would need to be modified to account for the full pattern of data we report.

Method

We report three experiments; because the methods and results are similar, we present the methods and results for all experiments together. Table 1 compares the design of the three experiments.

Participants

Participants were University of Alberta students. Experiment 1 was carried out in a third-year class of a Neurobiology of Learning and Memory course in which 79 students participated voluntarily without compensation. A total of 90 and 60 participants took part in Experiments 2 and 3, respectively, for partial fulfilment of course requirements of introductory psychology courses.

Materials

Study sets were drawn from a 126-word pool constructed using the MRC Psycholinguistic Database (Wilson, 1988). All stimuli were nouns that were two syllables in length and composed of six letters. Kucera−Francis frequency was constrained between 7−52 per million; imageability and concreteness were both constrained to 240−560 out of 700. As Experiment 1 was conducted in a classroom setting, the same study set was used for all subjects. For Experiments 2 and 3, words were drawn at random. In all three experiments, equal numbers of single-function and double-function pairs were presented. Single-function pairs did not share items with other pairs in the list (i.e., A−B, C−D, E−F, G−H). Double-function pairs shared items with other pairs in the list (i.e., I−J, J−K, K−L, L−I), forming a complete ring structure.

Procedure

Participants in all three experiments participated in a practice round (excluded from analyses) followed by experimental sets (Figure 2). Each round in the task consisted of several phases: study, distractor, cued recall (Test 1), distractor, and cued recall (Test 2).

Participants were presented with study sets comprising both single-function and double-function pairs in a randomized order with the restriction that double-function pairs with a common item could not be temporally contiguous at study or test. Each pair was simultaneously presented and was separated from the next by a blank screen that appeared for 150 ms. Distractor tasks were as follows. In Experiment 1, participants had to count the total number of a given digit in an 8 × 6 matrix made up of two digits. They were given 20 s to count and record their responses. In Experiments 2 and 3, participants had five addition questions involving addition of

Table 1. Three experiments with slight modifications were conducted with this paradigm

Experiment	Setting	Number of subjects	Mode of response	Number of rounds	Number of pairs per round	Number of single-function/double-function pairs	Presentation rate of stimuli (s)
1	In class	79	Written	3	8	4/4	4
2	Individual	90	Typed	7	12	6/6	4
3	Individual	60	Typed	7	12	6/6	2

Note: In Experiment 1 the distractor task was 20 s long, and in Experiments 2 and 3 it was 5 s per question. In all tasks, subjects were given 10 s per probe to respond.
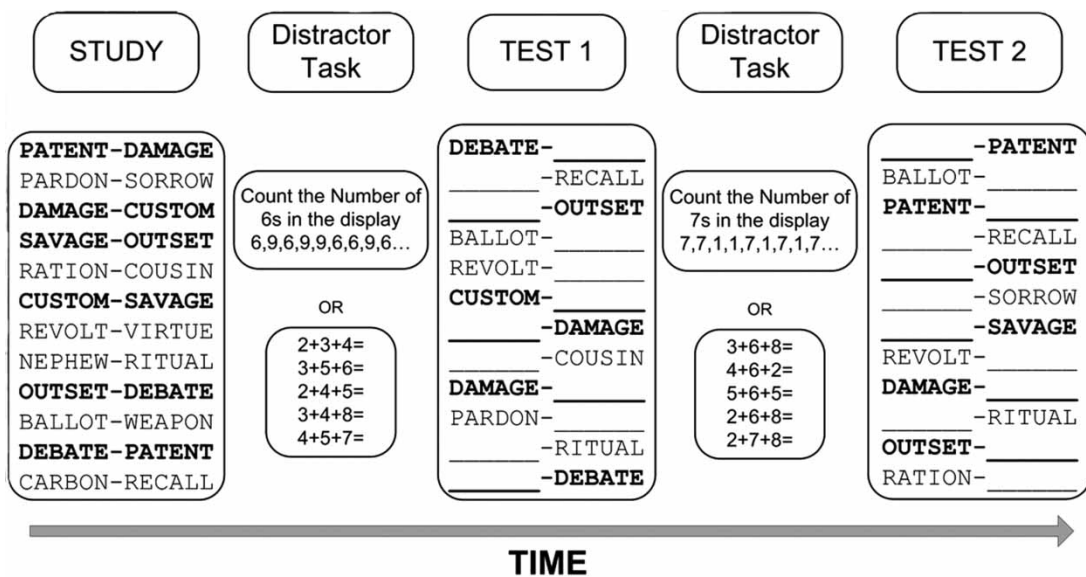

Figure 2. *A single round of the task. In the study/test phases, the order of presentation of single-function and double-function (bolded for illustration only) pairs/probes was randomized. In the test phases, the order of forward and backward probes was also randomized.*

three digits (from 2 to 8). Each distractor question was separated from the next by a blank screen that appeared for 150 ms. Participants were given 5 s to type their responses.

Cued recall consisted of a probe word and a blank line, either to the left or to the right of the word. In this way the participants were tested in either forward or backward direction for the pairs just studied. Each probe was separated from the next by a blank screen that appeared for 250 ms. In all three experiments, participants were given 10 s to recall and record their responses. Both response initiation (first key press) and

termination ("ENTER" key press) times were logged in Experiments 2 and 3. Because termination times exhibited the same qualitative effects as initiation times but with less sensitivity, we report only analyses of initiation times.

After each pair in the study set was tested exactly once ("Test 1"), a second distractor task followed, and then a second complete set of tests was conducted ("Test 2"). The direction of test on Test 1 and Test 2 could be the same (forward–forward or backward–backward) or different (forward–backward or backward–forward) as depicted in Table 2.

Table 2. Conditions of successive testing

Successive tests	Test 1 probe	Test 2 probe	Correlation contributed to
Forward−forward	REVOLT−?	REVOLT−?	Same
Backward−backward	?−VIRTUE	?−VIRTUE	Same
Forward−backward	REVOLT−?	?−VIRTUE	Different
Backward−forward	?−VIRTUE	REVOLT−?	Different
Bootstrap (different pairs)	ABSENCE−?	REVOLT−?	Control

Note: Participants can be tested in four possible conditions over two tests. Two of the conditions comprise the same probe in both tests, and two comprise different probes in both tests. The examples assume that the participant studied the pairs REVOLT−VIRTUE and ABSENCE−HOLLOW in the same study set.

In Experiment 3 only, participants were given a final free recall test at the end of the session; Participants were given 3 minutes to recall and record words from the task in any order that they remembered them (data not reported here). In all experiments each participant was given a strategy questionnaire at the end of the session (data not reported here).

Correlation of accuracy on successive tests

Our measure of correlation between Test 1 and Test 2 probes was Yule's Q, a measure of association related to the odds ratio, which is appropriate for dichotomous data. The "same" Q value estimates the correlation due to test−retest reliability alone and will be the highest correlation when successive testing is done in the same direction (forward−forward and backward−backward). The "different" Q value is the correlation between forward and backward directions, when Test 1 and Test 2 were in different directions (forward−backward and backward−forward) and is our measure of interest. The "control" Q value is the lowest possible expected correlation and measures the correlation between unrelated pairs within the same round. It is calculated by correlating one pair from Test 1 with a different pair from Test 2 (a bootstrap) and is included to control for subject and study set variability. Note that for the subject-level approach, we include all such pairs, but for the aggregate approach we include only a random pairing of each pair on Test 1 with exactly one different pair on Test 2 to avoid overestimating degrees of freedom. Thus, the

"same" and "control" correlations demonstrate the effective range of the "different" correlation. Yule's Q was evaluated using the log-odds ratio transform (Bishop, Fienberg, & Holland, 1975; Hayman & Tulving, 1989).

Statistical analysis

Analyses of variance (ANOVAs) were performed with the Greenhouse−Geisser correction for non-sphericity. Effects are considered significant based on an alpha level of .05. Effects not reported were not significant.

Results

Mean accuracy

Prior to our measure of interest, the correlation between accuracy on Test 1 and Test 2, we asked whether mean performance varied as a function of the factors in our design. We performed a three-way repeated measures ANOVA on each data set, on Test (2) × Function (2) × Direction (2). In Experiment 1 (Figure 3a), single-function pairs were recalled more accurately than double-function pairs [main effect of function, $F(1, 78) = 64$, $MSE = 0.33$, $p < .0001$], suggesting that interference presented a challenge to participants. Test 2 was more accurate than Test 1 [main effect of test, $F(1, 78) = 17$, $MSE = 0.082$, $p < .0001$], suggesting a small amount of output encoding. Forward probes were recalled more accurately than backward probes [main effect of direction, $F(1, 78) = 14$, $MSE = 0.40$, $p < .0001$]. Finally, the interaction Test × Function
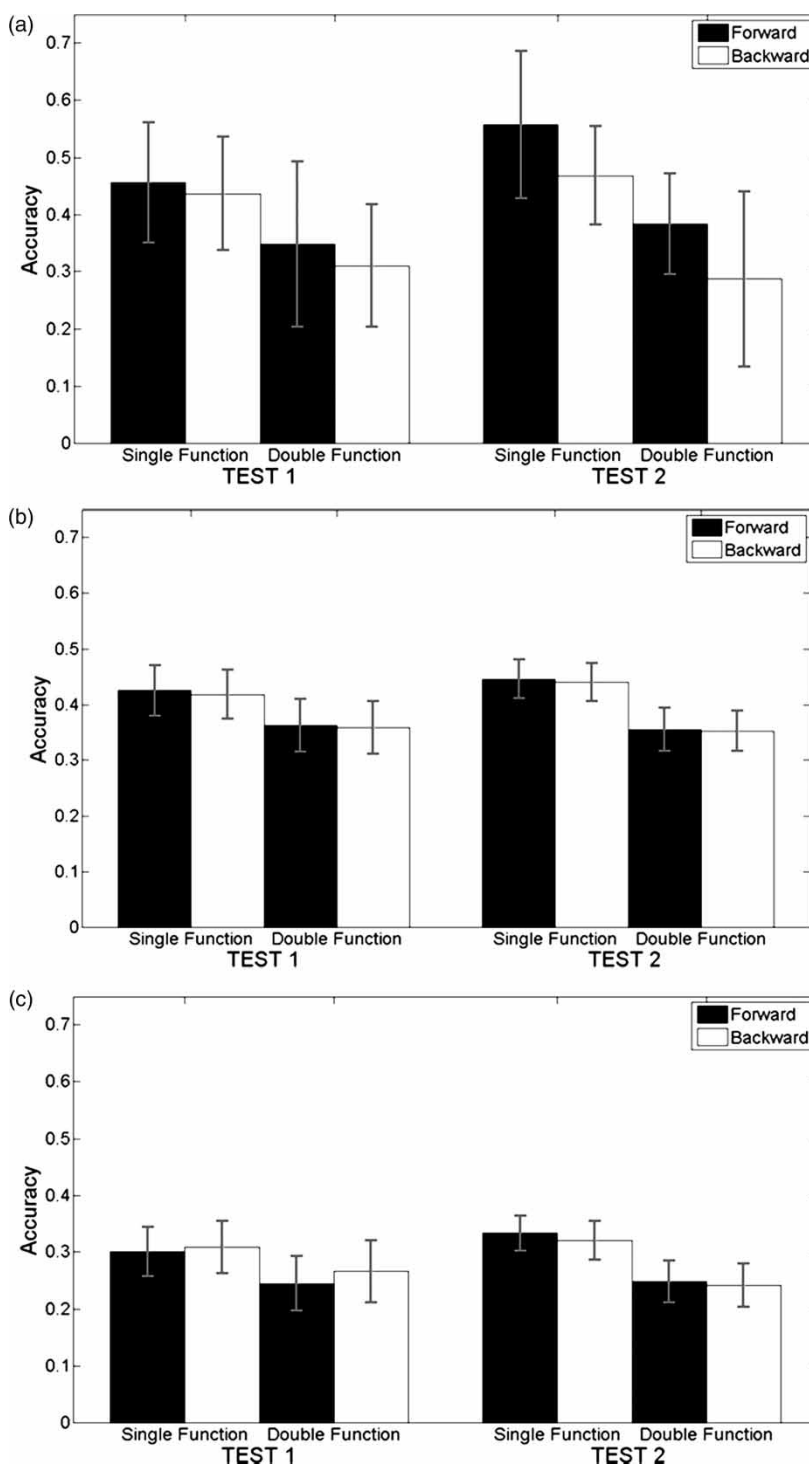
Figure 3. *Mean accuracy in (a) Experiment 1, (b) Experiment 2, and (c) Experiment 3 for single-function and double-function probes in forward and backward directions. Error bars are 95% confidence intervals.*

was significant, $F(1, 78) = 13$, $MSE = 0.072$, $p <$ .001, but simple effects revealed a single-function advantage over double-function pairs for both tests, so the interaction was quantitative, not qualitative.

Experiments 2 (Figure 3b) and 3 (Figure 3c) showed a subset of significant effects with the same pattern as that found in Experiment 1. For these experiments, the main effects of direction and test were not significant. However, the following were significant: Experiment 2, main effect of function, $F(1, 89) = 28$, $MSE = 0.037$, $p <$ .0001; and Test × Function interaction, $F(1, 89) = 12$, $MSE = 0.003$, $p < .001$; Experiment 3, main effect of function, $F(1, 59) = 21$, $MSE = 0.025$, $p < .0001$; and Test × Function interaction, $F(1, 59) = 12$, $MSE = 0.03$, $p < .001$. As in Experiment 1, these interactions were quantitative, not qualitative.

Response time

We analysed mean response times for correct responses with ANOVAs with the same design as that used for accuracy primarily to test for possible speed−accuracy trade-offs, which could have complicated interpretation of the accuracy-correlation analyses; however, all response time effects were congruent with effects of accuracy, suggesting that no significant trade-off was present in our data. Response times were not collected for Experiment 1 because responses were handwritten. Both Experiment 2 (Figure 4a) and Experiment 3 (Figure 4b) had significant main effects of function [Experiment 2: $F(1, 89) = 16$, $MSE = 3.8 \times 10^5$ ms^2, $p < .0001$; Experiment 3: $F(1, 59) = 9.8$, $MSE = 2.9 \times 10^5$ ms^2, $p < .01$] and of test [Experiment 2: $F(1, 89) = 46$, $MSE = 2.0 \times 10^5$ ms^2, $p < .0001$; Experiment 3: $F(1, 59) = 14$, $MSE = 3.1 \times 10^5$ ms^2, $p < .0001$], where correct responses on single-function probes were significantly faster than correct responses to double-function probes, and responses were faster on Test 2 than on Test 1. Finally, for Experiment 2 only, forward probes were recalled significantly faster than backward probes: main effect of direction, $F(1, 89) = 5.0$, $MSE = 2.1 \times 10^5$ ms^2, $p < .05$.

Correlation of accuracy on successive tests

The correlation (Yule's Q) between forward and backward probes was our measure of interest because it is more diagnostic of model mechanisms (see introduction). We also calculated the "same" and "control" correlations (see Method) to estimate the possible range of measurable Q values. We calculated Yule's Q in two ways, but the findings did not differ qualitatively between methods. In the first method—"aggregate"—the values in the contingency table that were used to calculate the correlation were collapsed across all subjects in a given experiment, and one correlation value was calculated for "same", "different", and "control". We calculated correlation this way for all experiments. Confidence intervals and pairwise comparisons were done using the log-odds ratio transform and standard error values, which can be computed directly for log-odds ratios as discussed by Bishop et al. (1975) and Hayman and Tulving (1989). In the second method—"subject level"—we computed values in the contingency table for each subject individually and calculated a correlation value for each subject individually. This removes correlation due to subject variability, which can inflate correlations (Simpson's paradox; cf. Hintzman, 1972). We calculated correlation this way for Experiments 2 and 3 (but not for Experiment 1 due to insufficient data per participant).

As illustrated in Figure 5, the correlation of "different" double-function probes was lower than that of the "different" single-function probes as predicted by convolution and positional-coding models due to interference in double-function pairs. Pairwise comparisons (aggregate method, via log-odds transform) revealed that, for all three experiments, the differences between "same" and "different" within single-function and double-function probes were significant and that the differences between "same" and "different" between single-function and double-function were significant in all three experiments ($p < .01$). All Q values were significantly greater than the control Qs ($p < .01$).

Calculating Q using the subject-level method allowed us to carry out ANOVAs. This yielded consistent qualitative results as the aggregate
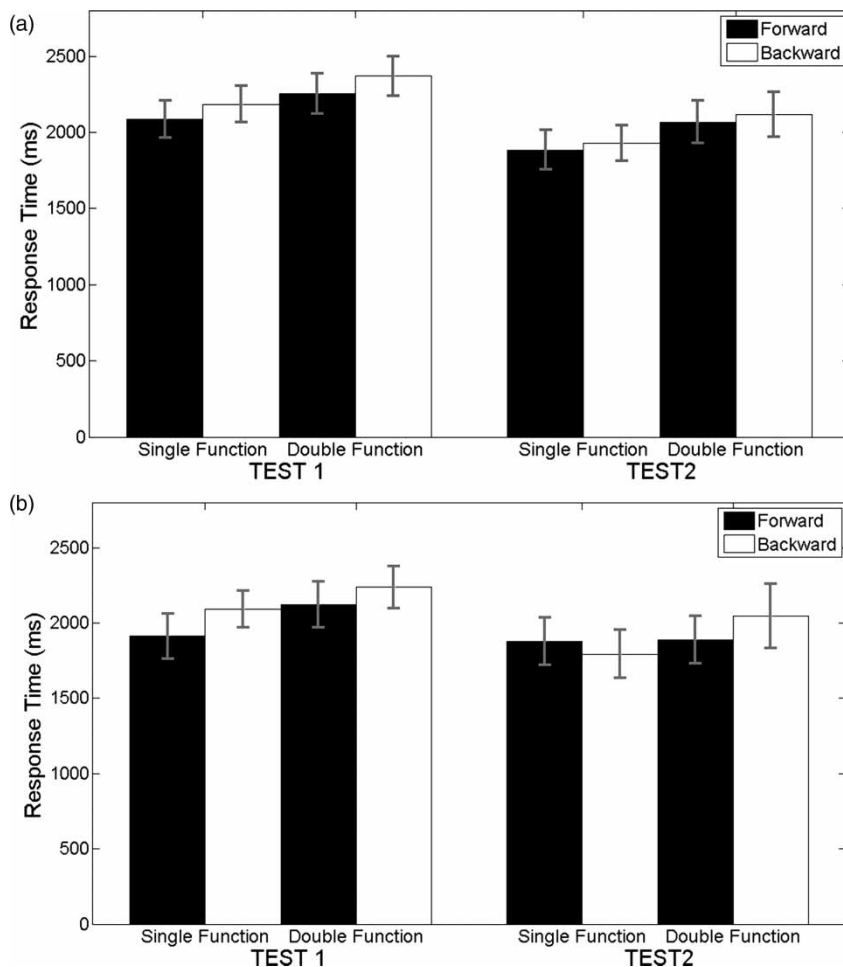
Figure 4. *Mean response times in (a) Experiment 2, and (b) Experiment 3 for single-function and double-function probes in forward and backward directions. Error bars are 95% confidence intervals.*

method but better summarizes our 2×2 design. We performed a two-way repeated measures ANOVA on Type (2) × Function (2). The levels of Type were "same" correlation and "different" correlation. Both Experiments 2 and 3 exhibited both significant main effects and a significant interaction—Experiment 2: main effect of type, $F(1, 89) = 31$, $MSE = 0.353$, $p < .0001$, with "same" correlations greater than "different"; main effect of function, $F(1, 89) = 140$, $MSE = 0.25$, $p < .0001$, with single-function pairs more accurate than double-function pairs; and a Type × Function interaction, $F(1, 89) = 24$, $MSE =$

0.21, $p < .0001$; Experiment 3: main effect of type, $F(1, 59) = 77$, $MSE = 0.28$, $p < .0001$; main effect of function, $F(1, 59) = 41$, $MSE = 0.29$, $p < .0001$; and Type × Function interaction, $F(1, 59) = 19$, $MSE = 0.22$, $p < .0001$. The interaction in both experiments was quantitative, not qualitative.

Discussion

Accuracy

Measures of mean performance show effects expected based on prior research. Single-function
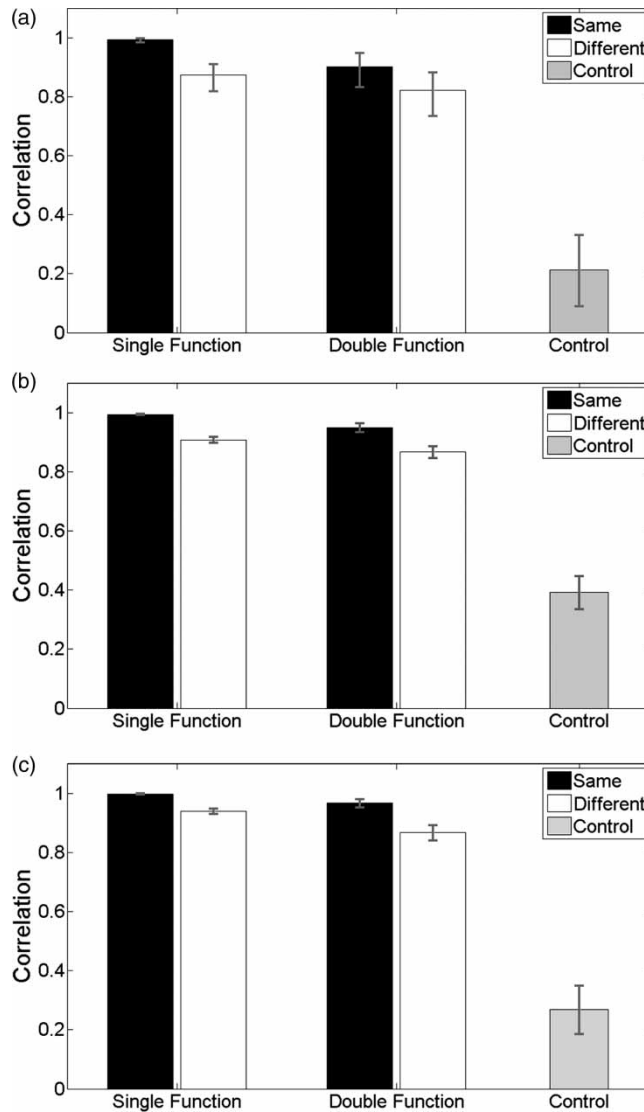
Figure 5. *Correlation of accuracy on successive testing (computing a single correlation value for each testing type for all subjects) in (a) Experiment 1, (b) Experiment 2, and (c) Experiment 3 for single-function and double-function probes. For a description of "same", "different", and "control" refer to Table 2. Error bars are 95% confidence intervals.*

pairs were significantly better recalled than double-function probes (Figure 3). This finding has been replicated in many studies since Primoff (1938): Howard et al., 2009; Slamecka, 1976; Young, 1959, 1961. Effects of direction were also not significant in accuracy, replicating numerous findings (for a review see Kahana, 2002) except

for Experiment 1 in which materials were not randomized across participants, so the overall forward-probe advantage in that experiment may be due to idiosyncrasies of the specific study sets used. Critically, these overall asymmetries in mean performance in Experiment 1 nonetheless yielded a correlation pattern consistent with

Experiments 2 and 3, underlining the independence of mean performance measures and Test 1–Test 2 correlation measures data (Madan et al., 2010). Interestingly, a response-time advantage for forward probes was observed in Experiment 2, suggesting that sometimes even when the association is symmetric (in both senses—mean accuracy and correlation measures), access can be asymmetric (Waugh, 1970). Regarding the model classes we consider:

1. The finding that accuracy for double-function pairs exceeds half the accuracy for single-function pairs, a replication of numerous prior results, presents a challenge to convolution-based models because these models have no way of distinguishing forward versus backward associations. The model would at best have to guess between two retrievable alternatives. Convolution-based models clearly need an additional mechanism by which to retrieve within-pair order information better than chance.
2. The lower accuracy of double-function than that of single-function pairs presents a challenge to matrix/concatenation-based models but only a mild one because, as mentioned in the introduction, such models could be very simply modified to assume that double-function pairs are studied less well due to repeated items being confusing or distracting to the participant.
3. Positional-coding models can accommodate the moderately reduced accuracy of double-function pairs because they encode within-pair position (along with pair position) but positional confusability ensures that this order information is imperfect. No major modification of positional-coding models is necessary to accommodate the mean-accuracy data.

Forward−backward correlation

The correlation between forward and backward probes is the measure that speaks most directly to the models we considered. All the "different" correlations are quite high despite being lower than possible ("same"), replicating numerous prior findings (Caplan et al., 2006; Kahana, 2002; Madan et al., 2010; Rizzuto & Kahana, 2001). As predicted by the interference-theory account (Caplan, 2005), the "different" correlation for double-function pairs was lower than that for "different" single-function pairs. Thus, interference from other items in the study set may have lowered the correlation. This finding relates to models as follows.

1. *Convolution models*, as laid out in the introduction, predict the reduced correlation for double-function pairs because within-pair order ambiguity means that pairs with a common item will compete with one another and introduce independent sources of competition depending on probe direction.
2. *Matrix and concatenation-based models*—as they are currently developed—are challenged by this finding for the reasons laid out in the introduction. Namely, because within-pair order is encoded perfectly (as long as the association itself is stored), an item in the left position is effectively entirely dissimilar to the same item in the right position. Because this prevents the very reversal errors that cause other classes of models to exhibit a reduced forward−backward correlation for double-function compared to single-function pairs, current matrix and concatenation models would need to be enhanced to accommodate this novel property of our findings—even if they were already modified to account for the numerous preexisting findings of reduced mean accuracy for double-function pairs as suggested above.
3. *Positional-coding models* naturally accommodate the reduced forward−backward correlation for double-function pairs for a similar reason as argued for convolution models—namely, that the source of interference will differ depending on probe direction. This argument was previously applied to cued recall of serial lists (Caplan, 2005; Caplan et al., 2006).

Test/retest reliability: A challenge to all models

A finding we had not anticipated was that the correlation of the "same" double-function probes was

also significantly lower than that of the "same" single-function pairs. The reason this characteristic was surprising is that the mechanism by which interference reduces the correlation between probe directions is by introducing substantial sources of interference that *differ* in strength depending on probe direction. However, two tests in the same direction (which enter into the "same" correlation) should be susceptible to precisely the same source of interference on both tests. Thus, we had expected that the "same" correlation would be just as high for double-function pairs as for single-function pairs even while the "different" correlation should have been reduced for double-function pairs. This logic can be applied to all three classes of model.

Modifications needed to retain each model class
Models arguably yield the most theoretical value when they fail. From that perspective, the present findings are useful in that they challenge each of the three major classes of model we consider, but each class is challenged in a different way. Although we cannot reasonably select one model class over another, we can use the empirical challenges to each to lay out how each model class must be extended; the further constraints on each model will help guide future development of the models to account for further empirical data. To this end, we now discuss each model class in turn with respect to how one might enhance the model class to accommodate all the features of the present findings that current implementations would miss.

Convolution models. This class of model as currently developed fails at the earliest stage, being unable to account for prior evidence of greater than chance memory for within-pair order both in associative recognition with reversed lures (Green & Tussing, 2001; Kounios et al., 2003) and in previous examples of better than chance forward recall accuracy in double-function lists (Howard et al., 2009; Primoff, 1938; Provyn, Sliwinski, & Howard, 2007; Young, 1961), replicated here. A complete exploration of means of adding order information to convolution models

would be lengthy and beyond the scope of this manuscript but there are a few candidate mechanisms that have been suggested previously (Plate, 1995; Westlake, 1970). First, a convolution model could be combined with the concatenation mechanism we discussed in reference to the second class of models. Then, associations could be learnt by autoconvolution of the concatenated vector: $\mathbf{w} = (\mathbf{a} \oplus \mathbf{b}) * (\mathbf{a} \oplus \mathbf{b})$. This mechanism would then be susceptible to the limitations of Class 2, so further elaboration would be similar to Class 2. A second approach might be to devote a subset of item-vector features to an explicit order code. Thus, associations would be stored like $\mathbf{w} = (\mathbf{a} \oplus \mathbf{t}_{\text{left}}) * (\mathbf{b} \oplus \mathbf{t}_{\text{right}})$, where $\mathbf{t}_{\text{left}}$ and $\mathbf{t}_{\text{right}}$ denote vector representations of the left-handed and right-handed within-pair positions, respectively. The outcome with respect to accuracy on double-function relative to single-function pairs and the correlation between forward and backward probes of double-function pairs would depend on the relative dimensionality and encoding reliability of the position-code features. if the $\mathbf{t}$ codes were too dominant in encoding and retrieval operations, they might undermine the desirable associative-symmetry properties of convolution models too much, losing the ability to fit the magnitude of the reduction in double-function accuracy as well as the high forward−backward correlation in cued recall of single-function pairs. A third strategy might be to return to an old suggestion that the convolution mechanism could be applied in the temporal domain, originally termed "temporal holography" (Gabor, 1968; Longuet-Higgins, 1968), but may be susceptible to similar limitations as perfect-order models (matrix and concatenation-based models). Following such an extension of convolution models, the reduction of test/retest reliability would still have to be addressed, probably in a manner that we suggest for the other two model classes below.

Matrix and concatenation-based models. This model class missed findings that suggest that within-pair order is not perfectly reliable (given retrieval of the association). Thus, in contrast to convolution models, this class of model needs a way of

making within-pair order more ambiguous. Recall first that to fit prior findings of highly correlated forward and backward cued recall accuracy, these models needed modification. Matrix models needed either the ability to probe the same encoded heteroassociation terms from the left or right, or to store forward and backward associations separately but with correlated encoding strengths. Concatenation-based models needed the left- and right-handed item strengths (or forward- and backward-directed association strength) to be highly correlated. The most obvious modification that could introduce direction ambiguity to these models would be to assume that first, whatever within-pair order is stored will be perfectly retrievable given retrieval of the association and, second, the model makes errors in initial encoding of directionality, sometimes storing the reverse within-pair order. Thus, there is a probability, $p_{reversal}$, that the model will store $\mathbf{b} \oplus \mathbf{a}$ rather than $\mathbf{a} \oplus \mathbf{b}$ (Caplan, 2004). This would maintain the high correlation between forward and backward cued recall of single-function pairs since if $\mathbf{a} \oplus \mathbf{b}$ were stored, both forward and backward probes would be effective, whereas if $\mathbf{b} \oplus \mathbf{a}$ were stored, both forward and backward probes would fail. It would also result in a reduction of accuracy for double-function pairs for the following reason. Consider a list I–J, J–K, K–L, L–I. If I–J and J–K were encoded correctly, but K–L were erroneously stored as L–K, then a forward probe of K–L would be likely to be highly accurate because $\mathbf{j}$ in the left position is completely unconfusable with $\mathbf{j}$ in the right position. However, the backward probe of J–K would be susceptible to interference from K–L (stored as L–K) since probing with $\mathbf{k}$ from the right would retrieve both J–K and the (incorrect) K–L. Because in general, forward and backward probes could be susceptible to competition from independently stored associations, this should reduce the forward–backward correlation for double-function relative to single-function pairs. Thus, a simple assumption of fallibility in encoding of within-pair order could enable concatenation and matrix models to explain most of the features of the data reported here.

However, the reduced test/retest ("same") correlation for double-function pairs requires further modification. To understand why, consider the example just mentioned. On both Test 1 and Test 2, the same source of interference (or lack thereof) is present. Thus, if a double-function pair faces a high level of competition on Test 1, then (without further assumptions) it will face the exact same level of competition on Test 2. Thus, there is no mechanism by which the test/retest correlation ("same") could be lower for double-function than single-function pairs. One might think that a probabilistic retrieval rule such as Luce's choice rule (Luce, 1959) would solve this problem. However, to the extent that one makes retrieval probabilistic, one will reduce *all* correlations ("same" as well as "different") for single-function pairs as well as double-function pairs. It may be impossible (or require precisely tuned model parameters) to achieve a reduction in forward–backward correlation for double-function pairs without undermining the model's ability to fit the remaining correlation values. We suggest one plausible approach that can bypass these trade-offs—output encoding, which simply refers to learning during recall. The logic is similar for all classes of model, so we expand on this idea in a separate discussion section below.

Positional-coding models. Recently, positional/contextual coding models have been adapted from their original motivation, to describe serial list learning, to also explain data on paired-associates learning (Caplan, 2005; Caplan et al., 2006; Howard et al., 2009). Because this unified modelling framework can potentially provide a more parsimonious account of verbal memory than do distinct modelling frameworks, explaining two classes of paradigms together, it is important to consider whether positional-coding models could also account for the paired-associate learning phenomena reported here. Significantly, the positional-coding model applied here assumes statistically independent storage operations (with uncorrelated encoding variability) for each of the items within a pair. In this way, the positional-coding model can treat memory for pairs

identically to memory for serial lists, by only specifying that the within-pair positional separation be substantially less than the between-pair positional separation. This class of model can nonetheless produce the observed high forward–backward correlation for pairs due to how variability combines during the probe process and the presence of a high degree of positional isolation between pairs (Caplan, 2005; Caplan et al., 2006). Positional-coding models should be able to explain the reduced forward–backward correlation for double-function pairs for the same reason as that applied to cued recall of serial lists in which paired items were not isolated from other list items (Caplan, 2005; Caplan et al., 2006). Rather, the overlapping pair, due to nearby positions being similar to one another, produces some confusion about within-pair order. Thus, double-function pairs are susceptible to interference from competing pairs, and the source of competition depends on probe direction. Assuming that each list item is stored with an independently drawn encoding strength, this differential source of interference will serve to partially decorrelate the otherwise high forward–backward correlation. As with the other model classes, however, current implementations of positional-coding models are challenged by the reduced test/retest correlation we found for double-function pairs because, as discussed for matrix- and concatenation-based models, the source and level of interference on Test 1 and Test 2 should be identical. Furthermore, the addition of a probabilistic retrieval mechanism threatens to reduce the forward–backward and test/retest correlation for single-function pairs as well as double-function pairs, causing the model to potentially be unable to fit the full pattern of data. We next discuss a possible solution that could be applied to this and the other model classes: output encoding.

Output encoding
Output encoding, learning during retrieval, was suggested as an account of the high "different" correlation in memory for pairs by Rizzuto and Kahana (2001). However, in their Hopfield network model they found that output encoding

served to inflate correlations but could not quantitatively explain the high value of the correlation without assuming that the underlying forward and backward encoding strengths were also highly correlated. Sommer, Schoell, and Büchel (2008) also found that output encoding could not explain a related high correlation in memory for object-location pairs and noted that cued recall can also produce reminiscence (retrieval of previously unrecalled items on Test 2). We reason that output encoding might have the effect of slightly increasing the "same" correlation for single-function pairs whereas for double-function pairs, output encoding could occasionally strengthen the level of interference between Test 1 and Test 2 of a given double-function pair, thus reducing the correlation between cued recall tests even in the same direction. The assumption is very simply that when a response is given (whether correct or incorrect), the model encodes the recalled association as though it were an additional presentation of the probe item paired with the recalled item (in the recalled direction); if no response was given on a cued recall trial, one can assume no output encoding occurs (Rizzuto & Kahana, 2001). For matrix, concatenation, and convolution models, the output-encoded pair can be added in precisely the same way as in initial study. For a positional-coding model, the output-encoded pair can be straightforwardly assigned a new positional code, which starts immediately following the last studied list position. A minimum of one additional free parameter would be associated with the output encoding mechanism, such as mean strength of output encoding, C_e. This extension is theoretically efficient in that it extends the applicability of an existing model process rather than introducing a completely new process.

For single-function pairs, output encoding functions like an additional study trial, potentially resulting in a higher encoding strength, which could facilitate retrieval on Test 2 relative to Test 1 or reencoding an error. Either way, output encoding serves to reinforce the correlation between Test 1 and Test 2. For double-function pairs, in addition to "same", the additional

potential benefit from the repeated encoding event, a competing pair might be strengthened, thus increasing the level of competition during retrieval of a particular double-function pair. This can reduce the test/retest reliability for double-function pairs without undermining the remaining high forward–backward correlations.

Positional-coding models of association memory

The success of positional-coding models, in that they require the least amount of modification from their current forms relative to the other model classes, suggests that other positional-coding models, as well as more complex models with positional-coding-like properties, such as Conrad's box model (Conrad, 1965), the Perturbation model (Lee & Estes, 1977, 1981), Mensink and Raaijmakers's context-dependent Search of Associative Memory models (Mensink & Raaijmakers, 1988, 1989; Raaijmakers, 2003), the Start-End Model (Henson, 1998), Burgess and Hitch's phonological loop model (Burgess & Hitch, 1999), the Temporal Context Model (Howard & Kahana, 1999), Serial Order in a Box (Lewandowsky & Farrell, 2000), and Scale-Independent Memory, Perception, and Learning (Brown, Neath, & Chater, 2007), could be extended to explain paired-associate learning.

Associative interference from semantic memory

The account presented here could apply to interference from preexperimental associates from semantic memory. Kahana (2002) demonstrated that one could produce either a forward- or a backward-probe advantage, if the paired items differed in their number of preexperimental associates. For example, if all the "A" items of the studied pairs had more preexperimental associates than the "B" items, then cued recall should be better in the backward direction on average, due to the "A" items being more ambiguous as probes than the "B" items. Our findings suggest that orthogonal to the question of symmetry or asymmetry in mean performance measures, the forward–backward correlation should reduce monotonically

as the number of preexperimental associates increases, regardless of whether those associates were attached to the "A" items or the "B" items. Indeed, in Experiments 2 and 3, we obtain symmetric mean performance but lowered forward–backward correlation for our double-function pairs, which were contrived to have an equal number of (episodic-memory) associates (two each). Thus, we predict that the forward–backward correlation for cued recall of previously learned semantic-memory associations (e.g., SALT–PEPPER) should be reduced relative to novel (highly episodic) associations.

Conclusion

In sum, the combination of mean-accuracy and correlation measures of cued recall of double-function and single-function pairs posed major challenges to existing forms of three major classes of models of pair memory. First, this underlines the importance of modellers directly specifying how within-pair order information should be stored, rather than treating within-pair order as a tangential question. Second, specific modifications of each model class were suggested to accommodate the full set of findings and without losing the ability to account for prior findings. Finally, the further constrained and extended models suggest specific ways in which participants might use within-pair order information to resolve associative interference.

REFERENCES

Anderson, J. A. (1970). Two models for memory organization using interacting traces. *Mathematical Biosciences*, 8, 137–160.

Asch, S. E., & Ebenholtz, S. M. (1962). The principle of associative symmetry. *Proceedings of the American Philosophical Society*, 106(2), 135–163.

Bishop, Y., Fienberg, S. E., & Holland, P. W. (1975). *Discrete multivariate analysis: Theory and practice.* Cambridge, MA: MIT Press.

Borsellino, A., & Poggio, T. (1971). Holographic aspects of temporal memory and optomotor responses. *Biological Cybernetics*, *10*(1), 58–60.

Brown, G. D. A., Neath, I., & Chater, N. (2007). A temporal ratio model of memory. *Psychological Review*, *114*(3), 539–576.

Burgess, N., & Hitch, G. J. (1999). Memory for serial order: A network model of the phonological loop and its timing. *Psychological Review*, *106*, 551–581.

Caplan, J. B. (2004). Unifying models of paired associates and serial learning: Insights from simulating a chaining model. *Neurocomputing*, *58–60*, 739–743.

Caplan, J. B. (2005). Associative isolation: Unifying associative and order paradigms. *Journal of Mathematical Psychology*, *49*(5), 383–402.

Caplan, J. B., Glaholt, M., & McIntosh, A. R. (2006). Linking associative and list memory: Pairs versus triples. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *32*(6), 1244–1265.

Conrad, R. (1965). Order error in immediate recall of sequences. *Journal of Verbal Learning and Verbal Behavior*, *4*, 161–169.

Ebbinghaus, H. (1913). *Memory: A contribution to experimental psychology*. New York, NY: Teachers College, Columbia University (Original work published 1885).

Gabor, D. (1968). Holographic model of temporal recall. *Nature*, *217*, 584.

Gillund, G., & Shiffrin, R. M. (1984). A retrieval model for both recognition and recall. *Psychological Review*, *91*(1), 1–67.

Green, R. L., & Tussing, A. A. (2001). Similarity and associative recognition. *Journal of Learning and Memory*, *45*, 573–584.

Hayman, C. A. G., & Tulving, E. (1989). Contingent dissociation between recognition and fragment completion: The method of triangulation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*(2), 228–240.

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

Hintzman, D. L. (1972). Theoretical notes on testing the independence of associations. *Psychological Review*, *79*(3), 261–264.

Hintzman, D. L. (1984). MINERVA 2: A simulation model of human memory. *Behavior Research Methods, Instruments & Computers*, *16*(2), 96–101.

Hintzman, D. L. (1986). "Schema abstraction" in a multiple-trace memory model. *Psychological Review*, *93*(4), 411–428.

Howard, M. W., Jing, B., Rao, V. A., Provyn, J. P., & Datey, A. V. (2009). Bridging the gap: Transitive associations between items presented in similar temporal contexts. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *35*(2), 391–407.

Howard, M. W., & Kahana, M. J. (1999). Contextual variability and serial position effects in free recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *25*(4), 923–941.

Humphreys, M. S., Bain, J. D., & Pike, R. (1989). Different ways to cue a coherent memory system: A theory for episodic, semantic, and procedural tasks. *Psychological Review*, *96*(2), 208–233.

Jones, J., & Pashler, H. (2007). Is the mind inherently forward looking? Comparing prediction and retrodiction. *Psychonomic Bulletin*, *14*(2), 295–300.

Kahana, M. J. (2002). Associative symmetry and memory theory. *Memory & Cognition*, *30*(6), 823–840.

Kounios, J., Bachman, P., Casasanto, D., Grossman, M., Smith, R. W., & Yang, W. (2003). Novel concepts mediate word retrieval from human episodic associative memory: Evidence from event-related potentials. *Neuroscience Letters*, *345*, 157–160.

Lee, C. L., & Estes, W. K. (1977). Order and position in primary memory for letter strings. *Journal of Verbal Learning and Verbal Behavior*, *16*, 395–418.

Lee, C. L., & Estes, W. K. (1981). Item and order information in short-term memory: Evidence for multilevel perturbation processes. *Journal of Experimental Psychology: Human Learning and Memory*, *7*(3), 149–169.

Lewandowsky, S., & Farrell, S. (2000). A redintegration account of the effects of speech rate, lexicality, and word frequency in immediate serial recall. *Psychological Research*, *63*, 163–173.

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

Longuet-Higgins, H. C. (1968). Holographic model of temporal recall. *Nature*, *217*, 104.

Luce, R. D. (1959). Detection and recognition. In R. D. Luce, R. R. Bush, & E. Galanter (Eds.), *Handbook of mathematical psychology* (pp. 103–189). New York, NY: Wiley.

Madan, C., Glaholt, M., & Caplan, J. B. (2010). The influence of item properties on association-memory. *Journal of Memory and Language*, *63*, 46–63.

Mandler, G., Rabinowitz, J. C., & Simon, R. A. (1981). Coordinate organization: The holistic representation

of word pairs. *The American Journal of Psychology*, *94*(2), 209–222.

Mensink, G., & Raaijmakers, J. G. W. (1988). A model for interference and forgetting. *Psychological Review*, *95*(4), 435–455.

Mensink, G., & Raaijmakers, J. G. W. (1989). A model for contextual fluctuation. *Journal of Mathematical Psychology*, *33*(2), 172–186.

Metcalfe Eich, J. (1982). A composite holographic associative recall model. *Psychological Review*, *89*(6), 627–661.

Murdock, B. B. (1982). A theory for the storage and retrieval of item and associative information. *Psychological Review*, *89*(6), 609–626.

Plate, T. A. (1995). Holographic reduced representations. *IEEE Transactions on Neural Networks*, *6*(3), 623–641.

Primoff, E. (1938). Backward and forward associations as an organizing act in serial and in paired-associate learning. *Journal of Psychology*, *5*, 375–395.

Provyn, J. P., Sliwinski, M. J., & Howard, M. W. (2007). Effects of age on contextually mediated associations in paired associate learning. *Psychology and Aging*, *22*, 846–857.

Raaijmakers, J. G. W. (2003). Spacing and repetition effects in human memory: Applications of the SAM model. *Cognitive Science*, *27*, 431–452.

Rizzuto, D. S., & Kahana, M. J. (2000). Associative symmetry vs. independent associations. *Neurocomputing*, *32–33*, 973–978.

Rizzuto, D. S., & Kahana, M. J. (2001). An autoassociative neural network model of paired-associate learning. *Neural Computation*, *13*, 2075–2092.

Slamecka, N. J. (1976). An analysis of double-function lists. *Memory & Cognition*, *4*(5), 581–585.

Sommer, T., Rose, M., & Büchel, C. (2007). Associative symmetry versus independent associations in the memory for object location associations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *33*(1), 90–106.

Sommer, T., Schoell, E., & Büchel, C. (2008). Associative symmetry of the memory for object-location associations as revealed by the testing effect. *Acta Psychologica*, *128*, 238–248.

Waugh, N. C. (1970). Associative symmetry and recall latencies: A distinction between learning and performance. *Acta Psychologica*, *33*, 326–337.

Westlake, P. R. (1970). The possibilities of neural holographic processes within the brain. *Kybernetik*, *7*(4), 129–153.

Wilson, M. D. (1988). The MRC Psycholinguistic Database: Machine Readable Dictionary, Version 2. *Behavior Research Methods, Instruments and Computers*, *20*(1), 6–11.

Young, R. K. (1959). A comparison of two methods of learning serial associations. *The American Journal of Psychology*, *72*(4), 554–559.

Young, R. K. (1961). Paired-associate learning when the same items occur as stimuli and responses. *Journal of Experimental Psychology*, *61*, 315–318.