

Bridging the gap: Transitive associations between items presented in similar temporal contexts

Marc W. Howard, Bing Jing, Vinayak A. Rao, Jennifer P. Probyn, Aditya V. Datey
Department of Psychology
Syracuse University

In press, *JEP:LMC* September 17, 2008, do not quote without permission

In episodic memory tasks, associations are formed between items presented close together in time. The temporal context model (TCM) hypothesizes that this contiguity effect is a consequence of shared temporal context rather than temporal proximity *per se*. Using double function lists of paired associates (e.g. A-B, B-C) presented in a random order, we examined associations between items that were not presented close together in time but were presented in similar temporal contexts. For instance A and C do not appear in close temporal proximity, but both occur in the context of B. After learning long double-function lists, across-pair associations fell off with distance in the list, as if participants were able to integrate the disparate experiences with the items into a coherent memory structure. Within-pair associations (e.g. A-B) were strongly asymmetric favoring forward transitions; across-pair associations (e.g. A-C) showed no evidence for asymmetry. While this pattern of results presented a stern challenge for a heteroassociative mediated chaining model, TCM provided an excellent fit to the data. These findings raise the possibility that contiguity effects in episodic memory do not reflect direct associations between items but rather a process of binding, encoding and retrieval of a gradually-changing representation of temporal context.

Episodic memory refers to the ability to vividly remember specific events situated in a particular spatiotemporal context (Tulving, 2002, 1983). The question of how the disparate stimuli that constitute an episode are bound together into an episodic memory for that event is a fundamental question in the study of human memory. Temporal contiguity effects have been extensively studied in the lab using the free recall task, in which subjects recall a list of words in the order they come to mind. The conditional response probability as a function of lag (lag-CRP) developed by Kahana and colleagues (e.g. Howard & Kahana, 1999; Kahana, 1996; Kahana, Howard, Zaromb, & Wingfield, 2002; Klein, Addis, & Kahana, 2005) estimates the probability of making a recall transition from a just-recalled words to other words in the list as a function of their temporal distance in the list. For instance if the 10th word in a list of 20 has just been recalled in a delayed free recall experiment, then, all other things being equal, the next word the subject recalls is more likely to

come from a nearby serial position (e.g. serial position 9 or 11) than a remote serial position (e.g. 5 or 15). The lag-CRP falls off gradually with distance, or lag, over several list positions.¹ In addition, lag-CRP functions in free recall show an asymmetry such that transitions forward in the list (e.g. a transition from item 10 to item 11) are more likely than transitions backward in the list (e.g. a transition from item 10 to item 9). Graded contiguity effects have been observed in a wide variety of conditions in free recall and a wide variety of episodic memory tasks (for a review see Kahana, Howard, & Polyn, 2008), suggesting that they reflect a very general property of episodic memory.

Explanations of the contiguity effect can be classified into two broad classes which we refer to here as heteroassociative and retrieved context accounts. Heteroassociative distributed memory models share the assumption that item representations are directly associated to one another (e.g., Chappell & Humphreys, 1994; Murdock, 1982; Humphreys, Bain, & Pike, 1989; Raaijmakers & Shiffrin, 1980; Lewandowsky & Murdock, 1989). An important question in formulating a heteroassociative model is how the representations of items presented at different times are simultaneously available so that they can be bound together. Classic notions of a limited capacity short-term store serve just this function (Atkin-

Supported by MH069938-01 from the NIH. The authors acknowledge the generous support of Michael Hasselmo and Chantal Stern in developing pilot versions of this task when the senior author was at Boston University. Mike Kahana and Mark Steyvers provided helpful comments on earlier versions of the paper. Thanks to Diana Hobbins, Radha Modi, and Madhura Phadke for helping collect and score data. R code to run the version of TCM described here can be found at <http://memory.syr.edu/simulationsw.html>. Udaya Jagadisan and Karthik Shankar worked on an earlier version of the R code.

¹Note that the existence of a contiguity effect does not preclude the existence of primacy or recency effects. If a primacy or recency effect also affects which word is recalled, then at extreme lags the it could overwhelm the contiguity effect, leading to a non-monotonicity in the lag-CRP at extreme lags (Farrell & Lewandowsky, in press; Howard, Sederberg, & Kahana, submitted).

son & Shiffrin, 1968; Jensen & Lisman, 2005; Raaijmakers & Shiffrin, 1980; Kahana, 1996). In the search of associative memory (SAM, Raaijmakers & Shiffrin, 1980) model the strength of the association between items in long-term memory is incremented when they are coactive in short-term memory, enabling SAM to describe contiguity effects in free recall (Kahana, 1996; Sirotin, Kimball, & Kahana, 2005).

In the temporal context model (TCM, Howard & Kahana, 2002; Howard, Fotedar, Datey, & Hasselmo, 2005; Sederberg, Howard, & Kahana, in press), temporal contiguity effects are not the consequence of direct heteroassociative connections, but rather a consequence of items' ability to be bound to and recover a gradually-changing representation of temporal context. Put another way, rather than relying on item-to-item associations, as heteroassociative models do, TCM produces temporal contiguity effects as a consequence of item-to-context-to-item associations. Suppose a pair of words ABSENCE—HOLLOW is presented. During study, the encoding context for each pair is composed of elements retrieved by each member of the pair. During study, the item representation for HOLLOW is associated with the encoding context of the pair. When ABSENCE is repeated as a cue, part of the encoding context of the pair is recovered. Because the context retrieved by ABSENCE at test overlaps with the encoding context of HOLLOW, the result is a behaviorally observed association between ABSENCE and HOLLOW.

Heteroassociative and contextual recovery accounts of episodic association cannot be distinguished by their ability to describe the form of temporal contiguity effects in standard episodic memory tasks.² In order to differentiate the predictions of these accounts of contiguity effects in episodic memory, we need to go beyond simple temporal proximity. According to TCM items presented close together in time become associated to each other not because of their temporal proximity *per se*, but because of the similarity between the encoding context of one item and the context recovered by the other. As long as there is similarity between these states of context, items should become associated to each other regardless of whether they were actually presented in close temporal proximity.

In double-function lists of paired associates (Primoff, 1938, see also Figure 1), items from different pairs do not co-occur, but there is an overlap in the temporal contexts in which they are studied. A double-function list might include the pairs ABSENCE—HOLLOW and HOLLOW—PUPIL. Note that the order in which the pairs are presented to the participant does not correspond to the structure of the list (Figure 1a). To the extent the subject learns the across-pair structure (Figure 1b), it must be inferred, or integrated, from temporally distinct events. In learning these pairs, ABSENCE and PUPIL are never presented as part of the same pair. Moreover, although a pair including ABSENCE and a pair including PUPIL may occasionally by chance be presented in close temporal proximity, HOLLOW is no more likely to occur close in time to ABSENCE than any other item from the list that is not part of a pair involving ABSENCE. However, ABSENCE and PUPIL are both presented in the context of HOLLOW. The contextual retrieval hypothesis predicts that ABSENCE and



Figure 1. Schematic of presentation of double-function pairs and the corresponding linked-list. Colors are shown for explanatory purposes and were not shown to the participants. Words were presented one at a time on the screen. Participants were able to distinguish members of the same pair by means of a longer delay between pairs than within-pair. On the left are a sample of pairs that the subject might be shown in order. On the right is the linked list that would be formed by stringing together pairs with overlapping items. The numbers on the right indicate the *linked-list lag* associated with several recall transitions. For instance, given PUPIL as a cue, a transition to the correct response, RIVER would have a linked-list lag of +1. Given PUPIL as a cue, a backward intrusion, HOLLOW in this example, would have a linked-list lag of -1. Remote intrusions, for instance ABSENCE or DARLING, would be associated with linked-list lags with absolute value greater than one.

PUPIL should become associated due to this commonality of the contexts in which they were experienced. In this paper, we present the results of a study in which participants learn a long list of double-function pairs. We compare the associative strength between the items that compose the pairs as a function of their “distance” in the double-function list by examining intrusions and transition probabilities from a final free recall session. The existence and form of these across-item associations, along with quantitative modeling, support a retrieved context account of these findings.

Associations that bridge across double-function pairs have been observed previously (Popper, 1959; Slamecka, 1976; Bunsey & Eichenbaum, 1996). However, the existence of associations that bridge across pairs in a double function list is not itself sufficient to rule out a heteroassociative account of temporal contiguity effects. Associations among items that were never presented together could be reconciled with heteroassociative models using mediated chaining (Slamecka, 1976). One can explain a behaviorally-observed association between ABSENCE and PUPIL if the item representation for HOLLOW is activated as an intermediate step. Mediated chaining could result in a bridging association at retrieval if the subject surreptitiously retrieves a link in the chain but, for whatever reason, withholds the response waiting until the next link in the chain is traversed. Heteroassociative medi-

² It is possible to distinguish fixed-capacity buffer models from TCM on the basis of their ability to describe contiguity effects across a delay, as in continuous-distractor free recall. However, this discrepancy is not a consequence of the basic assumptions about heteroassociation *versus* contextual recovery. For example, a version of SAM in which items did not drop out of the short-term buffer in an all-or-none fashion but rather decayed gradually would almost certainly be capable of describing at least the existence of the temporally-defined associations observed in continuous-distractor free recall.

ated chaining accounts predict a strong dependence between bridging associations and the links in the chain that must be traversed to bridge between the items.

Heteroassociative models may either predict associative symmetry between members of a pair, such that after learning a pair $A - B$, the backward association $B - A$ is just as strong as $A - B$, or asymmetry within pair, presumably reflecting stronger $A - B$ associations than $B - A$ associations. However, if a heteroassociative model has asymmetric associations, then it would tend to predict that the asymmetry observed within a single pair is amplified for remote associations as multiple asymmetric links must be traversed. To anticipate the results of the present experiment, we observed a strong asymmetry in adjacent associations between members of a pair, coupled with a much smaller asymmetry between items from different pairs. As a matter of fact, we found no positive evidence for asymmetry for remote transitions. If the heteroassociations are symmetric, then the mediated chaining model has difficulty accounting for the dramatic within-pair asymmetry we observed. If the heteroassociations are asymmetric, as one would expect from the asymmetry of the adjacent associations, then the model has difficulty accounting for the much smaller, not significantly different from zero, asymmetry across-pairs.³ Explicit modeling of a heteroassociative mediated chaining model will illustrate the difficulty that mediated chaining models have in accounting for this pattern of findings.

Utilizing Caplan's (2005) isolation principle to adapt a model of list memory to pairs (see also Caplan, Glaholt, & McIntosh, 2006), we simulated the pairs as a series of completely isolated two-item serial lists. TCM provides an excellent account of the data, including the divergence the strong asymmetry observed in the adjacent transitions and the much-reduced asymmetry observed in the remote transitions. Asymmetry in the adjacent transitions is not amplified in the remote transitions because of the way associations work in TCM. Associations in TCM are not formed directly between items but are mediated by the effect items have on context, which then cues other items for recall. More specifically, there are two components of retrieved context which have different properties; one of these components is essential for remote associations among items (Howard et al., 2005). The description of the model and its properties are explained in more detail in the Modeling section after presentation of the experiment.

Experiment

Subjects were presented with a list of paired associates containing both double-function and single-function pairs (see Figure 1). Single-function pairs, e.g. POLICE-WINDOW, are just standard, non-overlapping, paired associates. The double-function pairs give rise to a linked-list (Figure 1, right). *Linked-list lag* provides a convenient means to describe the distance between two items in the linked-list. For instance, consider the word PUPIL in the linked-list in Figure 1. The correct response to the stimulus PUPIL is RIVER, which is one step forward in the chain at a linked-list lag

of +1. If PUPIL was given as a probe and the subject responded HOLLOW, this would be a backward intrusion associated with a linked-list lag of -1. Items with a linked-list lag of ± 1 were actually presented together as part of the same pair. Larger absolute values of linked-list lag separate double-function items that were not presented together. For instance, the linked-list lag between PUPIL and DARLING is +2 and the linked-list lag from PUPIL to ABSENCE is -2. Linked-list lag allows us to compare the probability of intrusions or recall transitions as a function of their recall direction by comparing linked-list lags with the same absolute value (e.g. comparing +1 to -1). We can also look for the presence of bridging associations by looking for an effect of the absolute value of linked-list lag on associations between pairs that were not presented together (absolute value of linked-list lag ≥ 2).

We will examine the associations formed between double-function items, both within and across pair, in two ways. One is to observe the intrusions subjects make during paired associates testing. Subjects also performed surprise final free recall (FFR) of all items at the end of the experimental session. Transitions in this final free recall period presumably reflect the organization of information in memory in the relative absence of strategies to edit out intrusions during cued recall testing. The CRP of recall transitions as a function of *linked-list lag* during FFR provides a means to estimate the strength of the associations between those items. Insofar as this measure reflects associations between items that were not presented close together in time, this can be seen as a measure of the extent to which participants' memories were able to extract the linked-list structure from the collection of double function pairs.

Methods

Participants. Two-hundred-sixteen participants participated for course credit in an introductory psychology class at Syracuse University.

Materials. Study words were chosen from the noun subset of the Toronto word pool (Friendly, Franklin, Hoffman, & Rubin, 1982). Study lists were composed of 35 double-function pairs consisting of 36 distinct words and 8 single-function pairs consisting of 16 distinct words. The double-function pairs were formed by choosing 36 words randomly without replacement from the pool for each subject. The first two words were assigned to the first pair. The second and third word were assigned to the second pair and so on.

Procedure. Participants learned the pairs over four study-test trials. On each study-test trial each pair was presented for study three times and then, after a delay, each pair was tested once. This makes a total of twelve presentations of each pair over the course of the experiment.

³ Of course almost any model could trivially generate symmetric across-pair associations by making the magnitude of the remote associations arbitrarily small, thus causing the ratio of the two to go to one if there is a non-zero baseline of recall.

The order in which pairs were presented in each study-test trial was randomized separately for each participant subject to the constraint that double-function pairs from successive linked-list positions were never presented sequentially. This was done to ensure that any associations across double-function pairs “adjacent” in the linked-list could not be due to temporal contiguity. Words were presented one at a time in an uppercase font for 1000 ms. Words were also presented auditorally in a female voice. Presentations of items within a pair were separated by a delay of 100 ms, whereas pairs were separated by a delay of 1800 ms. Prior to each test, subjects performed a true-false arithmetic distractor for 30 s.

After the completion of the distractor task, each pair was tested one at a time. The order of tests was randomized, again subject to the constraint that adjacent pairs in the linked-list were never tested successively. On each test, the stimulus element of each pair was presented visually and auditorally as before. Subjects were instructed to recall the word that followed the probe in a study pair. Participants were given 5 s to respond verbally to each probe.

Following the last study-test trial, participants were administered a surprise FFR test. The delay between the completion of the last test item and the beginning of the FFR test was typically a few minutes, which consisted of time to notify the experimenter that the study-test cycles were complete, time for the experimenter to set up the FFR test and provide the participant with instructions for the FFR period. The length of this delay was large relative to the spacing between the presentation of items—and each pair was presented twelve times over several tens of minutes—so that recency effects in the FFR test would be expected to be negligible. Participants were given five minutes to recall as many study words as possible without regard to order. An experimenter encouraged participants to continue attempting to recall for the entire five minute period.

Analyses.

We measured the proportion of remote (absolute value of linked-list lag ≥ 2) double-function intrusions in response to a double-function paired-associate probe item that came from each linked-list lag. Because pilot testing demonstrated edge effects in recall of the pairs such that the first and last pairs in the linked-list were better recalled (presumably due to reduced associative interference from other double-function items), we omitted all intrusions in which the probe word or the remote intrusion was from the first or last pair in the linked-list.

In order to evaluate the associative structure of the double-function lists revealed by FFR transitions, we calculated an analogue of the lag-CRP developed to measure temporally-defined associations (Kahana, 1996; Howard, Addis, Jing, & Kahana, 2007). Given a pair of successively-recalled double-function items, this measure calculates the number of transitions at each linked-list lag divided by the number of correct recall transitions that would have been possible at that linked-list lag given the actual items that were recalled (see Howard et al., 2007, for details). We also re-ran this analysis restricting our attention to pairs of successively-recalled

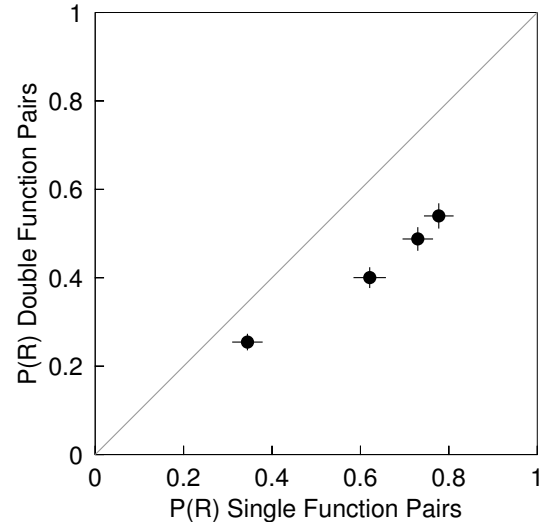


Figure 2. Probability of recall for double function pairs is lower than for single function pairs. The four points are probability of a correct cue recall after each of four trials. Trial 1 is in the lower left. The other trials are in sequence with monotonically increasing probabilities of recall. Error bars are 95% confidence intervals.

double-function items that were not part of the same pair (absolute value of linked-list lag ≥ 2). We ignored all transitions and potential transitions where either of the items was a member of the very first or very last pairs in the double-function list to avoid any possible edge effects.

Results and Discussion

We examine results from the paired associates testing before moving our attention to the final free recall data.

Paired associates testing.

Examination of probability correct on the paired associate tests revealed that single function pairs were learned better than double function pairs. Figure 2 shows the probability of a correct recall for double-function pairs as a function of the probability of recall for single-function pairs for each of the four study-test cycles. The left-most point reflects the performance on the test after the first presentation of the list. If there were no difference across pair types, the points would lie on the diagonal. In fact, the probability of recall on the first trial was significantly greater for single function pairs $.33 \pm .03$ than double function pairs, $.24 \pm .02$, $t(96) = 4.84$, $p < .001$. Pairwise comparisons for the other learning trials were also highly significant. Moreover, the difference between single and double-function pairs grew with learning, at least over the first couple of trials. A repeated measures ANOVA showed significant main effects of pair type (single vs double), $F(1, 672) = 391$, $MSe = 6.83$, $p < .001$, trial number, $F(3, 672) = 240$, $MSe = 4.20$, $p < .001$, as well as a significant interaction of pair type and trial $F(3, 672) = 11.5$, $MSe = .20$, $p < .001$.

Most authors attribute the disadvantage for double function pairs relative to single function pairs to associative interference from competing responses, especially the backward response (Primoff, 1938; Young, 1961). That is, if a subject is given HOLLOW as a probe, the correct response is PUPIL. However, the backward association between HOLLOW and ABSENCE will create a tendency to make a backward intrusion (i.e. a response at linked-list lag -1) which could interfere with the correct response. Because single function pairs are not subject to interference from a backward association, this should create an advantage for performance on single function pairs compared to double function pairs. Similarly, remote bridging associations compete with correct recalls for double-function but not single-function pairs.

Table 1 shows the proportion of responses of various types to the double-function probes. As can be seen from inspection of the table, there was a relatively large proportion of backward intrusions that persisted across trials. Nonetheless, the proportion of correct (forward) responses was larger at each trial, indicating an asymmetry. At each of the four trials in Table 1, there is a highly significant pairwise difference between the probability of a correct response and the probability of a backward intrusion (paired $t(215)$ from 9.87 to 19.7). To further quantify this apparent asymmetry, we ran a repeated measures ANOVA on forward and backward response types (the first two columns of Table 1) with trial as a factor. We found highly significant main effects of both response direction (forward vs backward), $F(1, 1505) = 1573.7$, $MSe = 25.9$, $p < .001$, and trial $F(3, 1505) = 100.1$, $MSe = 1.64$, $p < .001$, as well as a highly significant interaction of trial and direction $F(3, 1505) = 106.93$, $MSe = 1.76$, $p < .001$.⁴ We conclude that there was a robust asymmetry between responses of linked-list lags $+1$ and -1 during paired associate testing.

Table 1 also appears to show relatively large numbers of remote linked-list intrusions. To estimate the size of the asymmetry in these remote bridging associations we calculated the proportion of these remote intrusions that were observed at each linked-list lag. If the remote intrusions were a consequence of bridging associations, then we would expect that linked-list lag would have a significant effect on the probability of a remote intrusion probability. Figure 3a shows the result of these analysis collapsed over all test trials. First we note that there was a significant effect of linked-list lag on remote intrusion proportion. We conducted a repeated-measures ANOVA with absolute value of linked-list lag (2 through 7 inclusive) and recall direction (forward vs backward) as factors. We found a main effect of linked-list lag, $F(5, 2567) = 24.8$, $MSe = 0.16$, $p < .001$, providing strong evidence for the presence of remote bridging associations during double-function learning. There was no effect of direction, $F(1, 2567) = 1.56$, $MSe = 0.010$, $p > .2$ nor an interaction of intrusion direction and linked-list lag $F(5, 2567) = 1.41$, $MSe = 0.009$, $p > .2$. Despite the strong asymmetry in adjacent responses (linked-list lags ± 1 , see Table 1) there was no positive evidence for a corresponding asymmetry in the remote intrusions. The apparent discrepancy between the asymmetry observed in adjacent and

remote recalls here could be due to the fact that backward adjacent responses are errors whereas forward adjacent responses are not. Because of this difference in the instructions, we postpone directly comparing adjacent and remote asymmetries until we discuss the FFR data.

FFR transitions.

Figure 3b shows the results of the linked-list CRP calculated from the FFR transitions. The first notable feature is a strong asymmetry between forward and backward adjacent linked-list recall transitions. The linked-list CRP at a linked-list lag of $+1$ was significantly greater than that observed at linked-list lag -1 , paired-sample $t(215) = 6.39$, $p < .001$. The asymmetry was not only significant, but it was large when examined in terms of the size of the difference between the means and also examined as an effect size. The difference between the linked-list CRPs at linked-list lags ± 1 was .67 of the value at linked list lag -1 . The difference between the linked-list CRPs at linked-list lags ± 1 was almost equal to one standard deviation of the linked-list CRP at lag -1 (.97). As with the responses to paired-associate testing, FFR transitions showed a strong asymmetry among items that were presented as part of the same pair, i.e. items at adjacent linked-list positions with lag ± 1 .

Figure 3c shows a version of the linked-list CRP that only considers remote transitions that bridge across double-function pairs. That is, only remote transitions were considered, thus eliminating adjacent transitions from both the numerator and denominator of the measure calculated. The curve appears peaked in the center, with transitions to linked-list lags of ± 2 more likely than more remote linked-list lags and a gradual decline across several linked-list lags. To confirm that the linked-list CRP was elevated for more than just the lags ± 2 , we compared the linked-list CRP for small remote lags to those from more remote lags. To do so, we calculated the average over more remote lags (absolute values from 6 to 12, inclusive). The mean value of the linked-list CRP for these more remote linked-list lags was $0.027 \pm .001$. We then compared this value for more remote linked-list lags to the linked-list CRP, averaged across forward and backward transitions, for specific values of linked-list lag. We found a significant difference for linked-list lags ± 2 (paired Wilcoxon $V = 17382$, $p < .001$) and for linked-list lags ± 3 (paired Wilcoxon $V = 11580$, $p < .05$). Linked list lags ± 2

⁴ One might want to take the parametric statistics from these ANOVAs with a grain of salt due to the fact that the response values are not independent of each other. If a response is correct, it is necessarily not a backward intrusion. Similar concerns about a lack of independence can be raised for the intrusion analyses and final free recall CRP analyses that follow. For instance, if a response is an intrusion at linked-list lag -2 , it cannot be an intrusion at linked-list lag -3 . As a consequence, the proportion of intrusions at linked-list lag -2 is not independent of the proportion of intrusions at linked-list lag -3 . However, the linked-list used in the present study is quite long, so that the dependence between recall at any two lags is likely to be relatively weak. Moreover, the sample size is quite large so that any correction due to dependence is unlikely to affect our conclusions.

Table 1

Types of response to double-function probes for different trials. The column labeled $P(R)$ gives the probability correct. The column labeled “Bk” gives the probability of a backward intrusion (backward intrusions can also be described as intrusions with a linked-list lag of -1). The column labeled “Rm” gives the probability of a remote intrusion (remote intrusions are associated with linked-list lags with absolute value ≥ 2). The column labeled “SF” gives the probability of an intrusion that was part of the single-function pairs. The column labeled “XLI” gives the probability of an extra-list intrusion—a word that was not presented for study. The column labeled “NR” gives the probability that the subject did not make a response. The numbers in parentheses are the standard error multiplied by the critical t .

Trial	$P(R)$	Bk	Rm	SF	XLI	NR
1	.25 ± .02	.17 ± .01	.20 ± .02	.03 ± .01	.08 ± .02	.27 ± .02
2	.40 ± .02	.20 ± .02	.15 ± .02	.03 ± .01	.06 ± .01	.16 ± .02
3	.49 ± .03	.18 ± .02	.12 ± .02	.02 ± .01	.06 ± .02	.12 ± .02
4	.54 ± .03	.16 ± .02	.11 ± .01	.02 ± .01	.05 ± .01	.11 ± .02

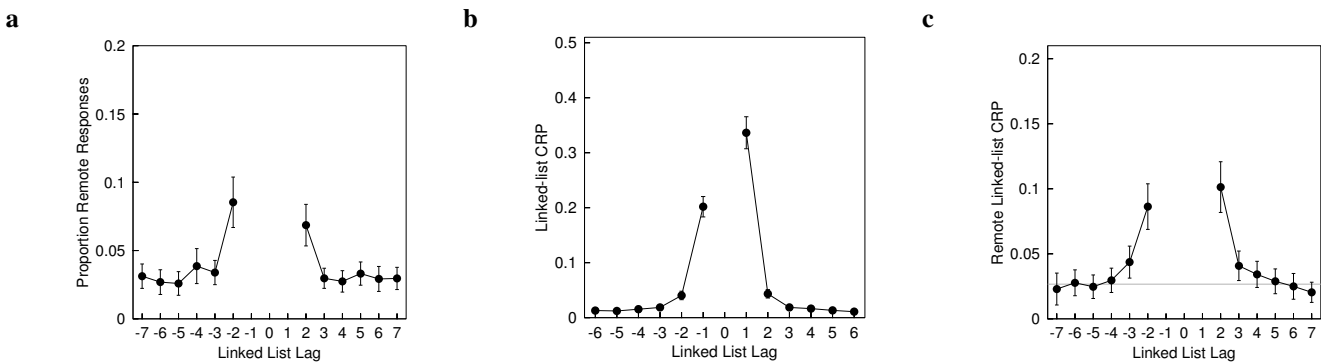


Figure 3. Remote bridging associations are symmetric. a. Remote intrusion analysis. Given that an intrusion from the double-function pairs is given as a response to a double-function probe, and the linked-list lag of the intrusion has an absolute value ≥ 2 , this graph shows proportion of such intrusions come from each value of linked-list lag. b. Final free recall (FFR) conditional response probability (CRP). Probability of an FFR transition is shown as a function of the linked-list lag (linked-list CRP) associated with the transition. c. Remote FFR linked-list CRP. Same analysis as in b, but restricted to remote FFR transitions. The error bars in all three panels reflect 95% confidence intervals.

and ± 3 also differed from each other (paired Wilcoxon $V = 8578$, $p < .001$). This quite conservative analysis demonstrates that remote bridging associations extend at least two pairs. That is, given B as a probe, there are not only significant B-D associations, but B-E associations as well. Although the overall number of responses are relatively small, the effect of linked-list lag on these remote transitions is relatively dramatic. Transitions of linked-list lag ± 2 are more than three times more likely than more remote transitions.

The remote transitions in FFR shown in Figure 3c do not show evidence for an asymmetry. A repeated-measures ANOVA on remote linked-list CRP taken for linked-list lags $+2$ to $+7$ and -7 to -2 with the absolute value of linked-list lag and direction of recall (forward vs backward) as factors found a significant effect of linked-list lag, $F(5, 2533) = 39.6$, $MSe = 0.31$, $p < .001$, but showed neither an effect of recall direction, $F(1, 2533) = 0.54$, $MSe = .0043$, nor an interaction between recall direction and lag, $F(5, 2533) = .66$, $MSe = .0052$. Three subjects failed to make a remote FFR transition and are excluded from these analyses. A paired-sample t -test comparing lag $+2$ to lag -2 failed to find a significant effect, $t(212) = 1.08$.

Here we directly compared the degree of asymmetry from adjacent transitions to the degree of asymmetry for remote transitions. To compare the difference between adjacent and remote transitions on a comparable footing, we divided the difference between forward and backward CRP values by the sum of the forward and backward CRP values. That is for the adjacent comparison, the asymmetry index was $[P(+1) - P(-1)]/[P(+1) + P(-1)]$. For the remote transitions, the index was $[P(+2) - P(-2)]/[P(+2) + P(-2)]$. To derive an estimate of the variability of this estimate we performed a bootstrap analysis. In each shuffle of the bootstrap, we sampled 216 subjects with replacement from our subject pool, generated the FFR linked-list CRP from this sampling of subjects and then recalculated the adjacent and remote asymmetry indices. We performed 1000 shuffles and report the standard deviation of this distribution as an empirical estimate of the standard error. Table 2 reports the results of these analyses. The much larger asymmetry index for the adjacent transitions ($.25 \pm .04$) than for the remote transitions ($.03 \pm .07$) provides strong positive evidence that the adjacent asymmetry index was much greater than the remote asymmetry index.

Table 2

Observed and model-generated asymmetry indices. Each asymmetry index is calculated by $(F - B)/(F + B)$, where F is the probability of a forward FFR transition and B is the probability of a backward FFR transition. The row labeled adjacent gives the index comparing +1 and -1 transitions. The row labeled remote gives the index comparing +2 and -2 transitions. The number in parentheses is an empirical estimate of the standard error of the mean (see text for details). Med. chain.=mediated chaining model. TCM=temporal context model. For each model, the predicted values and a z -score relative to the actual value is reported.

	Observed (SE)	Med. chain.	z	TCM
Adjacent	.25 (.04)	.19	-1.75	.24
Remote	.03 (.07)	.25	3.42	.12

These findings present a challenge to heteroassociative models based on mediated chaining. It remains possible, however, that the reduced asymmetry in the remote transitions is the result of a floor effect or from an asymptotic level of recall not subject to the assumptions of mediated chaining. We further illustrate the degree to which these findings challenge the heteroassociative framework by explicitly simulating a reasonable heteroassociative mediated chaining model.

Modeling

In this section we compare the ability of two models to describe the central results of this experiment. We restricted our attention to the CRP as a function of linked-list lag in final free recall. This frees us from having to explicitly model the recognition processes that presumably edit out intrusions during cued recall, which would itself be a formidable challenge. One model is a formulation of TCM. The other is a heteroassociative model that implements mediated chaining. The heteroassociative model is roughly inspired by the search of associative memory model (SAM, Raaijmakers & Shiffrin, 1980; Mensink & Raaijmakers, 1988), but it is primarily included here as a way to make quantitative our assumptions about mediated chaining in the empirical discussion above rather than as a serious candidate for how to model data from linked lists of paired associates. We start by describing the heteroassociative mediated chaining model.

A heteroassociative mediated chaining model

In order to formulate a mediated chaining account of the remote associations in Figure 3, the first problem that must be solved is how to traverse the links in the chain without recalling the items in between. For instance, starting from item A, how do we make it to C without recalling B beforehand? The distinction between sampling and recovery in SAM offers a suggestion if we are willing to suppose that items that are sampled, but not recovered, can be used as cues for subsequent retrieval attempts. Because B need not be recovered,

it can be omitted from recall but still serve as a cue for recall of C. We also assume, for simplicity, that the search cannot double back on itself. That is, an item cannot be used as a cue if it previously resulted in an unsuccessful recovery attempt.

Assume that we have a model in which there is a large chain of paired associates and we start from item i in the chain. We assume that the probability of sampling the item at position $i + 1$ is f and the probability of sampling the item at position $i - 1$ is $1 - f$. We assume for simplicity that the probability of successfully recovering the sampled item is r . Note that this does not necessarily imply that the probability of recovery is the same for forward and backward items. If we had assumed different recovery probabilities, this could be folded into the f parameter in the expression below. If the item is recovered it is said to be recalled and the process stops. If it is not recovered, the next item sampled comes one step forward in the linked-list with probability f and one step backward in the linked-list with probability $1 - f$. If the first failed retrieval attempt was in the forward direction, then a backward sample on the second retrieval attempt would result in resampling an item that was a cue for a failed recovery and recall stops (see Figure 4). Because the CRP analysis only includes successful retrievals, these events can be ignored in fitting the data and we find that the relative probability of eventually sampling and recovering an item at lag $j > 0$ from the initial item i is given by:

$$P_{SR}(i + j) \propto \frac{r}{1 - r} [(1 - r)f]^j \quad (1)$$

$$P_{SR}(i - j) \propto \frac{r}{1 - r} [(1 - r)(1 - f)]^j \quad (2)$$

The r comes from one recovery success and the factor of $(1 - r)^{j-1}$ comes from $j - 1$ recovery failures. Notice that the asymmetry here is expressed as a factor of f^j in the case of forward traversals of the chain compared to a factor of $(1 - f)^j$ in the case of backward traversals of the chain, such that the ratio increases as a function of j . In our simulations, we started from an item in the middle of the list, $i = 18$ and renormalized the probabilities of sampling and recovering an item by dividing by the sum across potential recalls to get the final value of P_{SR} (Howard & Kahana, 1999; Raaijmakers & Shiffrin, 1980).

To add an asymptotic level of recall aside from the sampling and recovery process, we also added a guessing probability g . If a guess was adopted, this was evenly distributed across all possible recallable items.

$$P_R(i + j) = g/N + (1 - g)P_{SR}(i + j) \quad (3)$$

$$P_R(i - j) = g/N + (1 - g)P_{SR}(i - j) \quad (4)$$

where N is the number of candidate items for recall (35 in our simulations). Note that the probability of a "guess" here also includes retrieval of semantic associates, or retrievals based on fixed list context as a cue.

This simple mediated chaining model has three parameters. The asymmetry at each of the links in the chain is controlled by f . Values of f above $\frac{1}{2}$ give rise to a forward asymmetry; values below $\frac{1}{2}$ give rise to stronger backward

associations than forward associations. The probability of stopping at any given linked-list lag is given by r . Finally, the guessing parameter controlling the asymptotic level of recall is given by g .

We implemented the simple heteroassociative mediated chaining model described above in R. We found the best-fitting parameters, as measured by minimizing the chi-squared between the model and the FFR CRP over linked-list lags -6 to $+6$, using a modified quasi-Newton method (method L-BFGS-B for R function `optim`) with bounds of zero and one on all three of the parameters.

The best-fitting parameters of this model were $f = .60$, indicating a forward asymmetry, $r = .43$, which was relatively low to allow substantial amounts of remote mediated associations and $g = .44$. This solution provided a poor fit to the empirical data, with a chi-squared of 20.05 with eight degrees of freedom (11 independent data points minus three free parameters) $p < .02$. Figure 4b illustrates how the model failed. The error bars in Figure 4b are the standard error times the critical t . The heteroassociative mediated chaining model overestimated the probability of an adjacent forward transition and underestimated the probability of an adjacent backward transition. Conversely, at remote linked list lags, e.g. ± 2 , the pattern was reversed. The heteroassociative mediated chaining model systematically underestimated the forward asymmetry in the adjacent transitions and overestimated the asymmetry in the remote transitions (see also Table 2).

A key goal in developing this model was to establish whether the substantial asymmetry observed in the adjacent associations and the not-different-from zero adjacent asymmetry observed in the remote associations were consistent with a mediated chaining account of the task. To determine whether this mediated chaining model was consistent with the data, we took the probabilities of a transition at each lag and then generated sets of 1000 responses using these probabilities. We analyzed these distributions of responses to calculate adjacent and remote asymmetry factors, as we did with the data above. The number of responses used here can be argued to be quite conservative—there were more than 4,000 responses that contributed to the analyses in Figure 3b. The use of many fewer responses, although it neglects variability due to participants, should overestimate the variability assigned to the model and thus overestimate the probability of observations far from the mode. The distribution of adjacent asymmetry factors was approximately symmetric with a mean of .188 and a standard deviation of .04 across 1000 samples. The adjacent asymmetry factor for the simulated data was less than that observed in the data for 937/1000 samples.

For the remote asymmetry factor, the model produced a distribution with mean .496 and standard deviation .05. This value was greater than that observed in the experimental data for 981/1000 samples. These results indicate that the pattern of observed results, and in particular the very small asymmetry observed for the remote asymmetry factor is very unlikely to result from this mediated chaining model.

While we cannot evaluate all possible mediated chaining

models, the data are sufficient to overwhelmingly reject this particular one. This illustrates the challenge that our findings present to heteroassociative models. In the general discussion we discuss some possible approaches whereby it may be possible to reconcile our findings with a heteroassociative framework, if not mediated chaining.

Temporal context model

According to TCM (Howard & Kahana, 2002; Howard et al., 2005; Howard, Kahana, & Wingfield, 2006; Rao & Howard, 2008; Sederberg et al., in press), the cue for episodic recall at all times is the current state of a temporal context vector \mathbf{t} . Context cues lexical representations of items, described as vectors \mathbf{f} via a matrix \mathbf{M} . TCM describes a set of rules for how context is updated from moment to moment, encoded and retrieved to enable this model to account for recency and contiguity, basic properties of episodic recall tasks. The following section provides a tutorial introduction to TCM. The only important difference between the implementation described here and previously-published versions of the model is the use of a different rule for contextual retrieval—readers already familiar with TCM may wish to skip to the subsection labeled “contextual retrieval”.

Temporal context changes gradually. In TCM, the cue for episodic recall the state of temporal context at time step i , \mathbf{t}_i , is formed from the previous state of context \mathbf{t}_{i-1} and an input pattern \mathbf{t}_i^N according to

$$\mathbf{t}_i = \rho \mathbf{t}_{i-1} + \beta \mathbf{t}_i^N \quad (5)$$

The input pattern \mathbf{t}_i^N is caused by the current item presented. Considerable attention will be paid to describing this input pattern later. For now we note that Eq. 5 results in a gradually-changing state of temporal context. The rate of change, ρ is constrained to be between zero and one. Equation 5 means that \mathbf{t}_i resembles to some extent \mathbf{t}_{i-1} . If the input patterns are uncorrelated, then in addition \mathbf{t}_i will more closely resemble \mathbf{t}_{i-1} than it does \mathbf{t}_{i-2} .

More formally, the previous state of context \mathbf{t}_{i-1} persists, multiplied by a scalar ρ , chosen such that the norm of \mathbf{t}_i is always unity.⁵ This is added to an input pattern \mathbf{t}_i^N weighted by a parameter $\beta = 1 - \rho$ that, together with the similarity of the input pattern to the previous state of context, controls the rate at which \mathbf{t}_i varies from time-step to time-step. The property that temporal context changes gradually is indispensable for the ability of TCM to explain recency and temporal contiguity effects.

⁵ In the simulations reported in this paper, we normalize vectors using the \mathcal{L}^1 norm rather than the more standard Euclidean \mathcal{L}^2 norm. The \mathcal{L}^1 norm of a vector \mathbf{v} is the sum of the absolute values of its components $|\mathbf{v}| = \sum_i |v_i|$. The reasons for adopting this have to do with computational simplicity and have no bearing on the conclusions reached in this paper. This change means that rather than describing points on a hypersphere, the vectors \mathbf{t} describe points on a simplex. Also, we assume that the input patterns \mathbf{t}^N are always \mathcal{L}^1 normalized.

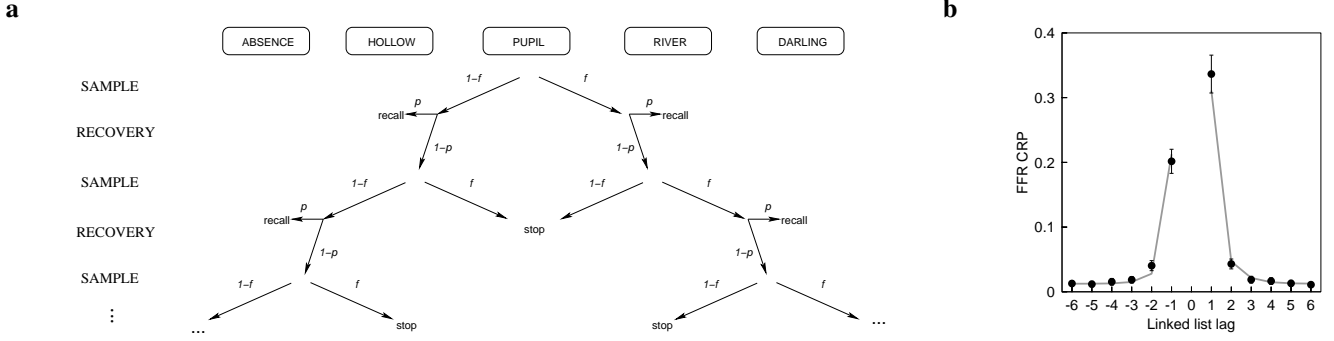


Figure 4. A heteroassociative mediated chaining model. a. The model proceeds in sampling and retrieval phases. Starting with HOLLOW as the cue, the item one step forward in the chain is sampled with probability f and the item one step back is sampled with probability $1 - f$. To enable remote associations, the sampled item may either be recovered, with probability p or fail to be recovered with probability $1 - p$, in which case another item is sampled. If an item that has already been the cue for a failed recovery attempt is re-sampled, the recall process stops. In addition, a guessing process allows an asymptotic level of recall to be generated (not shown). **b.** The symbols and error bars replot the data in Figure 4b. The solid grey line shows the best-fitting values generated by the mediated chaining model illustrated in a. The mediated chaining model underestimates the adjacent asymmetry and overestimates the remote asymmetry (see also Table 2).

Contextual encoding.

In the development that follows, we will adopt the Dirac notation by denoting column vectors as a “ket,” e.g., $|\mathbf{v}\rangle$ and row vectors as a “bra,” $\langle \mathbf{v}|$. The primary advantage of this notation is that it becomes intuitive to visually identify the inner product as a “bracket” $\langle \mathbf{v}|\mathbf{u}\rangle$ and the outer product as a “ketbra”: $|\mathbf{u}\rangle\langle \mathbf{v}|$. That is, $\langle \mathbf{v}|\mathbf{u}\rangle$ “compresses” to a scalar whereas $|\mathbf{u}\rangle\langle \mathbf{v}|$ “expands” to a matrix. The vectors discussed thus far, \mathbf{t} , \mathbf{t}^{IN} , and \mathbf{f} are all kets, e.g., $|\mathbf{t}\rangle$, $|\mathbf{t}^{\text{IN}}\rangle$, $|\mathbf{f}\rangle$. In circumstances where a plain vector is presented, this should be interpreted as a column vector, or ket.

In TCM, an item is encoded in the states of temporal context in which it is presented. As a consequence, a particular state of temporal context $|\mathbf{t}\rangle$ cues the lexical representation of an item $|\mathbf{f}\rangle$ to the extent it overlaps with that item’s encoding context. A matrix \mathbf{M} serves as the mediator for this cuing. For simplicity, we initialize \mathbf{M} to be zero before presentation of the list and then have \mathbf{M} change after presentation of item i such that it is encoded in the context that obtained before it was presented:

$$\Delta\mathbf{M} = |\mathbf{f}_i\rangle\langle \mathbf{t}_{i-1}| \quad (6)$$

Note that this just describes a linear associator between the states of context and the items that follow them in time. When \mathbf{M} is cued with a state of context by multiplying from the right, the result is a superposition of the items in the list, each weighted by the inner product of their encoding context to the cue context:

$$\mathbf{M}|\mathbf{t}\rangle = \sum_i |\mathbf{f}_i\rangle\langle \mathbf{t}_{i-1}|\mathbf{t}\rangle \quad (7)$$

Following previous treatments, we assume for simplicity that the item vectors $|\mathbf{f}\rangle$ are orthonormal. As a consequence, in a list in which each item is only presented once, we find that the degree to which item j is activated by a cue state $|\mathbf{t}\rangle$ is

just

$$\langle \mathbf{f}_j|\mathbf{M}|\mathbf{t}\rangle = \langle \mathbf{f}_j|\left(\sum_i |\mathbf{f}_i\rangle\langle \mathbf{t}_{i-1}|\right)|\mathbf{t}\rangle \quad (8)$$

$$= \sum_i \langle \mathbf{f}_j|\mathbf{f}_i\rangle\langle \mathbf{t}_{i-1}|\mathbf{t}\rangle \quad (9)$$

$$= \sum_i \delta_{ij}\langle \mathbf{t}_{i-1}|\mathbf{t}\rangle \quad (10)$$

$$= \langle \mathbf{t}_{j-1}|\mathbf{t}\rangle, \quad (11)$$

where δ_{ij} is the Dirac delta function, which is one if $i = j$ and zero otherwise. Note that the last line states that an item is activated by a context cue to the extent it overlaps with that item’s encoding context.

If a particular item $|\mathbf{f}_\alpha\rangle$ is repeated at multiple positions in the list, then the degree to which it is activated is given by the similarity of the cue context to the sum of the context states in which it was presented:

$$\langle \mathbf{f}_\alpha|\mathbf{M}|\mathbf{t}\rangle = \sum_i \langle \mathbf{f}_\alpha|\mathbf{f}_i\rangle\langle \mathbf{t}_{i-1}|\mathbf{t}\rangle \quad (12)$$

$$= \sum_{i:\mathbf{f}_i=\mathbf{f}_\alpha} \langle \mathbf{t}_{i-1}|\mathbf{t}\rangle, \quad (13)$$

where the notation in the sum means that the sum runs over all indices i where the item α was presented at that time step. One can think of this expression as the similarity of probe context to the average context in which item $|\mathbf{f}_\alpha\rangle$ was presented times the frequency with which it was presented.

Contextual retrieval.

TCM, contiguity effects are predicted to be a consequence of the effect items have on the state of context. The \mathbf{t}_i^{IN} at each time step is *caused* by the item presented at time step i . If an item causes a consistent input pattern across multiple presentations, then the context cue after the second presentation will resemble to some extent the contextual states surrounding the first presentation of the item. Because other items are

cued to the extent that the context cue resembles their encoding contexts (Eqs. 11, 13), repetition of an item gives rise to contiguity effects. That is, when an item is repeated, because the input pattern it causes resembles contextual states during list presentation, it provides an effective cue for the neighbors of its original presentation.

In TCM, there are two relevant sources of contiguity effects. First, the input pattern when an item is repeated can resemble the input pattern it caused when it was initially presented. Because this input pattern was part of the contextual states for items that followed the initial presentation of the item, this gives rise to a forward asymmetry, such that the items that followed the original presentation of the repeated item are cued (open symbols Figure 5b). In addition, if the repeated item manages to retrieve the state of context present when it was originally presented, this provides a symmetric retrieval cue (filled symbols Figure 5b, see also Howard & Kahana, 2002).⁶ This latter component reflects a change in the input pattern caused by an item across repeated presentations.

Thus far the model we have described does not diverge from previously-published versions of TCM in any important way. Those versions of TCM were designed to describe temporal contiguity effects between items presented once in a randomly-assembled series of items, as would typically be the case in single-trial free recall (Howard & Kahana, 2002; Howard et al., 2006; Sederberg et al., in press). The model was not designed to address situations in which items are repeated multiple times. In fact, the previously published versions of TCM are subject to collapse of the input space (but see Rao & Howard, 2008). Here we describe this problem with the old formulation and describe a new rule for contextual retrieval that avoids this problem.

In TCM, the space spanned by the context vectors \mathbf{t} is determined by the space spanned by the input vectors \mathbf{t}^{IN} (Eq. 5). In previous treatments of TCM (e.g. Howard & Kahana, 2002; Howard et al., 2005), contextual learning was described by the following expression for the input pattern caused by an item presented at time step i and then repeated at time step r :

$$\mathbf{t}_r^{\text{IN}} = \alpha_o \mathbf{t}_i^{\text{IN}} + \alpha_n \mathbf{t}_{i-1}. \quad (14)$$

That is, the context retrieved by the item at time step r is a combination of the input pattern caused by the item when it was originally presented, \mathbf{t}_i^{IN} , and retrieved temporal context, \mathbf{t}_{i-1} . This expression, however, has undesirable properties that become apparent when items are repeated many times. Briefly, the input pattern of each item ends up becoming similar to the input pattern of the items that preceded it. With a sufficient number of presentations of a small pool of items, the input vectors can collapse to a point (see page 106, Howard et al., 2005). That is, with runaway contextual learning, the \mathbf{t}^{IN} s corresponding to all the words in the experiment become identical. Another way of characterizing this problem is that in previous formulations, the consistent part of \mathbf{t}^{IN} isn't sufficiently consistent across multiple presentations of the item.

Here we use an extension of TCM that avoids runaway

contextual learning (Rao & Howard, 2008). This extension retains all of the key features of previously published versions of TCM. The important differences between this model and previous treatments come when items are repeated multiple times. In this extension, the input patterns \mathbf{t}^{IN} caused by an item are composed of a fixed component we will refer to as \mathbf{c} and a changing component we will refer to as \mathbf{h} . Each \mathbf{c}_i and each \mathbf{h}_i are caused by the item presented at time step i and depend only on the identity of that item and its previous history. The \mathbf{c} vectors for each item are fixed throughout the simulation. This solves the problem with learning in the previous treatments because the fixed component “anchors” the space so that it never collapses. This is because the consistent component of \mathbf{t}^{IN} never changes no matter how many times an item is presented.

If item α is presented at time step i , then

$$\mathbf{t}_i^{\text{IN}} = (1 - \gamma) \mathbf{c}_\alpha + \gamma \hat{\mathbf{h}}_\alpha. \quad (15)$$

The hat in the second term indicates that \mathbf{h} is normalized (using the \mathcal{L}^1 norm) prior to entering this expression. In the simulations that follow, we fix the \mathbf{c}_{α} s to be unit vectors that serve as the bases for the \mathbf{t} space. With learning, \mathbf{h}_α changes from one presentation of item \mathbf{f}_α to another as

$$\Delta \mathbf{h}_\alpha = \mathbf{t}_{i-1}. \quad (16)$$

Comparing these last two expressions with Eq. 14 we can see that \mathbf{h} supports contextual retrieval and that γ is analogous to α_n in the previous formulations of TCM. Note that if $\gamma = 0$, this model behaves much like a heteroassociative model with no backward and no across-pair associations.

Selecting an item for retrieval.

TCM describes a set of rules for how context changes over experience. While this is sufficient to provide a qualitative account of recency and contiguity phenomena—as well as transitive associations—an additional step is necessary to quantitatively describe probability of recall. The retrieval rule we use here is the Luce choice retrieval rule used in some previous treatments of the model (e.g. Howard & Kahana, 2002; Howard et al., 2005, 2006). This retrieval rule suffers from a number of shortcomings (Sederberg et al., in press). Sederberg et al. (in press) have presented a much more detailed retrieval rule using competing accumulators (Usher & McClelland, 2001, see also Polyn, Norman, & Kahana, revised). That model has numerous advantages over Luce choice rule, but is considerably more complex. Despite its limitations, the Luce choice rule is sufficient for the present purposes.

Given a particular context cue $|\mathbf{t}\rangle$, we define the activation of item i as

$$a_i = \langle \mathbf{f}_i | \mathbf{M} | \mathbf{t} \rangle \quad (17)$$

⁶ The symmetry comes because item j is activated by $|\mathbf{t}_{i-1}\rangle$ as a probe to an extent given by $\langle \mathbf{t}_{j-1} | \mathbf{t}_{i-1} \rangle$. This quantity is symmetric with respect to interchange of j and i due to a basic property of the inner product.

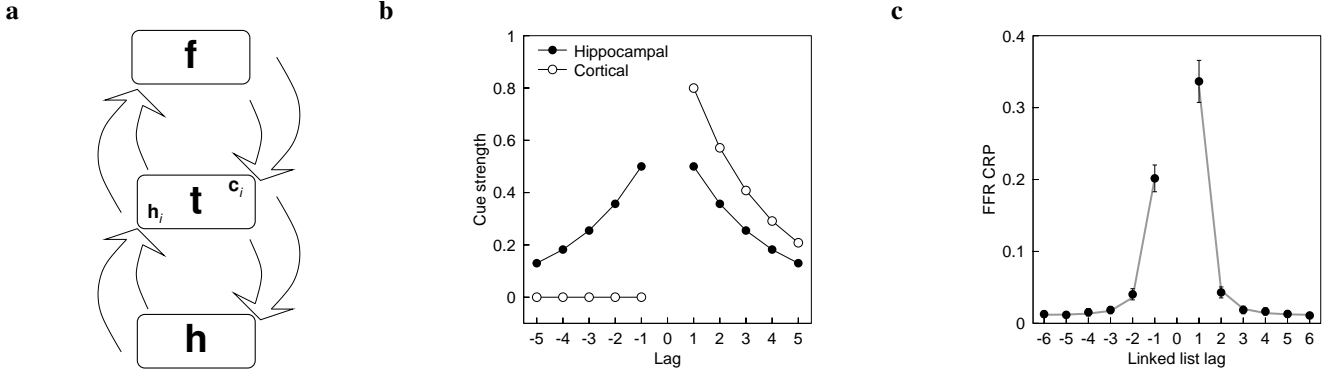


Figure 5. The temporal context model produces transitive associations. **a.** Architecture of the temporal context model used here. In the model, an item layer (f), corresponding roughly to a word recognition system, is connected to a temporal context layer (t). The temporal context layer is connected to a “hippocampal” layer (h) that functions to recover prior states of temporal context in which the item was previously encoded. When an item is presented on f , it provides two inputs to the context layer in sequence: first a constant input c that is assumed not to change over the scale of minutes and then a “hippocampal” input h that rapidly changes to recover states of temporal context. Both components contribute to the state of t that cues recall of items on f . **b.** Two components of contextual retrieval contribute to contiguity effects in TCM. One is the tendency for the input pattern caused when an item is repeated to resemble the input pattern the item caused when initially presented. This component, labeled “cortical,” and plotted with open symbols is an effective retrieval cue for items that followed the initial presentation of the item. The other component, which we refer to as contextual retrieval, enables the repeated item to recover the contextual state that was present when the item was studied. This component, labeled “hippocampal” and plotted with filled symbols, provides a symmetric retrieval cue for the item’s neighbors. **c.** Best-fitting solution from TCM to the final free recall linked-list lag CRP. The symbols with error bars represent the experimental data with 95% confidence intervals (same as Figure 4b). The thick grey lines that appear to connect the data points are the best-fitting values from TCM.

Given these activations, the probability of recalling item i from the list is given by the Luce choice rule:

$$P_R(i) = \frac{\exp 2a_i/\tau}{\sum_j \exp 2a_j/\tau}. \quad (18)$$

The sum in the denominator runs over all potential recalls. In Eq. 18, the parameter τ controls the sensitivity of the retrieval rule. As $\tau \rightarrow \infty$, the probability of recalling each item becomes equal. As $\tau \rightarrow 0$, the probability of recalling the most activated item becomes one. For this reason, Eq. 18 is referred to as a softmax rule in the connectionist literature.

Methods.

In simulating TCM, we chose the f s to be basis vectors for the f -space and the c s to be the basis vectors for the t -space. That is, the c for the first item was a one in the first component followed by all zeros; the c for the second item was a one in the second component with all other components zeros and so on. We initialized h_α to be zero for each item. The first time each item was presented γ was set to zero (because there is nothing to retrieve). We presented the model with a sequence of items corresponding to the pairs presented to a subject one at a time. To avoid any possibility that associations across pairs could be attributable to inadvertent simple temporal contiguity during study, we put an “infinitely long distractor interval” between each pair. That is, we presented a distractor—an input pattern orthonormal to every previous input pattern—with $\beta = 1$ before the first member of each pair. From Equation 5, this means that the effect of the preceding pair on context is completely lost when the first

member of the subsequent pair is presented. By presenting a different distractor between each pair, we ensure that the last word of a pair, and the first word of the next pair do not have any contextual overlap. The addition of the orthogonal distractor after each pair requires the dimensionality of the space to be quite large by the end of learning. Relaxing the degree of isolation between pairs results in diffuse across-pair associations comparable in effect to increasing τ . An infinite delay is not necessary to describe backward and transitive associations.

We randomly chose the lists and presentation order presented to ten different subjects and averaged the results across our simulated subjects. After presentation of the pairs, we averaged over the linked-list lags calculated with the ninth through twenty-seventh double-function items as probes. For each probe we presented an infinite distractor, then calculated the activations and probability of recall for the other items as a function of linked-list lag using Eq. 17 and the Luce choice rule Eq. 18. We kept track of results out to linked list lags ± 17 .

Three parameters were allowed to vary in fitting TCM to the FFR data. One was ρ , which controlled the degree of contextual overlap within pair (Eq. 5). The parameter γ controlled the degree to which the “hippocampal” component contributed to retrieved context (Eq. 15). Finally, τ controlled the sensitivity of the Luce choice rule (Eq. 18). We implemented TCM in R. We used the same methods to find the best-fitting parameters as were used for the mediated chaining model, including fitting only up to linked-list lags ± 6 .

Results and Discussion.

The best-fitting parameters for the fit of TCM to the data were $\rho = .16$, $\gamma = .97$ and $\tau = .24$. The high value of γ reflects the fact that backward adjacent and transitive associations require the hippocampal component of retrieved context. Figure 5c shows the best-fitting solution of TCM to the final free recall linked-list CRP. As can be seen from the figure, the fit is excellent. The best-fitting chi-squared with eight degrees of freedom was 4.27, n.s., indicating that deviations from the model's values are not different than would be expected by chance.

To directly compare the fits of TCM to those of the mediated chaining model, we took the best-fitting probabilities for each model and generated 1000 samples of 4,000 responses chosen according to those probabilities. We then compared this simulated data to the actual data and asked which model provided a better fit. For 969/1,000 samples, TCM was closer to the data than the mediated chaining model. If the fits of the model were equivalent, the probability of this occurring by chance is less than 10^{-243} .

In addition to the excellent quantitative fit, TCM also described the qualitative features of the data—a larger asymmetry for the adjacent *vs* remote transitions as well as robust transitive associations as manifest by the greater probability of recall for the items at ± 2 than for those at more remote linked-list lags (see Table 2). The adjacent asymmetry index for the best-fitting values generated by TCM was quite close to the observed empirical value, $z = .28$. The remote asymmetry index generated by TCM was much smaller than the predicted adjacent asymmetry index and not-significantly different than the empirically-observed value, $z = 1.3$, $p > .05$. This pattern can be contrasted with the results from the mediated chaining model, which generated an adjacent asymmetry slightly smaller than the data, $z = 1.75$, $p < .05$, but a remote asymmetry index, .25, that was actually *larger* than the predicted adjacent asymmetry factor and much larger than the empirically-observed value, $z = 3.42$, $p < .001$.

Whereas the mediated chaining model failed to correctly capture the pattern of greatly reduced asymmetry in moving from adjacent to remote transitions, TCM was able to do so within the margin of error of the data. This implementation of TCM also produced a small asymmetry in the remote transitions. In TCM, transitive associations depend on a non-zero value of γ . The asymmetry is also a result of contextual retrieval. Consider a linked list consisting of items *A* through *E* that have been presented many times each. Because of the hippocampal component of contextual retrieval, the context in which item *E* has been presented includes \mathbf{h}_D . This pattern includes a term that goes like \mathbf{t}_C^N , which includes both \mathbf{h}_C and \mathbf{c}_C . This means that both \mathbf{h}_C and \mathbf{c}_C are an effective retrieval cue for item *E*. Although the contextual cue for item *A* resembles \mathbf{h}_C , it does not resemble \mathbf{c}_C , resulting in the asymmetry when *C* is presented as a probe. Put another way, if we think of the *cs* as defining the space, with repeated presentation the \mathbf{t}^N s corresponding each item move towards the *c* of preceding items. The optimal probe context for each item also moves in this direction. Because of this movement, which depends on \mathbf{h} and contextual retrieval, the *c* for an

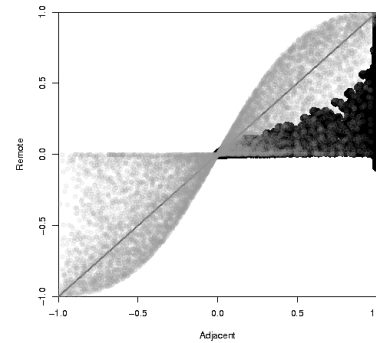


Figure 6. Model complexity analyses. The remote asymmetry index is shown as a function of the adjacent asymmetry index for a variety of parameter values from TCM (black symbols) and the heteroassociative mediated chaining model (grey symbols). There are 10,000 points for the heteroassociative model—the points are drawn with transparency to allow the perception of depth. Only 4,000 points were generated for the TCM simulation. For TCM, the vertical line of values with the adjacent factor at 1.0 are generated simulations with values of $\gamma < .2$. The dark grey line is the diagonal.

item ends up closer to the encoding context of successors in the double-function list than the encoding context of predecessors in the double function lists. This asymmetry would be altered if the rate at which \mathbf{h} and the encoding context move towards the *cs* of preceding items were not the same. With the rates the same, as in the application here, the ratio between $\langle \mathbf{f}_{i+j} | \mathbf{M} | \mathbf{t}_i^N \rangle$ and $\langle \mathbf{f}_{i-j} | \mathbf{M} | \mathbf{t}_i^N \rangle$ at steady state remains approximately constant as *j* increases. This is very different from the behavior of Eqs. 1 and 2, where the ratio grows with *j*. The non-linearity in the Luce choice retrieval rule (Eq. 18) magnifies the asymmetry at adjacent linked-list lags much more so than between more remote linked-list lags, accounting for the much larger asymmetry at adjacent linked-list lags than at more remote linked-list lags.

Model complexity analysis

In order to compare the range of qualitative behaviors demonstrated by TCM and the heteroassociative model we generated simulated data from a wide variety of parameter values and compared the adjacent and remote asymmetry indices for the simulated data. We sampled parameters uniformly over the complete range possible for the heteroassociative mediated chaining model—*f*, *r* and *g* were allowed to vary from zero to one. For TCM we also sampled uniformly. The range for ρ was zero to 0.999; the range for γ was zero to 1 and the range for τ was 0.1 to 1.0. The adjacent asymmetry index and the remote asymmetry index for each set of parameters for each model are shown in Figure 6.

In Figure 6, values from the heteroassociative chaining model are shown as light grey symbols; values from TCM are shown as black symbols. Unlike the heteroassociative chaining model, TCM is incapable of generating a back-

ward asymmetry. Moreover, the mediated chaining model can generate data in which the remote asymmetry is greater than the adjacent asymmetry. These are the points above the diagonal in the upper right quadrant and below the diagonal in the lower left quadrant. These regions are densely sampled by the heteroassociative model and include the best-fitting parameter values. Interestingly, both models are able to cover the region where the experimental data lie—with an adjacent asymmetry factor of .25 and a remote asymmetry factor of .03. However, the mediated chaining model reaches this low level of remote asymmetry by making the remote associations very small relative to the level of guessing. That is, the linked-list lag CRP curves generated by the mediated chaining model in this region are essentially flat for linked-list lags with absolute value of two or greater. We conclude that TCM's superior description of the data is not merely a consequence of greater model flexibility.

General Discussion

This paper is an attempt to shed light on the basis of the temporal contiguity effect in episodic memory tasks. Retrieved context accounts, exemplified by TCM, describe temporally-defined associations between items as resulting from shared temporal context rather than temporal proximity *per se*. After learning double-function lists of paired associates, TCM predicts across-pair transitive associations as a natural consequence of the mechanism hypothesized to underlie the contiguity effect. Heteroassociative accounts of the contiguity effect can explain across-pair associations by resorting to a mediated chaining hypothesis in which subjects step through the “links in the chain” to reach distant items.

Our empirical findings demonstrated robust associations formed both within- and across-pairs, using evidence from both patterns of intrusions and FFR transitions. The strength of these associations depended on the distance between items in the linked-list structure. It is as if participants were able to integrate isolated temporal events into a coherent linear structure that bridged across the pairs.

The specific form of these associations strongly favored TCM over the mediated chaining hypothesis. In both intrusion patterns and the linked-list CRP, we found a strong asymmetry among forward and backward adjacent associations, but no evidence for such an asymmetry in remote associations that bridge across pairs. While we cannot accept the null, which would imply a qualitative difference between adjacent and remote associations, we were able to establish that the adjacent asymmetry index was certainly greater than the remote asymmetry index. While we cannot rule out every possible mediated chaining model we simulated a straightforward mediated chaining model that failed to capture this pattern of results. In contrast, TCM provided an excellent quantitative fit to the pattern of results. This is not a consequence of a generally higher level of model flexibility, but rather a consequence of the structural assumptions of TCM. TCM's account of the contiguity effect, and in particular its principled account of the asymmetry observed in the contiguity effect, lead it to naturally predict graded transitive as-

sociations that depend on a process of contextual binding and retrieval.

According to TCM, two components contribute to temporally-defined associations and these components have different properties. Forward associations can be formed in the absence of contextual learning. In contrast, the ability of items to bind themselves to the temporal contexts in which they are presented and recover these associations when they are repeated enable the development of associations that bridge and integrate disparate experience.

Can the present findings be reconciled with a heteroassociative account? Although mediated chaining is the most straightforward way to reconcile the existence of across-pair associations with a heteroassociative source of contiguity effects, there are several other potential strategies that may be more successful at describing our data without abandoning a heteroassociative account of the contiguity effect. One possibility is that the across-pair associations result from surreptitious retrieval during study and/or testing of the pairs. That is, during testing of $A-?$, subjects may retrieve B and then surreptitiously use this as a cue to recover C . Despite the fact that they were not presented together, A and C would now have been experienced close together in time enabling them to be associated on the basis of temporal proximity. Several implementations of this idea were evaluated for earlier versions of this paper, but none were successful in fitting our pattern of results. The major problem was that they were never able to demonstrate sufficiently strong remote associations. This is because remote associations can only be learned if an adjacent association is expressed, and thus also learned.

Another approach would be to assume that there exists a symmetric heteroassociative mechanism and then focus one's attention on generating an explanation for the observed asymmetry within-pair. The advantage of this approach is that across a broad range of data, pairs can be argued to be symmetrically encoded (Kahana, 2002; Rizzuto & Kahana, 2001); perhaps the within-pair asymmetry in the present study is a consequence of consistently testing in only one direction. Indeed, (Rehani & Caplan, in preparation) did not observe within-pair asymmetry in double-function lists when testing proceeded in both directions⁷ Perhaps during training subjects learn to edit out, or perhaps suppress, backward responses during paired-associate testing because they are incorrect responses in that context. The challenge for this approach then becomes why the remote intrusions in the backward direction would not also be affected by this process.

Perhaps the most promising approach we have considered in reconciling our findings with a heteroassociative storage mechanism comes from the idea that retrieval follows a spreading activation process. That is, there is an associative matrix between items that is strictly heteroassociative, with associations only formed between members of the

⁷ It should be noted that there are a number of other methodological differences with the present study. Notably, they did not evaluate associations using FFR.

same pair. However, during retrieval activation spreads extensively across the associative links *prior* to the attempt to select an item for retrieval. The remote nodes may experience less asymmetry than adjacent nodes because there is less tendency for activation to spread away from them. That is, although forward associations receive more activation, they also send away more activation. If these activations are then sent through a sufficiently non-linear retrieval rule, the small level of asymmetry in the activations may be within the margin of error of the experimental data (not unlike the implementation of TCM presented here). This spreading-activation approach resembles TCM to some extent. When an item is presented as a probe, the effect is an implicitly-activated pattern of activation across items reflects the history of the item and takes into account the global history of stimulus relationships (not unlike implicitly-activated associates in the PIER model Nelson, Schreiber, & McEvoy, 1992; Nelson, McKinney, Gee, & Janczura, 1998). In TCM, when a probe is presented it generates a contextual state that also reflects the history of the item and the global history of stimulus relationships. It may be extremely difficult to distinguish predictions of such a spreading-activation model from those of TCM.

Perspectives on temporal contiguity

Heteroassociative accounts of contiguity effects in episodic recall tasks are quite intuitive when viewed from a certain perspective. They have the advantage of parsimony if the only goal to be accomplished is to provide a simple association between two stimuli. However, the retrieved context account may be more intuitively appealing if one starts from a slightly different perspective. Tulving (1983, 2002) characterizes episodic memory as “mental time travel” in which one remembers not only “what” but also “when” and “where” of a memory. Two-process theorists of item recognition describe recollection as the vivid recovery of contextual information associated with a probe item (e.g. Yonelinas, 2002). Starting from this perspective, the idea that episodic recall, even paired associate learning, is largely a consequence of binding and recovery to a representation of spatiotemporal context seems more natural than a simple heteroassociative account.

Nonheteroassociative accounts of contiguity motivated by neural considerations.

Similar arguments in favor of the intuitive appeal of a contextual recovery account of temporal contiguity effects in episodic memory can be made starting with considerations from neurobiology and cognitive neuroscience. The hippocampus, a medial temporal lobe structure widely believed to be essential for episodic memory, also maintains a representation of location, or spatial context, as an animal moves through its environment (e.g. O’Keefe & Dostrovsky, 1971; Wilson & McNaughton, 1993). Burgess and colleagues (Burgess, 2002; Burgess, Maguire, & O’Keefe, 2002) have argued that the primary function of the hippocampus in spatial navigation tasks is to enable binding of item representations to a representation of the spatial context in

which they were encountered (see also Howard et al., 2005; Smith & Mizumori, 2006). From this perspective, the contextual recovery account of temporal contiguity effects seems much more intuitive than a heteroassociative account.

Evidence that the hippocampus is responsible for recovery of a representation of temporal context can also be inferred from lesion studies conducted with rats. Bunsey and Eichenbaum (1996) used an analogue of paired associate learning using odors as stimuli to teach animals chained pairs of conditional discriminations $A - B$ and $B - C$. While animals with hippocampal lesions were able to learn $A - B$ and $B - C$ as well as sham-lesioned animals, they differed from the sham-lesioned animals in that they showed no evidence for bridging $A - C$ associations. Moreover, the animals with hippocampal lesions also were impaired at generalizing to backward $B - A$ associations. While it is very difficult to account for these findings using a heteroassociative account of the contiguity effect, TCM can create just this pattern of results if contextual recovery is impaired (Howard et al., 2005; see also Meeter, Myers, & Gluck, 2005; O’Reilly & Rudy, 2001; Wu & Levy, 1998).

Nonheteroassociative accounts of contiguity from animal behavior studies.

There is also evidence from animal learning paradigms that argues against a heteroassociative account of even what might appear to be simple conditioning. Secondary generalization, which has been experimentally observed in pigeons (e.g. Wasserman, DeVolder, & Coppage, 1992), refers to the learning setting in which two arbitrary stimuli are associated to a common response. Subsequently, pairing one of the stimuli to another response results in generalization such that the other stimulus also evokes the new response without explicit pairing. Secondary generalization is analogous to transitive association—the secondary generalization develops after the stimuli are experienced in the temporal context of a particular response. Rather than appealing to a heteroassociative account based on mediated chaining (Hull, 1947), Wasserman et al. (1992) attributed the finding of secondary generalization in pigeons to the development of conceptual categories.

Trace conditioning refers to a classical conditioning paradigm in which the offset of the conditioned stimulus (CS) and the onset of the unconditioned stimulus (UCS) are separated by a temporal delay. In their Experiment 2, Cole, Barnett, and Miller (1995) paired two auditory stimuli CS1 and CS2, each 5 s in duration, such that presentation of CS2 immediately followed the offset of CS1. Later CS1 was paired with an unconditioned stimulus (foot shock) such that presentation of the unconditioned stimulus followed the offset of CS1 by either zero or five seconds. The choice of delays and the duration of CS2 was such that in the five second condition, the onset of the UCS corresponded to the time at which CS2 would have terminated had it been presented as during the initial pairing of CS1 and CS2. Although conditioned responding to CS1 was greater after training in the zero second condition, the conditioned responding to CS2 *via* second order transfer was greater in the five second con-

dition. Cole et al. (1995) interpreted their results as evidence that the association between CS1 and CS2 reflected information about the temporal relationships implied by their pairing (see also Gallistel & Gibbon, 2000) rather than a heteroassociative mechanism. Moreover, it is necessary to integrate information about the relationship between CS1 and CS2 and the relationship between CS1 and the UCS to infer the relationship between CS2 and the UCS. This is analogous to the need to integrate information about $A - B$ and $B - C$ to create information about the $A - C$ association in the present experiment.

To Eichenbaum (2001), the effect of hippocampal lesion on learning of overlapping pairs of odors (Bunsey & Eichenbaum, 1996; Dusek & Eichenbaum, 1997) suggests that the hippocampus is responsible for maintaining and organizing information about the spatial and temporal relationships between items. In this view, the ability to organize stimuli into a “memory space” is considered to be a central property of the hippocampus’ function in support of declarative and episodic memory. This conception can be seen as integrating the positions of Wasserman et al. (1992) and Cole et al. (1995). TCM can be seen as a quantitative implementation of a memory space in which the patterns t^n evoked by the items reflect the temporal relationships between those items (Rao & Howard, 2008). These include not only temporal relationships that are explicitly trained, i.e. $A - B$ pairs, but also latent temporal relationships that are not explicitly presented (e.g. $A - C$).

Transitive association, contextual recovery and semantic spaces

Across-pair associations can be seen as evidence of subjects’ ability to extract the structure of experience from isolated events. Rao and Howard (2008) showed that contextual retrieval can be used to extract structures with a variety of topologies rather than just the linear structure of the double-function list used here. Moreover, in situations where the structure to be extracted is clustered, contextual retrieval allows for a more rapid learning of the structure. Such structures may be very common indeed. For example, if word A is a synonym of word B , and word B is a synonym of word C , then word A and C are much more likely to also be synonyms of each other than two words chosen at random. Clustering is a defining characteristic of small worlds networks, which have been shown to be ubiquitous not only in natural systems, but also in semantic structures (Steyvers & Tenenbaum, 2005), suggesting that contextual retrieval may facilitate learning in a great many situations.

Context has been a common feature of computational models of semantic memory (e.g. Landauer & Dumais, 1997; Griffiths & Steyvers, 2002; Jones & Mewhort, 2007). For instance, in latent semantic analysis (LSA, Landauer & Dumais, 1997), two words are judged to have similar meanings by virtue of having occurred in the same context. This property is analogous to pairwise $A - B$ associations. In addition, words are also judged to have similar meanings by virtue of having occurred in similar contexts, that is contexts that contain similar words (see for instance, Landauer

& Dumais, 1997, Figure A4). This latter property is closely analogous to the transitive $A - C$ associations observed here. Other computational models of semantic memory, including the probabilistic topics model (Griffiths & Steyvers, 2002, 2003), the hyperspace analogue of language (HAL, Lund & Burgess, 1996) and the BEAGLE model (Jones & Mewhort, 2007) also naturally predict transitive associations.

Dennis (2004, 2005) has explored the role of syntagmatic and paradigmatic associations (Nelson, 1977) to develop a model of verbal cognition. In free association norms, some associates appear to follow from simple temporal contiguity—for instance, RUN-FAST. Other associates appear to follow from a deeper level of meaning that is not necessarily extracted from simple cooccurrence, e.g., RUN-WALK. The former are referred to as syntagmatic associates; the latter are referred to as paradigmatic associates. Syntagmatic associations are analogous to the asymmetric part of the within-pair associations described here. Paradigmatic associations are analogous to across-pair associations that do not depend on simple temporal proximity. Indeed, learning of paradigmatic associations does not apparently depend on temporal contiguity (McNeill, 1966).

This striking convergence between computational models of semantic learning and contextual retrieval accounts of temporal contiguity effects in episodic recall suggests a deep relationship between semantic and episodic memory that may ultimately improve our understanding of each. In many of the semantic memory models, the notion of what constitutes a context is provided as an input to the model—collections of words are tagged as members of a single document. Moreover, temporal context in these models is often discrete—words either co-occur in a particular document or not. One of the central ideas of TCM is that temporal context changes gradually over time. It is intriguing to wonder about the implications of using a specific representation of temporal context specified by a model of episodic recall to construct semantic spaces.

Conclusions

We studied the associative structure induced by learning double-function lists of paired associates. We found evidence for associations between words that were never presented together but that were presented in similar temporal contexts by examining both intrusion probabilities and FFR transition probabilities. Whereas associations between adjacent members of the linked-list were strongly asymmetric, we observed no asymmetry in the remote bridging associations. This finding cannot easily be reconciled with mediated chaining, casting doubt on the heteroassociative account of temporal contiguity effects in paired associate learning. These transitive associations are, however, a perfectly natural consequence of accounts of contiguity effects in episodic memory that depend on contextual retrieval. TCM, a formal model of episodic association based on contextual retrieval, provided an outstanding quantitative fit to the key data describing pairwise and transitive associations. The analogy between contextual retrieval accounts of temporal contiguity

effects in episodic memory and computational models of semantic memory suggests that perhaps transitive associations will prove to be critical in understanding the transition between episodic and semantic memory.

References

- Atkinson, R. C., & Shiffrin, R. M. (1968). Human memory: A proposed system and its control processes. In K. W. Spence & J. T. Spence (Eds.), *The psychology of learning and motivation* (Vol. 2, p. 89-105). New York: Academic Press.
- Bunsey, M., & Eichenbaum, H. B. (1996). Conservation of hippocampal memory function in rats and humans. *Nature*, *379*(6562), 255-257.
- Burgess, N. (2002). The hippocampus, space, and viewpoints in episodic memory. *Quarterly Journal of Experimental Psychology*, *55*(4), 1057-80.
- Burgess, N., Maguire, E. A., & O'Keefe, J. (2002). The human hippocampus and spatial and episodic memory. *Neuron*, *35*(4), 625-41.
- Caplan, J. B. (2005). Associative isolation: unifying associative and order paradigms. *Journal of Mathematical Psychology*, *49*, 383-402.
- Caplan, J. B., Glaholt, M. G., & McIntosh, A. R. (2006). Linking associative and serial list memory: Pairs versus triples. *Journal of Experimental Psychology : Learning, Memory, and Cognition*, *32*(6), 1244-65.
- Chappell, M., & Humphreys, M. (1994). An autoassociative neural network for sparse representations: Analysis and application to models of recognition and cued recall. *Psychological Review*, *101*, 103-128.
- Cole, R. P., Barnet, R. C., & Miller, R. R. (1995). Temporal encoding in trace conditioning. *Animal Learning & Behavior*, *23*(2), 144-153.
- Dennis, S. (2004). An unsupervised method for the extraction of propositional information from text. *Proceedings of the National Academy of Science, USA*, *101 Suppl 1*, 5206-13.
- Dennis, S. (2005). A memory-based theory of verbal cognition. *Cognitive Science*, *29*, 145-193.
- Dusek, J. A., & Eichenbaum, H. (1997). The hippocampus and memory for orderly stimulus relations. *Proceedings of the National Academy of Science, USA*, *94*(13), 7109-7114.
- Eichenbaum, H. (2001). The hippocampus and declarative memory: cognitive mechanisms and neural codes. *Behavioural Brain Research*, *127*(1-2), 199-207.
- Farrell, S., & Lewandowsky, S. (in press). Empirical and theoretical limits on lag-recency in free recall. *Psychonomic Bulletin & Review*.
- Friendly, M., Franklin, P. E., Hoffman, D., & Rubin, D. C. (1982). The Toronto Word Pool: Norms for imagery, concreteness, orthographic variables, and grammatical usage for 1,080 words. *Behavior Research Methods and Instrumentation*, *14*, 375-399.
- Gallistel, C. R., & Gibbon, J. (2000). Time, rate, and conditioning. *Psychological Review*, *107*(2), 289-344.
- Griffiths, T., & Steyvers, M. (2002). A probabilistic approach to semantic representation. In W. Gray & C. Schunn (Eds.), *Proceedings of the 24th annual conference of the cognitive science society* (p. 381-386). Austin, TX: Cognitive Science Society.
- Griffiths, T., & Steyvers, M. (2003). Prediction and semantic association. *Advances in Neural Information Processing Systems*, *15*.
- Howard, M. W., Addis, K. A., Jing, B., & Kahana, M. (2007). Semantic structure and episodic memory. In T. K. Landauer, D. S. McNamara, S. Dennis, & W. Kintsch (Eds.), *LSA: A road towards meaning* (p. 121-141). Mahwah, NJ: Laurence Erlbaum Associates.
- Howard, M. W., Fotedar, M. S., Datey, A. V., & Hasselmo, M. E. (2005). The temporal context model in spatial navigation and relational learning: Toward a common explanation of medial temporal lobe function across domains. *Psychological Review*, *112*(1), 75-116.
- 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*, 923-941.
- Howard, M. W., & Kahana, M. J. (2002). A distributed representation of temporal context. *Journal of Mathematical Psychology*, *46*(3), 269-299.
- Howard, M. W., Kahana, M. J., & Wingfield, A. (2006). Aging and contextual binding: Modeling recency and lag-recency effects with the temporal context model. *Psychonomic Bulletin & Review*, *13*, 439-445.
- Howard, M. W., Sederberg, P. B., & Kahana, M. J. (submitted). Reply to Farrell & Lewandowsky: Changes in the shape of the lag-CRP predicted by TCM due to recency. *Psychonomic Bulletin & Review*.
- Hull, C. L. (1947). The problem of primary stimulus generalization. *Psychological Review*, *54*, 120-134.
- 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*, 208-233.
- Jensen, O., & Lisman, J. E. (2005). Hippocampal sequence-encoding driven by a cortical multi-item working memory buffer. *Trends in Neuroscience*, *28*(2), 67-72.
- Jones, M. N., & Mewhort, D. J. K. (2007). Representing word meaning and order information composite holographic lexicon. *Psychological Review*, *114*, 1-32.
- Kahana, M. J. (1996). Associative retrieval processes in free recall. *Memory & Cognition*, *24*, 103-109.
- Kahana, M. J. (2002). Associative symmetry and memory theory. *Memory & Cognition*, *30*(6), 823-40.
- Kahana, M. J., Howard, M., & Polyn, S. (2008). Associative processes in episodic memory. In H. L. Roediger III (Ed.), *Cognitive psychology of memory, Vol. 2 of learning and memory - a comprehensive reference* (J. Byrne, Editor) (p. 476-490). Oxford: Elsevier.
- Kahana, M. J., Howard, M. W., Zaromb, F., & Wingfield, A. (2002). Age dissociates recency and lag-recency effects in free recall. *Journal of Experimental Psychology : Learning, Memory, and Cognition*, *28*, 530-540.
- Klein, K. A., Addis, K. M., & Kahana, M. J. (2005). A comparative analysis of serial and free recall. *Memory & Cognition*, *33*(5), 833-839.
- Landauer, T. K., & Dumais, S. T. (1997). Solution to Plato's problem : The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, *104*, 211-240.
- Lewandowsky, S., & Murdock, B. B. (1989). Memory for serial order. *Psychological Review*, *96*, 25-57.
- Lund, K., & Burgess, C. (1996). Producing high-dimensional semantic spaces from lexical co-occurrence. *Behavior Research Methods, Instruments & Computers*, *28*(2), 203-208.

- McNeill, D. (1966). A study of word association. *Journal of Verbal Learning and Verbal Behavior*, 5(6), 548-557.
- Meeter, M., Myers, C. E., & Gluck, M. A. (2005). Integrating incremental learning and episodic memory models of the hippocampal region. *112*(3), 560-85.
- Mensink, G.-J. M., & Raaijmakers, J. G. W. (1988). A model for interference and forgetting. *Psychological Review*, 95, 434-55.
- Murdock, B. B. (1982). A theory for the storage and retrieval of item and associative information. *Psychological Review*, 89, 609-626.
- Nelson, D. L., McKinney, V. M., Gee, N. R., & Janczura, G. A. (1998). Interpreting the influence of implicitly activated memories on recall and recognition. *Psychological Review*, 105, 299-324.
- Nelson, D. L., Schreiber, T. A., & McEvoy, C. L. (1992). Processing implicit and explicit representations. *Psychological Review*, 95, 322-348.
- Nelson, K. (1977). The syntagmatic-paradigmatic shift revisited: A review of research and theory. *Psychological Bulletin*, 84(1), 93-116.
- O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map. preliminary evidence from unit activity in the freely-moving rat. *Brain Research*, 34, 171-175.
- O'Reilly, R. C., & Rudy, J. W. (2001). Conjunctive representations in learning and memory: principles of cortical and hippocampal function. *Psychological Review*, 108(2), 311-345.
- Polyn, S. M., Norman, K. A., & Kahana, M. J. (revised). Episodic and semantic organization during free recall: The control of memory search. *Psychological Review*.
- Popper, J. (1959). Mediated generalization. In R. R. Bush & W. K. Estes (Eds.), *Studies in mathematical learning theory* (p. 94-108). Stanford, CA: Stanford University Press.
- Primoff, E. (1938). Backward and forward associations as an organizing act in serial and in paired-associate learning. *Journal of Psychology*, 5, 375-395.
- Raaijmakers, J. G. W., & Shiffrin, R. M. (1980). SAM: A theory of probabilistic search of associative memory. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 14, p. 207-262). New York: Academic Press.
- Rao, V. A., & Howard, M. W. (2008). Retrieved context and the discovery of semantic structure. In J. Platt, D. Koller, Y. Singer, & S. Roweis (Eds.), *Advances in neural information processing systems 20*. Cambridge, MA: MIT Press.
- Rehani, M., & Caplan, J. B. (in preparation). The influence of associative interference on cued recall of word pairs. <http://count.psych.ualberta.ca/cml/>.
- Rizzuto, D. S., & Kahana, M. J. (2001). A neural network model of paired associate learning. *Neural Computation*, 13, 1835-1852.
- Sederberg, P. B., Howard, M. W., & Kahana, M. J. (in press). A context-based theory of recency and contiguity in free recall. *Psychological Review*.
- Sirotin, Y. B., Kimball, D. R., & Kahana, M. J. (2005). Going beyond a single list: Modeling the effects of prior experience on episodic free recall. *Psychonomic Bulletin & Review*, 12, 787-805.
- Slamecka, N. J. (1976). An analysis of double-function lists. *Memory & Cognition*, 4, 581-585.
- Smith, D. M., & Mizumori, S. J. (2006). Hippocampal place cells, context, and episodic memory. *Hippocampus*, 16(9), 716-29.
- Steyvers, M., & Tenenbaum, J. (2005). The large scale structure of semantic networks: statistical analyses and a model of semantic growth. *Cognitive Science*, 29, 41-78.
- Tulving, E. (1983). *Elements of episodic memory*. New York: Oxford.
- Tulving, E. (2002). Episodic memory: from mind to brain. *Annual Review of Psychology*, 53, 1-25.
- Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. *Psychological Review*, 108(3), 550-92.
- Wasserman, E. A., DeVolder, C. L., & Coppage, D. J. (1992). Non-similarity-based conceptualization in pigeons via secondary or mediated generalization. *Psychological Science*, 3(6), 374-378.
- Wilson, M. A., & McNaughton, B. L. (1993). Dynamics of the hippocampal ensemble code for space. *Science*, 261, 1055-8.
- Wu, X. B., & Levy, W. B. (1998). A hippocampal-like neural network model solves the transitive inference problem. In J. M. Bower (Ed.), *Computational neuroscience: Trends in research* (p. 567-572). New York: Plenum Press.
- Yonelinas, A. P. (2002). The nature of recollection and familiarity: A review of 30 years of research. *Journal of Memory and Language*, 46(3), 441-517.
- Young, R. K. (1961). Paired-associate learning when the same items occur as stimuli and responses. *Journal of Experimental Psychology*, 61, 315-8.