in press, *Psychology of Learning and Motivation, Vol. 49*

Short-Term Memory: New Data and a Model

Stephan Lewandowsky

University of Western Australia

and

Simon Farrell

University of Bristol

Running head: Model of Short-Term Memory

Address correspondence to: Stephan Lewandowsky

School of Psychology

University of Western Australia

Crawley, W.A. 6009, AUSTRALIA

lewan@psy.uwa.edu.au

URL: http://www.psy.uwa.edu.au/user/lewan/

Abstract

We review the extensive benchmark data on short-term memory for order, and discuss the prominent computational theories accounting for serial order memory. On the basis of recent diagnostic results, we identify four explanatory constructs that we suggest must be instantiated in a model in order to provide an adequate account of those data. We present one such model, called C-SOB, and show that it handles existing benchmark data as well as recent diagnostic results. We conclude by exploring some of the model's novel predictions.

1. Short-Term Memory: A Wealth of Data and Theories

1.1. The Engine of Cognition

What could be simpler than reading a few items, such as the digits in a phone number, and recalling them in the right order a short while later? Simple as this short-term serial recall task may appear, it has been shown to contribute to language abilities such as vocabulary acquisition (Baddeley, Gathercole, & Papagno, 1998) and utterance production (e.g., Adams & Gathercole, 1996). Short-term memory for serial order can also be a critical element in mental arithmetic (Fürst & Hitch, 2000) and, when a task combines short-term storage (e.g., retention of digits) with a processing component (e.g., mental arithmetic), the resultant "working memory span" measure is highly predictive of higher-level cognitive functioning including fluid intelligence and reasoning ability (e.g., Oberauer, Süß, Wilhelm, & Sander, 2007). Accordingly, short-term memory is an acknowledged core component of cognition and there is some consensus that extending our understanding of short-term memory will ultimately contribute to solving other puzzles of cognitive functioning. The importance of short-term memory (STM) is also reflected in the wealth of existing data and the varieties of existing quantitative theories.

In this article, we first survey benchmark results in STM research and introduce some of the current computational theories that vie for their explanation. We then explore a new generation of results to identify and select the most appropriate explanatory constructs for a theory of STM. Following this, we present a computational model that instantiates those constructs, known as C-SOB, and we conclude by exploring some of the model's novel predictions.

1.2. Wealth of Data

We focus on situations in which people are presented short lists (usually no more than 5-7 items) for study and immediate (or nearly immediate) retrieval. Memory for the order among items—in addition to memory for their identity—is considered crucial, and thus most of the relevant data were generated using serial recall tests or other order-sensitive tasks such as serial reconstruction. In serial recall, people must report the list items in the order in which they were presented, either from the beginning of the list (forward recall) or from the end (backward recall). In a reconstruction task, by contrast, all list items are shown again at retrieval in a random sequence and the participant's task is to re-arrange the items into their original presentation order.

With some 13,000 papers published on short-term memory to date[1], the rich data base necessarily escapes concise summary; fortunately, however, there is some agreement on a set of findings that are considered benchmarks. Table 1 lists those benchmarks and additionally shows which of the currently available theories can handle the results at a quantitative level of description. We discuss the various theories later and, for now, invite the reader to focus on the rows, which report the core data on short-term memory. In order to keep the quantity of information manageable, the table only considers findings to which more than one computational theory has been applied. On the basis of that criterion, we excluded the interaction between articulatory suppression and phonological similarity plus a number of interesting neuropsychological findings that have been explained by the model by Burgess and Hitch (1999). We also excluded the detailed account of the Hebb repetition effect—that is, the memorial benefit associated with surreptitious repetition of a list on every third study-test trial—provided by a later variant of the same model (Burgess & Hitch, 2006). The fact that no other model has been applied to those important areas of research attests to the power of the theory developed by Burgess and Hitch and must be borne in mind during the remaining discussion. We also excluded data gathered with probed recall, recognition, and judgments of recency, which only the OSCAR model (Brown, Preece, & Hulme, 2000) accommodates. Finally, we exclude latency data because only the SOB model (Farrell & Lewandowsky, 2002) provides response time predictions as a consequence of its architecture; however, we show later that response latencies may distinguish between rival explanatory constructs underlying theories of short-term memory.

Even after exclusion of data that are unique to the domain of a single theory, the table shows that there is no shortage of phenomena that have been quantitatively explained. Before turning to a summary of the competing theories, we highlight a subset of benchmark results that are well replicated and have shaped the development of models of short-term memory.

1.2.1. *Serial position curve.* The serial position curve for forward serial recall is the sine qua non of modeling in short-term memory; all theories accommodate the extensive primacy (i.e., superior performance for early list items) and limited recency (i.e., advantage for terminal items). This pattern reverses with backward recall, in which case steep and extensive recency but little primacy is observed (e.g., Li & Lewandowsky, 1993). Symmetry of primacy and recency is achieved with a reconstruction task (e.g, Lewandowsky, Nimmo, & Brown, in press; Nairne, 1992). Perhaps somewhat surprisingly, few models (if any) accommodate backward recall and reconstruction, and we therefore restrict consideration to the serial position curve for forward recall.

1.2.2. *Error patterns*. The forward serial position curve is accompanied by a highly regular pattern of different types of errors. When a list item is recalled in an incorrect position, this is considered a transposition error (it is the only error possible with a reconstruction task). Most transpositions involve neighboring list positions, such that the third item might be recalled in the second or fourth output position (rather than in the first or seventh). This property of transpositions is known as the "locality constraint" (Henson, Norris, Page, & Baddeley, 1996), and it obeys the orderly property that increasing displacements of list items are increasingly unlikely.

A further analysis of transposition errors reveals several reliable subtleties in the data that are not apparent from consideration of pairwise transpositions alone. Suppose an anticipation error has been committed, for example if the second list item is recalled first (i.e., "B" is recalled first from the list "ABCD"). This error can be followed by report of the first item (i.e., "BA…"), an error known as a "fill-in," or it can be followed by report of the third item (i.e., "BC…"), an error known as "infill", which preserves the relative ordering of the second and third item. Fill-in errors are known to be roughly twice as frequent as infill errors (e.g., Surprenant, Kelley, Farley, & Neath, 2005).

On occasion, people will mistakenly introduce an extra-list item into their recall. These "intrusions" frequently involve items from the immediately preceding list (Drewnowski & Murdock, 1980), in which case they are also known as "protrusions" (Henson, 1998a, 1999) and are likely to occur at the same serial position as on the original list (Henson, 1999). The probability of intrusions appearing in people's recall is known to increase across output positions (e.g., Henson, 1998a). When people are permitted to "skip" items, the resulting omission errors also become more likely across output positions. (Only models that capture this pattern across output positions are identified in Table 1 as explaining intrusions and omissions.)

Finally, erroneous repetitions of an item—for example, reporting the first list item twice during recall—occur notably infrequently. Henson (1996) reported that erroneous repetitions constituted 2% of all responses, and Vousden and Brown (1998) cited a figure of 5%. Despite their infrequency, repetition errors have a distinct distribution, with most repetitions involving early list items that are reported a second time late in recall. In consequence, repetition errors are typically separated by 3 or 4 output positions (e.g., Henson, 1996, observed an average separation of approximately 3.5 output positions). The spacing of repetition errors can give rise to a subtle violation of the locality constraint for transpositions involving the last output position. Specifically, when a moderate number of repetition errors are present, the last output position may involve more reports of the first item than of the second or third if repetitions are not excluded from analysis (Henson et al., 1996).

1.2.3. *Phonological similarity effect.* The phonological similarity effect refers to the ubiquitous finding that lists composed of similar-sounding items are less accurately recalled in the correct order than lists in which items do not sound alike (e.g., Baddeley, 1966, 1968; Conrad, 1964; Henson, et al., 1996; Wickelgren, 1965a, b). The effect is of considerable generality and occurs with consonants (Baddeley, 1968) as well as with words (Baddeley, 1966; Coltheart, 1993; Henry, 1991) and for the most part is due to an increase in transposition errors. With increased retention interval, the detrimental effects of phonological similarity are either reduced (e.g., Farrell, 2006) or reversed (e.g., Nairne & Kelley, 1999).

The phonological similarity effect also occurs when phonologically similar (e.g., B, P, T) and dissimilar (e.g., K, Q, R) items are mixed together on a single study list. Mixed lists are particularly diagnostic because performance on dissimilar items can differentiate between competing theories of memory. We consider mixed lists in detail in a later section.

1.2.4. *Grouping effects.* When a list is explicitly grouped, for example by inserting a temporal pause after every third item, recall performance improves considerably overall (Ryan, 1969) and there are within-group recency and primacy effects, thus creating a slightly scalloped appearance of the overall serial position curve (Hitch, Burgess, Towse, & Culpin, 1996). The size of the grouping effect does not appear to depend on the duration of the pause between groups (0.9 s vs. 3.4 s; Ryan, 1969), and grouping effects may even emerge spontaneously without any objective cue to group (Henson, 1996, Madigan, 1980), indicating grouping to be a natural manner in which information is organized in short-term memory.

Grouping of a list also affects the pattern of errors, primarily by reducing the frequency of transpositions overall. When transpositions do occur between groups, they tend to preserve within-group position, in which case they are referred to as “interpositions” (Henson, 1999).

1.2.5. *Effects of word length.* There is no doubt that lists of words that take longer to articulate (e.g., “hippopotamus”, “retromingent”, “plethysmograph”) are recalled less accurately in correct order than lists composed of shorter words such as “bat” or “putt” (Baddeley, Thomson, & Buchanan, 1975), and this pervasive result has been taken to support the view that information in short-term memory is subject to time-based decay (e.g., Mueller, Seymour, Kieras, & Meyer, 2003; though see Brown & Hulme, 1995). This conclusion has been subject to considerable controversy and alternative interpretations (e.g., Service, 1998). In particular, unless the number of syllables is kept constant between short and long words, the word length effect likely reflects differences in complexity of the material rather than the net effects of the passage of time (Service, 1998). When the number of syllables is kept constant but words of differing pronounciation directions are compared (e.g., “platoon” vs. “racket”), there is a small effect of word length but it appears to be confined to selected sets of stimuli (Lovatt, Avons, & Masterson, 2000).

Moreover, although the sheer volume of interest in the word length effect mandates its inclusion in our list of benchmarks, it is important to bear in mind that word length is a mere surrogate manipulation for the real variable of interest—viz. the effects of time on memory; we therefore suggest that examinations of word length are best replaced by other, more direct manipulations of encoding or retrieval time (e.g., Lewandowsky, Duncan, & Brown, 2004). We take up the role of time in STM in a later section.

1.2.6. *Data and theory.* Although our survey of benchmark results was necessarily brief and selective, it should nonetheless be apparent that the existing data base poses an enormous theoretical challenge. Can we find a tractable and coherent explanation for the overall pattern of results, given that no single model handles all phenomena listed in Table 1 (although some handle the list of benchmarks just reviewed)? Is it worth pursuing the possibility of an over-arching computational model or might it be preferable to strive for broad verbal explanations?

In our view, the precision and reliability of the data demands equally precise and exacting theorizing which can only be achieved by computational modeling (see Hintzman, 1991; Lewandowsky, 1993; Lewandowsky & Heit, 2006, for discussions of the benefits of computational modeling). Without such modeling, even seemingly clear-cut empirical dissociations defy unambiguous interpretation (Brown & Lamberts, 2003). We therefore restrict the remaining discussion to computational models. By implication, the influential working memory model of Baddeley and Hitch (e.g., 1974) is considered only in its explicit computational instantiations; that is, the primacy model (Page & Norris, 1998) and the models by Burgess and Hitch (e.g., 1999).

1.3. Varieties of Theories

A complete review of the available theories goes beyond the scope of this article; nonetheless, to place our own theoretical development into context, we find it helpful to at least provide a broad classification of models according to their architecture and underlying explanatory constructs. Table 2 provides this classification using 10 attributes that we find useful. (The last digits of the section numbers from here on refer to column numbers in the table to facilitate cross-referencing; the role of the shading of various cells in the table is discussed later.)

1.3.1. *Nature of representations.* Concerning representations, Page (2000) has provided arguments in favor of localist storage—that is, dedicating a unique place in memory to each item—whereas others (e.g., Hinton, McClelland, & Rumelhart, 1986) have provided arguments in favor of the opposite notion, namely that each item is represented by multiple units of information and that the same units contribute to storage of all items. The latter approach imlements distributed representations.

1.3.2. *Locus of similarity effects.* The locus of similarity effects, and in particular the need for a second output stage to explain phonological confusions, is likewise subject to debate (e.g., Page & Norris, 1998 vs. Lewandowsky & Farrell, in press) and we address this issue further below.

1.3.3. *Type of associations.* After the fading from prominence of chaining models (e.g., Lewandowsky & Murdock, 1989), the nature and type of associations is now considerably less controversial, and Table 2 shows the (near) consensus towards what we call "item marking;" viz. associations or "binding" between items and some independent representation of order such as time, temporal context, or ordinal list position. For example, in OSCAR (Brown et al., 2000), items are associated with a temporal "context" signal that autonomously evolves over time driven by a set of oscillators. The temporal distinctiveness mechanism in SIMPLE (Brown, Neath, & Chater, in press) can be understood in a similar manner. An alternative class of markers involves ordinal position rather than time, as instantiated in SEM (Henson, 1998a) and in the model we present below.

1.3.4. *Role of time at encoding.* It is important to underscore that item marking can equally involve time-based encoding—as in the oscillator-based OSCAR model or SIMPLE—or event-based encoding, as in SEM or the model we present below. In the former case, items are associated to a continuously changing signal whose temporal evolution is unaffected by encoding or retrieval events but depends on the passage of chronological time alone. In the latter case, items are associated to a signal that is unaffected by the passage of time per se but is assumed to evolve with retrieval and/or study events.

1.3.5. *Role of time during forgetting.* Within a given model, the role of time at encoding is not always mirrored by the role of time during retention; for example, the primacy model (Page & Norris, 1998) relies on temporal decay to explain forgetting whereas time plays no role during encoding. Conversely, even though OSCAR associates events to a temporal signal during encoding, forgetting is not driven by temporal decay.

1.3.6. *Primacy gradient.* Another property that characterizes many, but not all, models is the presence of a primacy gradient. A primacy gradient embodies the assumption that the quality of information available for retrieval of an item decreases across serial positions. In some cases, this is achieved by one or two weighting parameters that reduce encoding strengths across successive items (e.g., Brown et al., 2000; Houghton & Hartley, 1996; Lewandowsky, 1999; Lewandowsky & Li, 1994; Lewandowsky & Murdock, 1989; Page & Norris, 1998). By contrast, in SOB (e.g., Farrell & Lewandowsky, 2002) the primacy gradient arises as the natural consequence of weighting the encoding strength of any new item in proportion to its novelty.

1.3.7. *Response selection mechanism.* Response selection mechanisms fall into two broad classes, depending on whether or not they involve item marking and, hence, cueing by context. In context-free competitive cueing models ("CQ" in Table 2), items compete with each other for recall through mutual inhibition between items and self-excitation. In the most basic version of such models, this competition is based only on the intial strengths of the items (e.g., Page & Norris, 1998) whereas in other architectures the items' activations enter into several cycles of excitation and inhibition until they settle into a final pattern (e.g., Farrell & Lewandowsky, 2004). In contrast, models that rely on item marking typically postulate a CQ response-selection mechanism that processes the activations elicited by the context cue (e.g., Burgess & Hitch, 1999), with items whose stored contexts' most closely match the cueing context being more highly activated. Finally, in more abstract instantiations of the response selection mechanism, an item is selected not only in proportion to its own activation, but in proportion to its own activation compared to those of all other items, as described by Luce's choice rule (Luce, 1963).

1.3.8. *Output intereference.* There is considerable evidence that the act of recalling an item interferes with the accessibility of other items yet to be recalled (see Anderson & Neely, 1996, for a review). Accordingly, when input and output order are dissociated in serial recall, a contribution of output interference to primacy can be empirically identified (Cowan, Saults, Elliott, & Moreno, 2002; Oberauer, 2003). For example, Oberauer (2003) randomized the temporal and spatial input order of items by presenting them—separated in space and time—in a spatial array of boxes. Output order was dissociated from input order by randomly and successively cueing individual boxes for report of the presented item. Oberauer found that primacy was particularly pronounced when performance was plotted as a function of output position, aggregating across all input and spatial positions, suggesting that output interference contributes to the primacy effect. Given the evidence for output interference, it is perhaps surprising that this does not form a core assumption of any models of memory for serial order (though see Brown et al., 2000; Farrell & Lewandowsky, 2004).

1.3.9. *Response suppression*. A nearly universal property of models is the presence of response suppression. Response suppression refers to the assumption that each recalled item is temporarily suppressed and unavailable for further report. In any pure competitive cueing model with a primacy gradient (e.g., Page & Norris, 1998), in which the strongest or most active item is chosen for report without cueing by context, suppression is essential for maintaining serial reproduction. Without suppression, only the first (and strongest) item would be produced at each successive retrieval attempt. Suppression naturally accounts for the relative infrequency of erroneous repetitions noted earlier. Response suppression is also implicated by the fact that people are reluctant to repeat themselves during recall, even when a list does contain repetitions (e.g., Henson, 1998b). Because this deficit for report of a repeated item persists even in conditions in which people can detect repetitions, if instructed to do so, with more than 85% accuracy (Henson, 1998b), response suppression is typically understood to be an obligatory and automatic process.

1.3.10. *Energy-gated encoding.* Finally, Table 2 identifies SOB as the only theory to incorporate a process known as energy-gated encoding. Energy-gated encoding instantiates the idea that the strength of encoding into memory is a function of the novelty of the incoming information: Novel or surprising events are encoded strongly whereas familiar and unsurprising events that are similar to existing information in memory are encoded with reduced strength. Energy-gated encoding thus resembles the feedback signal instantiated in some other distributed memory models, such as the "closed-loop" version of TODAM (e.g., Lewandowsky & Murdock, 1989; Murdock, 1982).

1.4. Model Selection

Where do we go from here? How might one choose a preferred model from the set shown in the first two tables? By what criteria might one best propose a new model? The most obvious possibility, quantitative model comparison, suffers from the constraint that some results can be handled equally by different models, as is obvious from Table 1. A more specific illustration of this problem is provided in Figure 1, which shows the predicted serial position curves (left panel) and transposition gradients (center panel) of several models that are based on a variety of quite different explanatory constructs (whose identity is irrelevant for now but will be revealed later). Although there is some heterogeneity between predictions, they do not qualitatively diverge, which may make identification of the preferred model difficult.

We therefore pursue an alternative approach by seeking direct evidence for specific explanatory constructs. Although Tables 1 and 2 show that there are a variety of ways in which constructs can be combined to account for (largely) the same set of results, we now show that evidence has become available recently that favors some constructs over others.

2. Explanatory Constructs and New Data

We review four recent findings that can help identify the preferred explanatory constructs of short-term memory, either by providing new constraints or by re-evaluating earlier benchmark results in a new light. Specifically, (1) we present empirical support for energy-gated encoding and show that this evidence also favors a primacy gradient and speaks against the presence of a separate phonological confusion stage. (2) We then show that encoding and retention in memory are best understood without assuming that time plays a causal role; (3) we present evidence for the existence and role of response suppression; and (4) we underscore the importance of item marking (i.e., item-to-context associations). The constructs that are directly supported by evidence are shaded gray in Table 2.

2.1. Energy-Gated Encoding

Although all viable models of short-term order memory accommodate the basic phonological similarity effect, their predictions differ for lists on which similar and dissimilar items are inter-mixed. A number of early studies (e.g., Baddeley, 1968; Bjork & Healy, 1974; Henson et al., 1996) showed that dissimilar items on mixed lists were recalled with the same accuracy as those on pure dissimilar lists. The immunity of dissimilar-item recall to list composition compelled several theorists to propose that serial recall involves two independent stages of processing, with order errors occurring between positionally-encoded tokens in a primary stage, and with similarity-based confusions occurring in a separate secondary stage that only affects similar items but lets dissimilar items pass unhindered (e.g., Burgess & Hitch, 1999; Henson, 1998a; Page & Norris, 1998).

However, the absence of a mixed-list advantage is arguably counter-intuitive. Consider an extreme case of a single dissimilar item being embedded in a similar-sounding list (e.g., the letter *X* in the list B D G *X* T P); on any notion of distinctiveness, the dissimilar item should be recalled more accurately from this list than when it is surrounded by other mutually dissimilar items. Several recent studies have re-evaluated the mixed-list effect and have concluded that when guessing is controlled (unlike in the earlier experiments), dissimilar items on mixed lists are in fact recalled more accurately than their counterparts on pure lists (Farrell, 2006; Farrell & Lewandowsky, 2003; Lewandowsky & Farrell, in press). One implication of these new, revised mixed-list effects is that all models that predict identical performance for dissimilar items regardless of list context—i.e., those identified as handling the "old" mixed-list effect in Table 1—are challenged by the empirical re-evaluation of that phenomenon.

In support, a comparison of several theories by simulation by Lewandowsky and Farrell (in press) and Farrell (2006) showed that the mixed-list advantage for dissimilar items cannot be accommodated by two-stage models (the primacy model by Page & Norris, 1998, and the SEM by Henson, 1998a) and instead supports an alternative view based on a single memorial stage involving energy-gated encoding (i.e., SOB; Farrell & Lewandowsky, 2002, Lewandowsky & Farrell, in press).

Energy-gated encoding refers to the weighting of the encoding strength of new list items on the basis of their novelty compared to already-encoded information. We use the term "energy-gated" because within the SOB model, novelty of an item is measured by its "energy" with respect to the weight matrix of previously stored associations (Farrell & Lewandowsky, 2002). Items that differ from earlier memorized information are encoded more strongly than items that are similar to the contents of memory. In consequence, dissimilar items on mixed lists receive a memorial advantage because the mutual similarity of the list neighbors implies that they are encoded less strongly. On a pure dissimilar list, by contrast, list neighbors are all dissimilar to each other and hence encoded with greater strength, thus providing more competition to each other. To illustrate, consider the expected gradients of encoding strengths for two three-item lists: one consisting only of dissimilar items (referred to as “D” in abstract notation; hence, DDD) and one that has a dissimilar item sandwiched between two items that are similar to each other (SDS). The first item is stored with the same strength on both lists, as there are no previous items which could affect its novelty. The second item is also stored with the same strength in both cases, as it is (likely) equally novel with respect to the first item. However, the third item on the SDS list receives a smaller weight than its sibling on the DDD list because its similarity to the first item reduces its encoding strength by the energy-gated encoding scheme. It follows that the isolated dissimilar item on the list SDS will be recalled more accurately than its counterpart on the DDD list because the third item on the SDS list, having been encoded with less strength, will offer reduced competition during recall. Because the recent data (Farrell & Lewandowsky, 2003; Lewandowsky & Farrell, in press) conform to this prediction, we identify energy-gated encoding as a preferrred explanatory principle in short-term memory.[2]

Energy-gated encoding, in turn, naturally gives rise to other explanatory constructs, including the primacy gradient. Although the primacy gradient is often motivated by the need to predict a recall advantage for early list items (e.g., Brown et al., 2000; Lewandowsky, 1999; Lewandowsky & Murdock, 1989), there are independent reasons for its existence. For example, Brown et al. (2000) justified the primacy gradient by appealing to the “intuition that each successive item ... is progressively less ‘surprising’ or attention-demanding than the previous one” (p. 151), an intuition that Brown et al. considered to be consonant with the demands on an adaptively rational organism. This intuition is formalized by energy-gated encoding, which naturally yields a primacy gradient (Farrell & Lewandowsky, 2002).

2.2. No Role for Time

Do short-term memory representations involve a temporal component? Are memories temporally organized? Do they inexorably decay over time? Or are our memories built and altered on the basis of events, irrespective of the time that lapses between events? Few issues have been debated as intensely as this fundamental question concerning the nature of our memories. Turning first to the role of time at encoding, we consider whether temporal distinctiveness is a useful principle to characterize performance in short-term memory. According to temporal distinctiveness theories (e.g., Bjork & Whitten, 1974; Brown et al., in press; Crowder, 1976; Glenberg & Swanson, 1986), the temporal separation of events at encoding is a crucial determinant of memory performance. All other things being equal, distinctiveness models predict that the memorability of an event increases with its temporal separation from neighboring events. In consequence, people are more likely to remember details about widely separated events (e.g., biennial holidays) than events that are more closely spaced (e.g., frequent trips to conferences).

Contrary to this expectation, there have been numerous recent reports showing that temporal isolation does not facilitate serial retrieval from short-term memory. Across a number of studies employing a variety of retrieval tasks (serial recall, probed order recall, serial recognition) and presentation modality (auditory vs. visual), it has been demonstrated that items closely preceded or followed in time by other items are no worse recalled than items that are temporally isolated (e.g., Lewandowsky, Brown, Wright, & Nimmo, 2006; Nimmo & Lewandowsky, 2005, 2006). There are, however, two known exceptions to this general conclusion: First, isolation can benefit memory if the presentation schedule is predictable and people can use it for strategic encoding of the list (Lewandowsky, Wright, & Brown, in press). Second, isolation can benefit memory if output order is unconstrained; that is, when people are free to choose which item to report (by selecting a serial position for report in a reconstruction task), isolated items are recalled better than temporally crowded items (Lewandowsky, Nimmo, & Brown, in press).

We acknowledge those exceptions, and we acknowledge that they mandate a mechanism by which people can switch attention to time and use it to govern retrieval (even retroactively, after list presentation; see Farrell & McLaughlin, in press; Lewandowsky, Nimmo, & Brown, in press). Nonetheless, we set aside those exceptions here because our declared focus is exclusively on forward serial retrieval. Within those constraints, we identify the *absence* of temporal representations as the preferable explanatory principle of STM. By implication, we reject the idea that temporal distinctiveness presents a viable explanation for short-term serial recall, although we acknowledge the power of the approach in other domains of memory such as free recall (see Brown et al., in press; Brown, Morin, & Lewandowsky, 2006).

Turning to the role of time during retention, we have already noted the limitations of the word length effect. Lewandowsky, Duncan, and Brown (2004) introduced a paradigm in which retention time was experimentally manipulated while simultaneously blocking rehearsal (to avoid potential refreshing of memory traces) without introducing extensive interference. Participants orally recalled a list of letters while repeating an irrelevant word ("super") aloud in between retrievals, thus preventing articulatory rehearsal. Lewandowsky et al. found that recall performance was unaffected by the number of times people articulated the distractor in between memory retrievals, implying that time per se did not cause forgetting in serial recall. We conclude that processes other than the passage of time are responsible for forgetting in STM.

2.3. Response Suppression

Whereas output interference refers to a deleterious effect of retrieval on other items that have yet to be recalled, response suppression refers to a temporary inability to access an item just recalled. Notwithstanding the widespread appeal to response suppression, it has proven difficult to find direct evidence for its existence. Farrell and Lewandowsky (2007) tackled this issue by noting that many models rely (at least in part) on response suppression to explain recency (e.g., Brown et al., 2000; Henson, 1998a; Farrell & Lewandowsky, 2002; Lewandowsky & Murdock, 1989). When items must be reported in serial order, any competitive response selection mechanism predicts recency if the number of potential response candidates decreases across output position—a decrease that is achieved very naturally by response suppression. A causal link between response suppression and recency entails the prediction that the fewer list items are recalled (thus leaving more list items unsuppressed), the less recency will occur. Because a list item will remain unsuppressed if it is replaced by an intrusion (i.e., recall of an extra-list item), the extent of recency should decline with the number of intrusions during recall. Importantly, this prediction can be tested while controlling absolute accuracy: if two items are transposed during recall (i.e., two items are recalled in the wrong position), accuracy—as measured by proportion correct recall—is identical to the case in which two intrusions are reported; however, in the former case two more list items are suppressed than in the latter.

Figure 2 shows the results of a conditional re-analysis of 7 published and unpublished serial-recall experiments conducted by Farrell and Lewandowsky (2007). For all studies, only those lists were considered on which people committed exactly two errors during recall of all items bar the last one. Performance on the last item was then examined as a function of how many list items were suppressed during commission of the two errors, which ranged from 0 (two intrusions) to 2 (two transpositions). The results, shown by the line labeled “forward” in the figure, are quite straightforward: Performance on the last item increased with the number of list items that were reported previously, exactly as predicted by the suppression account of recency.

One potential problem with this analysis is that it might reflect a correlation between two measures of memorial information; even though accuracy is identical between the different mixtures of errors, one might argue that intrusions occur only when both item and positional information is lost whereas transpositions reflect loss of positional information in the presence of intact item information. On that view, two intrusions might be associated with lower terminal-item performance than two transpositions not because of response suppression but because of overall poorer memory for those lists.

This problem can be addressed by reversing the directionality of conditionalization and considering performance on the early part of the list as a function of errors during the later part. The line labeled “backward” in Figure 2 shows performance on the first three serial positions as a function of the mixture of error types observed during recall of the remaining list items (again keeping accuracy constant by considering only those lists on which exactly two errors were made). In contrast to the forward conditionalization, the mixture of late-list errors was not predictive of accuracy for early-list items. This result allays fears that the forward conditionalization merely revealed a correlation between two measures of memorial quality.

We therefore identify response suppression as a viable explanatory construct in short-term memory. We additionally conclude that response suppression contributes to the occurrence of recency.

2.4. Item Marking

The first two panels in Figure 1 demonstrated that the serial position curve and underlying transposition gradients can be explained by several rival constructs. Farrell and Lewandowsky (2004) showed that when response latency is considered in addition to response proportions, the various competing explanations can be empirically differentiated on the basis of their predicted “latency-displacement functions.”

The right-most panel of Figure 1 shows a family of predicted latency-displacement functions, which plot response latency as a function of the displacement of the recalled item. A displacement of zero refers to a correct response, a negative displacement refers to an anticipation (e.g., the third list item recalled first would be a displacement of -2), and a positive displacement constitutes a postponement (e.g., the first item recalled in the second position would be a +1 displacement). It is clear from the figure that different explanatory constructs can predict qualitatively different latency-displacement functions.

The predictions shown in the panel were obtained using a common architecture, based on localist representations within an iterative competitive activation network. The network permitted the instantiation of various explanatory constructs in a common framework and generated both the probability and latency of all possible responses (see Farrell & Lewandowsky, 2004, for more detail). The panel shows predictions from four explanatory constructs; namely, a primacy gradient (labeled PR in the figure), response suppression (RS), output interference (OI; instantiated by increasing the amount of noise added to the network across output positions); and an item marker (IT; instantiated by activating items according to the match between the position currently being cued and each item’s list position). The predictions differ considerably between the various combinations of constructs instantiated in the network: Whereas an item marker on its own produces completely symmetric and non-monotonic latency-displacement functions (with the fastest responses reserved for correct retrievals), the presence of a primacy gradient in conjunction with response suppression qualitatively alters the shape of the function and renders it monotonically negative. Addition of an item marker to the combination of a primacy gradient and response suppression flattens the slope of the function for postponements, without however removing monotonicity.

Farrell and Lewandowsky (2004) reported three experiments involving timed keyboard recall of lists of letters or digits. Across all experiments, the observed latency-displacement functions were consistently monotonic (or nearly so) with a negative slope, and additionally exhibited a reduction in slope for postponements compared to anticipations. As seen in Figure 1, these data are most compatible with a model that includes a primacy gradient, response suppression, and an item marker.[3] The data clearly compromise models that are based on an item marker alone.

2.5. Summary

We argue that the evidence just discussed has considerably narrowed down the preferred set of explanatory constructs summarized in Table 2. With the exception of Columns 1 and 7, which are neutral with respect to the available evidence, only the shaded cells in Table 2 are compatible with the set of recent results on STM. We next present a computational model of serial recall in short-term memory that is based on those constructs. The model incorporates a set of principles that we have been elaborating for several years. The formal instantiation we present here brings together a simpler serial recall model that handled basic phenomena in a dynamic connectionist network (Farrell & Lewandowsky, 2002), and a generalized extension that we have recently implemented to account for similarity effects (Farrell, 2006; Lewandowsky & Farrell, in press).

3. C-SOB: A Model of Serial Recall in Short-Term Memory

Farrell and Lewandowsky (2002) presented a distributed model of serial ordering known as SOB, or "Serial-Order in a Box," a name that acknowledges the model's reliance on the Brain-State-in-a-Box algorithm (e.g., Anderson, Silverstein, Ritz, & Jones, 1977). Briefly, SOB assumes that items are represented by vectors of features that are encoded into memory by adding their auto-associations to a common weight matrix. That is, the core architecture of the model consists of the continuous super-imposition of new information onto items already presented. A crucial property of SOB is that encoding strengths are a direct function of the novelty, or "energy," of incoming items which naturally gives rise to a primacy gradient.

Retrieval from SOB involves non-linear iterative dynamics. Memory is probed by presenting a cue vector to the weight matrix. In contrast to many other networks, the first response of the model is not taken to be its final answer: Instead, the output is fed back into the weight matrix across multiple iterations until a stables state, known as an “attractor,” is reached. If the model correctly recalls an item, that final state will be identical to the target item, whereas in the case of an erroneous recall, the model will reach an attractor that differs from the target. This iterative dynamic deblurring mechanism very naturally gives rise to latency predictions because each response in SOB takes a measurable amount of time to emit.

Farrell and Lewandowsky (2002) showed that SOB could explain a number of benchmark phenomena, including the basic shape of the serial position curve; the pattern of errors during recall, including the balance between transpositions, omissions, intrusions, and erroneous repetitions; the effects of list length on the distribution of errors, the overall level of recall and response latency; and the effects of natural language frequency on recall performance.

In addition to providing an account of some benchmarks, the initial SOB resolved two problems that had previously been ascribed to distributed representations (e.g., Page, 2000); namely, the presumed inability to disambiguate responses and selectively to suppress items. The first problem, disambiguation, refers to the fact that except in special circumstances, distributed models of memory do not reproduce an exact copy of a studied item but a “blurry” approximation. SOB solves this problem because its dynamic deblurring mechanism is guaranteed to converge onto an unambiguous and identifiable response. The second problem, concerning suppression, pertains to the difficulties that can be associated with the selective removal of specific information from neural networks. SOB solves this problem by the use of anti-learning, which is instantiated by removing the auto-association of each recalled item from the weight matrix by “relearning” it with a negative learning rate (see Anderson, 1991, for an introduction to anti-learning).

The core assumptions of the initial SOB model have remained unchanged to date and are embodied in the C-SOB model being presented here. There are two major differences between SOB and C-SOB: First, whereas all items in SOB were necessarily orthogonal, this constraint is relaxed in C-SOB and inter-item similarity is free to vary. Second, whereas SOB initiated retrieval by a random cue, in C-SOB items are associated to an item marker that is used as a retrieval cue. In recent applications of C-SOB in which the selection stage was not considered crucial to the predictions of the model, the computationally demanding iterative deblurring mechanism was replaced by a simpler response selection method using Luce’s choice rule. Here we focus on the complete version of C-SOB that incorporates the fully specified dynamic deblurring stage.

3.1. Architecture and Specification

3.1.1. *Network architecture*. The model consists of two layers of units; an input layer (N = 16) used to represent item markers, and an output layer (N = 150) representing list items, where N refers to the number of units in each layer. The two layers are fully interconnected by a weight matrix **C** (weights unidirectionally projecting from the input layer to the output layer), and an additional weight matrix **W** representing full interconnectivity between units in the output layer. The weight matrix **C** stores associations between positional markers and items, whereas the weight matrix **W** stores auto-associations (i.e., associations of items to themselves) that drive the dynamic disambiguation. The matrix **W** and its dynamic disambiguation properties are isomorphic to the core architecture of the initial SOB model (Farrell & Lewandowsky, 2002).

3.1.2. *Positional markers*. C-SOB incorporates a marker for each list position to which items are associated (cf. Henson, 1998a). The similarity between any two positional markers is an exponential function of their absolute separation in list positions; that is,

$$\cos(\mathbf{p}_i, \mathbf{p}_j) = t_c^{(|i-j|)}, \tag{1}$$

where i and j are the positions of the ith and jth items, $\mathbf{p}_i$ and $\mathbf{p}_j$ are distributed vectors representing positional markers, and t_c is a constant that was set to .8 in the majority of simulations below. Markers were generated from a weighted combination of orthogonal Walsh vectors of dimensionality 16, with the weights set to satisfy the similarity constraint embodied in Equation 1. Note that the positional markers are not affected by the passage of time but are identified by position alone.

3.1.3. *Pre-experimental learning.* C-SOB assumes that participants come to serial recall experiments armed with prior knowledge of the experimental vocabulary, which plays a fundamental role in disambiguating noisy traces. In the original version of SOB, pretraining was accomplished by storing autoassociations in the weight matrix using Hebbian learning. In C-SOB, knowledge about items is captured by auto-associations that are pre-trained using Widrow-Hoff learning:

$$\Delta\mathbf{W}_i = \eta_P \mathbf{v}_i (\mathbf{v}_i - \mathbf{o}_i)^T, \tag{2}$$

where $\mathbf{v}_i$ is the item being learned, $\mathbf{o}_i$ is the current output of the network obtained when cueing with $\mathbf{v}_i$ (that is, $\mathbf{W}_{i-1}\mathbf{v}_i$) , and η_P is the learning rate. Items are represented by binary vectors whose elements are set to +1 or –1. For each simulation below, training lasted for 200 cycles; with each cycle k involving training of all items with learning rate η_P =.03/k. (Reducing the learning rate across pretraining cycles was found to facilitate learning). The use of Widrow-Hoff learning is required when lists contain similar items; if all items are assumed to be orthogonal, Hebbian learning may be employed instead (Farrell & Lewandowsky, 2002).

3.1.4. *Encoding of study list.* List items are associated with successive positional markers using standard Hebbian learning (see, e.g., Anderson, 1995):

$$\Delta \mathbf{C}_i = \eta_e(i)\mathbf{v}_i\mathbf{p}_i^T \,, \tag{3}$$

where $\mathbf{C}$ is the matrix of positional marker-item weights, $\mathbf{v}$ is the vector representing the *i*th presented item, and $\mathbf{p}$ is a positional marker for the *i*th serial position. The learning rate η_e for the *i*th association, $\eta_e(i)$, was determined anew for each item using the energy between the association to be learned and the information in the weight matrix up to that point:

$$\eta_e(i) = \begin{cases} 1, & i = 1 \\ -\phi_e / E_i, & i > 1 \end{cases}, \tag{4}$$

where ϕ_e is a free parameter, and E_i, the energy of the *i*th association, is given by

$$E_i = -\mathbf{v}_i^T \mathbf{C}_{i-1}\mathbf{p}_i \,. \tag{5}$$

3.1.5. *Retrieval*. Retrieval consists of stepping through the positions using the markers to cue for their associated items, and then using pre-trained knowledge to deblur the retrieved ambiguous information. Retrieval at position *i* is cued by placing the positional marker $\mathbf{p}_i$ across the input layer, and computing the item unit activations:

$$\mathbf{v}_i' = \mathbf{C}_i\mathbf{p}_i \,. \tag{6}$$

The retrieved vector, $\mathbf{v}_i'$, is a "noisy" version of $\mathbf{v}_i$, containing a blend of the target $\mathbf{v}_i$ and the other items on the list, according to the overlap between the positional markers for those other items and the marker for the target. This noisy output is then disambiguated using dynamic iterative deblurring as follows.

The pattern of activations $\mathbf{x}$ across units in the item layer is intialised to $\mathbf{v}_i'$ and scaled to unit length to yield an intial state $\mathbf{x}(0)$. The activations in $\mathbf{x}$ are then iteratively updated according to:

$$\mathbf{x}(t+1) = \mathrm{G}[\gamma\mathbf{x}(t) + \alpha\mathbf{W}\mathbf{x}(t)] \,. \tag{7}$$

The two terms inside the function G in Equation 7 are the activation patterns from the previous iteration *t* (weighted by γ; included as an inertia term, Anderson, 1995), and the activations resulting from passing the previous state $\mathbf{x}(t)$ through the weight matrix $\mathbf{W}$. This crucial second term uses the knowledge in $\mathbf{W}$ to reconstruct items; effectively, each item that contributes to the "blend" $\mathbf{v}_i'$ cues for itself using the stored autoassociation in $\mathbf{W}$. Generally, the extent to which an item is successful in cueing for itself will depend on the extent of its storage in $\mathbf{W}$ and the weight of its contribution to the blend $\mathbf{v}_i'$. Because all items are pre-trained with an equal weight, the only systematic contribution to the updating will result from the weighting of items in $\mathbf{v}_i'$.

To prevent this iterative updating process from continuing indefinitely, the function G in Equation 7 “squashes” the updated activations to lie between –1 and +1. Updating terminates when all units are saturated (that is, all units equal +1 or –1), thus ensuring that deblurring ends in a final state that corresponds to a vertex (i.e., corner) in a zero-centered hyperspace. (It is for this reason that the algorithm is referred to as “Brain-State-In-a-Box.”)

Although all list items, by definition, are verteces, the reverse is not true, and the model can therefore generate extra-list intrusions. To model omissions, we assume a temporal criterion for convergence; if I_{max} iterations have passed without achieving convergence, a “pass” response is emitted. As in SOB, because each iteration is assumed to take a fixed amount of time, the time to convergence can be treated as recall latency, thus allowing the model to predict latencies for each response.

Once an item is recalled, it is suppressed by adjusting the weights between the positional marker layer and the item layer according to

$$\Delta\mathbf{C}_j = \eta_s(j)\mathbf{v}_{o,j}\mathbf{p}_j^{\ T}, \tag{8}$$

where j is the output position, and $\mathbf{v}_{o,j}$ is the recalled item; Equation 8 is identical to the Hebbian learning used at encoding, except that the learning rate is negative, thus causing “anti-learning” of the weights and attenuation of item representations (e.g., Anderson, 1991). To ensure that the extent of response suppression approximately matches that of learning, the learning rate for suppression, η_s, is also determined from the energy of the recalled item with respect to the association matrix $\mathbf{C}$ and the positional marker $\mathbf{p}$

$$E_j = -\mathbf{v}_{o,j}^{\ T}\mathbf{C}_{j-1}\mathbf{p}_j. \tag{9}$$

The weighting of response suppression is given by

$$\eta_s(j) = \frac{-E_j}{\phi_s E_1}, \tag{10}$$

where ϕ_s is a model parameter, and E_1 is the energy of the first recalled item (see Farrell & Lewandowsky, 2002); this term reduces to $-\phi_s^{\ -1}$ for the first output position.

Response suppression is followed by generalized output interference due to retrieval of the item. Following Brown et al. (2000), output interference is implemented by adding Gaussian noise to all weights in $\mathbf{C}$ with mean 0 and standard deviation σ_O.

3.2. Account of Benchmarks

To establish the basic viability of C-SOB, we now show that it accounts for the usual set of benchmark results. For continuity with recent applications (Farrell, 2006; Lewandowsky & Farrell, in press), the items in all simulations were assumed to be moderately dissimilar letters, and were constructed using the known multidimensional phonological similarity structure of the stimuli (see Lewandowsky & Farrell, in press).

All simulations used a constant set of parameter values that were not estimated from the data but were based on previous applications of the model. The encoding parameter ϕ_e was set to 720, and ϕ_s was set to 0.8 (i.e., in the centre of the range of previously used values; viz. .45; Lewandowsky & Farrell, in press, to 1.2; Farrell, 2006). The convergence parameters γ and α were set to .2 and 1.1, respectively, similar to the values used in Lewandowsky and Farrell (2000). The omission threshold χ was arbitrarily set to 100, and the amount of output interference σ_O was set to 1.

3.2.1. *Serial position curves*. Figure 3 shows the predicted serial position curves for serial recall across list lengths 3 through 7. The left panel shows the accuracy predictions and the right panel the associated predicted latencies. The model captured the essential aspects of the accuracy data: It exhibited "fanning" from a common origin with increasing list length, it produced extensive primacy, and it showed limited recency.

Likewise, the model captured the known relationship between list length and cumulative response latencies (e.g., Dosher & Ma, 1998), particularly the fanning in the cumulative latencies, which reveals that making lists longer increases the time to recall individual items; a serial process in which items took a fixed amount of time to recall would predict the curves to be exactly overlaid.

3.2.2. *Transposition gradients*. The transposition gradients underlying the preceding serial position curves are shown in Figure 4 (averaged across all serial positions). The transposition gradients exhibit the usual "locality constraint"; that is, most errors involve items from neighboring list positions. Unlike the data, however, they show a slight asymmetry, with more anticipations than postponements at most displacement distances.

The predicted relative frequency of fill-in ("BA…") and infill ("BC….") errors was found to vary considerably with list length, with the ratio of fill-ins to infills decreasing from 53 to 10, 2, .8, and .5, respectively, for list lengths 3, 4, 5, 6, and 7. We know of no corresponding behavioral data that relate fill-ins to list length, with all available reports limited to list lengths 6 to 8 and citing a ratio of around two (Henson, 1996; Henson et al., 1996; Surprenant et al., 2005). It is clear, though, that C-SOB under-predicts the ratio of fill-ins to infills for the list lengths for which data are available.

3.2.3. *Latency-displacement functions*. Figure 5 shows the predicted latency-displacement functions for a number of different list lengths (averaged across serial positions in each case). The functions mirror the (roughly) monotonically negative trend in the data. The figure also confirms the modeling by Farrell and Lewandowsky (2004), which suggested that a primacy gradient in conjunction with response suppression is necessary to achieve negative monotonicity in the latency-displacement functions.

Moreover, the reduction in slope associated with postponements (i.e., positive displacements) mirrors the effects of introducing an item marker in addition to a primacy gradient and response suppression (see the right-hand panel of the earlier Figure 1), thus confirming that Farrell and Lewandowsky's analysis of a generic network architecture transfers to a completely specified instantiation.

3.2.4. *Item and repetition errors*. The predicted proportions of the various error types across output positions are shown in Figure 6 for list lengths 5 (left panel) and 6 (right panel). In accord with the data, item errors (intrusions and omissions) increase across output positions whereas transpositions show an inverted U-shaped function, with a reduction in transpositions for the last one or two positions.

C-SOB also accounts for the proportion of repetition errors; for 6-item lists, 2.7% of all responses were repetitions, compared to values between 2% and 5% reported in the literature (Henson, 1996; Vousden & Brown, 1998). However, C-SOB fails to account for the separation of repetitions: Whereas in the data repetitions tend to be separated by 3 or 4 items (e.g., Henson, 1996), C-SOB predicted a considerable number of immediate repetitions. This result may reflect the fact that as recall proceeds, the current context marker will increasingly differ from the context that generated the first occurrence of a possible repetition, making remote repetitions unlikely.

3.2.5. *Similarity effects*. Lewandowsky and Farrell (in press) provided a detailed analysis of SOB's ability to handle the effects of phonological similarity, and we briefly summarize their findings here. (Note that their model did not include dynamic deblurring but was identical to C-SOB in all other respects.)

Figure 7 is adapted from their paper and shows the predicted effects of similarity for pure lists (i.e., DDDDDD and SSSSSS) and various mixed lists, including two in which only a single dissimilar items was present in position 2 (SDSSSS) and 4 (SSSDSS), respectively. The figure clarifies that SOB predicts the detrimental effects of similarity and it also underscores the mixed-list advantage for dissimilar items—discussed earlier—that is a necessary consequence of energy-gated encoding.

3.2.6. *Effects of grouping.* To model the effects of grouping, the item markers used in the preceding simulations were modified to a two-dimensional representation with the addition of a further 16 units which coded the position of the item within a group (as well as the position of the item in the list, carried over from previous demonstrations; cf Brown et al., 2000; Burgess & Hitch, 1999). The same parameter t_c was used to determine contextual similarity for both dimensions of representation, and was set to .6. To account for the increased number of weights projecting from the context layer to the item layer, ϕ_e was set to the larger value of 1200.

The resulting predictions are shown in Figure 8. The left-most panel shows that C-SOB predicts the scalloping of the serial position curves that is characteristic of grouped lists; the center panel shows the associated latency predictions; and the rightmost panel shows the underlying transposition gradients. The grouped lists exhibit the classic peaks in the transposition gradients that reflect the increased likelihood of interpositions (i.e., transpositions of items from two groups in the same within-group positions).

3.2.7. *Summary.* Our simulations confirmed that C-SOB can handle the conventional set of benchmark results. As shown in Table 1, this confirms that C-SOB ranks among competing alternatives in its ability to handle core results. Morevover, we showed that the model handles the particularly diagnostic mixed-list similarity effects. We next show that C-SOB makes novel predictions by turning to a further analysis of the effects of time on forgetting from short-term memory.

4. Experimental Predictions

4.1. Forgetting: Time vs. Interference

We noted earlier that the word length effect, frequently cited in support of temporal decay, is fraught with problems of interpretation and replicability and we pointed to alternative means of examining the issue, as in the studies of Lewandowsky et al. (2004) in which increasing a distractor-filled delay at retrieval did not have a deleterious effect on performance. Oberauer and Lewandowsky (2007) extended the methodology of Lewandowsky et al. (2004) by comparing the effects of the irrelevant task to a quiet control condition and by adding an attention-demanding speeded choice task (with arbitrary stimulus-response mappings) on top of the articulatory suppression in between memory retrievals. Notwithstanding this added irrelevant load, manipulating the time in between retrievals—by varying the number of articulations and choices that had to be performed—had virtually no effect on memory performance.

The data from two of Oberauer and Lewandowsky’s (2007) experiments are shown in the two panels of Figure 9, together with the best-fitting predictions of C-SOB (Oberauer and Lewandowsky omitted the dynamic deblurring but their model was identical to C-SOB in all other respects). The figure makes two noteworthy points: First, as shown in the left panel, the addition of an irrelevant articulatory task depressed performance considerably compared to a quiet control (labelled B in the figure), but there was little indication that extending the duration of articulation (3R vs. 1R) led to a further decrement in performance.[4] Second, as shown in the right panel, even if articulation is accompanied by an additional speeded-choice task, the absolute duration of those two tasks (1 vs. 4 AS+CRT) has no effect on memory performance. The figure also shows that C-SOB captures all those effects; namely, the dramatic decline of performance with introduction of an irrelevant task and the time-invariance of performance across varying durations of the irrelevant activity (or activities).

The procedures of Oberauer and Lewandowsky (2007) and of Lewandowsky et al. (2004) are modelled by assuming that the distractor item (i.e., the irrelevant to-be-articulated word) is encoded into memory by being associated to a derivative of the current context marker. To model the differences in distractor duration between conditions, the distractor word was added either once or several times in between retrievals, mirroring the way in which people articulate once or multiple times. (For further details of these simulations, see Oberauer & Lewandowsky, 2007.) When combined with energy-gated encoding, these representational assumptions give rise to the best-fitting predictions in Figure 9. Specifically, the first time a distractor is encountered, it receives a large encoding weight because C-SOB identifies it as novel. In consequence, the update to the weight matrix is quite distruptive of existing memories. If there are additional repetitions of the same distractor, those repetitions receive minimal encoding because they are identified as having little novelty. In consequence, there is little additional interference compared to the single-distractor condition.

4.2. Distractor Structure and Interference: An Experimental Test

The energy-gated encoding of distractors yields strong predictions involving the similarity structure of the to-be-articulated material. To illustrate, consider the extreme case in which every articulatory distractor consists of a different word; that is, people might say “table – horse – truck” after recall of the first item and “orange – zucchini – car” after the second item and so on. Within the energy-gated encoding framework, every distractor in this case is relatively novel and would be encoded quite strongly, hence maximizing the amount of interference and giving rise to considerable forgetting that increases in proportion to the number of articulations in between retrievals.

Figure 10 shows the complete set of predictions for all possible combinations of novelty within and between “bursts” of articulatory distractors. (A burst refers to the articulation sequence in between retrievals and thus refers to either 1 or 3 words being spoken aloud after recall of each list item.) The two lines shown in each panel represent one and three distractors, respectively, in between each retrieval. The top panels show predictions for "simple" bursts; that is, those involving repetition of the same word (as used by Oberauer & Lewandowsky, 2007), whereas the bottom panels show predictions for "complex" bursts (consisting of three different words, not used in experiments to date). The left-hand panels show the predictions when bursts are steady across output positions; that is when people have to repeat the same set of (one or three) items after each retrieval. The right-hand panels, by contrast, show predictions when the identity of bursts changes across output positions. All predictions were obtained using the best-fitting parameter estimates and representational assumptions used by Oberauer and Lewandowsky (2007).

In summary, the top-left panel corresponds to people saying the same word either once or three times in between all retrievals; the top-right panel represents people saying a word either once or three times after a retrieval, with the identity of that word changing across output positions. Finally, the bottom panels correspond to people articulating three different words at each output position, with the identity of those three words either repeating across output positions (left-hand panel) or changing after each retrieval (right-hand panel). In addition to predicting no fanning for simple bursts (cf. Oberauer & Lewandowsky, 2007), C-SOB predicts only limited additional impairment when the distractor changes across serial positions (top left vs top right). However, changing the distractor within a burst introduces marked fanning (top vs bottom panels), with this fanning being most pronounced when the distrctor bursts are repeated across output positions (bottom left vs bottom right).

C-SOB's predictions are subject to fairly straightforward experimental tests. In a first, as-yet unpublished experiment, conducted by Sonja Geiger and the first author, people were given a predictable and well-rehearsed sequence of distractors; namely, the months of the year. People studied lists of 5 letters for immediate forward serial recall. During recall, each oral report of a list item was followed by bursts of articulation of 1 or 3 names of months. The nature of the bursts (simple; "January-January-Janurary" vs. complex; "January-February-March") was manipulated between subjects, whereas the remaining two variables (number of distractors and whether or not the identity of distractors changed across output positions) were manipulated within subjects.

The results of this experiment are shown in Figure 11 using the same layout of panels as for the C-SOB predictions. Comparison of this figure with the predictions in Figure 10 confirms that the study in large part supported the predictions of the model. There was no fanning for the simple bursts, irrespective of whether or not they changed across output positions (mean difference between 1 and 3 distractors across serial positions 2 through 5 was a negligible .008, averaged across steady and changing bursts); there was fanning for complex bursts, and the extent of that fanning was greater for the steady condition (mean difference between 1 and 3 distractors across serial positions 2 through 5 was .1337) than the changing condition (mean difference .0382). The one prediction that failed to be supported concerned the differences between steady and changing simple bursts: Whereas SOB expected the latter to lead to slightly more interference than the former, the data showed no such effect.

To place these results into a wider theoretical context, it must be noted that a temporal view of forgetting would expect *all* types of articulation to engender the divergence between 1 and 3 distractors that in the experiment was limited to complex bursts. That is, because three repetitions of the word "January" take (roughly) as much additional time (compared to a single repetition) as does articulation of the sequence "January – February – March," a time-based model must expect equal forgetting in each case. The data clearly contradict this expectation.

We conclude that C-SOB's predictions differentiate it from theories that assign a role of time during forgetting (e.g., the primacy model, Page & Norris, 1998; SIMPLE, Brown et al., in press), and that the successful test of those predictions (the results of the study just discussed were replicated and extended in two additional experiments conducted by Sonja Geiger and the first author) further supports one of the constructs on which C-SOB is built; namely, that the passage of time per se plays no causal role in short-term memory. This encouraging outcome suggests that further exploration of C-SOB is warranted.

5. General Discussion

5.1. Limitations of C-SOB

Although we have shown that C-SOB can accommodate acknowledged benchmark effects in serial recall, and have demonstrated that the theory's predictions can be successfully tested, these successes were accompanied by two design decisions that in turn engendered two potential limitations: First, C-SOB explicitly rejects the notion that elapsed time plays a causal role in short-term memory and, second, the theory as stated cannot be applied to multi-trial effects. We discuss those limitations before placing C-SOB into a wider theoretical context.

5.1.1. *The role of time in short-term memory.* By rejecting a causal role of time, C-SOB is an unequivocally event-based theory, and thus differs from several other current contenders, such as the primacy model (Page & Norris, 1998), SIMPLE (Brown et al., in press), or the model by Burgess and Hitch (e.g., 1999, 2006). Is this strong theoretical commitment justified in light of the available data? Our response is twofold and considers the effects of time at encoding and during retention separately.

Concerning retention, we propose that there is now ample evidence to suggest that forgetting in short-term memory is not caused by the passage of time per se. We presented some of the relevant studies earlier, and there is now clear evidence that if retrieval is delayed by distractor activity, recall performance is largely unaffected. At a quantitative level, this is best illustrated by considering the average “time-loss rate” across the 4 experiments reported by Oberauer and Lewandowsky (2007). The time-loss rate captures the loss in accuracy for each second by which recall is delayed, and for the studies reported by Oberauer and Lewandowsky this value ranged from −.0012 (i.e., 1/10th of a percent loss per second) to −.0073 (less than one percent loss per second), with an average of −.004. In other words, no forgetting was observed despite delaying recall by as much as 12.5 seconds (for later list items). While it might be argued that the absence of forgetting reflected the operation of compensatory rehearsal that exactly canceled out the effects of decay or loss of temporal distinctiveness, we find it implausible to assume that rehearsal can take place when people simultaneously articulate distractors and perform a speeded choice task while also trying to recall a list.

Turning to the effects of time at encoding, the picture is somewhat more complex. We noted earlier that although there is considerable evidence that temporal isolation does not facilitate serial recall (e.g., Lewandowsky et al., 2006), there are two exceptions to this general finding that occur in known circumstances. Clearly, a comprehensive theory of memory must account for those exceptions as well as the general rule, and we acknowledge that C-SOB falls short in this regard. The first exception involves lists in which temporal isolation increases or decreases predictably across serial positions, as for example the list A.B..C…D….E, where each “.” represents a unit of time (e.g., Neath & Crowder, 1996). Lewandowsky, Nimmo, and Brown (in press) showed that the benefit of temporal isolation that is observed under those conditions can be explained by assuming that people selectively focus attention at encoding on those items that they know will be widely separated. No existing theory of memory, C-SOB included, captures and describes strategies of this type. The second exception involves situations in which people are free to report items in any order, as in free recall (Brown et al., 2006) or in an unconstrained reconstruction task (Lewandowsky, Nimmo, & Brown, in press). The latter study is particularly relevant because it showed that isolation benefited performance even when output order was statistically controlled—that is, isolation affected performance directly, rather than indirectly by facilitating selection of isolated items for early report, thereby protecting them against the deleterious consequences of output interference. The study by Lewandowsky et al. thus constitutes strong evidence that even when controlling for strategic factors, temporal isolation can causally and directly determine STM performance provided report order is unconstrained (the study included a comparison condition in which reconstruction was enforced to be in forward order; in replication of numerous earlier studies, no isolation effects were observed in that condition).

Because C-SOB currently does not explain performance in tasks in which output order is unconstrained, we defer an account of these isolation effects, although we acknowledge that C-SOB might ultimately have to acknowledge a role of time during encoding. This might be most readily achieved by augmenting the context markers with a temporal component that evolves with elapsed time rather than being driven by study or retrieval events.

5.1.2. *Prior knowledge and multi-trial effects*. The present instantiation of C-SOB assumed that participants had access to pre-existing knowledge about stimulus items, without however exploring the precise effects of that pre-existing knowledge on short-term recall. This does not present an in-principle limitation of the theory: In earlier work, Lewandowsky and Farrell (2000) and Farrell and Lewandowsky (2002) accounted for the beneficial effects of word familiarity on short-term recall by giving more frequent items additional pretraining. Recent evidence suggests a more complicated relationship between word frequency and serial recall performance. Hulme, Stuart, Brown, and Morin (2003) have shown that the high-frequency advantage disappears when high and low frequency words are mixed together on the same lists, suggesting that it is the associations that are formed between high-frequency items at encoding that drive the effect (see also Stuart & Hulme, 2000). It is unclear whether C-SOB would handle those effects.

Other evidence suggests that knowledge about sequential statistics also determines serial recall performance (e.g., Baddeley, Conrad & Hull, 1965; Botvinick & Bylsma, 2005; Thorn & Frankish, 2005). For example, Botvinick and Bylsma (2005) tested participants' immediate serial recall over many sessions with lists that contained sequential dependencies (e.g., *fie* might often be immediately followed by *kay* but less often by *tee*). Botvinick and Bylsma found that participants more accurately remembered sequences containing high-frequency transitions, and that they produced many regularization errors (i.e., erroneously replacing a low-frequency transition with a higher-frequency transition at recall). It is doubtful that C-SOB could handle those effects without further development.

Finally, learning can also be witnessed during experimental sessions in the form of the Hebb repetition effect, in which repeating sequences of items during an experiment cumulatively enhances recall for those lists relative to unrepeated control lists (e.g., Cumming, Page, & Norris, 2003; Hitch, Fastame, & Flude, 2005). In the light of claims that the fundamental role of verbal short-term memory is to support the laying down of long-term phonological representations (Baddeley et al., 1998), phenomena such as the Hebb effect obviously represent valuable targets for models of serial recall. The revised model of Burgess and Hitch (2006) can provide a detailed account of the fairly intricate pattern of results surrounding the Hebb effect, and we therefore anticipate that any future development of C-SOB must involve its application to multi-trial situations; not only to capture the Hebb effect but also more basic patterns such as the occurrence of protrusions (i.e., intrusions from the previous list which tend to preserve their original position). We expect that the inevitable retention of context-item associations across trials may lead to interesting interactions with similarity via the energy-gated encoding in C-SOB, perhaps giving rise to a possible mechanism for additional phenomena such as proactive interference.

5.2. C-SOB: Relationship to Other Theories

5.2.1. *SOB and C-SOB.* It is informative to compare C-SOB to its immediate antecedent, SOB (Farrell & Lewandowsky, 2002), both at an architectural level and in terms of its predictive power. There are two crucial differences between C-SOB and the original SOB: First, unlike SOB, C-SOB is not restricted to modeling memory for unrelated dissimilar items. Second, whereas SOB relied on a random endogenous cue to initiate retrieval, C-SOB includes a set of context markers. It turns out that those two changes necessarily go together. Within a distributed architecture, lists of similar items cannot be retrieved by a random cue alone because their attractors partially overlap, thus creating particularly strong basins of attraction that correspond to a blend of all similar items. Without context markers to cue retrieval, the model would only retrieve those strong blends which do not correspond to any of the list items.

The inclusion of context markers is also necessary for empirical reasons: First, recall that the analysis of latency-displacement functions points to the need for item markers. Second, item markers were required to enable the model to account for grouping effects. We explored numerous avenues to produce grouping effects without the use of context markers, all of which remained unsuccessful. The fact that the primacy model (Page & Norris, 1998), which includes no item markers, also cannot handle grouping effects supports the conclusion that grouping cannot be modeled without some type of external markers that are associated to the list items and that demarcate the various groups from each other.

These two architectural changes from the earlier SOB have yielded considerable pay-off in terms of predictive scope: Not only does the model now handle similarity and grouping effects, but the presence of context markers additionally permitted modeling of distractor activity during retrieval (see the earlier discussion of the data by Oberauer and Lewandowsky, 2007, and of the new experiment that explored the effects of varying the nature of distractors). The use of context markers does, however, entail a cost: As in many other models (e.g., SEM; Henson, 1998), the structure of the markers across positions is assumed rather than explained by the model. That is, although it is entirely plausible to postulate that the contexts of adjacent items are more similar to each other than the contexts of items separated by intervening events, the precise form of their similarity relationship is not derived from the model's architecture. Are there any candidate mechanisms on the horizon that might permit a more principled derivation of context markers? One model that contains a principled—albeit entirely time-based—mechanism for the evolution of context markers is OSCAR (Brown et al., 2000). OSCAR postulates that context is provided by a bank of independent oscillators that are operating at different frequencies and that, when considered together, uniquely identify any discrete moment in time. Accordingly, in situations in which list items are separated by constant temporal intervals, the similarity structure among the oscillator-based context vectors in OSCAR is nearly indistinguishable from that assumed for C-SOB. Indeed, if the presumed bank of oscillators is replaced by a set of nested gears, each of which is advanced to differing degrees

5.2.2. *C-SOB and other theories.* C-SOB can be compared to alternative theories by once again considering Tables 1 and 2. It is apparent that C-SOB matches or exceeds the predictive capabilities of most other contenders when the benchmark results in Table 1 are considered. However, to gather a more complete picture, we must consider two additional issues that were omitted from the table because, by our criteria, they fell outside the set of benchmarks: First, Burgess and Hitch (2006) presented an extension to their model that accommodated a wide range of data on the Hebb repetition effect. Second, the table does not include the results of the forgetting study presented ealier.

The first point reinforces our belief, stated at the outset, that the sequence of models by Burgess and Hitch (e.g., 1999, 2006) present a powerful accomplishment and a formidable target for rival theories. That said, we must note that this accomplishment comes at the cost of a panoply of theoretical constructs: The architecture of the Burgess and Hitch model relies on item marking as well as competitive queueing; it contains two sets of weights with differing decay characteristics (slow vs. fast); it contains two stages of retrieval, one of which involves feedback between two layers of phonemes; and its mechanism for response suppression is accompanied by decay of that inhibition. Although support can be adduced for each of those architectural features, their sum total can hardly be considered parsimonious. A desirable target for future theorizing must therefore be the attempt to match the theory's explanatory power within a more parsimonious rival framework.

The second point runs somewhat counter to the first one, because the results of the new forgetting study presented in this article are difficult to reconcile with a model in which forgetting is primarily due to time-based decay. In particular, neither the Burgess and Hitch models nor other time-based approaches (e.g., the primacy model of Page & Norris, 1998, or SIMPLE; Brown et al., in press) can explain why performance in the absence of rehearsal is unaffected by the passage of time per se, and why a rather subtle change in the nature of the to-be-articulated distractors changes the outcome so profoundly. Similarly, unlike C-SOB, none of the existing models—with the possible exception of the model by Botvinick and Plaut (2006)—can accommodate the mixed-list advantage for phonologically dissimilar items reported by Farrell and Lewandowsky (2003), Farrell (2006), and Lewandowsky and Farrell (in press).

The comparison between C-SOB and other theories presented here is primarily qualitative; given that one strength of computational models is their ability to make quantitative predictions of the type presented throughout this article, a necessary next step in theoretical development in STM must involve quantitative comparison of models. By way of precedent, Lewandowsky and Farrell (in press; see also Farrell, 2006) showed that although some models could qualitatively account for the mixed-list advantage for dissimilar items, only C-SOB accurately captured the quantitative profile of the data. Similarly, Oberauer and Lewandowsky (2007) compared C-SOB to the primacy model and SIMPLE at a quantitative level in its ability to accommodate the effects of distractor activity during recall. We suggest that future work in STM should build on those precedents and should engage in quantitative comparisons of a number of models on the same data, rather than demonstrating that a single model can account for some data (see Navarro, Pitt, & Myung, 2004).

5.2.2. *Theoretical conclusions.* What, then, is the current state of theoretical affairs in STM? Beginning with the specifics, we suggest that time-based forgetting is no longer viable as an explanatory construct. Instead, in the light of Oberauer and Lewandowsky's (2007) findings and the new data being collected in the laboratory of the present first author, we suggest that forgetting from short-term memory is primarily due to interference, both in the cross-talk between associations inherent in C-SOB's assumption of composite distributed representations, and in the novelty-sensitive encoding that reduces the quality of encoding of successive similar items.

Turning to the more general level, we conclude that there is strong evidence that implicates several architectural constructs in serial recall performance. Both response suppression and a primacy gradient are almost universally assumed in models of short-term memory; when these assumptions are not incorporated into models, the ability of those models to account for the data is restricted (for example, witness the lack of a primacy effect in the Burgess & Hitch model, which does not incorporate a primacy gradient). C-SOB arguably moves beyond other models in providing an endogenous account of the generation of the primacy gradient, and in specifying a learning mechanism responsible for response suppression. Another nearly ubiquitous assumption in Table 2 is that of item marking; given the empirical evidence in the form of latency-displacement functions and its apparent necessity to account for grouping and similarity effects in C-SOB, we suggest that the multi-dimensional contexts assumed by most models of serial recall are an essential representational device.

Finally, we point to some failings of current models, including C-SOB, as important areas for theoretical development. We suggest that theoretical development should focus on accounting for the rich set of constraints provided by the data on the Hebb repetition effect (e.g., Cumming et al., 2003; Hitch et al., 2005) and the role of long-term knowledge in short-term serial recall (e.g., Baddeley, Conrad & Hull, 1965; Botvinick & Bylsma, 2005; Thorn & Frankish, 2005). We also suggest that models should ultimately account for working memory performance beyond short-term serial recall. Given the importance of processing + storage tasks in predicting higher-level cognitive functioning (e.g., Oberauer et al., 2007), we see this a pathway for the current models to begin to contribute to a coherent and comprehensive theory of cognition.

6. Concluding Remarks

We have presented and explored a theory of short-term memory that we consider to be a welcome addition to the available theoretical contenders. We confirmed that the theory—with its assumptions of novelty-based encoding, response suppression, dynamic deblurring and positional marking—accounts for conventional benchmark data, and showed that, unlike some of its rivals, the model additionally accommodates recent results that either forced a re-evaluation of earlier data or provided new constraints on modeling. We also explored novel predictions of the theory that further differentiate it from rival models. We suggest that C-SOB—and in particular the energy-gated encoding that is central to its architecture—will provide a fruitful avenue for further exploration.

References

Adams, A., & Gathercole, S. E. (1996). Phonological working memory and spoken language development in young children. *Quarterly Journal of Experimental Psychology: Section A, 96,* 216-233.

Anderson, J. A. (1991). Why, having so many neurons, do we have so few thoughts? In W. E. Hockley & S. Lewandowsky (Eds.), *Relating theory and data: Essays on human memory in honor of Bennet B. Murdock* (pp. 477-507). Hillsdale, NJ: Erlbaum.

Anderson, J. A. (1995). *An introduction to neural networks*. Cambridge, MA: MIT Press.

Anderson, J. A., Silverstein, J. W., Ritz, S. A., & Jones, R. S. (1977). Distinctive features, categorical perception, and probability learning: Some applications of a neural model. *Psychological Review, 84*, 413-451.

Baddeley, A. D. (1966). Short-term memory for word sequences as a function of acoustic, semantic, and formal similarity. *Quarterly Journal of Experimental Psychology, 18,* 362-365.

Baddeley, A. D. (1968). How does acoustic similarity influence short-term memory. *Quarterly Journal of Experimental Psychology, 20A,* 249-263.

Baddeley, A. D., Conrad, R., & Hull, A. J. (1965). Predictability and immediate memory for consonant sequences. *Quarterly Journal of Experimental Psychology, 17*, 175–177.

Baddeley, A. D., Gathercole, S., & Papagno, C. (1998). The phonological loop as a language learning device. *Psychological Review, 105,* 158-173.

Baddeley, A. D., & Hitch, G. (1974). Working memory. In G. A. Bower (Ed.), *Recent advances in learning and motivation* (Vol. 8, pp. 647-667). New York: Academic Press.

Baddeley, A. D., Thomson, N., & Buchanan, M. (1975). Word length and the structure of short-term memory. *Journal of Verbal Learning and Verbal Behavior, 14*, 575-589.

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

Botvinick, M., & Bylsma, L. M. (2005). Regularization in short-term memory for serial order. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31,* 351–358.

Botvinick, M. M., & Plaut, D. C. (2006). Short-term memory for serial order: a recurrent neural network model. *Psychological Review,* 113, 201-233.

Brown, G. D. A., & Hulme, C. (1995). Modeling item length effects in memory span: No rehearsal needed? *Journal of Memory and Language, 34*, 594-621.

Brown, G. D. A., & Lamberts, K. (2003). Double dissociations, models and serial position curves. *Cortex, 39*, 148-152.

Brown, G. D. A., Morin, C., & Lewandowsky, S. (2006). Evidence for time-based models of free recall. *Psychonomic Bulletin & Review, 13,* 717-723.

Brown, G. D. A., Neath, I., & Chater, N. (in press). A temporal ratio model of memory. *Psychological Review*.

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

Bjork, R. A., & Whitten, W. B. (1974). Recency-sensitive retrieval processes in long-term free recall. *Cognitive Psychology, 6*, 173-189.

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

Burgess, N. & Hitch, G. (2006). A revised model of short-term memory and long-term learning of verbal sequences. *Journal of Memory and Language, 55*, 627-652.

Coltheart, V. (1993). Effects of phonological similarity and concurrent irrelevant articulation on short-term memory recall of repeated and novel word lists. *Memory & Cognition, 21,* 539-545.

Conrad, R. (1964). Acoustic confusions in immediate memory. *British Journal of Psychology, 55,* 75-84.

Cowan, N., Saults, J. S., Elliott, E. M., & Moreno, M. V. (2002). Deconfounding serial recall. *Journal of Memory and Language, 46*, 153-177.

Crowder, R. G. (1976). *Principles of learning and memory*. Hillsdale, NJ: Lawrence Erlbaum.

Cumming, N., Page, M. & Norris, D. (2003). Testing a positional model of the Hebb effect. *Memory, 11*, 43-63.

Dosher, B. A., & Ma, J.-J. (1998). Output loss or rehearsal loop? Output-time versus pronunciation-time limits in immediate recall for forgetting-matched materials. *Journal of Experimental Psychology: Learning, Memory and Cognition, 24*, 316-335.

Drewnowski, A., & Murdock, B. B. (1980). The role of auditory features in memory span for words. *Journal of Experimental Psychology: Human Learning and Memory, 6,* 319-332.

Farrell, S. (2006). Mixed-list phonological similarity effects in delayed serial recall. *Journal of Memory and Language, 55*, 587–600.

Farrell, S., & Lewandowsky, S. (2002). An endogenous distributed model of ordering in serial recall. *Psychonomic Bulletin & Review*, *9,* 59-79.

Farrell, S., & Lewandowsky, S. (2003). Dissimilar items benefit from phonological similarity in serial recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 29,* 838-849.

Farrell, S., & Lewandowsky, S. (2004). Modelling transposition latencies: Constraints for theories of serial order memory. *Journal of Memory and Language, 51,* 115-135.

Farrell, S., & Lewandowsky, S. (2007). *Response suppression contributes to recency in serial recall.* Unpublished manuscript, University of Western Australia.

Farrell, S., & McLaughlin, K. (in press). Short-term recognition memory for serial order and timing. *Memory & Cognition*.

Fürst, A. J., & Hitch, G. J. (2000). Separate roles for executive and phonological components of working memory in mental arithmetic. *Memory & Cognition, 28,* 774-782.

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.

Henry, L. A. (1991). The effects of word length and phonemic similarity in young children's short-term memory. *Quarterly Journal of Experimental Psychology, 43A,* 35-52.

Henson, R. N. A. (1996). *Short-term memory for serial order*. Unpublished doctoral dissertation, University of Cambridge, England.

Henson, R. N. A. (1998a). Short-term memory for serial order: The Start-End Model. *Cognitive Psychology, 36*, 73-137.

Henson, R. N. A. (1998b). Item repetition in short-term memory: Ranschburg repeated. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 24,* 1162-1181.

Henson, R. N. A. (1999). Positional information in short-term memory: Relative or absolute? *Memory & Cognition, 27*, 915-927.

Henson, R. N. A., Norris, D. G., Page, M. P. A., & Baddeley, A. D. (1996). Unchained memory: Error patterns rule out chaining models of immediate serial recall. *Quarterly Journal of Experimental Psychology, 49A,* 80-115.

Hinton, G. E., McClelland, J. L., & Rumelhart, D. E. (1986). Distributed representations. In D. E. Rumelhart, J. L. McClelland, & the PDP Research Group (Eds.). *Parallel distributed processing: Explorations in the microstructure of cognition, vol. 1: Foundations*. Cambridge, MA: MIT Press.

Hintzman, D. L. (1991). Why are formal models useful in psychology. In W. E. Hockley & S. Lewandowsky (Eds.), *Relating theory and data: Essays on human memory in honor of Bennet B. Murdock* (pp. 39-56). Hillsdale, NJ: Lawrence Erlbaum Associates.

Hitch, G. J., Burgess, N., Towse, J. N., & Culpin, V. (1996). Temporal grouping effects in immediate recall: A working memory analysis. *Quarterly Journal of Experimental Psychology, 49A,* 116-139.

Hitch, G. J., Fastame, M. C. & Flude, B. (2005). How is the serial order of a verbal sequence coded? Some comparisons between models. *Memory, 13*, 247-258.

Houghton, G., & Hartley, T. (1996). Parallel models of serial behaviour: Lashley revisited. *Psyche, 2(25)*. Symposium on implicit learning and memory (http://psyche.cs.monash.edu.au).

Hulme, C., Stuart, G., Brown, G. D. A. & Morin, C. (2003). High- and low-frequency words are recalled equally well in alternating lists: Evidence for associative effects in serial recall. *Journal of Memory and Language, 49*, 500-518.

Lewandowsky, S. (1993). The rewards and hazards of computer simulations. *Psychological Science, 4,* 236-243.

Lewandowsky, S. (1999). Redintegration and response suppression in serial recall: A dynamic network model. *International Journal of Psychology, 34,* 434-446.

Lewandowsky, S., Brown, G. D. A., Wright, T., & Nimmo, L. M. (2006). Timeless memory: Evidence against temporal distinctiveness models of short-term memory for serial order. *Journal of Memory and Language, 54,* 20-38.

Lewandowsky, S., Duncan, M., & Brown, G. D. A. (2004). Time does not cause forgetting in short-term serial recall. *Psychonomic Bulletin & Review, 11*, 771-790.

Lewandowsky, S., & Farrell, S. (2000). A redintegration account of word-length, lexicality, and articulatory suppression effects in short-term serial memory. *Psychological Research, 63,* 163-173.

Lewandowsky, S., & Farrell, S. (in press). Phonological similarity in serial recall: Constraints on theories of memory. *Journal of Memory and Language*.

Lewandowsky, S., & Heit, E. (2006). Some targets for memory models. *Journal of Memory and Language*, *55*, 441-446.

Lewandowsky, S., & Li, S.-C. (1994). Memory for serial order revisited. Psychological Review, 101, 539-543.

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

Lewandowsky, S. Nimmo, L. M., & Brown, G. D. A. (in press). When temporal isolation benefits memory for serial order. *Journal of Memory and Language*.

Lewandowsky, S., Wright, T., & Brown, G. D. A. (in press). The interpretation of temporal isolation effects. In N. Osaka, R. Logie, & M. D'Esposito (Eds.). *Working memory-Behavioral and neural correlates*. Oxford University Press.

Li, S.-C., & Lewandowsky, S. (1993). Intralist distractors and recall direction: Constraints on models of memory for serial order. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 19,* 895-908.

Lovatt, P., Avons, S. E., & Masterson, J. (2000). The word-length effect and disyllabic words. *Quarterly Journal of Experimental Psychology, 53A*, 1-22.

Luce, R. D. (1963). Detection and recognition. In R. D. Luce, R. R. Bush & E. Galanter (Eds.), *Handbook of mathematical psychology* (pp. 103-189). New York: Wiley.

Madigan, S. (1980). The serial position curve in immediate serial recall. *Bulletin of the Psychonomic Society, 15*, 335-338.

Mueller, S. T., Seymour, T. L., Kieras, D. E., & Meyer, D. E. (2003). Theoretical implications of articulatory duration, phonological similarity, and phonological complexity in verbal working memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 29*, 1353-1380.

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

Nairne, J. S. (1992). The loss of positional certainty in long-term memory. *Psychological Science, 3*, 199-202.

Nairne, J. S., & Kelley, M. R. (1999). Reversing the phonological similarity effect. *Memory & Cognition, 27*, 45–53.

Navarro, D. J., Pitt, M. A., & Myung, I. J. (2004). Assessing the distinguishability of models and the informativeness of data. *Cognitive Psychology, 49*, 47-84.

Neath, I., & Brown, G. D. A. (2006). SIMPLE: Further applications of a local distinctiveness model of memory. *The Psychology of Learning and Motivation, 46*, 201-243.

Neath, I., & Crowder, R. G. (1996). Distinctiveness and very short-term serial position effects. *Memory, 4*, 225-242.

Nimmo, L. M., & Lewandowsky, S. (2005). From brief gaps to very long pauses: Temporal isolation does not benefit serial recall. *Psychonomic Bulletin & Review, 12*, 999-1004.

Nimmo, L. M., & Lewandowsky, S. (2006). Distinctiveness revisited: Unpredictable temporal isolation does not benefit short-term serial recall of heard or seen events. *Memory & Cognition, 34,* 1368-1375.

Oberauer, K. (2003). Understanding serial position curves in short-term recognition and recall. *Journal of Memory and Language, 49*, 469-483.

Oberauer, K., & Lewandowsky, S. (2007). *Forgetting in immediate serial recall: Decay, temporal distinctiveness, or interference?* Manuscript submitted for publication.

Oberauer, K., Süß, H.-M., Wilhelm, O., & Sander, N. (2007). Individual differences in working memory capacity and reasoning ability. In A. R. A. Conway, C. Jarrold, M. J. Kane, A. Miyake, & J. N. Towse (Eds.), *Variation in working memory* (pp. 49-75). Oxford: Oxford University Press.

Page, M. P. A. (2000). A localist manifesto. *Behavioral and Brain Sciences, 23,* 443-467.

Page, M. P. A., Madge, A., Cumming, N., & Norris, D. G. (2007). Speech errors and the phonological similarity effect in short-term memory: Evidence suggesting a common locus. *Journal of Memory and Language, 56*, 49-64.

Page, M. P. A., & Norris, D. (1998). The primacy model: A new model of immediate serial recall. *Psychological Review, 105,* 761-781.

Ryan, J. (1969). Temporal grouping, rehearsal and short-term memory. *Quarterly Journal of Experimental Psychology, 21,* 148-155.

Service, E. (1998). The effect of word length on immediate serial recall depends on phonological complexity, not articulatory duration. *Quarterly Journal of Experimental Psychology, 51A*, 283-304.

Stuart, G. & Hulme, C. (2000). The effects of word co-occurrence on short-term memory: Associative links in long-term memory affect short-term memory performance. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26*, 796-802.

Surprenant, A. M., Kelley, M. R., Farley, L. A., & Neath, I. (2005). Fill-in and infill errors in order memory. *Memory, 13*, 267-273.

Thorn, A. S. C. & Frankish, C. R. (2005). Long-term knowledge effects on serial recall of nonwords are not exclusively lexical. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31*, 729-735.

Vousden, J. I., & Brown, G. D. A. (1998). To repeat or not to repeat: The time course of response suppression in sequential behaviour. In J. A. Bullinaria, D. W. Glasspool, & G. Houghton (Eds.), *Proceedings of the fourth neural computation and psychology workshop: Connectionist representations* (pp. 301-315). London: Springer Verlag.

Wickelgren, W. A. (1965a). Acoustic similarity and retroactive interference in short-term memory. *Journal of Verbal Learning and Verbal Behavior, 4,* 53-61.

Wickelgren, W. A. (1965b). Short-term memory for phonemically similar lists. *American Journal of Psychology, 78,* 576-574.

Author Note

Preparation of this paper was facilitated by various Discovery Grants from the Australian Research Council to the first author and by a Linkage International Grant from the Australian Research Council to both authors and Gordon Brown. Address correspondence to the first author at the School of Psychology, University of Western Australia, Crawley, W.A. 6009, Australia. Electronic mail may be sent to lewan@psy.uwa.edu.au. Personal web page: http://www.psy.uwa.edu.au/user/lewan/.

Footnotes

[1] According to an online search of PsychLit.

[2] Page, Madge, Cumming, and Norris (2007) recently suggested that the mixed-list advantage observed by Farrell and Lewandowsky (2003) was due to their method of presenting all lists of the same type in a contiguous block, which followed the precedent in of Henson et al. (1996). Page et al.'s (2007) explanation is unlikely given that the mixed-list advantage has also been observed when all trial types were randomly intermixed (Farrell, 2006; Lewandowsky & Farrell, in press).

[3] Farrell and Lewandowsky (2004) did not consider this model and instead focused only on the model including a primacy gradient and response suppression.

[4] The small divergence in the figure is not statistically significant and is also absent in a number of other experiments conducted in the first author's laboratory. If the effects of increasing the duration of articulation are expressed as the "loss of accuracy per second," in all experiments the 95% confidence intervals for those time-loss rates narrow and straddle zero. The exact values are given later in this paper.

Table 1. Summary of contemporary theories of short-term memory and the phenomena they can account for at a quantitative level. Note: B&H92 = Burgess & Hitch (1992); B&H99 = Burgess & Hitch (1999); B&P06 = Botvinick & Plaut (2006); OSCAR = Brown et al., (2000); PM = Primacy model, Page & Norris (1998); SEM = Start-End model, Henson (1998a); C-SOB = Farrell (2006), Lewandowsky & Farrell (in press), and the present article; SIMPLE = Brown et al. (in press), Neath & Brown (2006).

Phenomenon		*B&H 92*	*B&H 99*	*B&P 06*	*OSCAR*	*PM*	*SEM*	*C-SOB*	*SIMPLE*
List length		✓	✓	✓	✓	✓	✓	✓	?
SPC									
	Primacy	✓	✓	✓	✓	✓	✓	✓	✓
	Recency		✓	✓	✓	✓	✓	✓	✓
Errors									
	Transposition Proportion	✓	✓	?	?	✓	✓	✓	?
	Transposition gradient	×	✓	✓	✓	✓	✓	✓	✓
	Fill-in vs. infill	?	?	✓	?	✓	?	×	×
	Repetitions	×	✓	✓	×	?	✓	✓	×
	Omissions	✓	✓	×	✓	✓	✓	✓	✓
	Intrusions	✓	✓	×	✓	✓	✓	✓	✓
	Protrusions	×	✓	×	×	×	✓	×	✓
Phonological Similarity									
	Overall	✓	✓	✓	✓	✓	✓	✓	✓
	Transpositions	✓	✓	✓	✓	✓	✓	✓	✓
	Mixed-Lists (old)	×	✓	✓	×	✓	✓	×	?
	Mixed-Lists (new)	×	×	✓ **(?)**	×	×	×	✓	×
	Reversal with delay	×	×	×	×	×	×	✓	✓
Grouping									
	SPC	×	✓	×	✓	×	✓	✓	✓
	Interpositions	×	✓	✓	✓	×	✓	✓	✓
	Modality effect	×	✓	×	×	×	✓	×	×
Word Length									
	Basic effect	✓	✓	×	×	✓	×	×	✓
	AS abolish	×	✓	×	×	✓	×	×	×
Word frequency (pure lists)		×	×	×	×	✓	×	✓	×
Modality effect		×	✓	×	×	✓	✓	×	×
Backward recall		×	×	×	✓	✓	×	×	×

Table 2. Summary of contemporary theories of short-term memory and the explanatory constructs on which they rely. See text for explanation of column numbers and shading.

	1	2	3	4	5	6	7	8	9	10
Model	*Nature of Representations*	*Locus of Similarity Effects*	*Type of Associations*	*Role of Time at Encoding*	*Role of Time During Forgetting*	*Primacy Gradient*	*Response Selection Mechanism*	*Output Interference*	*Response Suppression*	*Energy-gated encoding*
Burgess & Hitch (1999)	Localist (except context)	Separate confusion stage	Item↔Item and Item↔Context	Temporal context	Decay	No	Context + CQ	No	Yes	No
Primacy Model (Page & Norris, 1998)	Localist	Separate confusion stage	None	None	Decay	Yes	CQ	No	Yes	No
SEM (Henson, 1998)	Localist	Separate confusion stage	Item↔Context (× 2 markers)	None	Decay (certain applications), interference otherwise (?)	Yes (+ recency gradient)	Context + CQ	No	Yes	No
SIMPLE (Brown et al., in press)	Localist	Primary stage	Item↔Temporal context	Temporal information crucial	Interference, but temporally based	No	Context + Luce	Possible	Possible	No
OSCAR (Brown et al., 2000)	Distributed	Primary stage	Item↔Temporal context	Temporal context	None (interference)	Yes	Context + CQ	Yes	Yes	No
C-SOB (Farrell, 2006; Lewandowsky & Farrell, in press)	Distributed	Primary stage	Item↔Context	None	None (interference)	Yes	Context + Dynamic deblurring (or Luce)	Yes	Yes	Yes
Botvinick & Plaut (2006)	Distributed	Primary stage	Recurrent	None	None	No	Recurrent + CQ	No	Optional	No

Figure Captions

Figure 1. Predicted accuracy serial position curves (left panel), latency serial position curves (center panel), and latency-displacement functions (right panel) for a number of explanatory constructs implemented in a common architecture. IT = item marking; RS = response suppression; OI = output interference; PR = primacy gradient.

Figure 2. Average performance across 7 experiments for the terminal list item (forward) and for the first three list items (backward) conditionalized on the types of errors committed on the remaining list positions. See text for further explanation.

Figure 3. Serial position curves for various list lengths predicted by C-SOB. The left panel shows predictions for accuracy and the right panel predictiosn for cumulative latency.

Figure 4. Transposition gradients for various list lengths (averaged across serial positions) predicted by C-SOB.

Figure 5. Latency-displacement functions for various list lengths (averaged across serial positions) predicted by C-SOB.

Figure 6. Pattern of errors across output positions predicted by C-SOB for list lengths 5 (left panel) and 6 (right panel).

Figure 7. Serial position curves for lists composed of similar (S) and dissimilar (D) items predicted by C-SOB.

Figure 8. Accuracy serial position curves (left panel), latency serial position curves (center panel), and latency-displacement functions (right panel) for grouped and ungrouped 9-item lists predicted by C-SOB.

Figure 9. Data (plotting symbols) and predictions from C-SOB (continuous lines) for two experiments by Oberauer and Lewandowsky (2007). The left panel shows an experiment that compared a quiet baseline (B) to two conditions involving one (1R) or three (3R) articulations of a suppressor before each retrieval. The right panel shows an experiment in which each retrieval was preceded by one (1 AS+CRT) or four (4 AS+CRT) interfering events, each of which involved articulation of a suppressor and a simultaneous speeded choice task.

Figure 10. Predictions from C-SOB for four different types of distractors during retrieval. The top row of panels refers to simple “bursts” of distractors, in which the same word is repeated once (lines labeled 1) or three times (3) after each retrieval. The bottom row refers to complex bursts, in which each articulation involves a different word. The left column of panels refers to the case in which bursts of distractors remain unchanged across output positions, and the right column refers to the case in which the identity of words in a burst changes across output positions.

Figure 11. Data from an unpublished experiment that compared four different types of distractors during retrieval. The top row of panels refers to simple “bursts” of distractors, in which the same word is repeated once (lines labeled 1) or three times (3) after each retrieval. The bottom row refers to complex bursts, in which each articulation involves a different word. The left column of panels refers to the case in which bursts of distractors remain unchanged across output positions, and the right column refers to the case in which the identity of words in a burst changes across output positions.

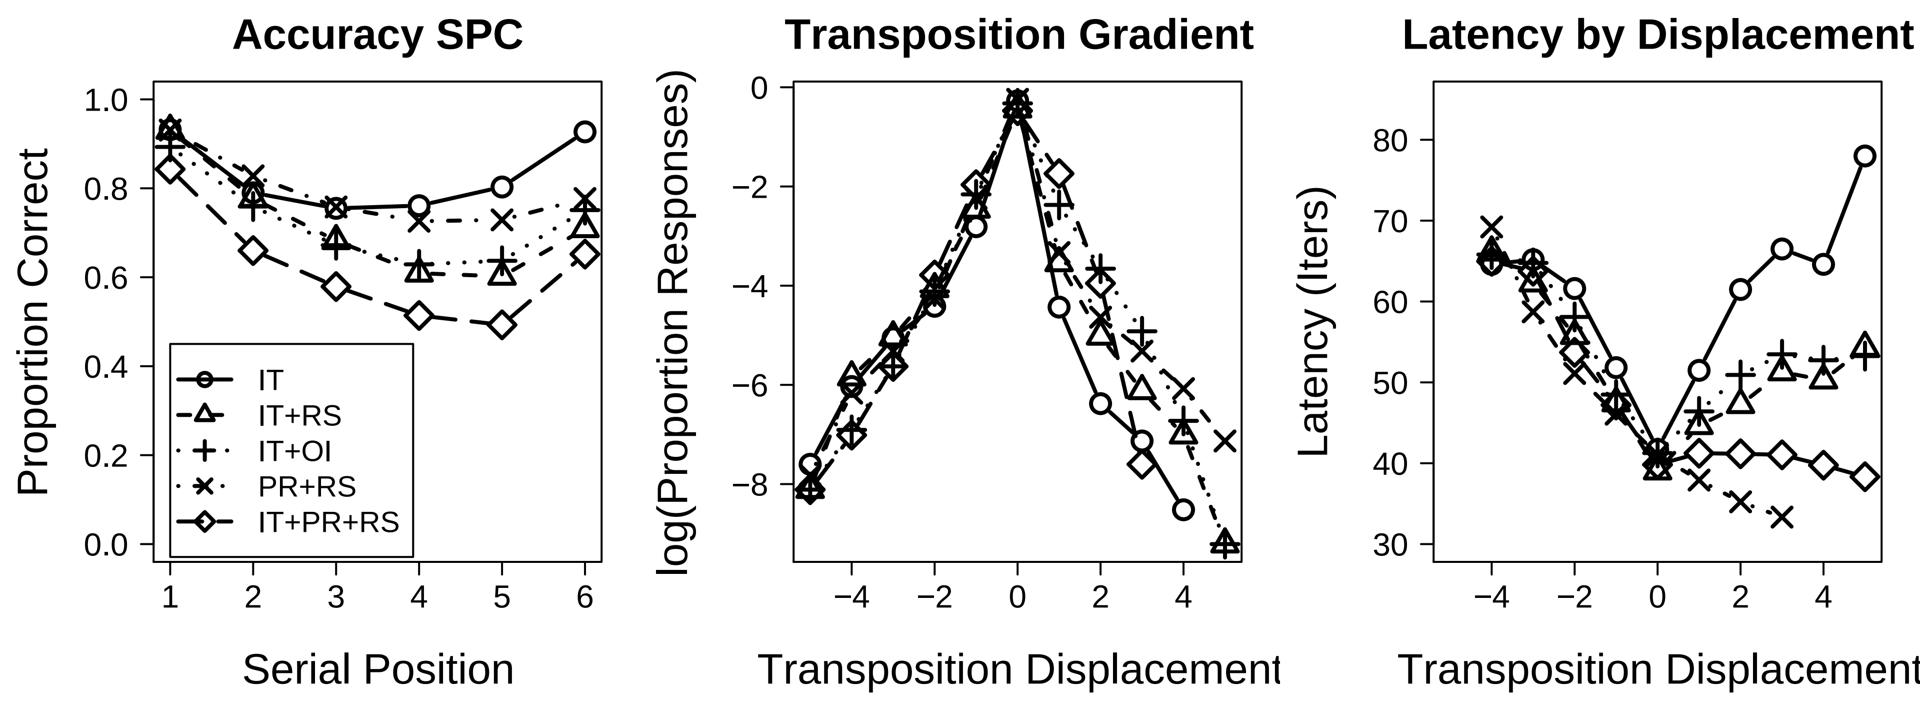
Accuracy SPC
Transposition Gradient
Latency by Displacement
Proportion Correct
Serial Position
log(Proportion Responses)
Transposition Displacement
Latency (Iters)
Transposition Displacement
IT
IT+RS
IT+OI
PR+RS
IT+PR+RS

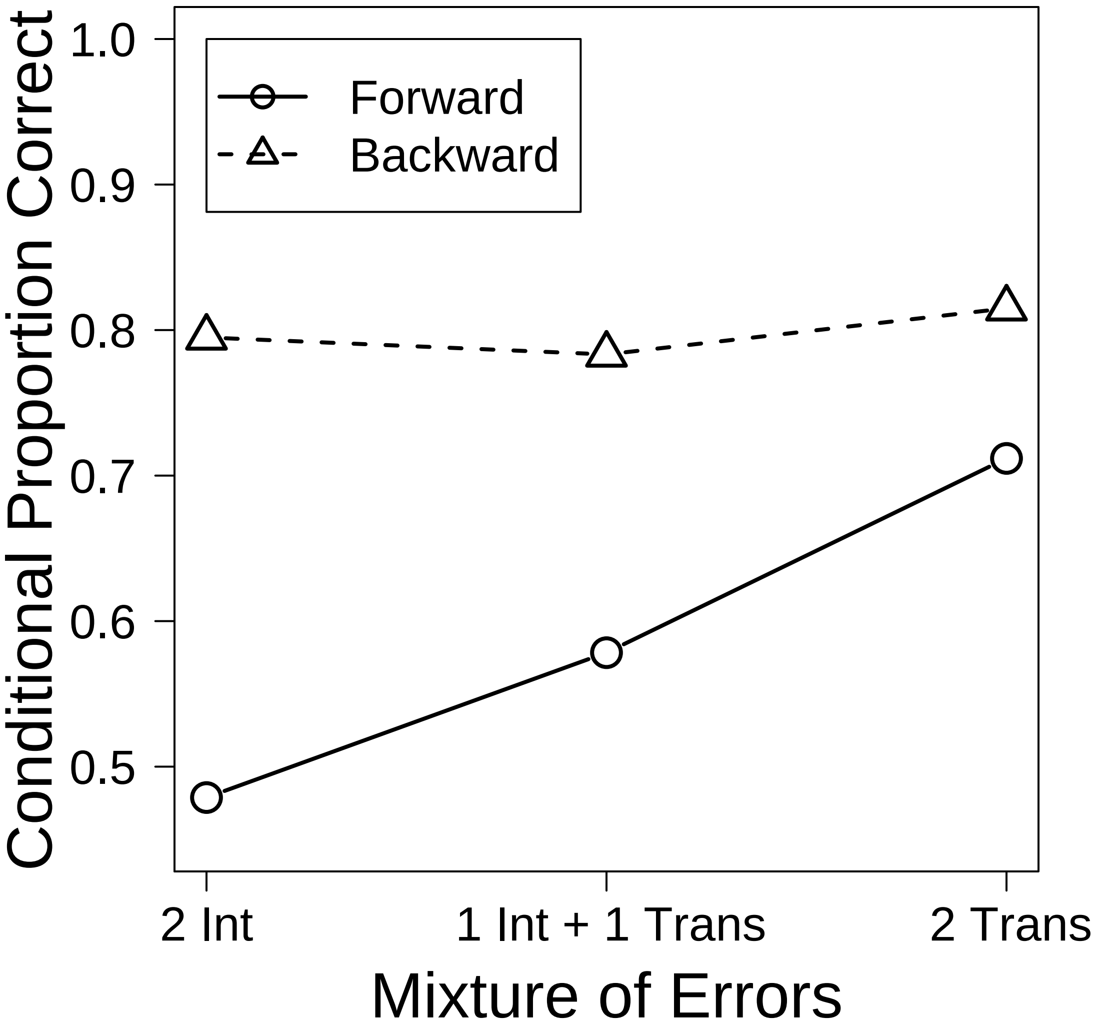
Conditional Proportion Correct
1.0
0.9
0.8
0.7
0.6
0.5
Forward
Backward
2 Int
1 Int + 1 Trans
2 Trans
Mixture of Errors

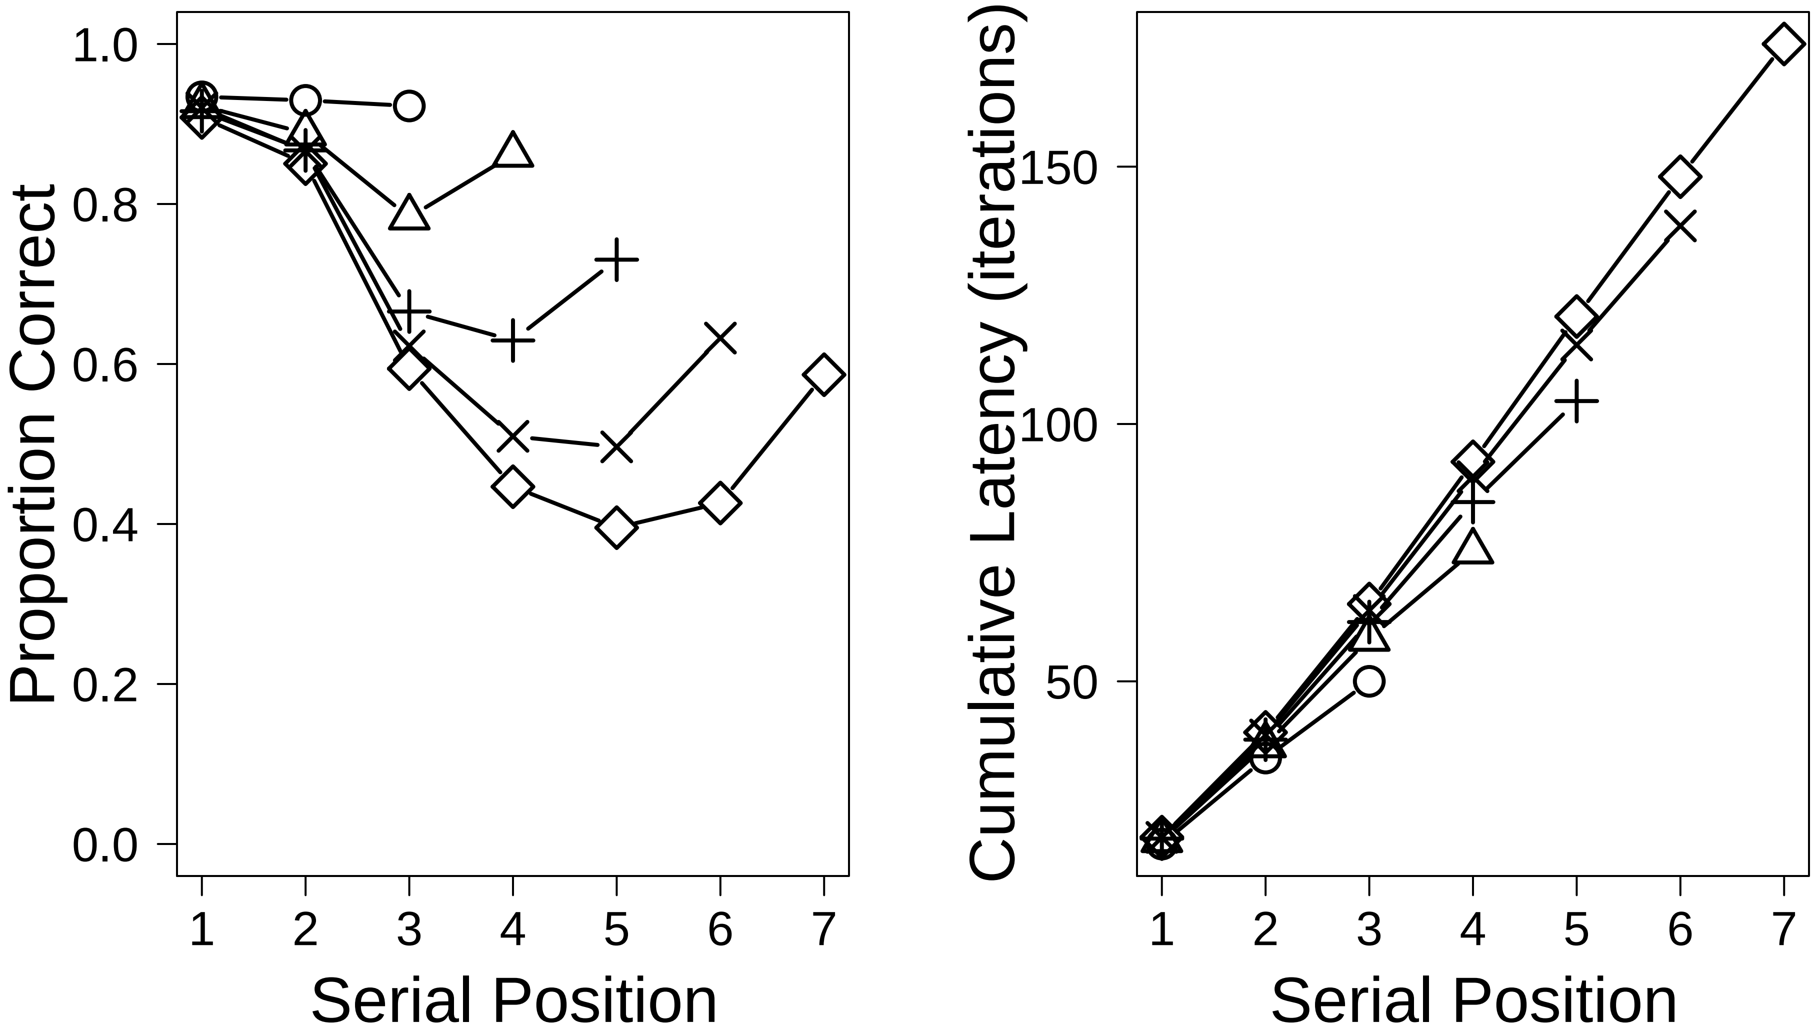

Proportion Correct
1.0
0.8
0.6
0.4
0.2
0.0
1
2
3
4
5
6
7
Serial Position
Cumulative Latency (iterations)
150
100
50
1
2
3
4
5
6
7
Serial Position

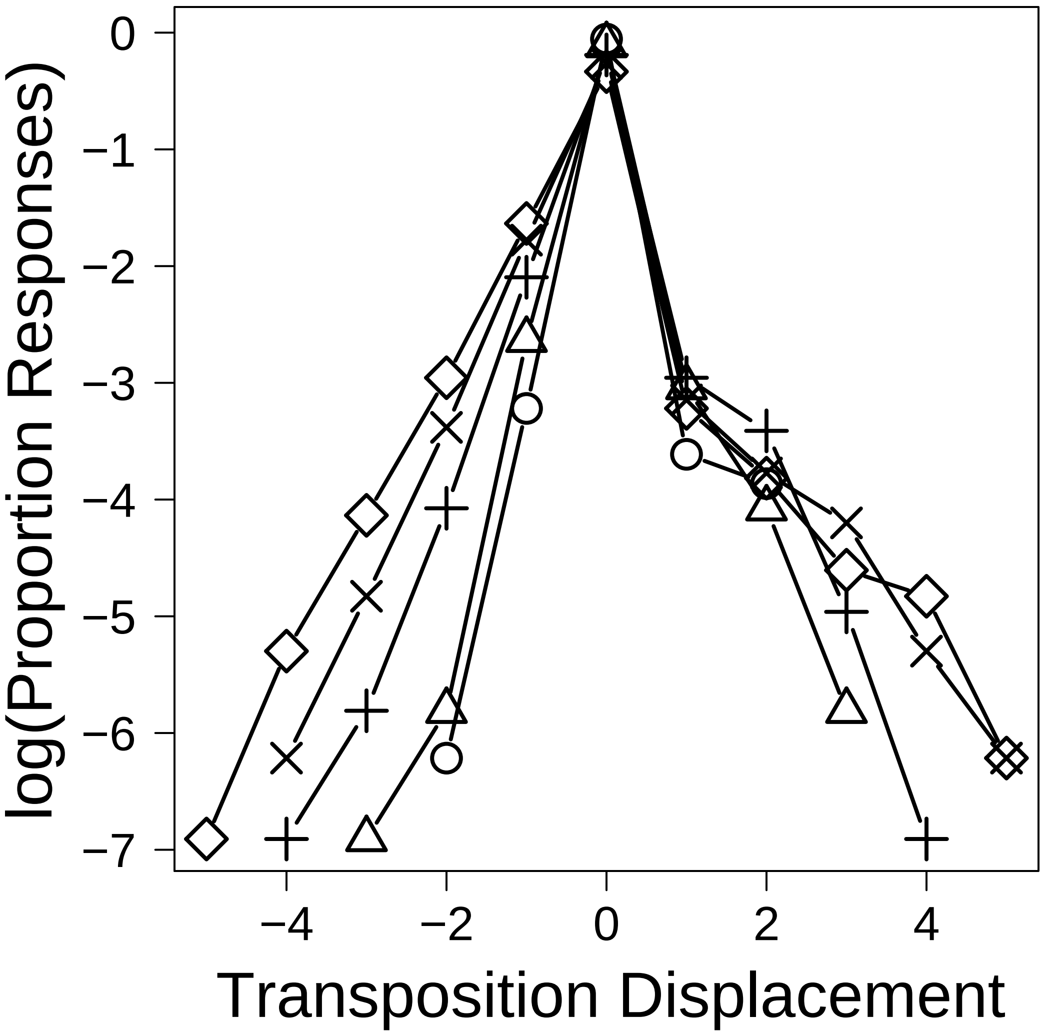

log(Proportion Responses)
0
−1
−2
−3
−4
−5
−6
−7
−4
−2
0
2
4
Transposition Displacement

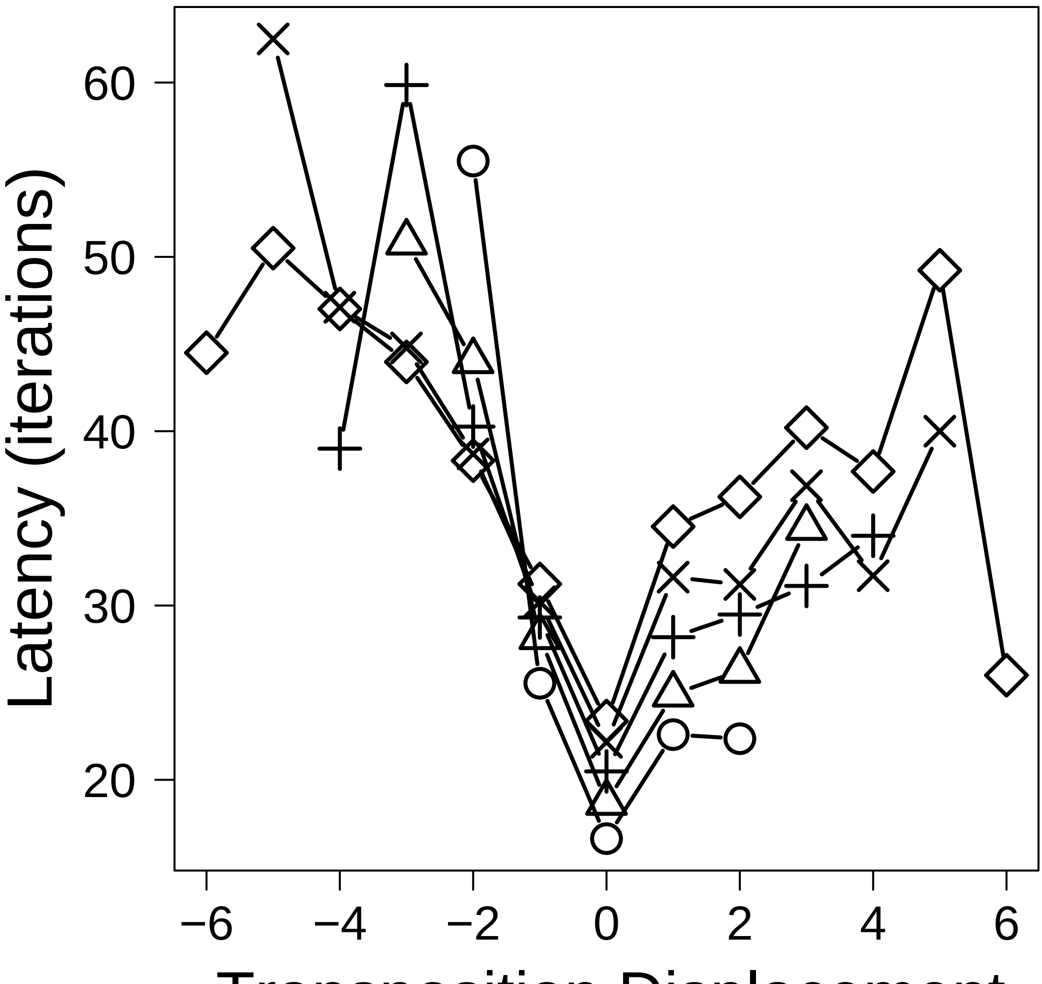

60
50
40
30
20
Latency (iterations)
−6
−4
−2
0
2
4
6
Transposition Displacement

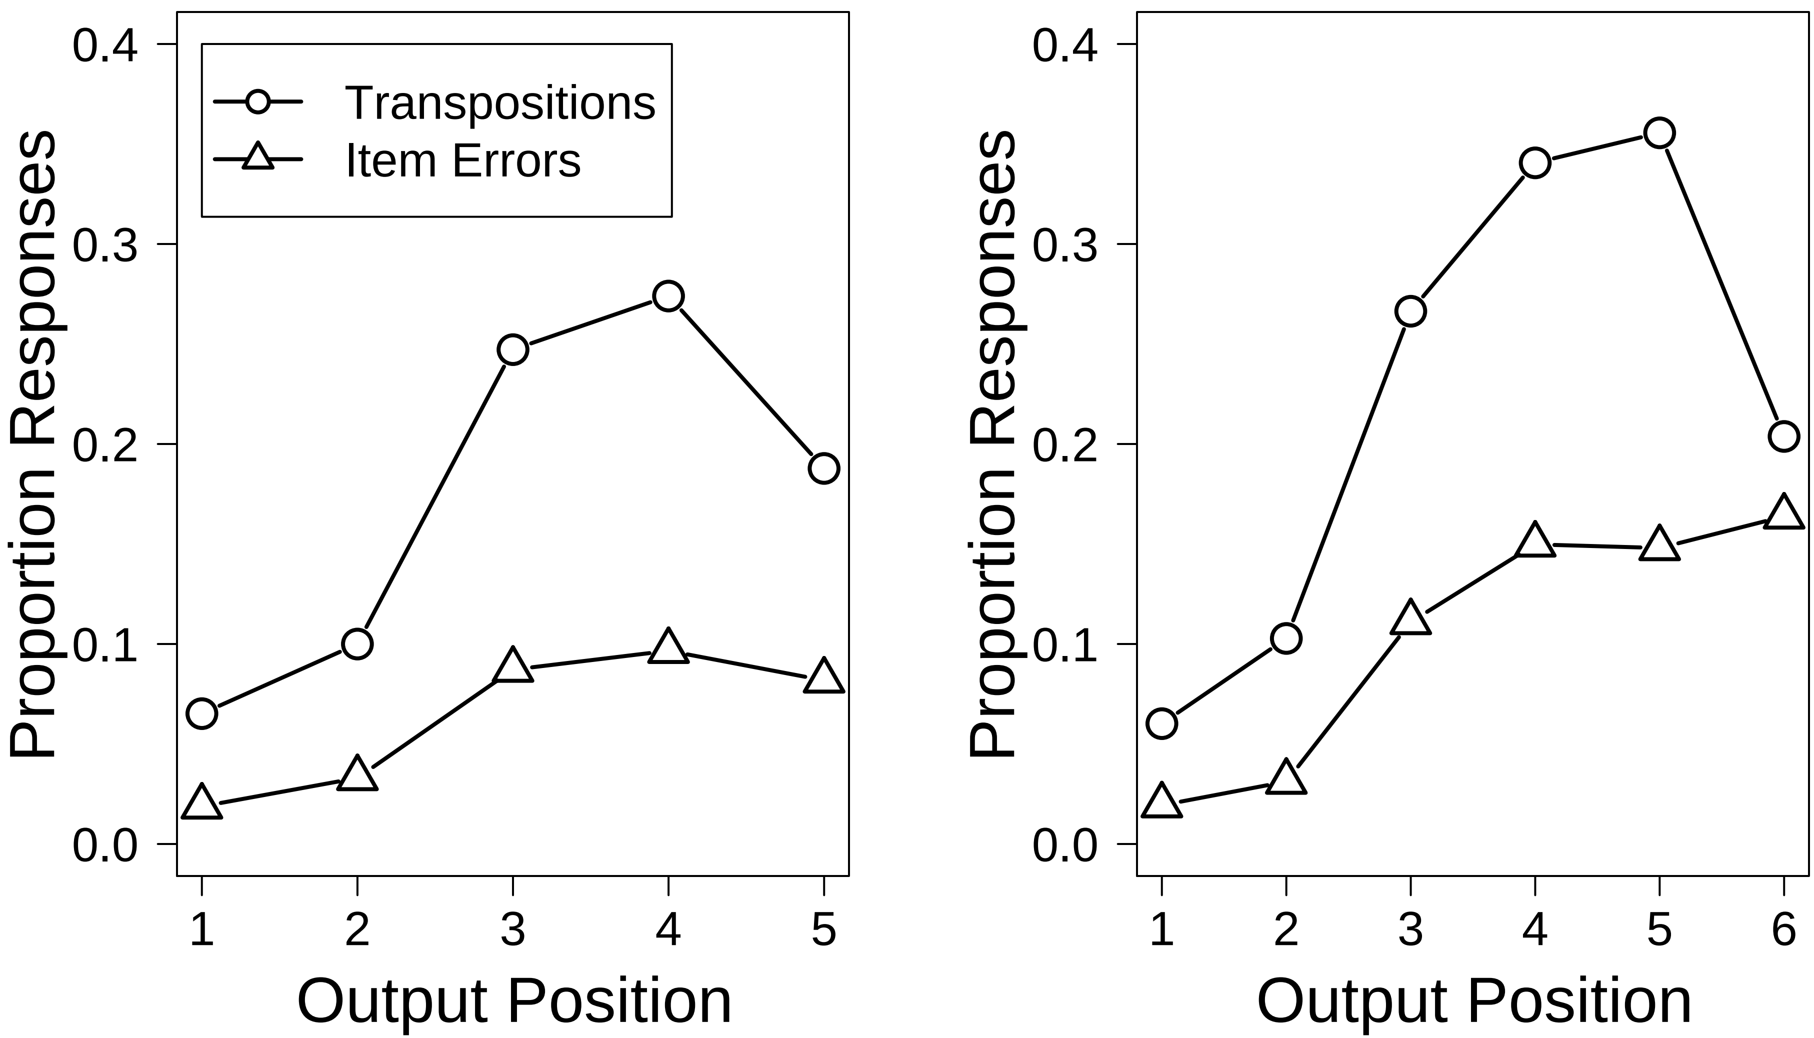
Transpositions
Item Errors
Proportion Responses
0.0
0.1
0.2
0.3
0.4
1
2
3
4
5
Output Position
Proportion Responses
0.0
0.1
0.2
0.3
0.4
1
2
3
4
5
6
Output Position

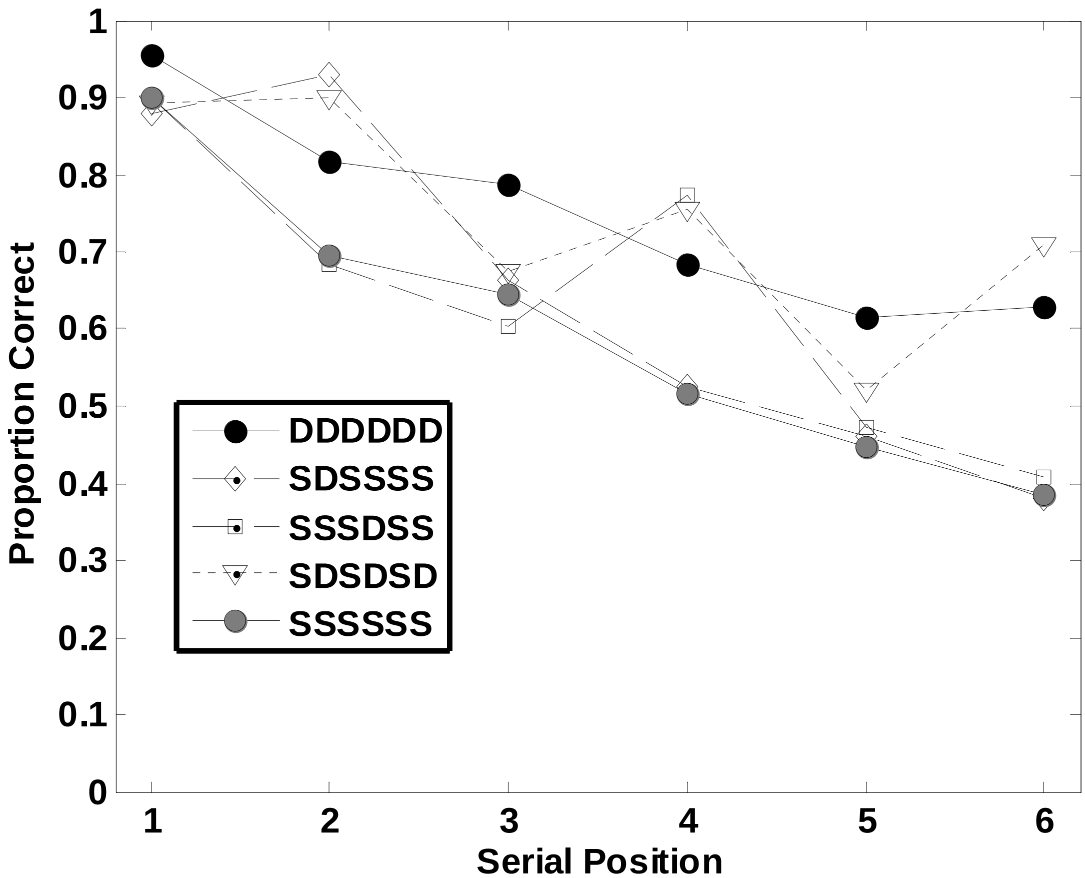

Proportion Correct
Serial Position
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
1
2
3
4
5
6
DDDDDD
SDSSSS
SSSDSS
SDSDSD
SSSSSS

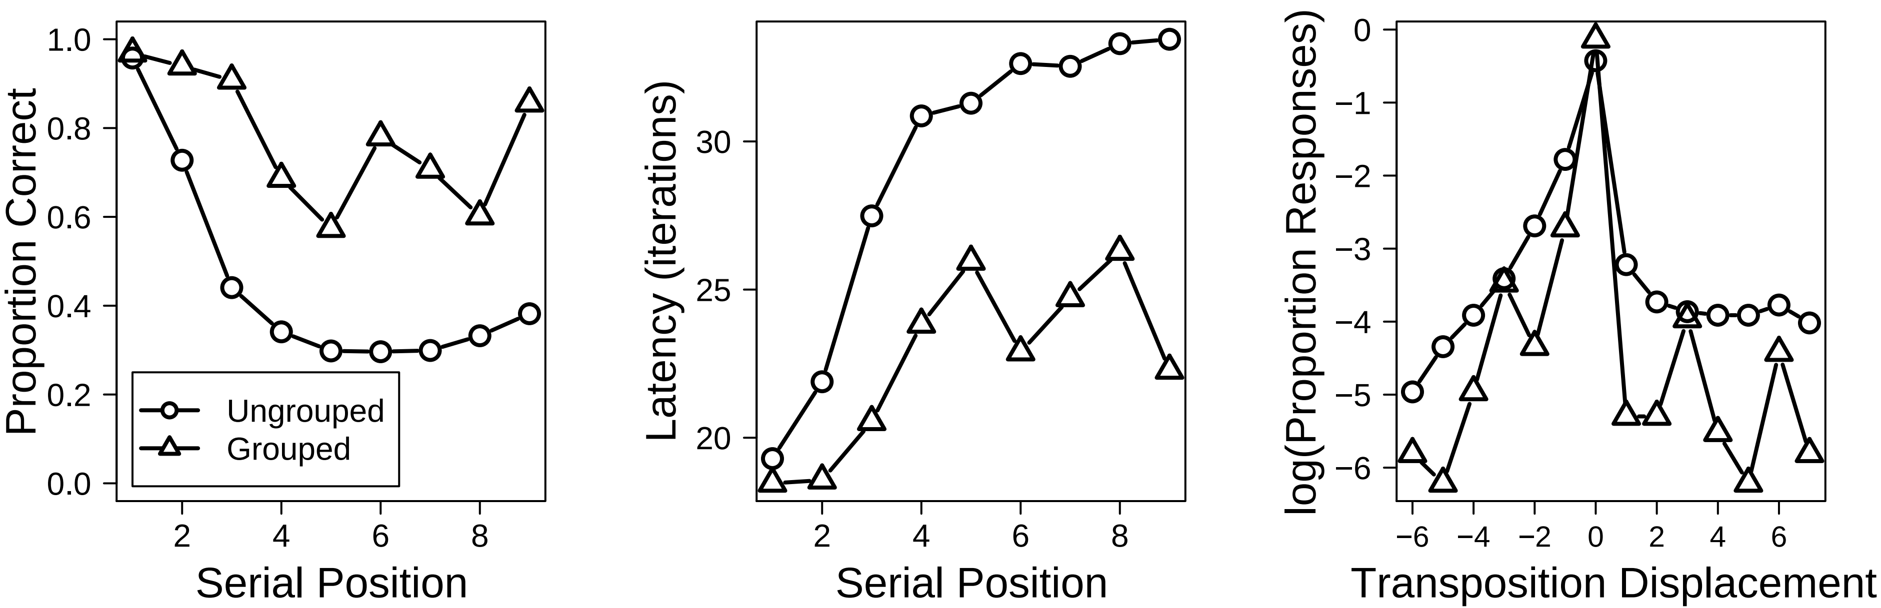
Proportion Correct
1.0
0.8
0.6
0.4
0.2
0.0
2
4
6
8
Serial Position
Ungrouped
Grouped
Latency (iterations)
30
25
20
2
4
6
8
Serial Position
log(Proportion Responses)
0
−1
−2
−3
−4
−5
−6
−6
−4
−2
0
2
4
6
Transposition Displacement

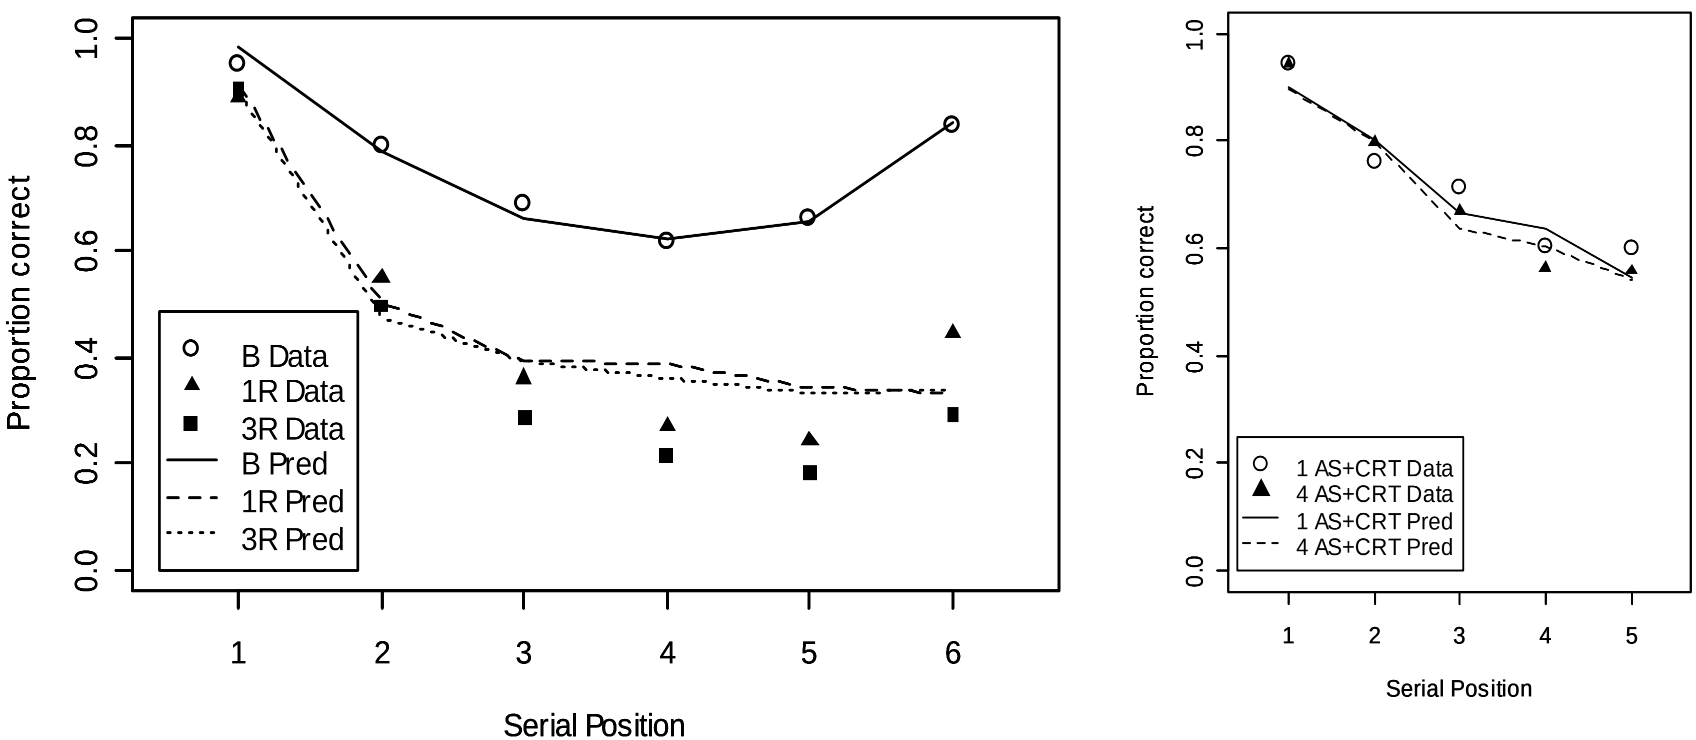

Proportion correct
1.0
0.8
0.6
0.4
0.2
0.0
B Data
1R Data
3R Data
B Pred
1R Pred
3R Pred
1
2
3
4
5
6
Serial Position
Proportion correct
1.0
0.8
0.6
0.4
0.2
0.0
1 AS+CRT Data
4 AS+CRT Data
1 AS+CRT Pred
4 AS+CRT Pred
1
2
3
4
5
Serial Position

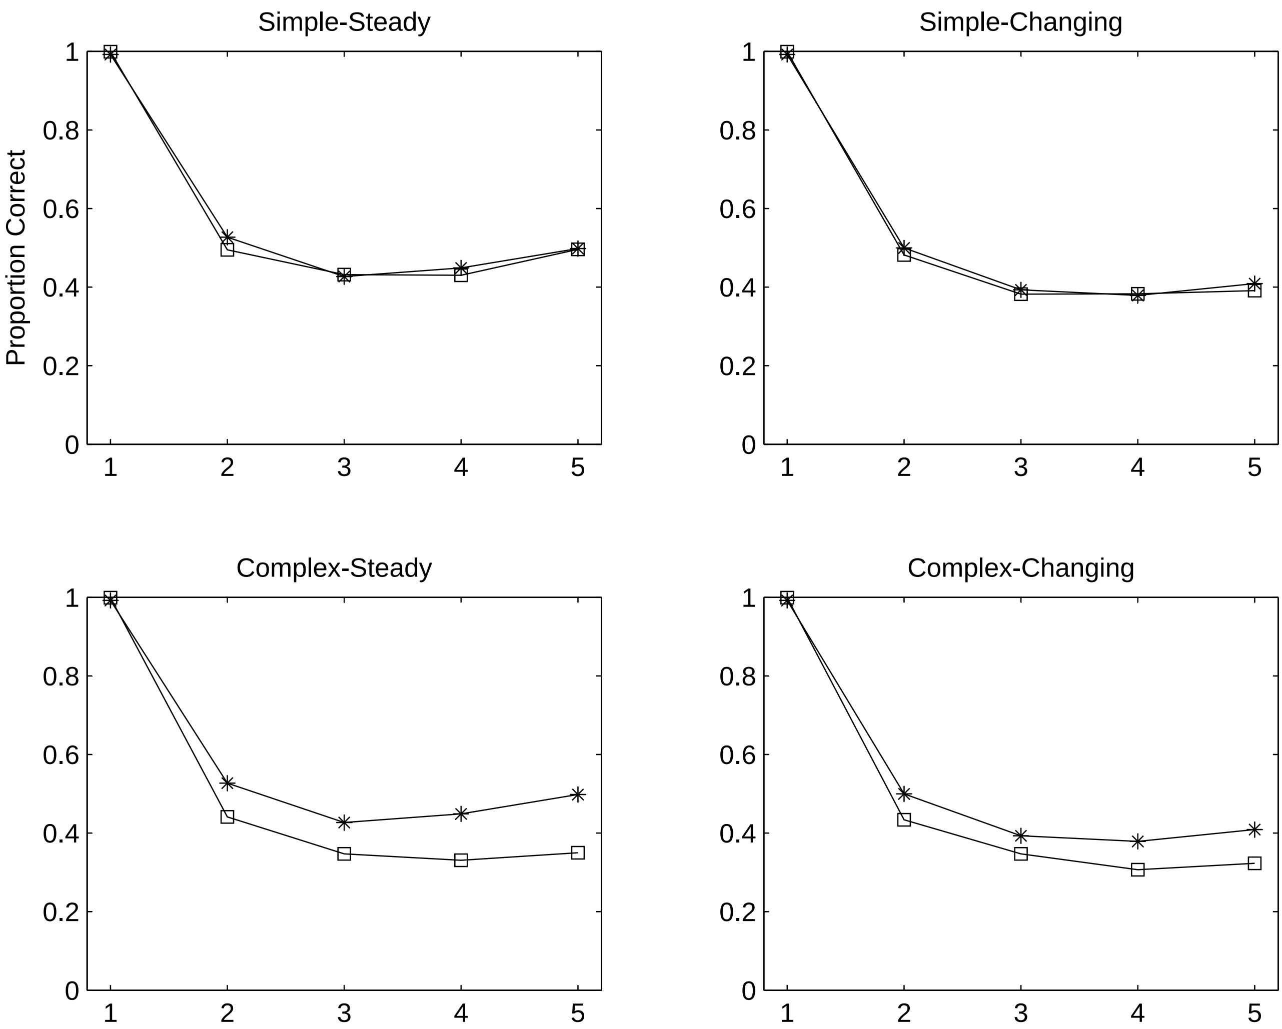
Simple-Steady
Simple-Changing
Complex-Steady
Complex-Changing
Proportion Correct
1
0.8
0.6
0.4
0.2
0
1
2
3
4
5

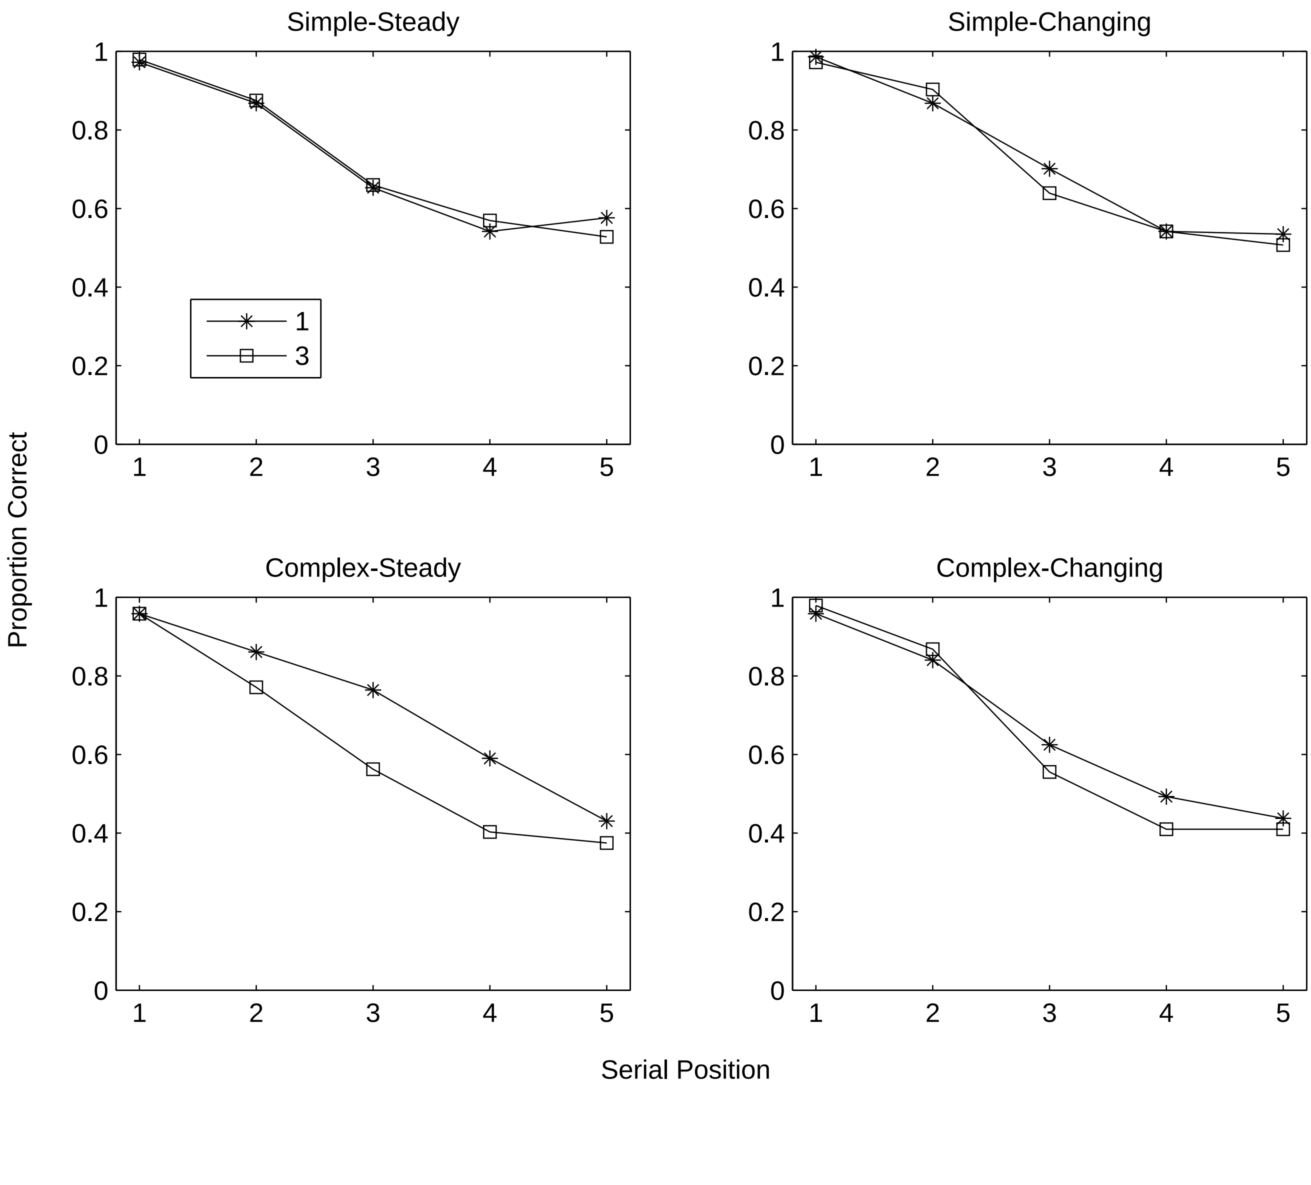

Simple-Steady
Simple-Changing
Complex-Steady
Complex-Changing
1
3
Proportion Correct
Serial Position
0
0.2
0.4
0.6
0.8
1
1
2
3
4
5