

The persistence of memory: Contiguity effects across hundreds of seconds

MARC W. HOWARD, TESS E. YOUKER, AND VIJAY S. VENKATADASS
Syracuse University, Syracuse, New York

A contiguity effect—the finding that stimuli that occur close together in time become associated to each other—is observed between words that are separated by several seconds. The traditional account of contiguity effects is that item representations become associated to each other while active in a *short-term memory buffer*—a limited-capacity store that can hold a small, integral number of items. Participants studied and free recalled 48 lists of words. At the end of the session, participants were given a surprise final free recall test on all of the items from all of the lists. In addition to a standard contiguity effect between items presented at nearby serial positions, we simultaneously observed a contiguity effect between items presented in different lists. This latter contiguity effect extended over several lists, or several hundred seconds, well beyond the range that can be attributed to a buffer holding a small, integral number of items.

Short-term memory buffers (Atkinson & Shiffrin, 1968; Davelaar, Goshen-Gottstein, Ashkenazi, Haarmann, & Usher, 2005; Raaijmakers & Shiffrin, 1980) provide a natural account of the *immediate recency effect*—the finding that items from the end of the list are better recalled than other items. Because the contents of the buffer are rapidly displaced by incoming information, buffer models predict immediate recency because items from the end of the list are more likely to remain activated in the buffer (Atkinson & Shiffrin, 1968; Davelaar et al., 2005; Raaijmakers & Shiffrin, 1980). Buffer models also naturally account for the *contiguity effect*—the finding that items that are presented close together in time become associated to one another. In free recall, participants recall the words in the list in whatever order they come to mind. The contiguity effect is manifest in free recall as an excess of recall transitions between words presented close together in the list (Kahana, 1996)—an operationalization of the strength of association. Buffer models predict contiguity effects if connections in long-term memory are built up between items that are simultaneously active in the buffer (Kahana, 1996; Raaijmakers & Shiffrin, 1980; Sirotin, Kimball, & Kahana, 2005). If a contiguity effect is caused by a short-term memory buffer, then the range over which it is observed constrains the range over which traces must remain active in the buffer, and hence, its capacity.

Previous studies have examined recency effects across a variety of time scales (e.g., Bjork & Whitten, 1974; Glenberg et al., 1980; Howard & Kahana, 1999), suggesting to some that a single scale-invariant mechanism accounts for recency effects (e.g., Brown, Neath, & Chater, 2007; Neath & Brown, 2006). There is considerably less information pertaining to contiguity effects across long time scales, although Howard and Kahana (1999) observed contiguity effects in continuous-distractor free recall when the list

items were separated by a distractor-filled delay. The goal of the present study is to examine recency and contiguity at multiple scales.

In this study, we presented participants with multiple lists of words for an immediate free recall test. After study and recall of 48 lists, subjects were given a surprise final free recall (FFR) test. We will evaluate recency at multiple scales by comparing within-list recency on the immediate free recall test with across-list recency on the FFR test. Moreover, during FFR, we can evaluate contiguity effects across time scales simultaneously by comparing transitions between items from the same list with transitions across different lists.

METHOD

Lists of 10 items were presented one at a time. Following each list, participants performed an immediate free recall test of the most recent list. After 48 lists were studied and recalled, participants were presented with a surprise FFR test in which they were instructed to remember all of the words from all of the lists in any order they came to mind.

Participants

A total of 294 participants participated for course credit in an introductory psychology class at Syracuse University.

Materials

Study lists consisted of 10 words from the noun subset of the Toronto word pool (Friendly, Franklin, Hoffman, & Rubin, 1982).

Procedure

Lists were presented both auditorially and visually under conditions that were designed to discourage rehearsal. For each study word, participants were required to press a key in order to indicate whether the word was concrete or abstract. If participants did not respond within 1,200 msec the computer would buzz and advance to the next word. The next word followed after a delay of 500 msec.

M. W. Howard, marc@memory.syr.edu

The screen remained blank for 500 msec after the last item. Next, a row of asterisks was displayed for 400 msec. Simultaneously, an auditory signal instructed participants to free recall the list. Participants were given 30 sec for verbal free recall. The total time between the initiation of one list and the initiation of the next, which included an opportunity to rest between lists, was 49 ± 6 sec. Half of the study lists contained a repeated item. The other half of the lists were control lists, without any repeated items. The effect of repetition on the immediate free recall data is described elsewhere (Howard, Venkatadass, Norman, & Kahana, 2007). In this report, lists with repeated items are omitted from all analyses that rely on serial position within a list.

At the very end of the session, participants were given 5 min for FFR. The delay between the completion of the final list and the beginning of the FFR period included an opportunity for a brief break, the time for the participant to notify the experimenter that he or she was finished with the main experiment, and verbal instructions for the FFR session. The total length of this delay was 250 ± 40 sec.

CRP Analyses

In order to measure contiguity effects within and across lists, we calculated conditional response probability functions (CRPs; Kahana, 1996; Kahana, Howard, Zaromb, & Wingfield, 2002). The CRP measures the probability of a recall transition, taking into account the possible transitions between correct items. We computed a within-list CRP in order to measure the contiguity effect between items in the same list, and an across-list CRP in order to measure the contiguity effect between items from different lists. Contiguity within list can be measured by calculating within-list lag, the difference between the serial positions of successive recalls from the same list. In calculating the within-list CRP, we only considered transitions between correct recalls in which both items were from the same list (and that list did not include repeated items). Following previous work, this calculation controls for the number of recall transitions that would be possible at a given recall attempt. For instance, if the 9th word in our 10-item list had just been recalled, then a lag of +2 would not be possible. A numerator and a denominator were retained for each within-list lag. The numerator was incremented when a transition of the appropriate lag was observed. For each observed transition, the denominator was incremented for each possible lag that would have been a correct (nonrepeated) recall. We only included participants in these analyses who had recalled at least 10 across-list FFR transitions.

We also generalized the CRP to describe associations across lists using across-list lag. If a participant recalled a word from List 23, and if the next recalled word was presented in List 31, this recall transition would be associated with an across-list lag of +8. Similarly, if the subject recalled a word from List 17 followed by a recall from List 12, this would be across-list lag -5. The calculation of the across-list CRP followed the methods of the within-list CRP, with a few exceptions. Rather than only including pairs of final free recalls from the same study list, only pairs that came from different study lists were included. In addition, we did not attempt to control for the number of items that would be available from a particular list at each output position.

Surrogate Data Set

In order to control for nonassociative tendencies that could contribute to output order in FFR, we constructed a surrogate data set from the participants' FFR protocols. We started with the set of pairs of correct FFRs that each participant contributed to the across-list CRP analysis. We then shuffled the first and second members of these pairs so that no items from the same list were paired with each other and that the original pairs were not reconstructed. In the surrogate data set, any causal relationship between the sequence of recalled words would be disrupted while leaving the marginal probability of recall unaffected.

RESULTS AND DISCUSSION

Figure 1 shows the recency effect during immediate free recall (panels A and B) and across lists during FFR

testing (panels C and D). Figure 1A shows the standard serial position curve—probability of recall as a function of serial position during immediate testing. The lack of a substantial primacy effect in Figure 1A is consistent with our intent to disrupt rehearsal.¹ Figure 1B shows the probability that the first word a participant recalled during the immediate recall period came from each position within the list. The large recency effect extending several items at the end of the list is consistent with previous work on immediate free recall (see, e.g., Murdock, 1962). If a buffer model is responsible for the recency effect, then the duration of the recency effect allows us to estimate buffer capacity. The capacity of the buffer in immediate free recall appears to be at most a few items. Figure 1C shows the probability that each item is final free recalled as a function of the number of the list in which it was presented; Figure 1D shows the probability that the first word recalled in the FFR session came from each list. As can be seen from both measures, there is a recency effect that extends over several lists—a *long-term recency effect* (Bjork & Whitten, 1974; Glenberg et al., 1980). Each list was separated by approximately 50 sec; thus, the five to eight lists that constitute the across-list recency effect correspond to tens of items and a delay of several hundred seconds—well beyond the capacity that we would estimate from immediate free recall. This result suggests either that separate mechanisms account for immediate and long-term recency (Davelaar et al., 2005; Raaijmakers, 1993) or that a common mechanism other than a fixed-capacity short-term memory buffer is responsible for recency effects over both time scales (Brown et al., 2007; Howard & Kahana, 2002).

Although the long-term recency effect has been well documented, less is known about contiguity effects over long time scales. The results of the within-list CRP analysis of the FFR data are shown in Figure 2B. The fact that the curve in Figure 2B peaks around zero demonstrates a contiguity effect—given that a participant just recalled a word, the next word recalled will tend to be from nearby positions within the list. This finding can be taken as evidence that an association was formed between nearby items in the list.² Even though the data in Figure 2B were collected from FFR rather than from recall of a single list, they are very similar to previous findings from single-trial free recall (Kahana, 1996; Kahana et al., 2002). The within-list contiguity effect extended for a few list positions, as did the immediate recency effect.

In this experiment, however, we also measured a contiguity effect across lists by examining transitions between words from different lists (see Figure 2A). Figure 2C shows the results of an across-list CRP analysis (see Method for details). One thing we can note from Figure 2C is that the across-list CRP appears to increase with higher values of across-list lag (toward the right side of the figure). This increase is a manifestation of a recency effect; transitions with large values of across-list lag tend to be transitions to more recent lists. In addition, there appears to be a peak near an across-list lag of zero, indicating an advantage for transitions to nearby lists, and suggesting a contiguity effect across lists.

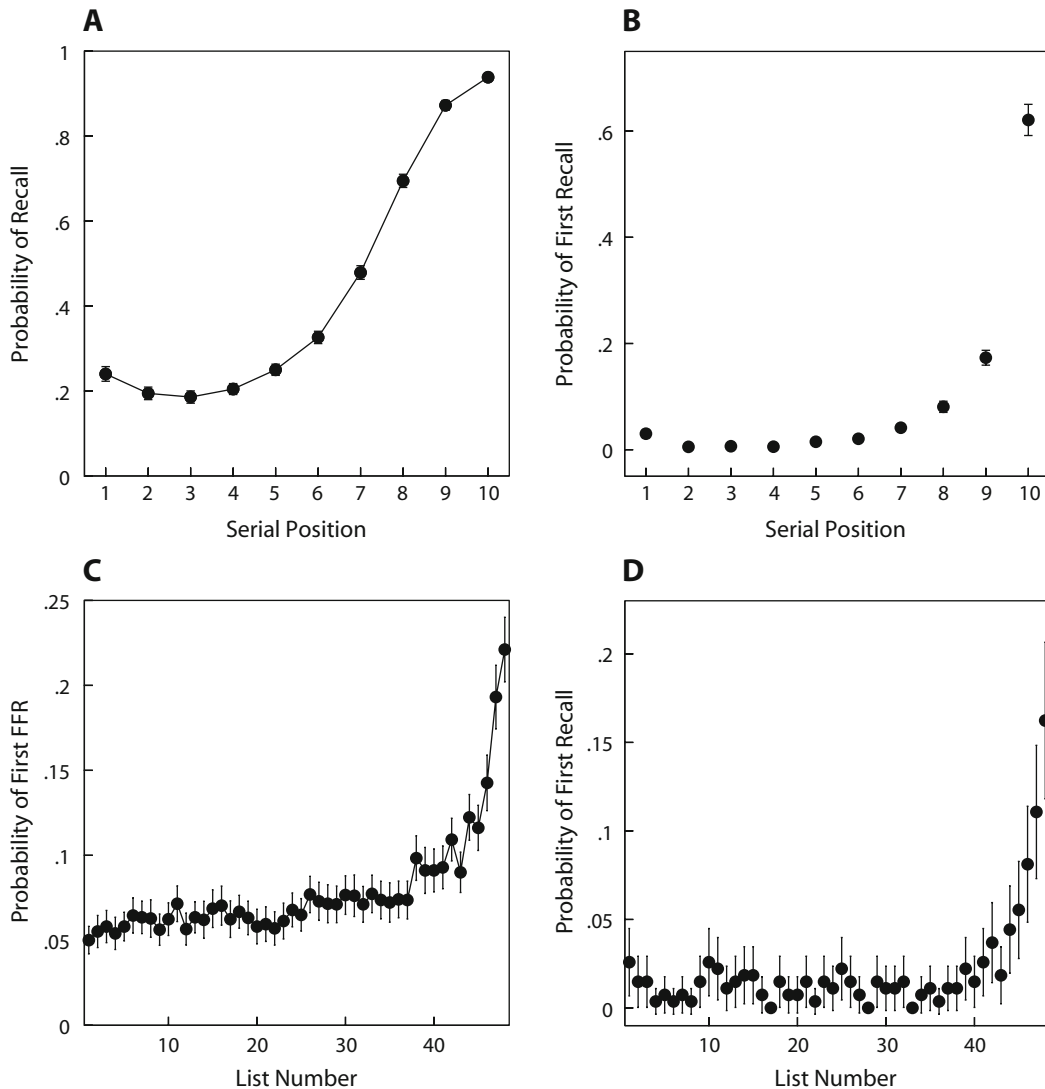


Figure 1. The recency effect across time scales. (A) Serial position curve from immediate free recall of control lists. (B) Probability of initiating immediate free recall as a function of position within the list. (C) Probability of final free recall as a function of list number. (D) Probability of initiating final free recall as a function of list number. Error bars reflect 95% confidence intervals.

There are several artifactual explanations of the apparent across-list contiguity effect that need to be ruled out before accepting the solid symbols in Figure 2C as evidence for genuine temporally defined association across lists. If the recency effect persists across all output positions, then both members of any pair of recalls will tend to be from the end of the list. In this case, we would expect to see an advantage for small across-list lags, despite a lack of real association between items. Another possibility is that correlated encoding across nearby lists accounts for the peak in Figure 2C. Imagine that a participant attends to Lists 20–25 and ignores all the others. In this case, we would observe small absolute values of across-list lags in the absence of any genuine associations between items. In order to statistically control for these possibilities, we generated a surrogate data set in which each pair of recalls that entered into the across-list CRP analysis was randomly shuffled (see also Method). We

constructed 10,000 shuffles and reexpressed the observed values of across-list CRP as a z score calculated with respect to the distribution of values obtained from the surrogate data set. The surrogate data set preserves the marginal probability of recall across lists (and, in fact, all other conceivable stimulus variables) for each participant, but disrupts any causal relationship between the pair of items that were actually recalled. The discrepancy must be attributable to correlated retrieval processes—or an association—between the recalled words. The gray line in Figure 2C provides the mean value across shuffles; the light gray region provides the mean, plus or minus one standard deviation.

Figure 2D shows the z score of the actual across-list CRP calculated with respect to the distribution of values generated from the surrogate data set. Because this analysis controls for nonassociative factors—such as the recency effect and other constant stimulus selection fac-

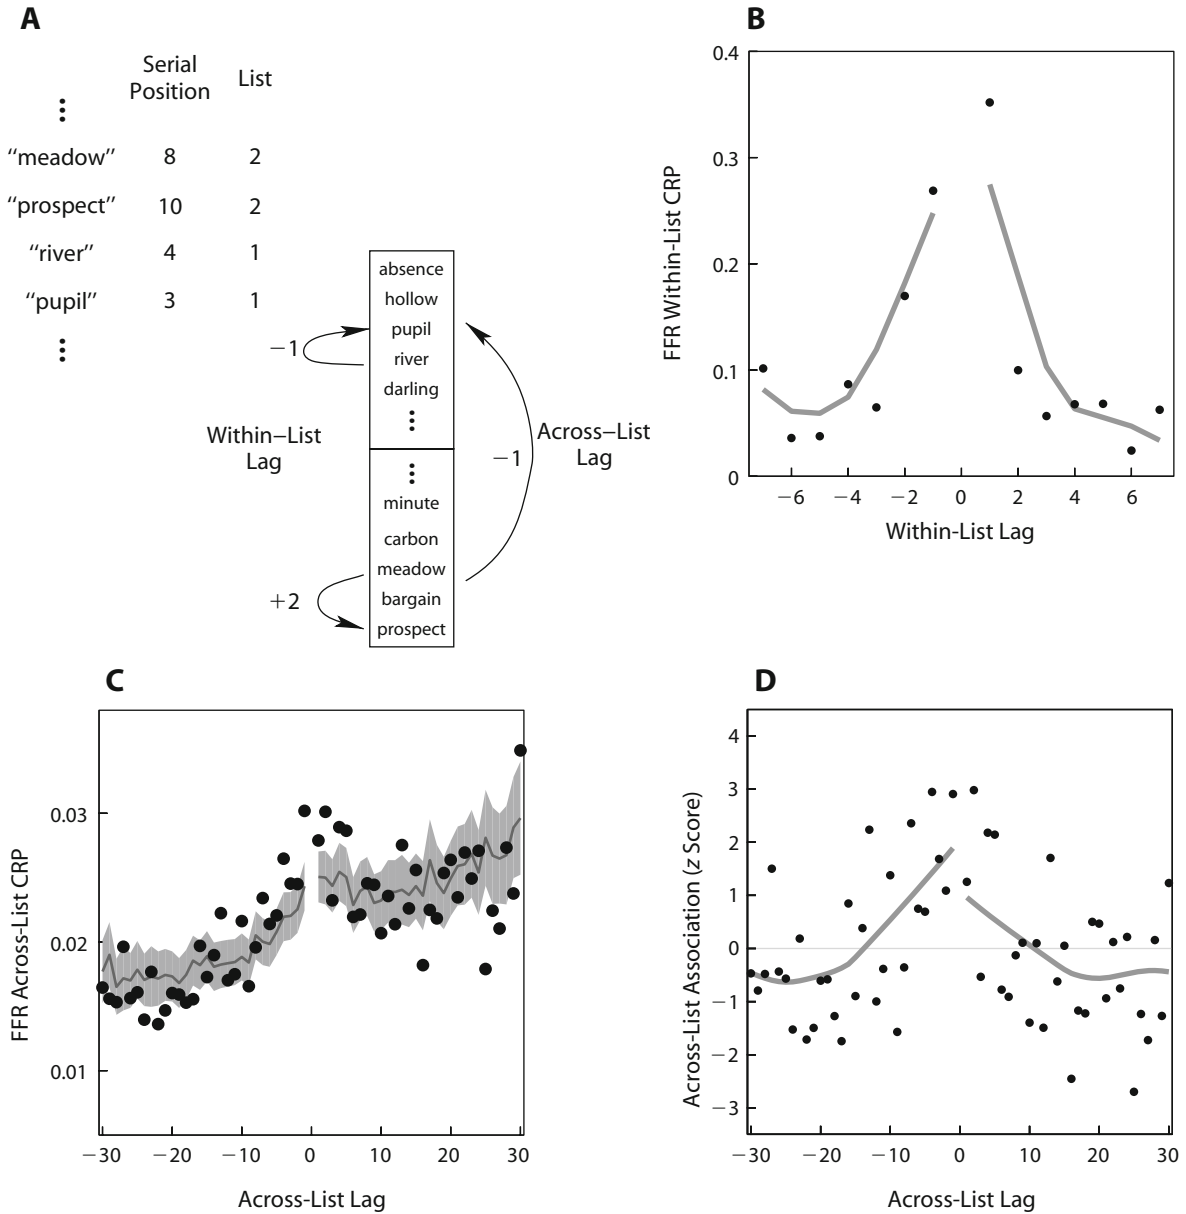


Figure 2. The contiguity effect is observed simultaneously across time scales. (A) By examining the order in which words are recalled in the final free recall (FFR) test, we can measure temporally defined associations both within a single list and across lists. For transitions in which the first and second recalled member are from the same list, we define within-list lag as the difference in their positions in the original list. Analogously, when a recall transition involves items from different lists, we can assess temporally defined associations by measuring the across-list lag of the transition. (B) Conditional probability of FFR transitions between members of the same list as a function of the within-list lag between their original presentations. Smooth curves come from a LOWESS fit to the data. (C) Conditional probability of FFR transitions between members of different lists as a function of across-list lag. The dark gray curve shows the probability generated from a simulated data set in which the pairs of recalled items were shuffled. The light gray region gives the mean \pm one standard deviation. (D) The across-list CRP in (C) converted to a z score calculated with respect to the distribution of values obtained from the surrogate data set. Smooth curves come from a LOWESS fit to the data.

tors—we will refer to the z score as a measure of across-list association between items. Examination of Figure 2D suggests that the strength of the association decreases in both the forward and backward directions over about 10 lists,³ which corresponds to a separation of approximately 100 other items and about 8 min of time.

We quantified the range of across-list associative strength using a variety of methods. We will only report the most conservative of these here. We divided the across-list lag values into different zones. If the strength of associations—that is, the average value of the associative z score shown in Figure 2D—decreases from one zone to

another, then the range of association must extend at least to the start of the more nearly adjacent zone. Our first zone consisted of the central 10 values of across-list lag—those with an absolute value ≤ 5 . The second zone consisted of the next 20 points, with absolute values of across-list lag being >5 and ≤ 15 . The third zone consisted of the next 30 values, with absolute value being >15 but ≤ 30 . We calculated the average across-list associative strength for each participant over each of the three ranges. We found a highly significant advantage for the associations observed in the central zone over both the second [$t(287) = 3.80$, $p < .001$] and third [$t(287) = 5.54$, $p < .001$] zones. This finding provides quantitative confirmation of the observation that there is a significant advantage in associative strength across lists. Moreover, we observed a significant advantage for the second zone over the third [$t(287) = 2.13$, $p < .04$]. This analysis places a lower limit on the range of associations at six lists—or about 300 sec—between items. Even with this conservative estimate, the range of the across-list contiguity effect exceeds that observed within list by about a factor of 100.

The formation of associations between items presented in different lists seems to require a mechanism to support temporally defined associations over long time scales. Perhaps, however, the items were actually experienced much closer together in time than their presentations would suggest. Perhaps items bridge across lists as intrusions. In this way, an item from, say, List 10 that was intruded during study of List 20 might become associated to the items studied in List 20 because of the temporal proximity of the intrusion to either study or retrieval of the List 20 items. Although overall levels of prior-list intrusions were relatively low, as in previous studies (Zaromb et al., 2006), we observed a recency effect for prior-list intrusions across several lists. We repeated the analyses shown in Figures 2C and 2D excluding all items that were ever recalled as a prior-list intrusion at any time throughout the experiment. Doing so eliminated .086 of the final free recalls. The results with the prior-list intrusions excluded were indistinguishable from those reported with the complete data set.

GENERAL DISCUSSION

Previous authors have described recency and contiguity effects as a consequence of persistent activation of item representations in a short-term memory buffer that holds a discrete number of items in an all-or-none fashion. If this is the case, then we should be able to estimate the capacity of the buffer by examining the range over which recency and contiguity effects are observed. We compared recency effects within a list in immediate free recall with recency effects across lists in final free recall. We observed similar functions relating memory to recency across scales; the similarity of the probability of first recall curves (Figures 1B and 1D) was particularly striking, despite the fact that the time scale over which recency was observed varied by about a factor of 100. We were also able to compare contiguity effects within and across lists by calculating lag-CRP functions describing FFR transitions between words from the same list or words presented across lists. Again,

we observed strikingly similar forms for the within- and across-list contiguity effects, despite the fact that the scale of the across-list effect was about a factor of 100 greater than that for the within-list effect. These findings suggest either that different models describe recency and contiguity across different time scales (Davelaar et al., 2005; Raaijmakers, 1993) or that at least some of our assumptions about the role of traditional models of working-memory maintenance in verbal learning are violated.

Although a previous study reported contiguity effects in continuous-distractor free recall (Howard & Kahana, 1999) between items separated by 15 sec of interitem distractor, the across-list contiguity effect that we observed in this study extends the range of contiguity effects by at least an order of magnitude. Moreover, the within- and across-list effects were observed simultaneously in this study, making it harder to attribute contiguity to an artifact of some strategy that varies across conditions in continuous-distractor free recall. The across-list contiguity effect we observed in the present experiment did not appear to show a dramatic asymmetry between forward and backward recall transitions. An asymmetry favoring forward recall transitions has been observed previously under a wide variety of conditions at shorter time scales (Kahana, Howard, & Polyn, in press).

If different mechanisms are responsible for recency and contiguity over different time scales (Davelaar et al., 2005; Raaijmakers, 1993), then our findings imply that these mechanisms have similar properties. In this case, recency and contiguity would be seen as general design principles that are advantageous for multiple memory systems to implement. It is also possible that a common mechanism accounts for recency and contiguity effects over both short and long time scales (for a computational neuroscience perspective on these issues, see Drew & Abbott, 2006; Miller & Wang, 2006). The striking similarity of the functional form of both recency effects (Figure 1) and contiguity effects (Figure 2) across time scales makes the position that short- and long-term recency and contiguity effects arise from a common mechanism more appealing. In this case, the most likely assumption to abandon in order to bring buffer models into line with the data is the assumption that the buffer holds information in an all-or-none fashion. The temporal context model (TCM; Howard, Fotedar, Datey, & Hasselmo, 2005; Howard & Kahana, 2002) proposes that incoming information is maintained by changing the current state of a temporal context vector. In this view, information decays gracefully instead of dropping out precipitously, enabling recency and contiguity effects that can be observed over long periods of time (Howard & Kahana, 2002). Although TCM describes the maintenance of recent information, it does not include the control or executive processes that play such an important role in the descriptive power of working memory models (Atkinson & Shiffrin, 1968; Baddeley, 1986), so that it cannot replace this functionality.

AUTHOR NOTE

We thank Aditya Datey and Bing Jing for programming support and Diana Hobbins, Michael Zaremba, and Donna Bridge for assistance with data collection. This article was supported by Grants 1-R01 MH069938

to M.W.H. and 2-R01 MH55687 to Michael J. Kahana, who provided valuable feedback throughout this project. Correspondence concerning this article should be addressed to M. W. Howard, Department of Psychology, Syracuse University, 430 Huntington Hall, Syracuse, NY 13244 (e-mail: marc@memory.syr.edu).

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, pp. 89-105). New York: Academic Press.
- BADDELEY, A. D. (1986). *Working memory*. Oxford: Oxford University Press, Clarendon Press.
- BJORK, R. A., & WHITTEN, W. B. (1974). Recency-sensitive retrieval processes in long-term free recall. *Cognitive Psychology*, **6**, 173-189.
- BROWN, G. D. A., NEATH, I., & CHATER, N. (2007). A temporal ratio model of memory. *Psychological Review*, **114**, 539-576.
- CRAIK, F. I. M. (1970). The fate of primary memory items in free recall. *Journal of Verbal Learning & Verbal Behavior*, **9**, 658-664.
- 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.
- DREW, P. J., & ABBOTT, L. F. (2006). Extending the effects of spike-timing-dependent plasticity to behavioral timescales. *Proceedings of the National Academy of Sciences*, **103**, 8876-8881.
- 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 & Instrumentation*, **14**, 375-399.
- GLENBERG, A. M., BRADLEY, M. M., STEVENSON, J. A., KRAUS, T. A., TKACHUK, M. J., & GRETZ, A. L. (1980). A two-process account of long-term serial position effects. *Journal of Experimental Psychology: Human Learning & Memory*, **6**, 355-369.
- 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**, 75-116.
- HOWARD, M. W., & KAHANA, M. J. (1999). Contextual variability and serial position effects in free recall. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **25**, 923-941.
- HOWARD, M. W., & KAHANA, M. J. (2002). A distributed representation of temporal context. *Journal of Mathematical Psychology*, **46**, 269-299.
- HOWARD, M. W., VENKATADASS, V., NORMAN, K. A., & KAHANA, M. J. (2007). Associative processes in immediate recency. *Memory & Cognition*, **35**, 1700-1711.
- KAHANA, M. J. (1996). Associative retrieval processes in free recall. *Memory & Cognition*, **24**, 103-109.
- KAHANA, M. J., HOWARD, M., & POLYN, S. (in press). Associative processes in episodic memory. In J. Byrne (Series Ed.) & H. L. Roediger III (Vol. Ed.), *Learning and memory—A comprehensive reference*. San Diego: Elsevier, Academic Press.
- 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, & Cognition*, **28**, 530-540.
- MILLER, P., & WANG, X.-J. (2006). Power-law neuronal fluctuations in a recurrent network model of parametric working memory. *Journal of Neurophysiology*, **95**, 1099-1114.
- MURDOCK, B. B. (1962). The serial position effect of free recall. *Journal of Experimental Psychology*, **64**, 482-488.
- NEATH, I., & BROWN, G. D. A. (2006). SIMPLE: Further applications of a local distinctiveness model of memory. In B. H. Ross (Ed.), *The psychology of learning and motivation* (Vol. 46, pp. 201-243). San Diego: Academic Press.
- RAAIJMAKERS, J. G. W. (1993). The story of the two-store model of memory: Past criticisms, current status, and future directions. In D. E. Meyer & S. Kornblum (Eds.), *Attention and performance XIV* (pp. 467-488). Cambridge, MA: MIT Press.
- 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: Academic Press.
- 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.
- ZAROMB, F. M., HOWARD, M. W., DOLAN, E. D., SIROTIN, Y. B., TULLY, M., WINGFIELD, A., & KAHANA, M. J. (2006). Temporal associations and prior-list intrusions in free recall. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **32**, 792-804.

NOTES

1. As an additional test for evidence of rehearsal, we also examined FFR probability as a function of original serial position (not shown). While there arguably was some evidence for a slight negative recency effect (Craik, 1970), there was no hint of a within-list primacy effect in FFR.

2. Although it is underestimated by the shape of the LOWESS curves, there was a reliable asymmetry between within-list lags +1 and -1, consistent with previous findings.

3. At extremely large and small across-list lags, the function started to increase again. However, there is less data at extreme values of across-list lag, and these values are less stable. Moreover, it may reflect an edge effect related to recalls either to or from the very first or last lists in the experiment.

(Manuscript received November 21, 2006;
revision accepted for publication May 21, 2007.)