Running head: CONTEXT IN TCM

TCM can explain recency and contiguity, but not at the same time

Simon Farrell
Department of Experimental Psychology
University of Bristol

Stephan Lewandowsky
School of Psychology
University of Western Australia

Word count: 3977 (approximate count due to use of LaTeX)

Dr. Simon Farrell
Department of Experimental Psychology
University of Bristol
12a Priory Road
Clifton, Bristol BS8 1TU
UK
e-mail: Simon.Farrell@bristol.ac.uk
fax: +44 (0)117 928 8588

Abstract

The temporal context model (TCM) is a model of the evolution of temporal context in episodic memory, and has been claimed to account for recency and contiguity effects in free recall (M. W. Howard & M. J. Kahana, 2002). However, previous demonstrations of the model have, for mathematical convenience, introduced assumptions that circumvent the core principles of the model. We show that, when TCM is instantiated as described with the assumption of a gradually changing context, the model is unable to capture the locality of transitions in free recall. We demonstrate that this restriction is not specific to particular parameter values, and note the need to incorporate additional, and as yet unspecified, assumptions into TCM to circumvent the problems that follow from the core assumption of a gradually evolving context.

TCM can explain recency and contiguity, but not at the same time

A canonical paradigm for examining long-term memory in humans is the free recall task, in which participants are required to recall as many items (e.g., words) as possible from a list in any order. Data from the free recall task have been central to theorising on human memory (e.g., Laming, 1999; Raaijmakers & Shiffrin, 1981; Tan & Ward, 2000; for an early review see Murdock, 1974). A particular debate surrounds the origin and nature of the recency effect in free recall, the greatly enhanced recall of items from the last few serial positions of a list (e.g., Postman & Phillips, 1965). Although findings such as the specifically disruptive effect on recency of post-list distractors is taken as evidence for the involvement of a limited-capacity short-term buffer (Glanzer & Cunitz, 1966; Postman & Phillips, 1965), others have argued that recency instead follows from the temporal distinctiveness of terminal items (e.g., ?, ?; Glenberg & Swanson, 1986; Neath, 1993). Evidence for this latter viewpoint comes from the finding that recency re-emerges when distractors are placed throughout the list in addition to after the list (e.g., Bjork & Healy, 1974).

Another form of recency apparent in the free recall task comes from examination of conditional response probabilities (CRPs). For each item i the probability of next recall of all other items may be conditionalised on recall of item i. The *lag-CRP* function plots these probabilities as a function of the lag j between the recalled item i and the successor recall $i \pm j$. This analysis of the contingencies in free recall reveals a further two consistent properties of free recall (Howard & Kahana, 1999; Kahana, 1996; Kahana, Howard, Zaromb, & Wingfield, 2002). First, recall of an item tends to be followed by recall of an item from nearby in the list, such that the lag-CRP function declines with increasing lag: this is referred to as *lag-recency*. Second, the lag-CRPs are asymmetric; although transitions with a smaller lag are more likely, for a given distance a positive or forward transition (i.e., a transition to a later list item) is more likely than a transition to an earlier item. Critically, the locality and asymmetry of transititions evident in the lag-CRP function are obtained even when accounting for the constraint that there are fewer opportunities to make transitions of greater distance in recall [e.g., a forward transition of the size of the list length can only happen one way, by recalling the first item and then the last item; there are many more opportunities to recall the immediate successor of a just-recalled item (a lag of +1) as all but one of the list items will have an item following].

One popular theory recently developed to account for such effects in free recall, and episodic memory generally, is the temporal context model (TCM; Howard & Kahana, 2002; Howard, 2004). As implied by the name, and in common with temporal distinctiveness models, TCM assumes a critical role for the association between items and the context in which they occur. A novel assumption in TCM is the notion of retrieved context: The current context is assumed to be a composite of both the prior context and the temporal context associated with items previously presented.

Although TCM is a useful framework for considering recency and contiguity effects in episodic memory, we argue that a single, coherent explanation of both these phenomena cannot be achieved within a constant set of assumptions in TCM. We show that assumptions about how context evolves across time, when actually incorporated into the model, lead to surprising behaviour not apparent in previous applications of the model.

Assumptions of context in recency and lag-recency

As presented (Howard & Kahana, 2002), TCM makes the central assumption that context evolves continuously. The temporal context at time step i is assumed to be the weighted sum of the context at time $i-1$ and the context retrieved by the item present at time i. This central assumption is formally captured in TCM's evolution equation for generating a new context $\mathbf{t}_i$:

$$\mathbf{t}_i = \rho_i \mathbf{t}_{i-1} + \beta \mathbf{t}_{IN}, \tag{1}$$

where $\mathbf{t}_{i-1}$ and $\mathbf{t}_{IN}$ are temporal context vectors that represent, respectively, the previous temporal context and the context retrieved by the previous item $i-1$ (weighted by the parameters ρ_i and β, with ρ_i chosen such that the length of the vector $\mathbf{t}_i$ equals 1). These two components are responsible for the model's account of the recency and lag-recency effects, respectively.

The recency effect in TCM follows from the relative overlap between the temporal context present at the end of recall and the temporal context used to initiate recall; if it is assumed that the list-final context (the context present at the end of presentation of the list) will carry over into recall, there will naturally be a close match between the recall context and the temporal context associated with items towards the end of the list (cf. non-reinstatable context in Brown, Preece, & Hulme, 2000). In contrast, the retrieved component $\mathbf{t}_{IN}$ allows the model to account for lag-recency effects: When item i is recalled and is used to evolve the temporal context, it will retrieve a context that will partially match those contexts bound to items $i-1$ and $i+1$, with increasingly poorer matches for increasing separation (lag) j between item i and item $i \pm j$.

In accounting for data at a quantitative level, Howard, Kahana, and colleagues make a simplifying assumption that, though subtle, has significant implications for the account offered by TCM for recency and lag recency. As noted above, it is assumed that list items and post-list distractors drive the evolution of temporal context to give that context that exists immediately prior to recall (Howard & Kahana, 2002); for simplicity we call this the *pre-recall context*. By contrast, when accounting for the lag-recency effect (i.e., lag-CRP functions), Howard and colleagues assume that an infinitely long time intervenes between the list and the to-be-modelled recall attempt; in the model, this is implemented by assuming that the temporal context that carries over from the previous recall is orthogonal to all contextual states from the list (Howard & Kahana, 2002). This means that the temporal context used to cue for an item in generating CRP functions is the weighted sum of the context retrieved by the previously recalled item, and a temporal context orthogonal to all previously presented contexts. The quantitative implications of this assumption, and its divergence from the principle of a gradually evolving context, are demonstrated in a set of simulations reported next.

Modelling recency and lag-recency in TCM

Replication of previous demonstrations

Representative behaviour of the model is shown in Figure 1. The simulations generating the figure were implemented as per the original description of the model in Howard and Kahana (2002), using representative parameter estimates($\beta = .4$; $\tau = .25$; $d = 7$ where d is the number of orthogonal items equivalent to 16s of distracting activity; these parameter values were also used for the following two simulations)[1]; further details about TCM are set out in Howard and Kahana (2002). The functions depicted in the left hand of Figure 1 show predicted first recall probability (FRP) functions; for each serial position, these depict the probability with which that item is recalled first, and give a measure of the recency effect (Howard & Kahana, 1999, 2002). As shown in Figure 1, when the pre-recall context is assumed to be used as a cue for the first recall, the model shows an extensive amount of recency and an absence of primacy; the recency accords with empirical FRPs obtained in a number of experiments (see the analysis of a number of experiments in Howard & Kahana, 1999), though the absence of primacy deviates from some empirical FRPs (see, e.g., Figures 1 and 3 in Howard & Kahana, 1999). The model also predicts effects of delaying recall that correspond to empirical data (Howard & Kahana, 1999, 2002): Delaying recall with some interfering activity reduces the extent of recency (middle left panel), with a recovery of recency when the distracting activity is presented after every item, as well as following the list (bottom left panel).

The right panels in Figure 1 show the model's account of conditional response probability (lag-CRP) functions that describe all recalls following the first one; these predictions were obtained under the standard assumption that an infinite amount of time has elapsed since list presentation. Hence, recall of the next item was cued by combining a context that was orthogonal to all preceding ones (a *novel* context) with the context retrieved by the item just recalled. Following precedent (Howard & Kahana, 1999, 2002; Kahana, 1996) the differential opportunity to make transitions of various distances was accounted for by dividing the predicted transition at each lag by the total number of opportunities to make that transition. As shown in the figure, the model predicts the two crucial characteristics of lag-CRPs: Transitions to items nearby on the input list are more likely than to items further way, and there is a bias to recall items following the item just recalled rather than items preceding it (asymmetry in lag-recency; Howard & Kahana, 1999; Kahana, 1996).

It is crucial that these predictions are obtained under two separate assumptions. When predicting FRPs, the pre-recall context is used as a cue to recall, with the first term in Equation 1 being $\mathbf{t}_{i-1}$; when predicting CRPs, a random context $\mathbf{t}_{i\pm\infty}$ is combined with retrieved context to provide a cue for the next item. This dissociation is justified by Howard, Kahana and colleagues' focus on conditions of delayed recall. They note that the CRP function changes across the first several output positions in immediate free recall, whereas in delayed recall it is fairly static (Howard & Kahana, 1999). This has led to a focus on delayed recall and on the continuous distractor paradigm when examining lag-CRP functions (Howard & Kahana, 2002; Howard, Kahana, & Wingfield, 2006). Restricting focus in this way allows Howard, Kahana, and colleagues (Howard & Kahana, 2002; Howard et al., 2006) to argue that the context at retrieval, following delay, will be mostly uncorrelated with the list-final context, and thus assume a recall context that is orthogonal to all list contexts.

However, this assumption creates a gap between the intended deep properties of the model and its implementation. On the one hand, the model is intended to capture the notion of a gradually evolving temporal context, with the evolution of context being driven by retrieval of context using the item just presented or recalled as a cue, in combination with the current contextual state; this is exactly what is described in Equation 1. On the other hand, the model is not implemented in this fashion when accounting for the lag-CRP function. In fact, to date there has been no published simulation that shows the model's prediction when it is permitted to behave as described by Equation 1; that is, let context evolve from the state associated with the end of list presentation to its final state after some retention interval.

TCM predictions when assuming continuously evolving context

What happens when the final list context (in the case of immediate recall), or a successively evolved context driven by post-list distractors (in the case of delayed recall) is used instead of a random context to predict the lag-CRP functions? In other words, what happens to the lag-CRP functions if the model is asked to transition from first-item recall to the recall of further items in the manner described in Equation 1? Figure 2 shows that the model produces behaviour that clearly does not correspond to that in Figure 1, which depicts the previous demonstrations of the model that are in accord with behavioral data. Of particular concern is that the model no longer predicts localised recall transitions. Instead, the lag-CRPs for recalls immediately following the first retrieval come to resemble the first recall probability functions depicted in Figure 1: The updating of temporal context by the context retrieved by the previously recalled item is swamped by overlap with the final list context $\mathbf{t}_{i-1}$. This is a noticeable problem even for delayed recall with and without a continuous distractor. Even after updating the temporal context with an additional 7 items, the model still shows a strong tendency to make extreme transitions (middle and right panels). Both qualitatively and quantitatively, this behaviour of the model clearly violates the principle of contiguity that is apparent in psychological data (Howard & Kahana, 1999; Kahana, 1996; Kahana et al., 2002). It also suggests that the stated core property of the model—namely, the evolution of context across retrieval events—is unable to give rise to realistic lag-CRP functions. Only when additional assumptions are made that negate this core property can the model predict lag-CRP functions.

TCM predictions when assuming a recall context orthogonal to list contexts

One might attempt to resolve this problem by assuming that a novel context is always used to generate the next context state during recall. This solution is undesirable for two reasons. First, although this assumption leads to correct predictions for the lag-CRP function as shown in Figure 1 and in earlier demonstrations, it simultaneously sacrifices another cornerstone of the model, namely its ability to predict the recency effect. This problem is illustrated by the FRPs in Figure 3 which are based on an orthogonal cueing context. It is evident that TCM then predicts a flat FRP function; since the cueing context is orthogonal to all list contexts, all items are equally accessible and hence equally likely to be recalled. Second, use of a uniformly orthogonal context is also undesirable because it negates the spirit of the model which puts a central role on a gradually evolving context; Howard and Kahana (2002) point to Equation 1 as a central component of the model because it allows temporal context to change gradually.

In summary, there is no unified set of assumptions that permit the model to predict both crucial phenomena to which it lays claim. If the reasonable assumption is made that the pre-recall context cues recall of the first item, then the model successfully predicts the correct FRP functions but it cannot, on its own, transition to a context that gives rise to the proper lag-CRP functions. Conversely, if it is assumed that context is unrelated to its state prior to recall—without however explaining how this comes about—then the model can handle the lag-CRP functions but it can then no longer predict a realistic FRP function. Regardless of which of those two assumptions is adopted, the stated core property of the model—namely, its evolving context—must be disabled or ignored for the predictions to be in accord with the data.

Parameter independence of the problems of TCM

Critically, this behaviour in the model is not due to a particular selection of parameter values. We conducted a further set of simulations in which context gradually evolved during recall; that is, we assumed that the pre-recall context was used to cue for the first recall, and that this context was carried over to the next recall to be combined with the context retrieved by the first recalled item. This model was fit to the representative patterns generated by the model in Figure 1 under the assumptions made by Howard and Kahana (2002) (that is, where pre-recall context is used to generate the FRP function and an orthogonal context is used to generate the CRP function). In other words, we sought to reproduce the model's known behaviour by bridging the two hitherto separate sets of assumptions with the mechanism specified by Howard and Kahana (2002) and with a freely estimated set of parameters.

Fitting of the free parameters β (the rate of contextual evolution) and τ (a scaling factor in the Luce choice rule used in the model; see Howard, 2004) was accomplished by minimizing an aggregate root mean squared deviation (RMSD) measure obtained by summing the individual RMSDs for the FRPs and CRPs; the Nelder-Mead minimization algorithm (Nelder & Mead, 1965) was used to find the best fitting parameter values. Figure 4 shows that the predictions under the best-fitting parameter estimates (immediate: $\beta = .97$, τ=.89; delayed: $\beta = .57$, τ=.15; continuous distractor: $\beta = .54$, τ=.09) still deviated systematically from the desired model behaviour: Although the model largely captured the behaviour of the FRP functions, the model still showed gross non-monotonicity in the CRP functions, with increasing probabilities for remote transitions. This behaviour appears to be an inevitable consequence of the assumption of a continuously evolving context.

The inevitable trade-off between recency and lag-recency in the model is most evident when examining the change in FRP and lag-CRP functions as the recall context is made less and less similar to the pre-recall context (cf. Howard, 2004). [2] The left panel of Figure 5 shows the change in recency and lag-recency as the degree of similarity between the list-final context $\mathbf{t}_{i-1}$ and the context used to initiate immediate free recall varies. Following Howard (2004), recency is measured by the probability of recalling the last list item on the first recall [FRP(last)], whereas the weighting of recency and lag-recency in the lag-CRP functions is measured by taking the ratio of the lag-CRP for lag +1 and lag $+(L-1)$, the largest positive lag (where L is the list length; cf. Howard, 2004). As shown in the panel, there is a continuous trade-off between recency and lag-recency in the model: As the temporal context used to initiate recall is rotated away from an orthogonal vector (cosine between list-final context vector and pre-recall context vector = 0) and towards the list-final context (cosine approaching 1), recency increases and the lag-CRP function shows fewer localised transitions and more extreme transitions. The middle and right panels of Figure 5 show this change in more detail, by tracing the effects on the FRP and lag-CRP functions. As can be seen, as the pre-recall context is made more similar to the list-final context, recency is more apparent in the FRP function (increased probability of first recall of later list items) and in the lag-CRP function (a prominent tendency to transition from the previous recalled item to an item towards the end of the list). Conversely, as the context becomes orthogonal to the list-final state, recency disappears but the lag-CRP functions take on their expected shape.

What implications do our results have for the TCM? We suggest that our results show that the TCM cannot model the gradual evolution of context and simultaneously account for recency and lag-recency in free recall. As Figure 4 shows, even under the most favourable parameter values, the model produces behaviour that is incommensurate with the pervasive lag-recency effect (Howard & Kahana, 1999; Kahana, 1996). This inability to handle contiguity effects in situations in which recency can be accounted for implies that the TCM can no longer be distinguished from random context models (e.g., Estes, 1955; Murdock, 1974).

To explain both FRP and CRP functions, the model must make additional assumptions about non-continuous changes in context in order to generate further contextual states at recall. One such assumption might be that the list-final context is used to cue the first recall, and that prior to the second recall the temporal context is reset—perhaps to some random signal—before combining it with the context retrieved by the item first recalled. However, this would predict a flat second recall probability function (similar to that in Figure 3), whereas the empirical recall probability functions display prominent recency well in to recall (i.e., up to the fourth or fifth output position in lists of 20 items; see Figure 1 of Howard & Kahana, 1999).

Given that the behaviour of the model is dependent on assumptions made about which prior context is driving contextual evolution at any point in recall, we stress caution in interpreting quantitative fits of free recall data under these assumptions (e.g., Howard & Kahana, 2002; Howard et al., 2006). In particular, we caution against attributing the quantitative fit of the model to the core principles of TCM (contextual evolution through retrieved context), when in fact auxiliary assumptions were required to give a full account of the data. In previous applications, the model as applied is not actually a model of gradually evolving context; assuming a gradually evolving context as per Equation 1 introduces problematic behaviour into the model that requires additional assumptions to be addressed; critically, these assumptions have received no motivation in previous presentations of the model, and have been made on the basis of convenience (e.g., Howard, 2004; Howard & Kahana, 2002).

In closing, we should emphasise that our simulations do not necessarily call into question all core mechanisms in the model. The recency observed in continuous distractor tasks certainly favours explanations based on temporal distinctiveness (?, ?; Glenberg & Swanson, 1986; Neath, 1993), and the lag-recency effect suggests some local associations between items on lists (Howard & Kahana, 1999; Kahana, 1996), both of which are embodied in TCM. However, in light of our simulations, the exact instantiations of these principles in the model may need to be revisited.

References

Bjork, E. L., & Healy, A. F. (1974). Short-term order and item retention. *Journal of Verbal Learning and Verbal Behavior*, *13*, 80-97.

Brown, G. D. A., Preece, T., & Hulme, C. (2000). Oscillator-based memory for serial order. *Psychological Review*, *107*, 127-181.

Estes, W. K. (1955). Statistical theory of spontaneous recovery and regression. *Psychological Review*, *62*, 145–154.

Glanzer, M., & Cunitz, A. R. (1966). Two storage mechanisms in free recall. *Journal of Verbal Learning and Verbal Behavior*, *5*, 351-360.

Glenberg, A. M., & Swanson, N. G. (1986). A temporal distinctiveness theory of recency and modality effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *12*, 3-15.

Howard, M. W. (2004). Scaling behavior in the temporal context model. *Journal of Mathematical Psychology*, *48*, 230-238.

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

Howard, M. W., & Kahana, M. J. (2002). A distributed representation of temporal context. *Journal of Mathematical Psychology*, *46*, 269-299.

Howard, M. W., Kahana, M. J., & Wingfield, A. (2006). Aging and contextual binding: Modeling recency and lag-recency effects with the temporal context model. *Psychonomic Bulleting & Review*, *13*, 439–445.

Kahana, M. J. (1996). Associative retrieval processes in free recall. *Memory & Cognition*, *24*, 103-109.

Kahana, M. J., Howard, M. W., Zaromb, F., & Wingfield, A. (2002). Age dissociates recency and lag recency effects in free recall. *Journal of Experimental Psychology:*

Learning, Memory, and Cognition, *28*, 530-540.

Laming, D. (1999). Testing the idea of distinct storage mechanisms in memory. *International Journal of Psychology*, *34*, 419-426.

Murdock, B. B. (1974). *Human memory: Theory and data.* Potomac, MD: Lawrence Erlbaum Associates.

Neath, I. (1993). Contextual and distinctive processes and the serial position function. *Journal of Memory and Language*, *32*, 820-840.

Nelder, J. A., & Mead, R. (1965). A simplex method for function minimization. *Computer Journal*, *7*, 308-313.

Postman, L., & Phillips, L. W. (1965). Short-term temporal changes in free recall. *Quarterly Journal of Experimental Psychology*, *17*, 132-138.

Raaijmakers, J. G. W., & Shiffrin, R. M. (1981). Search of associative memory. *Psychological Review*, *88*, 93-134.

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*, 1589–1625.

Author Note

Collaboration on this grant was assisted by an Australian Research Council International Linkage Grant to Stephan Lewandowsky, Gordon Brown, and Simon Farrell. We thank Marc Howard for his clarification of some details of operation of the model, and Anna Lelièvre for her assistance in proofreading the manuscript. Correspondence should be addressed to Simon Farrell, Department of Psychology, University of Bristol, 12a Priory Road, Clifton, Bristol BS8 1TU, UK; e-mail: Simon.Farrell@bristol.ac.uk.

Notes

[1] Code for the simulations is available from the first author's homepage at http://seis.bris.ac.uk/∼pssaf

[2] In his demonstrations Howard (2004) noted a positive correlation between recency and contiguity effects. However, that conclusion was based on the assumption of an infinite retention interval.

Figure Captions

Figure 1. Typical recency and lag-recency predictions from the TCM in its standard implementation. Predicted first recall probabilities (FRP; left) and lag-conditional response probability functions (lag-CRPs; right) are shown under conditions of immediate recall (top panels), delayed recall (middle panels), and continuous distraction (bottom panels). In the lag-CRP plots, each point shows the probability of recall of list item $i + j$ following recall of item i. Probabilities are corrected for chance.

Figure 2. CRPs predicted by TCM when the temporal context incorporated into the cue for recall of the next item is the context following presentation of the list and any distractors.

Figure 3. First recall probabilities predicted by TCM under conditions of immediate recall, delayed recall, and continuous distraction, where the temporal context used as a cue is a random context.

Figure 4. Fitted FRP and CRP functions. The circles reproduce the target pattern from Figure 1, whereas the solid lines show fits of the model when the final list context is used to generate both first recalls and subsequent recalls. See Figure 1 and text for more detail about FRP and CRP functions.

Figure 5. Left panel: Ratio of FRP for last two list items (circles), and ratio of CRP at lag 1 to CRP at lag $L - 1$ (triangles), as a function of the cosine between the list-final context vector and the context used to intiate recall. Middle panel: Change in the FRP function for increasing cosine between the list-final context vector and the context used to initiate recall (lighter to darker lines). Right panel: Change in the lag-CRP function for increasing cosine between list-final context vector and the context used to intiate recall (lighter to darker lines).

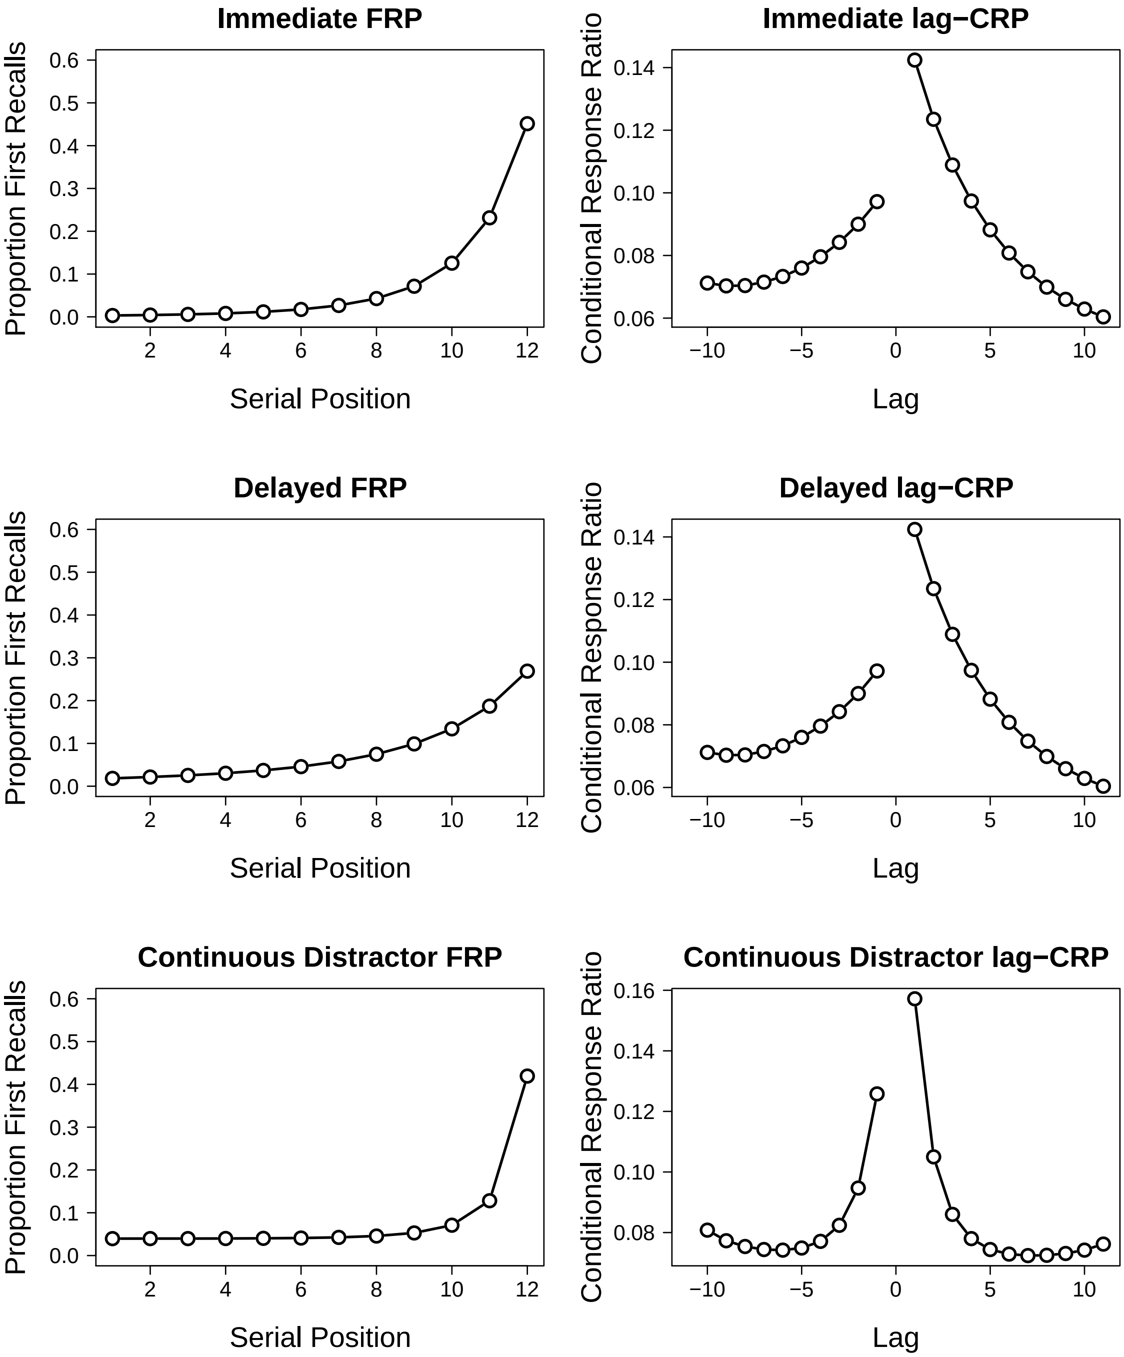
Immediate FRP
Proportion First Recalls
0.0
0.1
0.2
0.3
0.4
0.5
0.6
2
4
6
8
10
12
Serial Position
Immediate lag−CRP
Conditional Response Ratio
0.06
0.08
0.10
0.12
0.14
−10
−5
0
5
10
Lag
Delayed FRP
Proportion First Recalls
0.0
0.1
0.2
0.3
0.4
0.5
0.6
2
4
6
8
10
12
Serial Position
Delayed lag−CRP
Conditional Response Ratio
0.06
0.08
0.10
0.12
0.14
−10
−5
0
5
10
Lag
Continuous Distractor FRP
Proportion First Recalls
0.0
0.1
0.2
0.3
0.4
0.5
0.6
2
4
6
8
10
12
Serial Position
Continuous Distractor lag−CRP
Conditional Response Ratio
0.08
0.10
0.12
0.14
0.16
−10
−5
0
5
10
Lag

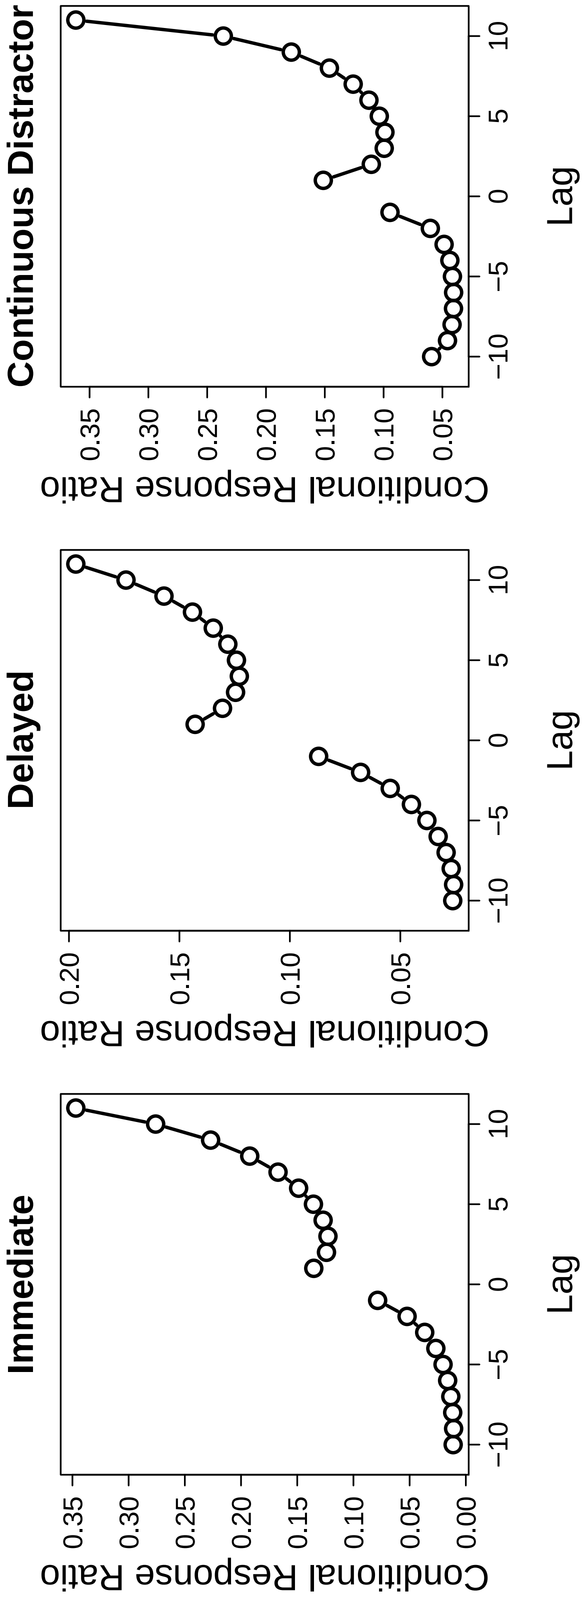
Immediate
Delayed
Continuous Distractor
Conditional Response Ratio
Lag
0.00
0.05
0.10
0.15
0.20
0.25
0.30
0.35
−10
−5
0
5
10

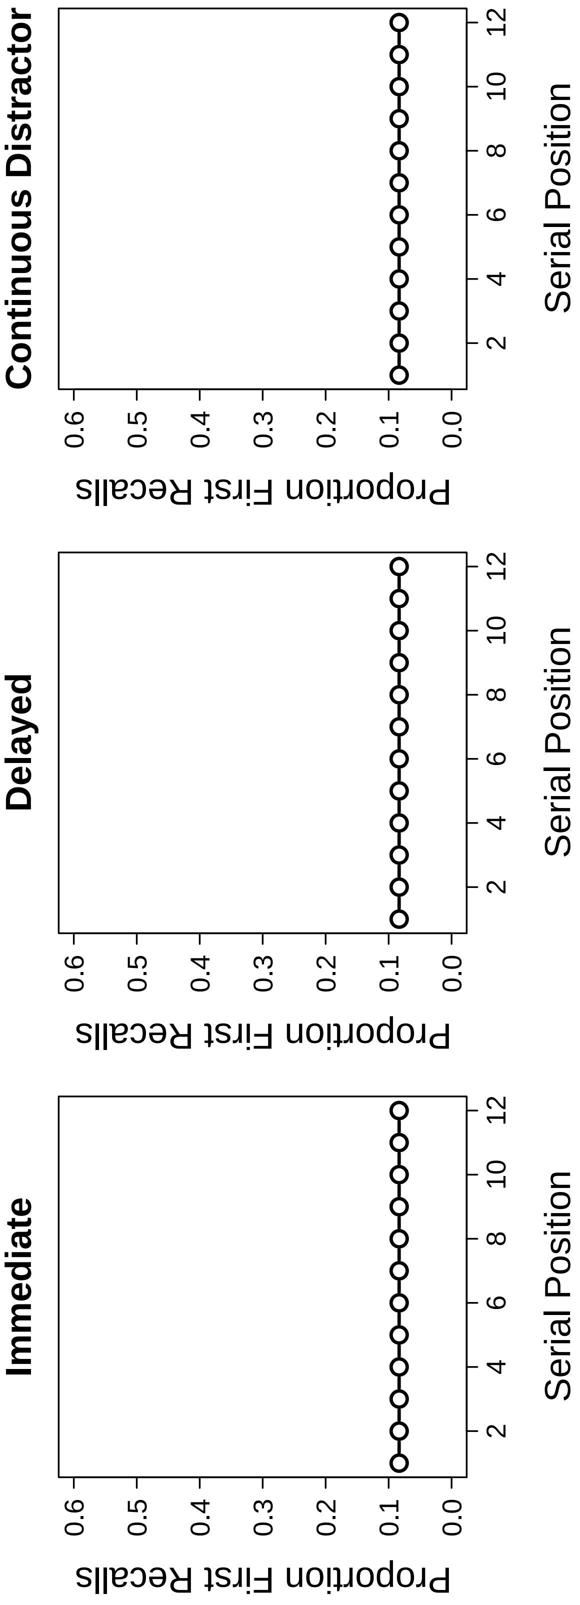
Immediate
Delayed
Continuous Distractor
Proportion First Recalls
0.0
0.1
0.2
0.3
0.4
0.5
0.6
Serial Position
2
4
6
8
10
12

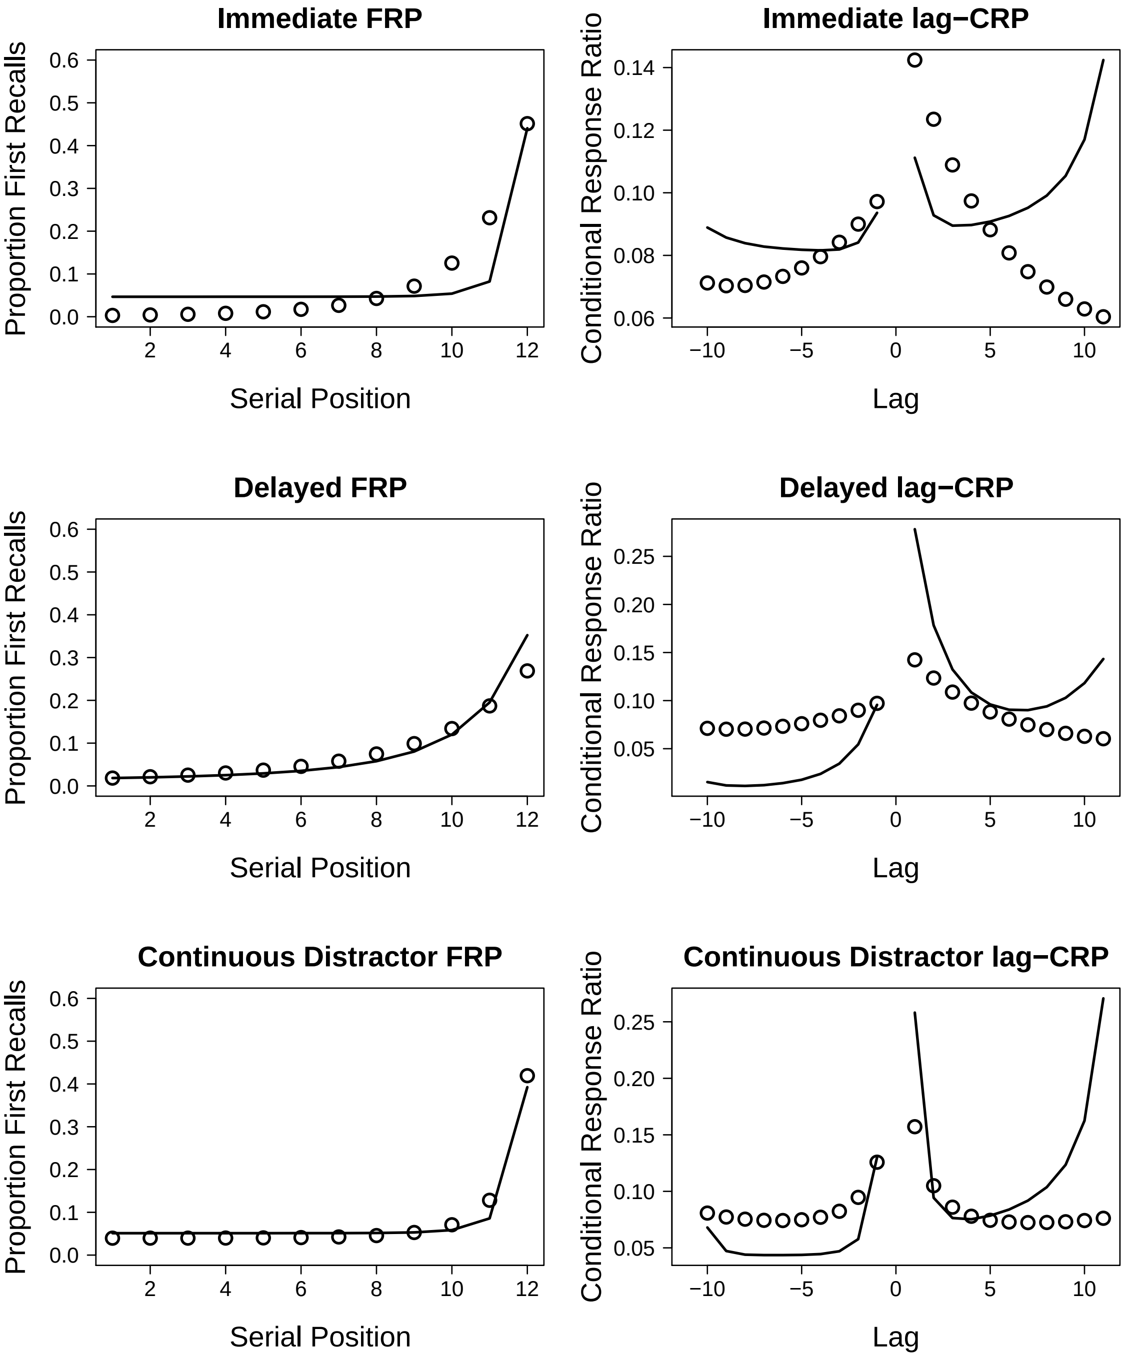
Immediate FRP
Proportion First Recalls
0.0
0.1
0.2
0.3
0.4
0.5
0.6
2
4
6
8
10
12
Serial Position
Immediate lag−CRP
Conditional Response Ratio
0.06
0.08
0.10
0.12
0.14
−10
−5
0
5
10
Lag
Delayed FRP
Proportion First Recalls
0.0
0.1
0.2
0.3
0.4
0.5
0.6
2
4
6
8
10
12
Serial Position
Delayed lag−CRP
Conditional Response Ratio
0.05
0.10
0.15
0.20
0.25
−10
−5
0
5
10
Lag
Continuous Distractor FRP
Proportion First Recalls
0.0
0.1
0.2
0.3
0.4
0.5
0.6
2
4
6
8
10
12
Serial Position
Continuous Distractor lag−CRP
Conditional Response Ratio
0.05
0.10
0.15
0.20
0.25
−10
−5
0
5
10
Lag

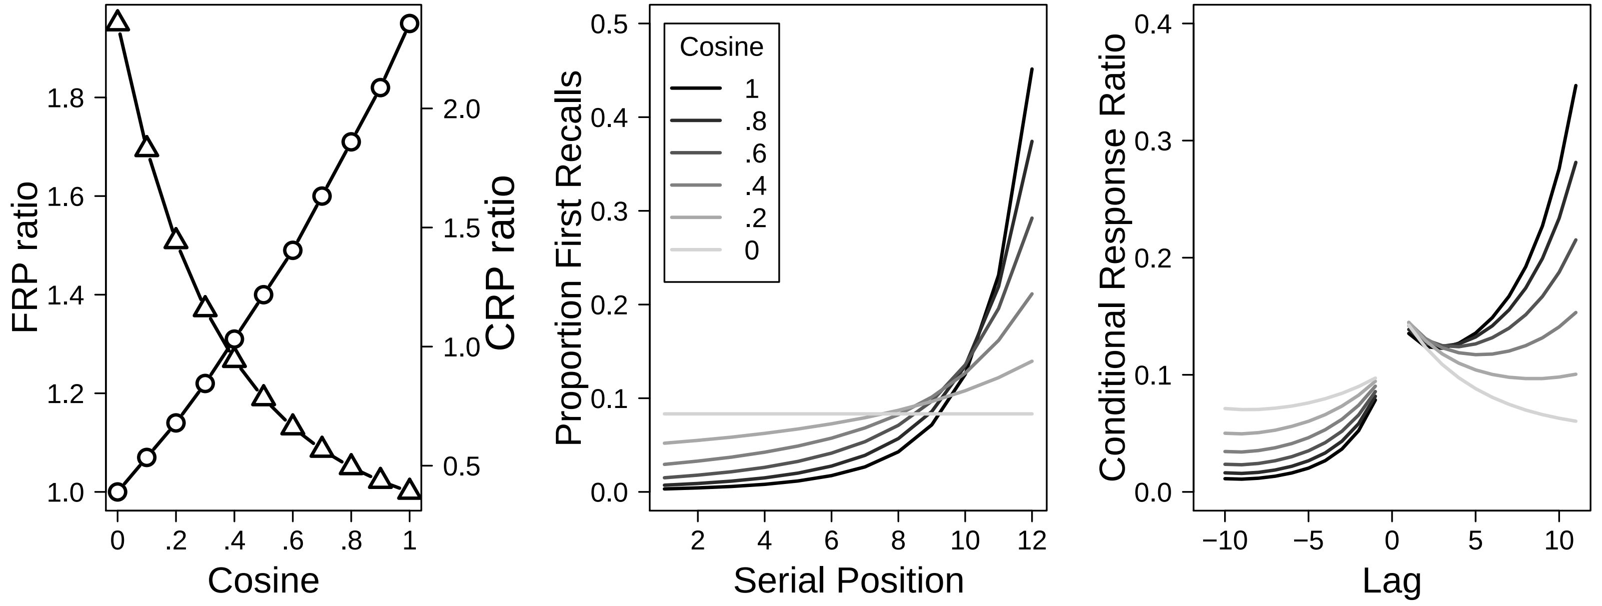
FRP ratio
1.0
1.2
1.4
1.6
1.8
CRP ratio
0.5
1.0
1.5
2.0
0
.2
.4
.6
.8
1
Cosine
Proportion First Recalls
0.0
0.1
0.2
0.3
0.4
0.5
Cosine
1
.8
.6
.4
.2
0
2
4
6
8
10
12
Serial Position
Conditional Response Ratio
0.0
0.1
0.2
0.3
0.4
−10
−5
0
5
10
Lag