

Compound Cuing in Free Recall

Lynn J. Lohnas and Michael J. Kahana
University of Pennsylvania

According to the retrieved context theory of episodic memory, the cue for recall of an item is a weighted sum of recently activated cognitive states, including previously recalled and studied items as well as their associations. We show that this theory predicts there should be compound cuing in free recall. Specifically, the temporal contiguity effect should be greater when the 2 most recently recalled items were studied in contiguous list positions. A meta-analysis of published free recall experiments demonstrates evidence for compound cuing in both conditional response probabilities and interresponse times. To help rule out a rehearsal-based account of these compound cuing effects, we conducted an experiment with immediate, delayed, and continual-distractor free recall conditions. Consistent with retrieved context theory but not with a rehearsal-based account, compound cuing was present in all conditions, and was not significantly influenced by the presence of interitem distractors.

Keywords: compound cuing, free recall, memory

Freely recalling a list of items is a path dependent process. If I've just recalled a list item, r_1 , it is likely to remind me of another item, r_2 , which reminds me of r_3 and so forth. Thus, an earlier response r_1 should influence a later response r_3 through its effect on the intermediate response r_2 . When considering the responses made within a given recall period (in contrast to responses across lists), a model is said to be stochastic or to have the first-order Markov property if r_1 influences r_3 only through its effect on r_2 . Such models have a simplicity and elegance that makes them very appealing. It is perhaps not surprising that they have been around for a long time and have been used to explain a variety of intralist recall phenomena (Anderson, 1972; Feigenbaum & Simon, 1962; Laming, 2010; Lewandowsky & Murdock, 1989; Raaijmakers & Shiffrin, 1981).

The assumption of a stochastic retrieval process is also implicit in many empirical analyses of the free recall paradigm. For instance, perhaps one of the most prominent forms of organization in free recall is the temporal contiguity effect, quantified by the tendency to follow recall of the item from serial position i with one of the available items $i + \text{lag}$ for smaller values of $|\text{lag}|$. The temporal contiguity effect appears robustly across manipulations of list length, presentation rate, and presentation modality (Kahana, 1996). The temporal contiguity effect appears to be relatively insensitive to the time scale of item presentation, as it is unaffected by a demanding interitem arithmetic distractor task

(Howard & Kahana, 1999), and it appears at very long time scales in both list memory tasks (Howard, Youker, & Venkatadass, 2008; Unsworth, 2008) and autobiographical memory tasks (Moreton & Ward, 2010). Another major form of organization in free recall of random word lists is semantic proximity, wherein participants tend to successively recall words with similar meanings (Howard & Kahana, 2002b).

Although temporal contiguity, semantic proximity, and other forms of organization in free recall assume that the retrieval cue is the just-recalled item, an alternative possibility is that multiple prior items combine to form a compound cue for the next response. As discussed below, this idea has parallels in empirical studies of other memory tasks and in theories of the temporal contiguity effect itself. According to retrieved context theory, recalling an item retrieves its previously associated contextual states, which combine with the current state of context to form the retrieval cue for the next response (e.g., Howard & Kahana, 2002a). Because context is a recency-weighted sum of previously activated cognitive states, retrieved context models embody a compound cue assumption and predict that the temporal contiguity effect should be enhanced when a sequence of previously recalled items was studied at neighboring list positions.

Here we aim to test the predicted compound cue effect in both a meta-analysis of previously published free recall studies and in a new experiment that allows us to examine the degree to which compound cuing reflects interitem rehearsal. We also compare the empirical compound cue results to predictions to three sets of model simulations. Whereas most of the organizational analyses in free recall suggest that recall is stochastic inasmuch as the just-recalled item substantially informs recall of the next item, the compound cue analyses would allow us to detect violations of this property, both qualitatively from the empirical results and quantitatively based on model predictions. First, we present simulations of a retrieved context model of free recall (the context maintenance and retrieval model; Polyn, Norman, & Kahana, 2009) which uses the recency-weighted sum of recalled items to cue retrieval. To contrast, we also compare the compound cue results to predictions

This article was published Online First August 19, 2013.

Lynn J. Lohnas and Michael J. Kahana, Department of Psychology, University of Pennsylvania.

This research was funded by National Institutes of Health Grant MH55687. We thank Jonathan Miller and Patrick Crutchley for assistance with designing and programming the experiment, and we thank Kylie Hower, Joel Kuhn, and Elizabeth Crutchley for help with data collection.

Correspondence concerning this article should be addressed to Michael J. Kahana, Department of Psychology, University of Pennsylvania, Philadelphia, PA 19104. E-mail: kahana@psych.upenn.edu

of a model that relies primarily on the just-recalled item as a cue (the episodic search of associative memory model [eSAM]; Sirotnin, Kimball, & Kahana, 2005). Such a model may help to determine whether seemingly second-order Markovian processes could be predicted by a first-order Markovian model. To confirm that the eSAM predictions reflect its retrieval cue assumption, we also present simulations of a SAM variant that relies on compound cues for recall (fSAM; Kimball, Smith, & Kahana, 2007).

In the literature on priming and item recognition tasks, compound cuing has frequently been used to refer to the joint influence of the current cue item and the preceding cue item or prime on task performance (Clark & Shiffrin, 1987; Doshier & Rosedale, 1989; Ratcliff & McKoon, 1988; Schwartz, Howard, Jing, & Kahana, 2005). In recall tasks, however, compound cuing usually refers to the influence of multiple prior items on recall performance; specifically, whether the response before the last combines with the last response to form a compound retrieval cue.

Studies of probed recall and serial learning suggest that participants can use compound cuing to facilitate retrieval (Chance & Kahana, 1997; Cohen & Sekuler, 2010; Kahana & Caplan, 2002; Posnansky, 1972). Posnansky (1972) examined the advantage of three previous items over one previous item to cue recall of the next item. On a small but significant number of trials, participants failed to recall the correct item using one previous item as a cue but could recall the item when additionally provided with the two items presented before the one-item cue. If participants relied exclusively on associations between adjacent items, we would not expect that providing additional items would aid recall. However, because the compound cue was given only when the single-item cue failed, this study leaves open the possibility that participants simply benefited from having a second recall attempt. In a more direct investigation of compound cuing, Kahana and Caplan (2002) had participants study lists of word triples or entire serial lists and then probed them for recall with either single item or double item cues. Although they found that participants exhibited superior recall accuracy in the double item cue condition irrespective of whether the cue was in the forward or backward direction, they found that participants exhibited greater recall accuracy and faster response times when probed in the forward direction than the backward direction.

Chance and Kahana (1997) examined a more endogenous form of compound cuing in a study of interresponse times (IRTs) in serial list learning. They instructed participants to learn two sequences of 15 common nouns that contained an overlapping subsequence of 1, 2, 4, or 8 items. According to classic chaining theories, which assume item i serves as the cue for item $i + 1$ (Lewandowsky & Murdock, 1989; Wickelgren, 1966), one would expect to find slower IRTs when exiting the overlapping subsequence due to associative interference from the other list. Chance and Kahana (1997) did find slower exit IRTs, but only when the overlapping segment had more than two elements. We would expect this result if participants were using a compound cue composed of more than two items: Because the compound cue contains more items than the overlapping subsequence, this helps to eliminate the ambiguity in the repeated subsequence, thus attenuating the associative interference seen in the exit IRTs. Cohen and Sekuler (2010) extended these findings to a motor sequence task. Participants learned a series of motor movements as either one complete sequence (control group) or as a set of subsequences

(experimental group). Consistent with a compound cue account, they found that the control group outperformed the experimental group.

A common feature of these previous studies is that the retrieval cue is largely determined by the experimenter. Even in serial recall, where there is no explicit cue, the requirement of ordered recall makes the sequence of prior responses a reliable cue for the next response. If, however, compound cuing is a general feature of the memory system, then one would expect to observe a compound cuing advantage even in free recall. Here we examine compound cuing in free recall: As an extension of traditional analyses of temporal contiguity, which examine how the just-recalled item influences the next transition, we investigate how the previous two items serve as a compound cue for the next transition. In the analyses reported here, we examine the conditional response probability as a function of serial position lag in the study list (lag-CRP). Given that a participant has just recalled the item from serial position i , the lag-CRP indicates the probability that the next recalled item comes from serial position $i + \text{lag}$.

Meta-Analysis

To examine compound cuing in free recall, we considered a set of free recall studies in which each participant contributed a large enough number of recall transitions such that we could estimate the lag-CRP conditional on the previous response.

Method

Studies.

Murdock (1962). Ninety participants performed immediate free recall of 80 lists. Lists were composed of words chosen randomly and without replacement from the Thorndike and Lorge (1944) word pool with G count > 20 . Using group testing, 15 participants contributed to each of six conditions as follows, where the first number indicates list length and the second number indicates presentation rate in seconds: 10–2, 15–2, 20–1, 20–2, 30–1, 40–1. Immediately following presentation of the last list item, participants had 90 s to write down as many words as possible from the just-presented list.

Sederberg et al. (2006). Forty-eight participants performed delayed free recall of 48 lists. Lists were composed of 15 high-frequency nouns, each presented visually for 1,600 ms with an 800–1,200 ms blank interstimulus interval (ISI). After the presentation of the last item, participants performed math problems of the form $A + B + C = ?$, where A , B , and C are positive, single-digit integers, for 20 s. Immediately following the distraction period, participants had 45 s to vocally recall as many words as possible from the just-presented list. Participants' responses were digitally recorded and processed offline to specify exact response onset times. The Sederberg et al. (2006) article reported analyses of the first 35 of the 48 participants included here.

Sederberg, Miller, Howard, and Kahana (2010). Twenty-seven participants performed delayed free recall of 48 lists. Lists were composed of 16 high-frequency nouns, each presented visually for 1,000 ms with a 300–700 ms blank ISI. After the presentation of the last item, participants performed math problems of the form $A + B + C = ?$, where A , B , and C are positive, single-digit integers, for 20 s. Immediately following the

distraction period, participants had 45 s to vocally recall as many words as possible from the just-presented list. Participants' responses were digitally recorded and processed offline to specify exact response onset times. In each list, noun pairs were divided into four groups of increasing semantic relatedness (Steyvers, Shiffrin, & Nelson, 2004) such that two pairs of items from each of four groups were selected without replacement and arranged such that one pair occurred at adjacent serial positions and the other pair was separated by at least two other items.

Polyn et al. (2009) and Lohnas, Polyn, and Kahana (2011). One hundred and six participants performed immediate free recall of 24–72 lists. Lists were composed of 24 words, each presented visually for 3,000 ms with an 800–1,200 ms blank ISI. Immediately following presentation of the last list item, participants had 90 s to vocally recall as many words as possible from the just-presented list. Participants' responses were digitally recorded and processed offline to specify exact response onset times. For each item, participants made either a size or animacy judgment, indicating their response during item presentation via a key press. On single-task lists, every word was judged with the same task. On task-shift lists, participants shifted back and forth between the two judgment tasks. We restricted our analyses to the single-task lists.

Measures

Conditional response probability as a function of lag (lag-CRP). For each participant, we initialized a set of numerators and denominators to 0, one for each possible transition lag (Figure 1A). For a list length of l , there are $2(l-1)-1$ possible lags: $l-1, l-2, l-3, \dots, 2, 1, -1, -2, \dots, -(l-2), -(l-1)$. We excluded transitions to and from items repeated within a recall period (e.g., transitions with lag = 0). We also excluded transitions to and from intrusions, as the lag is not defined for those items. For each list, we stepped through each recall transition, incrementing the numerator value matching the actual serial position lag of that transition and incrementing the denominators matching the set of all possible recall transitions. Specifically, possible transitions excluded lags that would be to an item outside of the list, as well as to items already recalled during the current trial. After incrementing the numerators and denominators for all of a participant's lists, the lag-CRP for that participant was simply the numerator divided by the denominator for each possible lag.

Compound lag-CRP. Most broadly, the compound lag-CRP can be considered as a lag-CRP that considers only a subset of transitions. Specifically, a compound lag-CRP considers transitions assuming that the previous transition lag_1 meets a particular criterion. That is, in addition to requiring that the transition lag_2 does not include a repetition or intrusion, we also require that it follows a transition of a particular value, e.g., $lag_1 = +1$. As with the lag-CRP, we initialized a set of numerators and denominators for each possible transition lag. We stepped through each recall transition, and for those transitions that met the specified criterion, we incremented the numerators and denominators matching the set of possible transitions. The compound lag-CRP for a participant was the numerator divided by the denominator at each lag.

An example of the lags contributing to the compound lag-CRP calculation on a single trial is shown in Figure 1C. Whereas in the standard lag-CRP analysis, each of the first two output transitions contribute separately to the conditional probabilities (e.g., lag =

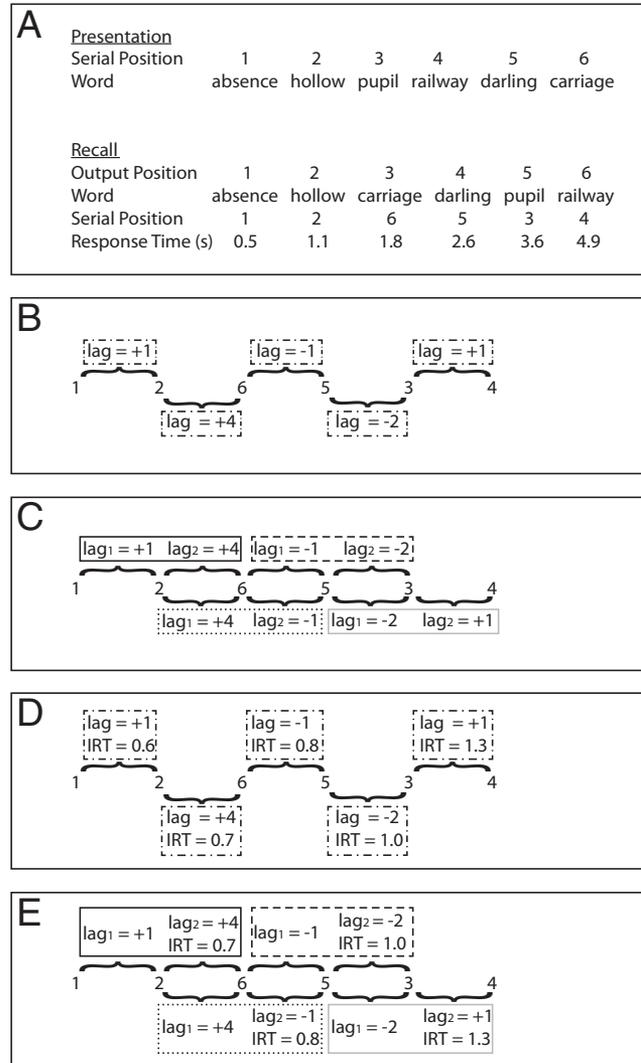


Figure 1. Calculation of conditional response functions. A. Sample list presentation and recall sequence. B. Lag calculation used for the conditional response probability as a function of lag (lag-CRP) based on the recall sequence provided in A. The boxes around each lag indicate that all lags contribute to the lag-CRP shown in Figure 2A. C. Calculation of lag_1 and lag_2 as used in the compound lag-CRP. Each box uses the same line scheme as the plots provided in Figures 2, 4, and 5 (i.e., conditional on a particular value of lag_1). The light gray box for the final pair of lag_1, lag_2 values indicates that this transition does not contribute to the compound lag analyses, as $lag_1 = -2$ does not fit any considered lag_1 criteria. D. Calculation of lag and conditional response latency (CRL) used for the lag-CRL analysis. E. Calculation of lag_1 and lag_2 as used in the compound lag-CRL. The line scheme is the same as Panel C. IRT = interresponse time.

+1, lag = +4), in the compound lag-CRP, these two transitions are considered together as $lag_1 = +1$ and $lag_2 = +4$. Based on the value of $lag_1 = +1$, the value of $lag_2 = +4$ increases the numerator of the actual transition made in the +1 compound lag-CRP. This is indicated by the solid box outline, corresponding to the function of the same line style in Figure 2. For the next considered transition, the transition from Item 2 to Item 6 becomes

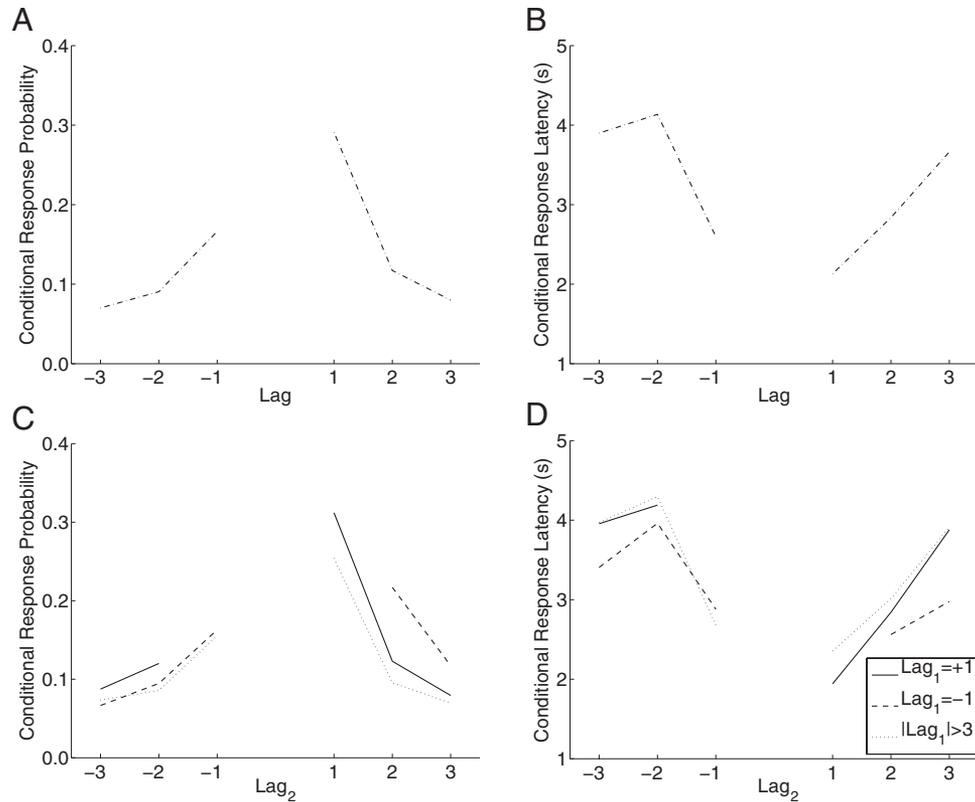


Figure 2. Single and compound cues in a meta-analysis of free recall studies. A. The contiguity effect, or tendency to successively recall neighboring items, is reflected in the conditional response probability as a function of lag (lag-CRP). B. Neighboring items recalled successively also exhibit faster interresponse times (IRTs), as exhibited in the conditional response latency as a function of lag (lag-CRL). C. Compound cuing and the contiguity effect. The conditional response probability as a function of the current lag (lag_2) illustrates the tendency to transition between neighboring items in recall. Here we show the lag-CRP separately for three cases of the lag of the previous transition, lag_1 : $\text{lag}_1 = +1$ (solid); $\text{lag}_1 = -1$ (dashed); $|\text{lag}_1| > 3$ (dotted). Included studies are Lohnas et al. (2011); Murdock (1962); Polyn et al. (2009); Sederberg et al. (2006) + replication; Sederberg et al. (2010). D. Compound cuing influences recall timing. The conditional response latency as a function of the current lag (lag_2) illustrates the tendency to transition between neighboring items in recall with faster IRTs. Here we show the lag-CRL separately for three cases of the lag of the previous transition, lag_1 : $\text{lag}_1 = +1$ (solid); $\text{lag}_1 = -1$ (dashed); $|\text{lag}_1| > 3$ (dotted). Data are from Lohnas et al. (2011); Polyn et al. (2009); Sederberg et al. (2006) + replication; Sederberg et al. (2010).

the lag_1 value of interest, and based on the subsequent transition the result is $\text{lag}_2 = -1$. Here, $|\text{lag}_1| > 3$, so this transition contributes to the corresponding compound lag-CRP and compound lag-CRL with dotted lines in Figure 2.

For the immediate free recall studies, we excluded the first two output transitions, as temporal contiguity is artifactually enhanced due to the pronounced recency effect (Davelaar, Goshen-Gottstein, Ashkenazi, Haarmann, & Usher, 2005; Sederberg, Howard, & Kahana, 2008). We excluded 11 participants because they did not contribute at least one observation to each of the lags considered in the t tests reported below. In total, 260 participants were included in this analysis.

Conditional response latency as a function of lag (lag-CRL).

For each participant, we initialized a set of elements to 0, one for each possible transition lag (Figure 1D). For a list length of l , there are $2(l-1) - 1$ possible lags: $l-1, l-2, l-3, \dots, 2, 1, -1,$

$-2, \dots, -(l-2), -(l-1)$. We excluded transitions to and from items repeated within a recall period, which includes transitions of $\text{lag} = 0$. We also excluded transitions to and from intrusions, as the lag is not defined for those items. For each list, we stepped through each recall transition, keeping track of the IRTs for the value matching the lag of each transition. The lag-CRL for a participant was the mean value of all IRTs made at a particular lag. For instance, as shown in the sample calculation in Figure 1D, both the first and last output transitions would contribute CRLs to $\text{lag} = 1$.

Compound lag-CRL. For the compound lag-CRL, the mean IRT is calculated conditional on the previous transition being of a particular lag value, termed lag_1 (Figure 1E). In addition to requiring that participants had at least one observation for each of the t tests reported below, we also required that they actually made each of those transitions at least once (i.e., it needed to be a possible transition and an actual transition). Whereas in the compound

lag-CRP, if a participant never makes an actual transition for a possible lag this corresponds to a CRP of zero, in the compound lag-CRL this corresponds to an undefined CRL. As a result, we excluded an additional 48 participants. This analysis also excluded participants from Murdock (1962), as response times were not recorded in that experiment. In total, 122 participants were included in this analysis.

Results

Given that a participant has successively recalled items from serial positions r_1 and r_2 , we sought to characterize the probability that the next recalled item would be from serial position $r_3 = r_2 + \text{lag}_2$, conditional on the availability of the appropriate lags (i.e., that r_3 falls within the list boundaries and has not been recalled on the current trial). We refer to this as the conditional response probability as a function of lag, or lag-CRP (Figure 2A). Because we further examined the dependence of the lag-CRP on the lag between two previously recalled items ($\text{lag}_1 = r_2 - r_1$), we refer to it as the compound lag-CRP to distinguish it from the traditional lag-CRP that only considers items r_2 and r_3 .

Figure 2B shows the compound lag-CRPs averaged across the included studies. The solid line shows the compound lag-CRP of each transition conditional on the previous transition satisfying $\text{lag}_1 = +1$. For this function, which we term the +1 compound lag-CRP, transitions to an item of $\text{lag}_2 = +1$ are more likely than any other transition. Similarly, the dashed line shows the compound lag-CRP conditional on the previous transition satisfying $\text{lag}_1 = -1$. For the -1 compound lag-CRP, the highest conditional probabilities of lag_2 are also in the forward direction and even exceed those of the +1 compound lag-CRP.

For these compound lag-CRPs, not all values of lag_2 are possible. For the +1 compound lag-CRP, the transition to the item of $\text{lag}_2 = -1$ is a repetition: Because the participant just transitioned from item of serial position r_1 to $r_2 + 1$, then $\text{lag}_2 = -1$ defines a transition from $r_2 - 1$ back to r_1 . Following the same logic, the transition of $\text{lag}_2 = +1$ in the -1 compound lag-CRP corresponds to a repetition as well.

The compound lag-CRP for $|\text{lag}_1| > 3$ (dotted line) closely approximates a standard lag-CRP and can be considered a control for the $|\text{lag}_1| = 1$ compound lag-CRPs. Although it is tempting to compare the +1 and -1 compound lag-CRPs to a standard lag-CRP function, the latter includes all possible compound lags, including the two conditions we wish to compare. As mentioned in the introduction, transitions at larger absolute lags are much rarer, and thus the control (compound) lag-CRP reflects compound transitions from items that are less influential on the current transition.

To quantify the compound cue advantage, we defined a compound cue score separately for the forward and backward recall transitions. For forward transitions, the compound cue score is the difference between the $\text{lag}_2 = +1$ point of the +1 compound lag-CRP and the $\text{lag}_2 = +1$ point of the control lag-CRP. For backward transitions, the compound cue score is the difference between the $\text{lag}_2 = -1$ point of the -1 compound lag-CRP and the $\text{lag}_2 = -1$ point of the control lag-CRP. Positive compound cue scores indicate a compound cue advantage, whereas negative scores indicate a disadvantage.

Analysis of the forward compound cue score across all experiments revealed a strong compound cue advantage, mean com-

pound cue score = 0.06; $t(259) = 6.48$, $p < .001$. Thus, a participant is more likely to recall the next item presented in the list if the previous two recalled items were presented successively.

In the studies analyzed here, we did not find reliable differences in the compound lag-CRPs for negative values of lag_2 , mean compound cue score = 0.01; $t(259) = 0.94$, $p > .3$. This suggests that participants benefit most from compound cuing when the next recall is in the forward direction. These results are consistent with the benefits of compound cuing in serial learning because in that task participants benefit from compound cues when recalling in forward order (Chance & Kahana, 1997; Cohen & Sekuler, 2010; Kahana & Caplan, 2002; Posnansky, 1972).

To further characterize the contribution of compound cues in free recall, we examined the IRTs of transitions. Analogous to the standard and compound lag-CRPs, Figures 2D and 2E plot the standard and compound lag-CRLs, respectively. The IRTs are faster for $\text{lag}_2 = +1$ in the +1 compound lag-CRL than the control lag-CRL, which we can again quantify as compound cue scores reliably greater than zero, $t(121) = 2.76$, $p < .01$. Thus, following a transition of $\text{lag}_1 = +1$, the $\text{lag}_2 = +1$ item is not only more likely to be recalled but will also be recalled with a faster IRT. Similar to response probabilities, this benefit in faster IRTs is only seen in the forward direction, as the backward compound cue score is not reliably greater than zero, $t(121) = 0.44$, $p > .5$.

Experiment

One possible explanation for the compound cuing effect observed in our meta-analysis is that participants use rehearsal to group items into “chunks” such that a sequence of three (or more) items may become bound together as a group (Laming, 2010). Overt rehearsal studies provide some support for this view, as participants often rehearse sequences of neighboring list items together (Rundus, 1971; Tan & Ward, 2000). Thus, one could argue that compound cuing arises because participants are recalling the sequences that they have rehearsed. To help rule out this possibility, we conducted a free recall experiment that included a continual-distractor condition designed to attenuate rehearsal.

Method

These data were collected as part of the Penn Electrophysiology of Encoding and Retrieval Study, involving three multisession experiments that were sequentially administered. The data reported here come from 108 young adults who took part in Experiment 2. These data exclude four additional participants who performed poorly in the distractor tasks (taking on average more than 4 s per arithmetic problem, as described below), suggesting that they might have used the distraction intervals for rehearsal. One session from one participant was excluded because a fire alarm sounded in the middle of the session.

Each session consisted of 12 lists of 16 words presented one at a time on a computer screen. Each study list was followed by a free recall test, and each session ended with a recognition test. Half of the six sessions were randomly chosen to include a final free recall test before recognition, but only the initial free recall test phase is considered here.

Words were presented concurrently with a task cue, indicating the judgment that the participant should make for that word. Each

word was presented with one of two encoding tasks: a size judgment (“Will this item fit into a shoebox?”) or an animacy judgment (“Does this word refer to something living or not living?”), and the current task was indicated by the color and typeface of the presented item. The task manipulation was not considered here.

Each word was drawn from a pool of 1,638 words. Lists were constructed such that varying degrees of semantic relatedness occurred at both adjacent and distant serial positions, although the manipulation of semantic relatedness is not considered here. Each item was on the screen for 3,000 ms, followed by a jittered 800- to 1,200-ms ISI. If the word was associated with a task, participants indicated their response via a key press. After the last item in the list, there was a 1,200- to 1,400-ms jittered delay, after which a tone sounded, a row of asterisks appeared, and the participant was given 75 s to attempt to recall any of the just-presented items.

In each distractor interval, participants solved math problems of the form $A + B + C = ?$, where A , B , and C were positive, single-digit integers, though the answer could be one or two digits. When a math problem was presented on the screen, the participant typed the sum as quickly as possible. The task was self-paced, such that a participant might have been presented with but not responded to a problem at the end of the distraction interval. Participants were given a monetary bonus based on the speed and accuracy of their responses. In the first two trials, participants performed free recall, with one trial having a distractor period following the last word presentation for 8 s. For the other of the first two trials, participants performed an 8-s distractor period prior to and following each word presentation. In the remaining 10 trials, participants performed free recall with five possible time durations for the between-item and end-of-list distractor tasks. As listed here, the first number indicates the between-list distractor duration and the second number indicates the end-of-list distractor, both in seconds: 0–0, 0–8, 0–16, 8–8, 16–16. A 0-s distractor refers to the typical, nonfilled duration intervals as described in the previous paragraph.

Participants who did not have valid recalls for all possible values of lag_1 in a particular condition were excluded from that condition only. To control for multiple comparisons, here we

consider the Bonferroni-corrected $p = .01$ threshold to indicate significance.

Results

Figure 3A shows the serial position curves for all five conditions: the two continual-distractor conditions (8 s or 16 s), the two delayed conditions (8 s or 16 s), and the immediate recall condition. Comparing the immediate and delayed conditions, the introduction of a distractor greatly reduces recall of recency but not primacy or midlist items. To quantify the recency effect, in each condition we calculated the correlation between the last three serial positions and the corresponding recall probabilities at those serial positions. Each of the recency correlations for the delayed conditions were reliably reduced from the immediate condition (both $p < .01$). However, participants do exhibit striking long-term recency in the continual-distractor conditions (Bjork & Whitten, 1974; Howard & Kahana, 1999), as the recency correlations were not reliably different between each of the continual-distractor conditions and the immediate condition (both $p > .05$). Nonetheless, primacy remained relatively constant across conditions. Analogous to our measure of the recency effect, we quantified the primacy effect as the correlation between the first three serial positions and the corresponding recall probabilities at those serial positions. Pairwise comparisons between each of the distractor conditions and the immediate condition revealed comparable levels of primacy (all $p > .2$). Consistent with Howard and Kahana (1999), participants exhibited strong long-range temporal contiguity, as seen in the lag-CRPs in both the delayed and continual-distractor conditions (Figure 3B).

In the continual-distractor conditions, participants performed arithmetic distractor tasks between each word presentation. A size or animacy encoding task helped to direct participants’ attention to the word during its presentation. These manipulations were expected to reduce rehearsal of prior items during and between presentations, thus providing a measure of compound cuing that is less susceptible to the influence of rehearsal processes. As shown

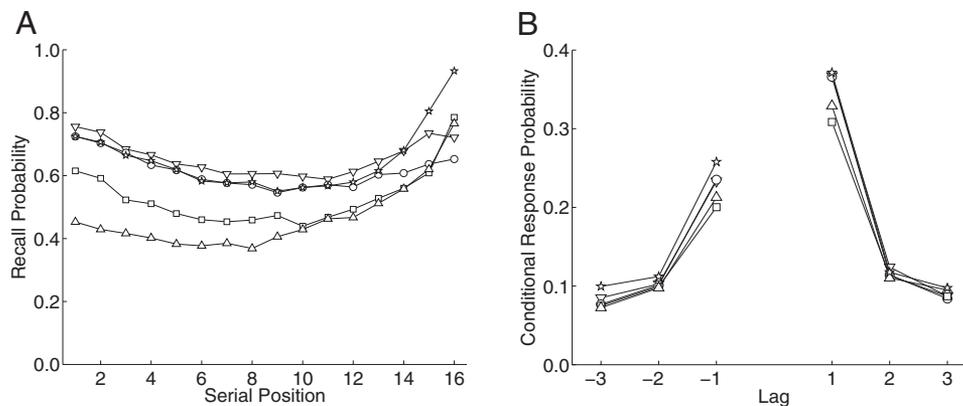


Figure 3. Recall probabilities in the experiment. In each panel, the five free recall conditions are plotted: continual-distractor 8 s (squares), continual-distractor 16 s (upward triangles), delayed 8 s (downward triangles), delayed 16 s (circles), immediate (stars). A. Serial position curves. B. Conditional response probabilities as a function of lag.

in the top two panels of Figure 4, participants exhibited strong forward compound cuing in the continual-distractor conditions: 8 s, $t(107) = 2.75$, $p < .01$; 16 s, $t(100) = 3.09$, $p < .005$. Replicating the major finding of our meta-analysis, the forward compound cue score is also reliable in the delayed conditions (middle panels of Figure 4): 8 s, $t(106) = 4.81$, $p < .0001$; 16 s, $t(105) = 4.67$, $p < .0001$, and in the immediate recall condition, $t(107) = 3.10$, $p < .005$. In addition, a t test between the 16-s conditions of delayed and continual-distractor free recall (presumably,

the two conditions most strongly attenuating rehearsal) revealed that the compound cue scores are not reliably different, $t(100) = 0.08$, $p > .5$, suggesting that the same underlying mechanism motivates compound cuing in the two conditions. For the backward compound cue score, the value was significant only in the continual-distractor 16-s condition, $t(100) = 2.78$, $p < .01$, although this distribution of values did not differ reliably from the distribution of backward compound cue scores in the delayed 16-s condition, $t(100) = 0.42$, $p > .5$.

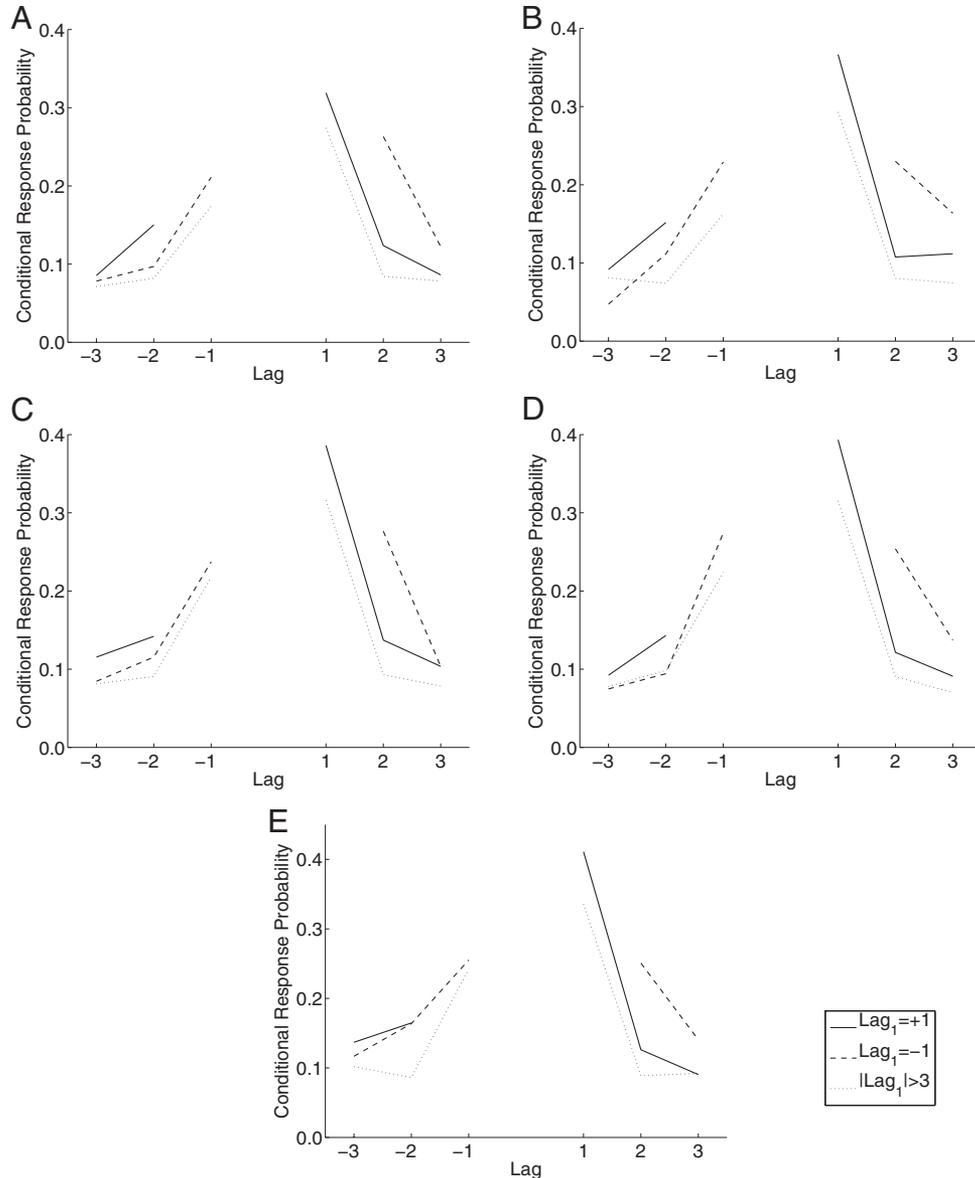


Figure 4. Compound cuing in the experiment. Within-subject comparisons revealed significant compound cuing across all five conditions of distraction intervals. In each panel, the lag-CRP is plotted separately for three cases of the lag of the previous transition, lag_1 : $\text{lag}_1 = +1$ (solid); $\text{lag}_1 = -1$ (dashed); $|\text{lag}_1| > 3$ (dotted). Top row: Continual-distractor free recall. A. Interstimulus and end-of-list distractor interval = 8 s. B. Interstimulus and end-of-list distractor interval = 16 s. Middle row: Delayed free recall. C. End-of-list distractor interval = 8 s. D. End-of-list distractor interval = 16 s. E. Immediate free recall. CRP = conditional response probability.

Simulation 1: Context Maintenance and Retrieval Model

The compound cuing results we observed in immediate, delayed and continual-distractor free recall are exactly what one would expect if the cue for recall included information not only from the just-recalled item but also from the sequence of prior recalls. This idea is embodied in retrieved context models of episodic memory, such as the temporal context model of Howard and Kahana (2002a) and its more recent variants (e.g., Gershman, Moore, Todd, Norman, & Sederberg, 2012; Howard, Kahana, & Wingfield, 2006; Polyn et al., 2009; Sederberg, Gershman, Polyn, & Norman, 2011; Sederberg et al., 2008). According to these models, the context cue used for recall of items contains a recency-weighted sum of previously activated cognitive states, and as such predict that the models predict that the temporal contiguity effect should be enhanced when a sequence of previously recalled items was studied at neighboring list positions. Here we examine whether a specific version of retrieved context theory—the context maintenance and retrieval model (CMR)—can predict our observed results at a quantitative level. After presenting simulations of the CMR model we also consider two recent variants of the SAM model that have been successful in explaining a wide range of free recall phenomena (Sirotin et al., 2005; Kimball et al., 2007).

CMR Model

According to CMR, the vector representation of the item presented at serial position i (denoted \mathbf{f}_i) and its corresponding context representation (denoted \mathbf{c}_i) interact through associative matrices (M^{FC} and M^{CF}) that are updated according to a standard Hebbian learning rule, such that $\Delta M^{FC} = \mathbf{c}_i \mathbf{f}_i^T$ and $\Delta M^{CF} = \mathbf{f}_i \mathbf{c}_i^T$. The context-to-item associative matrix, M^{CF} , is initialized to represent the semantic relations between items as determined using latent semantic analysis (LSA; Landauer & Dumais, 1997; see Polyn et al. 2009, for details). The item-to-context matrix, M^{FC} , is initialized to the identity matrix.

Context evolves according to the equation:

$$\mathbf{c}_i = \rho_i \mathbf{c}_{i-1} + \beta M^{FC} \mathbf{f}_i, \quad (1)$$

where ρ_i is defined such that $\|\mathbf{c}_i\| = 1$. Context is a weighted sum of the past several contextual states, and recent states dominate the representation. β is a model parameter that determines how much \mathbf{c}_i changes with each studied item.

To simulate delayed free recall, we assume that \mathbf{c} further evolves according to Equation 1 as participants engage in a distracting mental activity. In this case, we consider the distractor to be a single orthogonal item \mathbf{f}_D with the value of β_D determining the change in \mathbf{c} .

As a simplified implementation of the primacy effect, a primacy gradient ϕ_i scales the strengths of the item-to-context associations of early list items according to

$$\phi_i = \phi_s e^{-\phi_d(i-1)} + 1, \quad (2)$$

where ϕ_s and ϕ_d are model parameters.

At the time of recall, cuing with \mathbf{c}_i retrieves a vector $\mathbf{f}^{\text{N}} = M^{CF} \mathbf{c}_i$. To determine which item the model recalls, \mathbf{f}^{N} serves as the input to

a leaky, competitive accumulation process (Usher & McClelland, 2001), whose value at time step t is determined by

$$\begin{aligned} \mathbf{x}_t &= (1 - \tau\kappa - \tau\lambda\mathbf{N}) \mathbf{x}_{t-1} + \tau \mathbf{f}^{\text{N}} + \epsilon, \\ \epsilon &\sim \mathcal{N}(0, \eta), \\ \mathbf{x}_t &\rightarrow \max(\mathbf{x}_t, \mathbf{0}). \end{aligned} \quad (3)$$

This process runs iteratively until one of the accumulating elements crosses a threshold or until the recall period is over. Each element of \mathbf{x}_t corresponds to an element in \mathbf{f}^{N} . τ is a time constant, κ is a leak parameter, η is a noise parameter, and λ is a parameter that controls lateral inhibition, by scaling the strength of an inhibitory matrix \mathbf{N} that connects each accumulator to all of the others except itself.

The recall period is modeled as a series of competitions; each competition either produces a winning item or runs out of time, at which point the recall period is over and the next trial begins. When an item wins the recall competition, it is re-presented to the model, updating context according to Equation 1. The rate of context updating can differ between encoding and recall events (β_{enc} and β_{rec} , respectively). The updated state of context activates a different set of features on \mathbf{f}^{N} , and the recall competition begins again. Thus, the contextual cue for recall of the next item is a blend of the previous state of context and the input from the most recently presented item.

The foregoing description omits several model details that are not critical to explain our current findings, such as the parameters that govern the relative influence of preexperimental and experimental associations (for a full description, see Polyn et al., 2009).

Results

We fit the CMR model to the mean data from the 16-s delayed free recall condition of our experiment (Figure 4). Specifically, we were interested in whether CMR could produce the three compound lag-CRPs of Figure 4D, while maintaining accurate predictions of the standard lag-CRP that is not conditionalized on prior responses (Figure 3B, open circles), and the serial position curve (Figure 3A, open circles).

To find the best-fitting model parameters, we first performed a grid search of 10,000 parameter sets. Each parameter value was randomly selected on a uniform distribution of a predetermined range. Twenty generations of a genetic algorithm search were run from this initial grid search, where each successive generation took the most fit 20% of the previous generation and used these parent parameter sets to form 1,000 new parameter sets to simulate, by randomly repairing the parameters and adding random mutation to all values, using a normal distribution with mean zero and standard deviation 10% of the parameter range. The goodness-of-fit for a particular simulated data set was quantified using the root-mean-squared deviation between the model and the analyses described above.

CMR predicts compound lag-CRPs that are quite similar to those observed in all conditions of our experiment and in our meta-analysis. As shown in Figure 5A, transitions of $\text{lag}_2 = 1$ are more likely for the +1 compound lag-CRP than the control compound lag-CRP. In the case of the +1 compound lag-CRP, the two previously recalled items (r_1 and r_2) are associated with similar context representations. Because the context retrieval cue is most

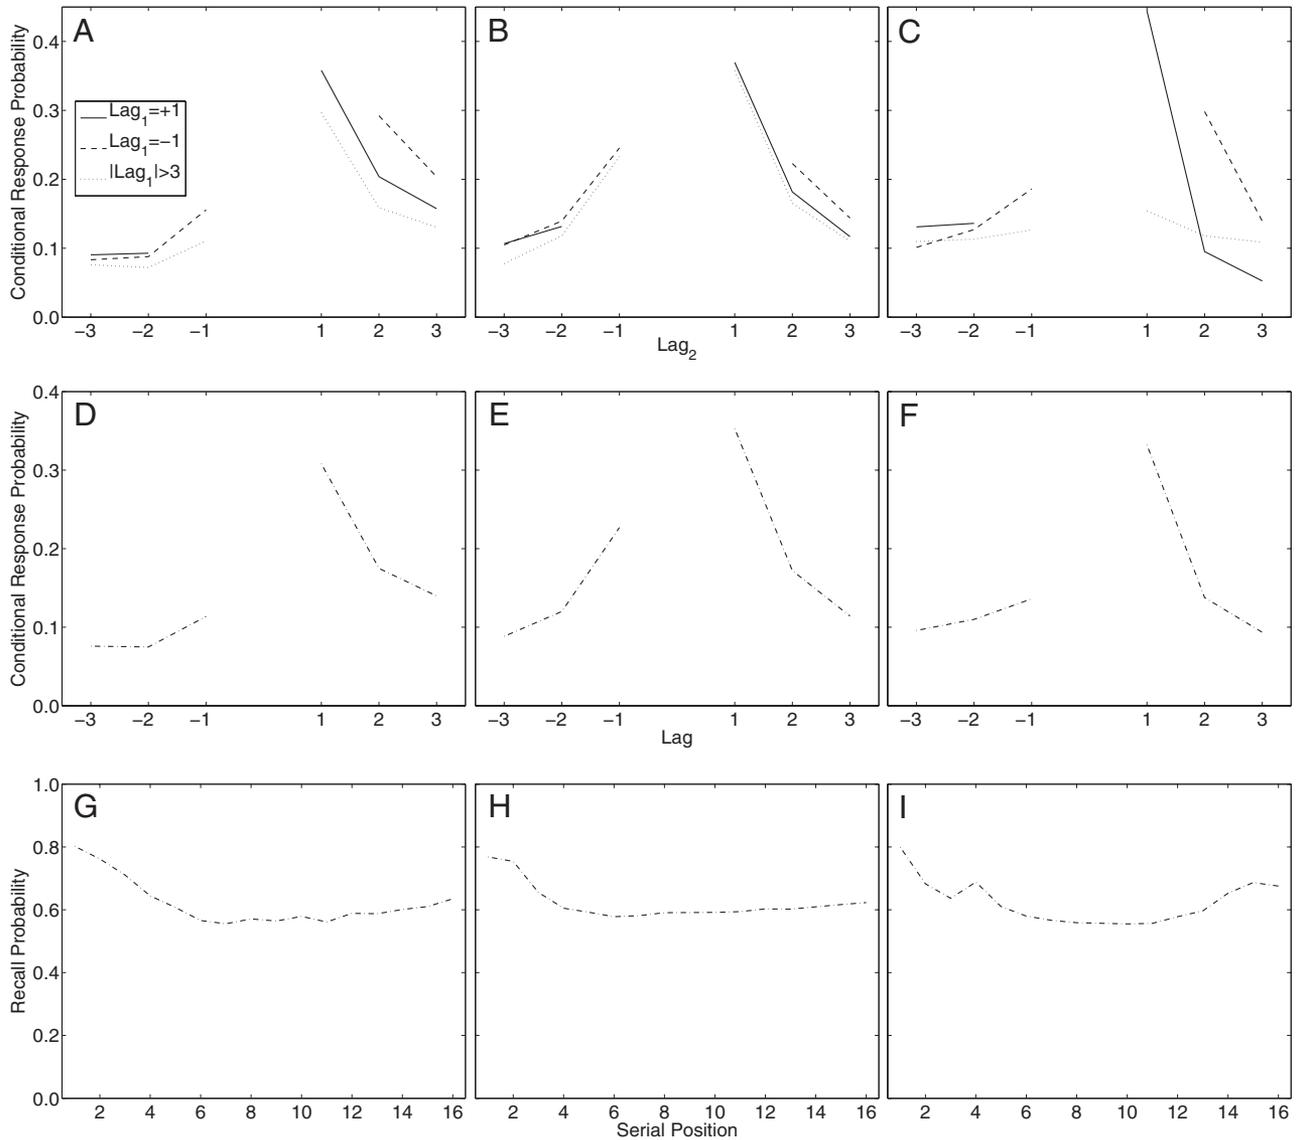


Figure 5. Model predictions of compound lag-CRPs, standard lag-CRPs, and serial position curves. The left column shows predictions of the context maintenance and retrieval model (CMR); the middle column shows predictions of the episodic search of associative memory model (eSAM); the right column shows predictions of a version of the search of associative memory model that assume a multiplicative rule for associations between items at encoding and retrieval (fSAM). A, B, C. The conditional response probability as a function of lag (lag-CRP) is plotted separately for three cases of the lag of the previous transition, lag_1 : $\text{lag}_1 = +1$ (solid); $\text{lag}_1 = -1$ (dashed); $|\text{lag}_1| > 3$ (dotted). D, E, F. Standard lag-CRP. G, H, I. Serial position curves. CRP = conditional response probability.

strongly influenced by the contexts associated with these two items, this more strongly encourages recall of items with contexts similar to r_2 , including its closest neighbor $\text{lag}_2 = +1$. The model parameters used in fitting these data are reported in Table 1.

Because our primary interest concerns each model's ability to fit the compound cue effect, here we report each model's goodness-of-fit based on the Akaike information criterion (AIC; Akaike, 1974) applied to the forward compound cue score. In comparison to the two models reported below, CMR's fit yields the lowest AIC

value, indicating the closest performance to the experimental data (AIC = 12).

Simulation 2: Search of Associative Memory

Although we have suggested that the context-integration mechanism in CMR is necessary to reproduce the compound cuing effect, the central role of this mechanism in the model makes it difficult to uniquely identify its role in generating the effect. As a

Table 1
Best-Fit Parameters of CMR

Parameter	Description	Value
β_{enc}	Context drift rate during encoding	.313
β_D	Context drift rate during distractor interval	.960
β_{rec}	Context drift rate during recall	.674
γ_{FC}	Forward asymmetry during encoding	.571
ϕ_s	Weight of primacy factor during encoding	.698
ϕ_d	Decay of primacy factor during encoding	.544
s	Retrieval weight for interitem semantic strength	2.98
κ	Decay weight in decision process	.100
λ	Lateral inhibition weight in decision process	.500
η	Standard deviation of noise in decision process	.218

Note. CMR = context maintenance and retrieval model.

comparison to the CMR model, we examine the compound cuing predictions of the eSAM model (Sirotnin et al., 2005). The eSAM model builds on the class of SAM models, which can account for a variety of memory findings in recall tasks (e.g., Gillund & Shiffrin, 1984; Raaijmakers, 2003; Raaijmakers & Shiffrin, 1980, 1981; Shiffrin, Ratcliff, & Clark, 1990). The eSAM model was introduced to account for numerous effects concerning the role of semantic associations in episodic recall tasks. As with previous versions of SAM, the strengths of associations in long-term episodic memory are governed by the operation of a short-term memory buffer and its control processes, as described below. Relevant to our current interests, the just-recalled item (along with a list-context representation) serve as the cue for the next response. Thus, if eSAM predicted the compound cue effects, this would imply that the context-integration mechanism of CMR is sufficient but not pivotal in explaining the compound cuing results.

eSAM

In eSAM there are two memory stores that contain item-to-item associations: a short-term memory (STM) buffer and long-term memory (LTM). LTM is composed of a semantic component (S_s) that represents the semantic associations between items (these values are set identically to those used for CMR, as specified in the previous section). LTM is also composed of two episodic components that store item-to-item associations (S_e) and item-to-context associations (S).

When item i is presented to the model, it enters a limited capacity STM, which contains four recently presented items except during presentation of the first three items (Raaijmakers & Shiffrin, 1980; Sirotnin et al., 2005). If the buffer is full, the newly presented item displaces the n^{th} buffer item with probability $\frac{q(1-q)^{n-1}}{1-(1-q)^{r_{cur}}}$, where r_{cur} refers to the number of items currently in the buffer and q is a free parameter that determines the degree of bias to displace older items versus displacing all items uniformly (Phillips, Shiffrin, & Atkinson, 1967).

Episodic interitem associations are initialized to 0.001. Forward and backward associations between any pair of items (i and j) that co-occupy the STM buffer are incremented, respectively, according to

$$S_e(i, j) = S_e(i, j) + \frac{b_{for}}{r_{cur}} \quad (4)$$

$$S_e(j, i) = S_e(j, i) + \frac{b_{bac}}{r_{cur}}$$

The association between items in STM and list context are similarly incremented as

$$S(i, c) = S(i, c) + \frac{a}{r_{cur}} \quad (5)$$

The weights of these associations are controlled by the free parameters a , b_{for} , b_{bac} .

In delayed free recall, recall begins with recalling items from long-term memory. An item i is recalled only if it is successfully sampled and recovered, which occur with probabilities P_s , P_r , respectively:

$$P_s(i|j, c) = \frac{S_s(i, j)^{W_s} S_e(i, j)^{W_e} S(i, c)^{W_c}}{\sum_{k=1}^{LL} S_s(k, j)^{W_s} S_e(k, j)^{W_e} S(k, c)^{W_c}} \quad (6)$$

$$P_r(i|j, c) = 1 - e^{-W_s S_s(i, j) - W_e S_e(i, j) - W_c S(i, c)}$$

where LL refers to list length, c refers to the current list context, and j refers to the previously recalled item. At the beginning of the recall period, only list context contributes to recall; specifically $W_e = W_s = 0$, and $W_c = 1$. For subsequent recalls that incorporate item information, W_e , W_s , W_c are free parameters.

A recalled item's associative strength to list context is incremented with the scalar parameter e_c , and item-to-item associations are also updated according to

$$S_e(j, i) = S_e(j, i) + f_{for} \quad (7)$$

$$S_e(i, j) = S_e(i, j) + f_{bac}$$

where f_{for} , f_{bac} are model parameters.

If an item-context cue is used L_{max} times but still fails to recover an item, then eSAM reverts to using only context as a cue. The recall process ends for a trial when K_{max} recovery failures have accumulated across all retrieval cues. K_{max} is a model parameter, and L_{max} is set such that $L_{max} = 0.1K_{max}$.

Results

Model parameters used to fit these data were determined with the same algorithm described in the CMR section and are reported in Table 2. As implemented here, the eSAM model does not predict any reliable differences in the compound lag-CRPs. That is, in eSAM recall of an item r_3 is minimally influenced by the recall r_1 (Figure 5B). This is most apparent in eSAM's inability to predict the forward compound cue effect: The transition of $lag_2 = +1$ is not markedly different conditional on $lag_1 = +1$ or $|lag_1| > 3$. Although the eSAM predicts that the difference between these two transition probabilities, termed the forward compound cue score, will be 0.01, in the experimental data the forward compound cue score is 0.08. Moreover, the goodness-of-fit of the eSAM model is less than that of the CMR model ($AIC = 17$).

Table 2
Best-Fit Parameters of SAM

Parameter	Description	eSAM	fSAM
b_{for}	Forward interitem strength, encoding	.687	.946
b_{bac}	Backward interitem strength, encoding	.527	.466
a	Item-context increment during encoding	.361	.300
f_{for}	Forward interitem strength during recall	.058	.731
f_{bac}	Backward interitem strength, recall	.528	0
e_c	Item-context increment during recall	.719	.035
W_e	Weight for interitem episodic strength	2.17	5.00
W_c	Weight for item-context strength	.010	4.12
W_s	Weight for interitem semantic strength	3.69	.350
K_{max}	Attempts after which recall stops	25	188
q	Bias to displace older buffer items	.165	.308
a_s	Semantic item-context increment, encoding		.410

Note. eSAM refers to the search of associative memory (SAM) model using classic item–item associations at encoding and retrieval. fSAM refers to the SAM model that assumes a multiplicative compound cue rule defines item–item associations formed during encoding and used during retrieval.

This result is not meant to imply that the SAM framework cannot account for these effects but rather that the most basic version of eSAM, which relies heavily on interitem associations, does not predict influences of compound cues. Whereas the CMR model, which embodies retrieved context theory, can explain the compound cue effects, the eSAM model that lacks explicit context representations cannot explain these effects. Thus, this simulation suggests that some nonstochastic process, which incorporates more than the just-recalled item into the retrieval cue, influences recall organization. In the next simulation we consider the fSAM model, a variant of the eSAM model that uses compound cues to guide retrieval.

fSAM

Kimball et al. (2007) introduced the fSAM model to generalize the SAM framework to explain the false memory paradigm. The major distinction of the fSAM model concerns item encoding and retrieval: whether each item is stored and retrieved individually, or whether all items in STM serve as a multiple-item cue. Kimball et al. (2007) compared performance of the fSAM model based on encoding and retrieval using single items, an additive combinatorial rule, or a multiplicative combinatorial rule. Here we consider only the version of the model that used the multiplicative multiple-item encoding and retrieval rules, as this was the only version that could account for behavioral performance in the false memory paradigm. Thus, our assumptions of the fSAM model are identical to the eSAM model, except for the following:

1. In addition to updating the item-context association of the just-presented item i , the item-context association for each list item k (including i) is updated according to

$$S(k, c) = S(k, c) + a_s \prod_j^{j \in M} S_s(k, j), \quad (8)$$

where a_s is a model parameter and M refers to the set of items currently in STM.

2. When an item-context cue is used during recall, item i is sampled with probability

$$P_s(i|j \in M, c) = \frac{(\prod_j^{j \in M} S_s(i, j))^{W_s} (\prod_j^{j \in M} S_e(i, j))^{W_e} S(i, c)^{W_c}}{\sum_{k=1}^{LL} (\prod_j^{j \in M} S_s(k, j))^{W_s} (\prod_j^{j \in M} S_e(k, j))^{W_e} S(k, c)^{W_c}}, \quad (9)$$

and recovered with probability

$$P_r(i|j \in M, c) = 1 - e^{-w_s \prod_j^{j \in M} S_s(i, j) - w_e \prod_j^{j \in M} S_e(i, j) - w_c S(i, c)}. \quad (10)$$

Results

The predictions of the fSAM concerning compound cuing are shown in Figure 5C. The fSAM model captures the compound cuing effect, in that a transition of $\text{lag}_2 = 1$ is much more likely following a transition of $\text{lag}_1 = +1$ than $|\text{lag}_1| > 3$. This suggests that the ability of the SAM framework to predict a forward compound cuing effect relies on the model's use of compound cues during encoding and retrieval. Consistent with the predictions of the CMR model, using a recency-weighted sum of recalled items to cue recall benefits successive recalls of items with $\text{lag}_1 = +1$. Although fSAM's goodness-of-fit is poorer than eSAM ($AIC = 21$), this reflects the fact that fSAM predicts a forward compound cue effect greater than the experimental data (fSAM = 0.21; data = 0.08) and thus captures qualitatively the main effect of interest.

Discussion

A model with a stochastic retrieval process assumes that recall of an item r_3 is only directly influenced by the prior recall r_2 . Our analysis of the dynamics of responses in free recall shows that people exhibit striking violations of the stochastic assumption. In a meta-analysis of five free recall studies, we found that participants were more likely to make a transition of $\text{lag}_2 = +1$ when the previous transition, lag_1 , was also $+1$, in comparison to when the previous transition was more remote, i.e., $|\text{lag}_1| > 3$ (Figure 2B). Participants were not only more likely to make such transitions but also made such transitions more quickly, as exhibited in their interresponse times (Figure 2D). We also reported this benefit in probability of recall in the forward compound cue in a new experiment that included a continual-distractor free recall condition (Figure 4). The presence of compound cuing in continual-distractor free recall (as well as in immediate and delayed free recall) suggests that compound cuing is not simply a product of recalling sequences of rehearsed items.

We also assessed the predictions of compound cuing in the context maintenance and retrieval model (CMR; Polyn et al., 2009). The CMR model assumes that context is a recency-weighted sum of presented and recalled items and uses this context to cue recall of the next item. The nature of this retrieval cue leads CMR to predict the patterns of compound cuing found in a delayed recall condition of the new experiment. A context-based framework such as CMR would also predict compound cuing in continual-distractor free recall for the same reason it predicts compound cuing in delayed free recall: The relative strengths of item representations in context will be highest for successively presented items.

In contrast, a variant of the search of associative memory model (eSAM; Sirotin et al., 2005) does not predict strong differences in lag_2 as a direct function of lag_1 . In eSAM, the retrieval cue is based on the

relatedness between r_2 and not-yet-recalled items; as such, recall of r_3 is only directly influenced by the prior recall r_2 . That eSAM cannot account for the compound cue advantage suggests that the retrieved context assumption, as embodied in CMR, may be necessary if not also sufficient to explain the compound cue result. In general, the retrieval process SAM framework is not restricted to be stochastic. Compound cues can be implemented in SAM by assuming that a set of recently recovered items contributes to the cue (e.g., Kimball et al., 2007; Ratcliff & McKoon, 1988, 1995). Here we considered a version of the fSAM model (Kimball et al., 2007) that shares most assumptions of the eSAM model, except that items are associated via a multiplicative rule at encoding and retrieval rather than only being directly associated with one another. This version of the SAM model did predict the advantage of the forward compound cue.

The critical distinction between CMR and fSAM, which predict compound cuing, and eSAM, which does not, is the retrieval cue. Even if context associations are formed at encoding, our results suggest that such context representations must be incorporated into the retrieval cue to produce compound cue effects. Consistent with this idea, Kimball et al. (2007) found that fSAM could only account for false-memory paradigm effects when fSAM assumed that compound associations were used both at encoding and retrieval. An open question concerns whether models that assume chunking or grouping processes at encoding but not at recall (e.g., Farrell, 2012) can produce the forward compound cue effect.

Nonetheless, we do not wish to make strong arguments defining the set of models that can or cannot capture the compound cue effect; our primary focus is to provide a set of novel findings as predicted by retrieved context models such as CMR. This model assumes that context is a recency-weighted sum of recalled items and thus assumes that r_3 is more strongly influenced by r_2 than r_1 . We found this to be the case in our experimental data as well: Although r_1 and r_2 form a compound cue for retrieval of r_3 , the influence of r_2 is clearly stronger than that of r_1 . In the extreme case where r_1 and r_2 equally informed the next recall r_3 , we would not expect a meaningful interpretation of r_3 when only considering r_2 . Yet previous analyses have shown that a large proportion of the variance in transitions can be explained when considering only the previous transition (Kahana, 1996). The stronger influence of r_2 can be intuited from the retrieved context account, as in the simulations of CMR. Because context is a recency-weighted sum of recalled items, r_2 is more strongly represented in the context used to cue recall than r_1 .

The reported compound cue advantage in free recall can be used to understand the part list cuing effect. In the part list cuing paradigm, participants are provided with a subset of presented items as cues immediately preceding the recall period. In initial studies, performance was worse than standard free recall (Slamecka, 1968). If participants benefit from a double or triple item cue over a single item cue in serial recall, it may seem counterintuitive that presentation of a subset of the list hinders recall. However, in part list cuing, the cues may not be the most beneficial to the participant. Indeed, studies in which the items comprised a randomly selected subset of presented items hindered recall (Roediger, 1973; Slamecka, 1968), but when cued stimuli are re-presented in their original presentation order, this can attenuate or eliminate the inhibitory effect (Allen, 1969; Basden, 1973; Sloman, Bower, & Rohrer, 1991). Whereas in free recall the units of organization imposed by the participant can be unclear, in serial recall the optimal units are implicit in the instruction to recall items in order. Thus, the fact that participants benefit from being

presented with compound rather than single cues in serial recall is also consistent with these ideas (Kahana & Caplan, 2002; Posnansky, 1972).

Our reported free recall results seem relatively consistent with the results in serial recall. In serial learning the benefit of compound cues is greater in the forward than in the backward direction (Kahana & Caplan, 2002). In our free recall meta-analysis and experiment, we always found a benefit of forward compound cues; in only one condition of our experiment did we find a benefit of a backward compound cue. Although in free recall participants may recall list items in any order, they nonetheless exhibit striking asymmetry in their transitions in free recall (Kahana, Howard, & Polyn, 2008); thus, it is not surprising that the compound cue advantage is more reliable in the forward direction.

Our finding of compound cuing demonstrates that intrainit transitions in free recall are not stochastic. These results are consistent with retrieved context theory, which assumes that the retrieval cue is a recency-weighted sum of prior recalls.

References

- Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, *19*, 716–723. doi:10.1109/TAC.1974.1100705
- Allen, M. M. (1969). Cueing and retrieval in free recall. *Journal of Experimental Psychology*, *81*, 29–35. doi:10.1037/h0027460
- Anderson, J. A. (1972). A simple neural network generating an interactive memory. *Mathematical Biosciences*, *14*, 197–220. doi:10.1016/0025-5564(72)90075-2
- Basden, D. R. (1973). Cued and uncued free recall of unrelated words following interpolated learning. *Journal of Experimental Psychology*, *98*, 429–431. doi:10.1037/h0034367
- Bjork, R. A., & Whitten, W. B. (1974). Recency-sensitive retrieval processes in long-term free recall. *Cognitive Psychology*, *6*, 173–189. doi:10.1016/0010-0285(74)90009-7
- Chance, F. S., & Kahana, M. J. (1997). Testing the role of associative interference and compound cues in sequence memory. In J. Bower (Ed.), *Computational neuroscience: Trends in research* (pp. 599–603). New York, NY: Plenum Press. doi:10.1007/978-1-4757-9800-5_93
- Clark, S. E., & Shiffrin, R. M. (1987). Recognition of multiple-item probes. *Memory & Cognition*, *15*, 367–378. doi:10.3758/BF03197727
- Cohen, N. R., & Sekuler, R. (2010). Chunking and compound cueing of movement sequences: Learning, retention and transfer. *Perceptual and Motor Skills*, *110*, 736–750. doi:10.2466/pms.110.3.736-750
- Davelaar, E. J., Goshen-Gottstein, Y., Ashkenazi, A., Haarmann, H. J., & Usher, M. (2005). The demise of short-term memory revisited: Empirical and computational investigations of recency effects. *Psychological Review*, *112*, 3–42. doi:10.1037/0033-295X.112.1.3
- Doshier, B., & Rosedale, G. (1989). Integrated retrieval cues as a mechanism for priming in retrieval from memory. *Journal of Experimental Psychology: General*, *118*, 191–211. doi:10.1037/0096-3445.118.2.191
- Farrell, S. (2012). Temporal clustering and sequencing in short-term memory and episodic memory. *Psychological Review*, *119*, 223–271. doi:10.1037/a0027371
- Feigenbaum, E. A., & Simon, H. A. (1962). A theory of the serial position effect. *British Journal of Psychology*, *53*, 307–320. doi:10.1111/j.2044-8295.1962.tb00836.x
- Gershman, S. J., Moore, C. D., Todd, M. T., Norman, K. A., & Sederberg, P. B. (2012). The successor representation and temporal context. *Neural Computation*, *24*, 1553–1568. doi:10.1162/NECO_a_00282
- Gillund, G., & Shiffrin, R. M. (1984). A retrieval model for both recognition and recall. *Psychological Review*, *91*, 1–67. doi:10.1037/0033-295X.91.1.1

- 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. doi:10.1037/0278-7393.25.4.923
- Howard, M. W., & Kahana, M. J. (2002a). A distributed representation of temporal context. *Journal of Mathematical Psychology*, *46*, 269–299. doi:10.1006/jmps.2001.1388
- Howard, M. W., & Kahana, M. J. (2002b). When does semantic similarity help episodic retrieval? *Journal of Memory and Language*, *46*, 85–98. doi:10.1006/jmla.2001.2798
- 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. doi:10.3758/BF03193867
- Howard, M. W., Youker, T. E., & Venkatadass, V. (2008). The persistence of memory: Contiguity effects across hundreds of seconds. *Psychonomic Bulletin & Review*, *15*, 58–63. doi:10.3758/PBR.15.1.58
- Kahana, M. J. (1996). Associative retrieval processes in free recall. *Memory & Cognition*, *24*, 103–109. doi:10.3758/BF03197276
- Kahana, M. J., & Caplan, J. B. (2002). Associative asymmetry in probed recall of serial lists. *Memory & Cognition*, *30*, 841–849. doi:10.3758/BF03195770
- Kahana, M. J., Howard, M. W., & Polyn, S. M. (2008). Associative retrieval processes in episodic memory. In H. L. Roediger III (Ed.), *Learning and memory: A comprehensive reference: Vol. 2. Cognitive psychology of memory*. Oxford, England: Elsevier.
- Kimball, D. R., Smith, T. A., & Kahana, M. J. (2007). The fSAM model of false recall. *Psychological Review*, *114*, 954–993. doi:10.1037/0033-295X.114.4.954
- Laming, D. (2010). Serial position curves in free recall. *Psychological Review*, *117*, 93–133. doi:10.1037/a0017839
- Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, *104*, 211–240. doi:10.1037/0033-295X.104.2.211
- Lewandowsky, S., & Murdock, B. B. (1989). Memory for serial order. *Psychological Review*, *96*, 25–57. doi:10.1037/0033-295X.96.1.25
- Lohnas, L. J., Polyn, S. M., & Kahana, M. J. (2011). Contextual variability in free recall. *Journal of Memory and Language*, *64*, 249–255. doi:10.1016/j.jml.2010.11.003
- Moreton, B. J., & Ward, G. (2010). Time scale similarity and long-term memory for autobiographical events. *Psychonomic Bulletin & Review*, *17*, 510–515. doi:10.3758/PBR.17.4.510
- Murdock, B. B. (1962). The serial position effect of free recall. *Journal of Experimental Psychology*, *64*, 482–488. doi:10.1037/h0045106
- Phillips, J. L., Shiffrin, R. J., & Atkinson, R. C. (1967). The effects of list length on short-term memory. *Journal of Verbal Learning & Verbal Behavior*, *6*, 303–311. doi:10.1016/S0022-5371(67)80117-8
- Polyn, S. M., Norman, K. A., & Kahana, M. J. (2009). A context maintenance and retrieval model of organizational processes in free recall. *Psychological Review*, *116*, 129–156. doi:10.1037/a0014420
- Posnansky, C. J. (1972). Probing for the functional stimuli in serial learning. *Journal of Experimental Psychology*, *96*, 184–193. doi:10.1037/h0033503
- Raaijmakers, J. G. W. (2003). Spacing and repetition effects in human memory: Application of the SAM model. *Cognitive Science*, *27*, 431–452. doi:10.1207/s15516709cog2703_5
- 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, pp. 207–262). New York, NY: Academic Press.
- Raaijmakers, J. G. W., & Shiffrin, R. M. (1981). Search of associative memory. *Psychological Review*, *88*, 93–134. doi:10.1037/0033-295X.88.2.93
- Ratcliff, R., & McKoon, G. (1988). A retrieval theory of priming in memory. *Psychological Review*, *95*, 385–408. doi:10.1037/0033-295X.95.3.385
- Ratcliff, R., & McKoon, G. (1995). Sequential effects in lexical decision: Tests of compound-cue retrieval theory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 1380–1388. doi:10.1037/0278-7393.21.5.1380
- Roediger, H. L. (1973). Inhibition in recall from cueing with recall targets. *Journal of Verbal Learning & Verbal Behavior*, *12*, 644–657. doi:10.1016/S0022-5371(73)80044-1
- Rundus, D. (1971). An analysis of rehearsal processes in free recall. *Journal of Experimental Psychology*, *89*, 63–77. doi:10.1037/h0031185
- Schwartz, G., Howard, M. W., Jing, B., & Kahana, M. J. (2005). Shadows of the past: Temporal retrieval effects in recognition memory. *Psychological Science*, *16*, 898–904. doi:10.1111/j.1467-9280.2005.01634.x
- Sederberg, P. B., Gauthier, L. V., Terushkin, V., Miller, J. F., Barnathan, J. A., & Kahana, M. J. (2006). Oscillatory correlates of the primacy effect in episodic memory. *NeuroImage*, *32*, 1422–1431. doi:10.1016/j.neuroimage.2006.04.223
- Sederberg, P. B., Gershman, S. J., Polyn, S. M., & Norman, K. A. (2011). Human memory consolidation can be explained using the temporal context model. *Psychonomic Bulletin & Review*, *18*, 455–468. doi:10.3758/s13423-011-0086-9
- Sederberg, P. B., Howard, M. W., & Kahana, M. J. (2008). A context-based theory of recency and contiguity in free recall. *Psychological Review*, *115*, 893–912. doi:10.1037/a0013396
- Sederberg, P. B., Miller, J. F., Howard, W. H., & Kahana, M. J. (2010). The temporal contiguity effect predicts episodic memory performance. *Memory & Cognition*, *38*, 689–699. doi:10.3758/MC.38.6.689
- Shiffrin, R. M., Ratcliff, R., & Clark, S. E. (1990). List-strength effect: II. Theoretical mechanisms. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *16*, 179–195. doi:10.1037/0278-7393.16.2.179
- Sirotny, 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. doi:10.3758/BF03196773
- Slamecka, N. J. (1968). An examination of trace storage in free recall. *Journal of Experimental Psychology*, *76*, 504–513. doi:10.1037/h0025695
- Sloman, S. A., Bower, G. H., & Rohrer, D. (1991). Congruency effects in part-list cuing inhibition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *17*, 974–982. doi:10.1037/0278-7393.17.5.974
- Steyvers, M., Shiffrin, R. M., & Nelson, D. L. (2004). Word association spaces for predicting semantic similarity effects in episodic memory. In A. F. Healy (Ed.), *Cognitive psychology and its applications: Festschrift in honor of Lyle Bourne, Walter Kintsch, and Thomas Landauer* (pp. 237–249). Washington, DC: American Psychological Association.
- Tan, L., & Ward, G. (2000). A recency-based account of the primacy effect in free recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *26*, 1589–1626. doi:10.1037/0278-7393.26.6.1589
- Thorndike, E. L., & Lorge, I. (1944). *The teacher's word book of 30,000 words*. New York, NY: Teachers College.
- Unsworth, N. (2008). Exploring the retrieval dynamics of delayed and final free recall: Further evidence for temporal-contextual search. *Journal of Memory and Language*, *59*, 223–236. doi:10.1016/j.jml.2008.04.002
- Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, *108*, 550–592. doi:10.1037/0033-295X.108.3.550
- Wickelgren, W. A. (1966). Associative intrusions in short-term recall. *Journal of Experimental Psychology*, *72*, 853–858. doi:10.1037/h0023884

Received August 17, 2012

Revision received March 21, 2013

Accepted May 19, 2013 ■