Time-Based Resource Sharing in ARCADIA

Kevin O'Neill, Will Bridewell, & Paul Bello
{**kevin.oneill, will.bridewell, paul.bello**}**@nrl.navy.mil**
Naval Research Laboratory
Washington, DC 20375 USA

Abstract

We provide a new computational model of working memory in the complex span task implemented in the ARCADIA cognitive framework. While there exist implementations of working memory successful enough to account for many of the benchmark findings in the working memory literature, we demonstrate that further progress requires the integration of these models with a rich conception of attention. ARCADIA provides this intersection, allowing for precise control of the focus of attention on a time scale fine enough to begin to disentangle the overlapping effects of interference, temporal decay, and attentional refreshing.

Keywords: working memory; time-based resource sharing; attention; cognitive architecture

Introduction

What causes forgetting in working memory? The question remains a matter of intense debate, because its answer has implications for why working memory is capacity limited and how it can be maximized. Two relevant paradigms have been utilized to study working memory and its limitations. In the serial recall or simple span task (SST), a number of memoranda are presented one-by-one, and the subject is asked to recall the memoranda in presentation order after a short delay. Data from the SST typically exhibit greatly increased recall for early list items (the primacy effect) and slightly increased recall for late list items (the recency effect) (Tan & Ward, 2008). The complex span task (CST) extends the SST by placing processing demands over a fixed interval after the presentation of each memorandum. This processing demand can be a set of distracting verbal (reading words, letters, or digits), arithmetic (incrementing or decrementing operations), or visuo-spatial (discriminating the location or size of an object) tasks which systematically induce forgetting. As such, the CST provides a deeper window into the relationship, and possible trade-off, between processing and storage in working memory.

To date, several mechanisms have been posited to account for forgetting in working memory; notably, the time-based resource sharing (TBRS) model of working memory relies on temporal decay and the serial-order-in-a-box complex span (SOB-CS) model on interference between memory representations. While these theories have been treated as mutually exclusive, differences in each theory's predictive gaps may instead indicate that both theories are generally correct but incomplete. Indeed there is evidence indicating that both temporal decay and interference contribute to the overall character of memory loss (Altmann & Schunn, 2012). Exacerbating this problem is the nature of modeling techniques currently used to test the time-sensitive predictions of decay and interference— while models of each mechanism rely on attentional restoration processes to counteract forgetting, none make a coordinated effort to account for attention (or distractor processing) in its own right. This simplification reduces the number of assumptions programmed into the models but causes imprecision with regard to the exact time-course of attentional processing, and hence to the exact predictions of the models themselves. Reliably capturing the effects of decay and interference then necessitates studying them within a process-level architecture with attention at its core. We present a new implementation of TBRS in ARCADIA, an attentional framework for cognitive modeling which meets this demand. Through this implementation we take the first steps towards a more integrated picture of working memory.

Background

Temporal Decay

Decay-based theories including the time-based resource sharing (TBRS) model declare that memories simply decay over time, requiring a rapid compensatory attentional refreshing process in order to maintain their representations (Barrouillet, Bernardin, & Camos, 2004). In a CST, the trade-off between distractor processing and attentional refreshing has been shown to create a cognitive load effect in which the complex span is inversely proportional to cognitive load (estimated by the distractor rate) (Barrouillet & Camos, 2007). Defined by the equation $CL = aN/T$, where N is the number of distractors to be processed in between memory encoding, T is the time available for distractor processing, and a is a parameter reflecting distractor difficulty, the cognitive load effect has the important implication that regardless of the nature of the distracting task, complex span is determined by the proportion of time that the distracting task requires attention.

TBRS has been implemented in TBRS*, a model of working memory consisting of a two-layer connectionist network (Oberauer & Lewandowsky, 2011). In TBRS*, decay is modelled as an exponential decrease in network weights which can be counteracted by attentional refreshing. Distractor processing is not modelled *per se*, rather the time required for distractor processing (as estimated by response time) is used to schedule refreshing, which follows a forward cumulative order from the beginning to the end of the list. TBRS* has been shown to qualitatively reproduce the cognitive load effect, as well as primacy and recency effects in the complex span task, proximity effects over errors, and error type distributions (Oberauer & Lewandowsky, 2011).

However, there remain issues with the predictions of TBRS* (and TBRS in general) noted by Oberauer and
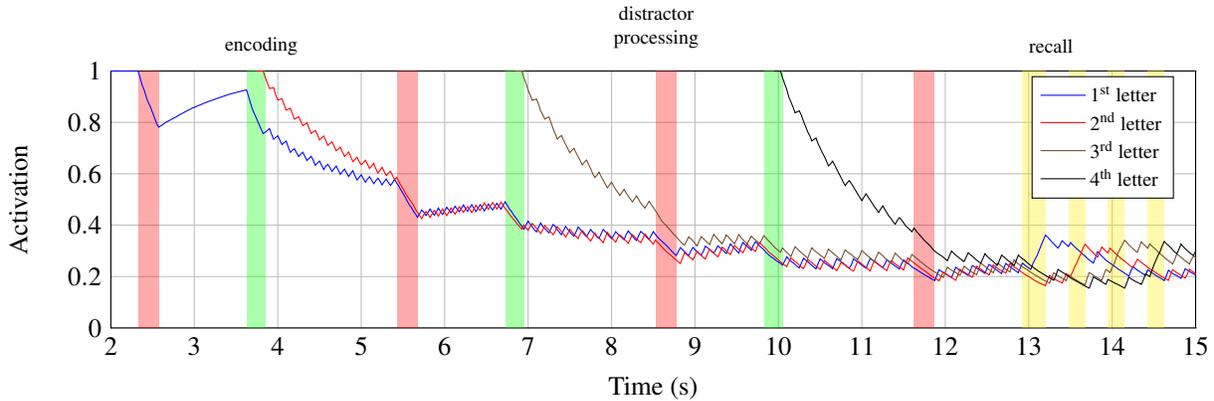
Figure 1: Example activation trace for 4 letters with 1 distractor over 1s intervals between letters

Lewandowsky (2011). Most prevalently, TBRS cannot account for the phonological similarity effect, intertrial interference effects, and feature overlap effects between memoranda and distractors. In addition, decay theories have difficulty predicting that while introduction of three identical distractors (a simple burst) does not impair memory more than introduction of a single distractor, the introduction of three unique distractors (a complex burst) significantly impairs memory (Lewandowsky, Geiger, Morrell, & Oberauer, 2010). These results have been called into question since complex bursts likely capture attention longer than simple bursts, but without having precise measurements of attentional capture it is difficult to determine what the TBRS model actually predicts (Barrouillet, Portrat, Vergauwe, Diependaele, & Camos, 2011). Here is just one example of results remaining in contention as a direct result of the lack of a rich model of attention. In any case, it is clear that TBRS leaves much to be explained in the realm of modality-specific effects.

Interference

In contrast, interference theories of forgetting reject the idea of memory decay in favor of representation-based interference between items in memory based on novelty or feature-sharing (Oberauer, Lewandowsky, Farrell, Jarrold, & Greaves, 2012). SOB-CS is an implementation of an interference account of forgetting in working memory adapted to account for the cognitive load effect found in Barrouillet et al. (2004). The core assumption of SOB-CS enabling its generalization to the CST is that distractors are involuntarily encoded into memory causing interference with memoranda, and therefore must be removed using an attention-demanding unbinding process. In contrast to decay-based theories, SOB-CS assumes active removal of irrelevant memory items rather than active maintenance of relevant memoranda, and default maintenance of memoranda rather than default decay (Oberauer et al., 2012). This dynamic makes detangling the predictions of SOB-CS from TBRS quite difficult— to do so, it is necessary to break the time-course of distractor bursts into operation duration, in which distractor processing occupies attention, and free time, in which memory representations are being attended to, and to analyze their effects in-

dependently. TBRS predicts a positive effect of free time, since free time is used to refresh memoranda, and a negative effect of operation duration, since memoranda decay during this time. While SOB-CS likewise predicts a positive effect of free time, since free time is used to unbind distractor representations, it predicts little effect of operation duration, since memoranda remain at constant activation over time. These predictions were put to the test, and the small effect of operation duration seems to favor SOB-CS (Oberauer et al., 2012).

It should be noted that there is little doubt that interference plays a role in working memory. The major discrepancy is whether it is the sole source of forgetting, or only plays a secondary role in modulating the behavior of a larger source of forgetting. Specifically, interference theories have difficulty explaining the cognitive load effect when distractors and memoranda are categorically different, or when the CST includes repetitive (but not predictable) distractors as in (Barrouillet et al., 2011).

TBRS in ARCADIA

In the following we use the reading digit span (RDS) task from Experiment 7 in (Barrouillet et al., 2004) as a case study of time-based resource sharing in ARCADIA. The RDS task is a CST in which the memoranda are letters and the distractors consist of digits to be read out loud. At the end of the trial, the letters are recalled in serial order. The reading digit span is then the maximum number of letters that can be recalled in serial order with reading digits as a distracting task; it is determined by incrementally increasing the number of letters to be recalled until recall falters.

With respect to the RDS task, our model makes specific predictions in line with TBRS about the qualitative behavior of activations of items in working memory due to momentary shifts in attention. When a new memory item is being encoded (see the green regions of Figure 1), the activation of that item should rise to an asymptote while the activations of all other items exponentially decrease. In the attention-demanding portions of recall (see the yellow regions of Figure 1), activations follow the same pattern; the item being recalled increases in activation while all other items decrease in activation. Within the attention-demanding portion of digit
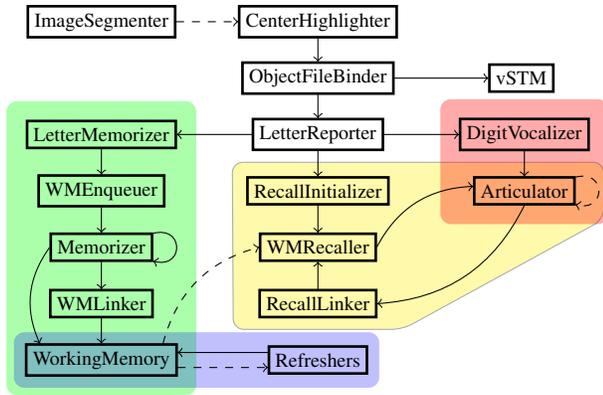
Figure 2: Information flow between components used during visual processing (unshaded), memory encoding (green), distractor processing (red), attentional refreshing (blue), and recall (yellow) in the TBRS model in ARCADIA.

vocalization (see the red regions of Figure 1) all memory items suffer from decay, since attention is momentarily diverted towards the distractor. During the rest of the CST, however, attention is freely available to maintain the activations of working memory representations.

ARCADIA is an attention-centric framework for the construction of cognitive models capable of receiving perceptual input from an environment (Bridewell & Bello, 2016). In every cycle of the model (corresponding to 25ms), an attentional strategy selects a unitary focus of attention from the outputs of numerous processing components (collectively called accessible content), which in turn gets broadcast globally along with the rest of accessible content back to the components to modify processing on the next cycle. This task-specific selection mechanism puts serial constraints on attentional processes while allowing for parallel execution of dissociable components. For the RDS task, maximal priority is assigned to ongoing vocalizations still requiring attention—this ensures that a vocalization will continue uninterrupted by visual processing or attentional refreshing. All other attention-demanding actions (including encoding, memorization, and recall) are prioritized next. In the absence of any action requests visual processing is preferred, so that processing of a new stimulus initiates directly after its presentation. Finally, any attentional refreshing requests are attended to. Since refreshing has the lowest priority in the task, ARCADIA will only refresh during free time.

The functional organization of the TBRS model in ARCADIA is depicted in Figure 2; the unshaded components are involved in visual processing, the green shaded components in encoding, the blue shaded components in attentional refreshing, the red components in distractor processing, and the yellow components (and *WorkingMemory*) in recall. Solid lines indicate information flow requiring serial attention, while dashed lines indicate information flow not strictly demanding attention. Note that all items in working memory are globally accessible pre-attentively (the outgoing arrows from working memory are dashed); as a result, the compo-

nents accessing working-memory representations may only access items which are sufficiently active relative to that component's function. For this reason ARCADIA uses two activation threshold parameters over working memory items; θ_r and θ_v are the minimum activation levels for working memory items to be retrieved and vocalized, respectively.

Visual Processing

Before being encoded into working memory, each stimulus must first be processed through ARCADIA's visual pipeline (see the unshaded area of Figure 2). This pipeline begins with the pre-attentive segmentation of the visual scene into contours defined by each visible stimulus (*ImageSegmenter*). The *CenterHighlighter* then pre-attentively produces a fixation request causing the model to attend to the stimulus. After the stimulus is attentively fixated, its properties are bound by the *ObjectFileBinder* into a single object representation for storage in visual short-term memory (*vSTM*). Finally, the semantic content of the stimulus, its character value, is extracted by the *LetterReporter* using optical character recognition. The visual pipeline thus consumes four cycles of attentional processing each lasting 25ms, making for a total of 100ms of attentional demand.
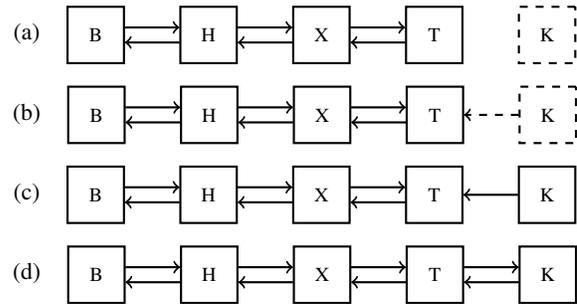

Figure 3: Memory Representation in ARCADIA

Encoding

Since letters must be recalled in serial order, they need to be encoded with positional information (see the green area of Figure 2). We use a doubly linked list representation in which each memorandum is assigned a unique identification number (UID), and is associated with the UIDs of neighboring items in the list (see Figure 3). After visual processing (Figure 3a), and as a preprocessing stage of encoding, the *WMEnqueuer* retrieves the preceding memorandum, associates its UID to the current item, and assigns the item its own UID (Figure 3b). Next, the stimulus is stored in working memory by the *LetterMemorizer* and the *Memorizer* at a baseline activation b; this memorization process takes t_m cycles (Figure 3c). In ARCADIA visual processing takes roughly 100ms and preprocessing takes 50ms, so we set t_m to two cycles (50ms) to place total encoding time squarely inside the current estimate of 150-300ms (Oberauer et al., 2012). After the new item is successfully encoded into working memory, the *WMLinker*

updates the preceding item to associate it with the new item (Figure 3d). All of the operations involved in encoding require attention to be processed.

Distractor Processing

The distractors in the RDS task are visually-presented digits which must be read out loud as quickly as possible. In ARCADIA, visual processing of digits follows the same pipeline as letters (see above). After the character is recognized by the *LetterReporter*, the *DigitVocalizer* extracts the digit's lexical representation and issues a vocalization request. Vocalization is carried out by the *Articulator*, which uses text-to-speech software to produce audible sound clips as well as articulation times; these times are used to synchronize speech execution with ARCADIA's own simulation time. Individual differences in articulation times are accounted for by multiplying the raw articulation time by the free parameter a_d, the vocalization duration factor.

It is often presumed that articulation of a stimulus does not occupy the attentional bottleneck for any lengthy duration, since articulatory suppression does not prevent attentional refreshing (Barrouillet et al., 2011). However, to our knowledge, there exists at present no estimate of how much time speech requires attention. For this reason we incorporated another parameter, a_a, which defines the proportion of articulation time that must be attended. Thus, the total amount of time attended to each distractor is $t_p + a_d a_a t_v$, where t_p is the visual processing time for the stimulus (approximately 100ms), t_v is the raw vocalization time (approximately 400ms), and $0 \leq a_a \leq 1$. After this time, the model engages in attentional refreshing.

Decay and Attentional Refreshing

By default, items in working memory decay exponentially every cycle. When an item is attended to, whether in service of refreshing, vocalization, or memory updating, the activation of that item is instead increased towards an asymptote of 1. The mean rates of decay and refreshing are determined by the free parameters D and R, but the actual rates d and r are drawn every update from Gaussian distributions with mean D and R, and with corresponding standard deviations $pD/4$ and $pR/4$. Thus, the single parameter p is used to control roughly the maximum variation of each distribution as a percentage of D and R. In alignment with TBRS*, the activation $A_i(c)$ of every item i in working memory at cycle c is updated at each cycle according to the following equation:

$$A_i(c+1) = \begin{cases} A_i(c) + r\,(1 - A_i(c)) & \text{if } i \text{ is being attended} \\ A_i(c) - d\,A_i(c) & \text{otherwise} \end{cases}$$

Current estimates of the duration of attentional refreshing are 40-50ms, corresponding to roughly two 25ms cycles in ARCADIA (Lemaire, Pageot, Plancher, & Portrat, 2017). Attentional refreshing then is initiated by strategy-specific components which retrieve the item for refreshing during the first cycle, and the *Refresher* carries out the refresh operation over

the next cycle. As theorized by Oberauer and Lewandowsky (2011), the exact strategy used to attentionally refresh memoranda determines the relative size of primacy and recency effects, because different strategies allocate different amounts of time for refreshing items in each serial position. We implemented three such strategies. The random strategy is implemented by the *RandomRefresher*, which randomly refreshes any above-threshold item in working memory. The cumulative strategy used in TBRS* begins with the first item, and sequentially moves forward throughout the list. Cumulative refreshing is implemented in two components; the *FirstRefresher* always proposes the first item for refreshing, and the *CumulativeRefresher* always proposes the following item. Finally, preventative refreshing is implemented by the *PreventativeRefresher*, which selects the item in working memory with minimum activation above the retrieval threshold. Although Oberauer and Lewandowsky dismissed this strategy as requiring a homunculus, preventative (or least activated first) refreshing can be implemented without an internal homunculus and in TBRS*, it matched the data significantly better than other candidate strategies (Lemaire et al., 2017).

Recall

Vocalization in recall follows roughly the same process as vocalization in distractor processing. The task-specific component *RecallInitializer* detects presentation of the recall cue, in this case the symbol +, and initiates recall of the first item. During recall the item must be retrieved by the *WMRecaller* and then vocalized by the *Articulator*; this means that the item's activation must be greater than both θ_r and θ_v to be correctly recalled. In the event where an item is forgotten, the model uses the substitute word "pass" and the response is categorized as an omission error. Thus, at the moment, ARCADIA will not commit any item errors (we will address this issue in the Discussion). Importantly, temporal decay ensues throughout recall, and attentional refreshing is utilized during the free time in recall. This allows for the possibility, for instance, that an above-threshold item decays below θ_r during recall of an earlier list item. On the other hand, it allows for the possibility that an item below θ_v is refreshed above threshold just before it is needed for recall.

Name	Description	Values
D	mean decay rate	0.025 / 0.05
R	mean refresh rate	0.025 / 0.15
p	proportional randomness	0.5 / 0.5
θ_r	retrieval threshold	0.1 / 0.15
θ_v	vocalization threshold	0.1 / 0.2
b	baseline activation	1.0 / 1.0
a_d	speech duration factor	1.0 / 1.0
a_a	speech attention factor	0.375 / 0.5
t_m	memorization time	2 / 2

Table 1: Parameters for ARCADIA for cumulative / preventative and random refreshing strategies

Results

After a directed grid search minimizing for RMSE on the cognitive load effect (bracketed numbers indicating [low bound, high bound, step size]) over D [0.025, 0.1, 0.025], R [D, 0.25, 0.05], θ_r [0.05, 0.25, 0.05], θ_v [θ_r, 0.25, 0.05], and a_a [0, 1, 0.125], we settled on the parameters listed in the last column of Table 1 for the cumulative and preventative/random refreshing strategies respectively. We use two different parameter settings because while preventative and random refreshing strategies always refresh items above θ_r, the cumulative strategy attempts to refresh all of the list items in forward order and accordingly sometimes fails to refresh items under θ_r. These parameters sufficed to replicate qualitatively the expected activation behavior predicted by TBRS, as well as the cognitive load effect, but they are not fixed.

As shown in Figure 1, the activation trace for a sample trial of the RDS task looks similar to that of TBRS* (the shaded regions indicate the attention-demanding portions of processes of identical color in Figure 2). In contrast to TBRS*, items are encoded at a fixed baseline activation (in this case at maximal activation). Additionally, refreshing can occur during recall, providing maintenance for items waiting to be recalled.
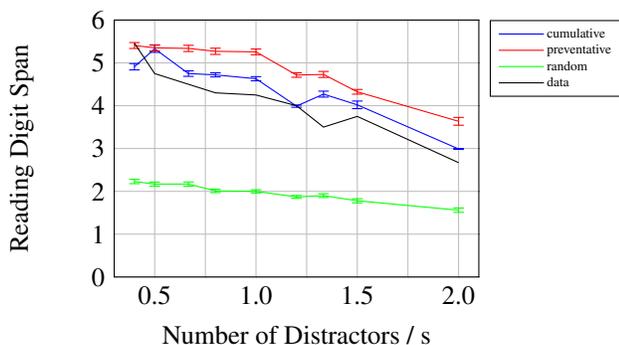

Figure 4: Cognitive Load Effect in ARCADIA

Cognitive Load Effect

In addition to capturing the behavior of activations in working memory predicted by TBRS, ARCADIA captures the cognitive load effect: as the rate of distractor processing increases, the reading digit span decreases linearly. Figure 4 shows the cognitive load effects for ARCADIA on trials with 4, 8, or 12 digits between each letter presented during 6s, 8s, or 10s intervals. The random (green line, $r = .94$, RMSE = .86, n=1000), preventative (red line, $r = .89$, RMSE = .52, n=1000), and cumulative (blue line, $r = .88$, RMSE = .39, n=1000) refreshing strategies all correlate significantly (ps $<$ .002) with data in Barrouillet et al. (2004) plotted in black, but only the cumulative and preventative strategies provide close approximations to the data.

Primacy and Recency Effects

As predicted by TBRS*, plotting recall accuracy as a function of serial position in the list reveals that the general shape of recall accuracy varies dramatically with respect to the refreshing strategy (see Figure 5). Under cumulative refreshing (plotted in blue), ARCADIA exhibits a large primacy effect with little to no recency effect. Random refreshing (plotted in green) generates the opposite behavior; little to no primacy, but extended recency. Finally, preventative refreshing (plotted in red) exhibits a much more even curve with little semblance of primacy or recency. This failure to capture both primacy and recency effects simultaneously may result from the choice of our particular parameter values, since TBRS* captures the serial position curve well without recourse to a primacy gradient (Oberauer et al., 2012).
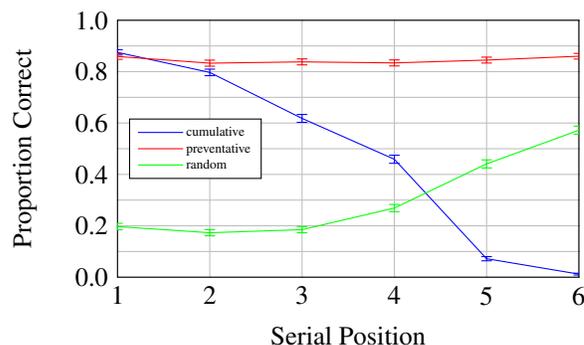

Figure 5: Serial position curves for cumulative, preventative, and random refreshing strategies

Discussion

These results demonstrate that our attention-centric model of working memory implements the TBRS theory closely. While our current set of parameters do not reproduce serial position curves observed in human data, they do highlight the expected effect of each refreshing strategy over time. Cumulative refreshing exhibits primacy since the first item is refreshed most often, random refreshing exhibits recency since early list items have more time to decay, and preventative refreshing exhibits even recall since each item is refreshed as needed. It is likely then that parameter fitting for serial position data would yield models capable of explaining both cognitive load and serial position effects.

Conclusion

The TBRS model of working memory, as implemented in ARCADIA, explains the cognitive load effect as the trade-off between attentional capture in processing and storage. This approach not only allows for our model to yield predictions about the effects of attentional capture on recall, but is also extensible to effects involving the privileged status of the focus of attention. In one study, reaction times were reduced for the last-attended item (Vergauwe & Langerock, 2017). Additionally, time costs associated with switching attention between memoranda and distractor processing affect cognitive load (Barrouillet et al., 2004). It is possible that interference

in working memory arises from a general center-surround inhibition over the focus of attention; under this interpretation working memory items and distractors interfere with one another in virtue of being attended, not strictly because of involuntarily encoding of distractors into working memory. Following this line of reasoning there is evidence that the competition between exogenous and endogenous attentional demands directly determines the extent to which working memory representations bias visual search (Kiyonaga & Egner, 2014). ARCADIA has built into its architecture a heightened status for the focus of attention; since all of the effects above presumably rely on this sort of heightened state, ARCADIA stands out as a desirable framework for their explanation.

While ARCADIA currently predicts the cognitive load effect, more could be done to produce realistic serial position curves and error distributions— minimally, incorporating this data in parameter fitting should improve our results. We have seen consistently that the refreshing strategy largely determines the shape of the serial position curve; if each strategy fails to reproduce primacy and recency on its own, then dynamic competition between different strategies could better account for the data. One possibility is that participants use a default refreshing strategy spontaneously but retain the ability to deliberately impose a different strategy with effort. It has been shown that people engage in spontaneous refreshing, but further research must address the degree to which deliberate strategies are utilized (Vergauwe & Langerock, 2017). In addition, articulatory rehearsal is often considered to be a complementary process to attentional refreshing (Lucidi et al., 2016); however, it has also been hypothesized to have little causal influence on recall (Lewandowsky & Oberauer, 2015). Although verbal rehearsal puts some demand on attention, limiting the short-term use of refreshing, it remains unclear how verbal rehearsal interacts with refreshing to generate primacy and recency.

As a simplifying assumption, this version of working memory in ARCADIA only commits omissive errors. The human data for CSTs, however, also include transposition errors (recall of the incorrect list item at a position) and extralist intrusions (recall of a letter not belonging to the memory list). Transposition profiles for each item follow a gradient which peaks at the correct serial position, and decreases nonlinearly as serial position increases or decreases (Oberauer et al., 2012). Reproducing this finding while maintaining a localist representation could require imposing similar probability distributions for each item upon retrieval. The incorporation of primacy, recency, transposition error, and other effects described above would provide a wealth of explanatory power for working memory research.

Acknowledgments

References

Altmann, E. M., & Schunn, C. (2012). Decay versus interference: A new look at an old interaction. *Psychological science*, *23*.

Barrouillet, P., Bernardin, S., & Camos, V. (2004). Time constraints and resource sharing in adults' working memory spans. *Journal of Experimental Psychology: General*, *133*(1), 83–100.

Barrouillet, P., & Camos, V. (2007). The time-based resource-sharing model of working memory. In N. Osaka, R. H. Logie, & M. D'Esposito (Eds.), *The cognitive neuroscience of working memory*. Oxford University Press.

Barrouillet, P., Portrat, S., Vergauwe, E., Diependaele, K., & Camos, V. (2011). Further evidence for temporal decay in working memory: reply to Lewandowsky and Oberauer (2009). *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *37*(5), 1302-17.

Bridewell, W., & Bello, P. (2016). A theory of attention for cognitive systems. In *Fourth annual conference on advances in cognitive systems.*

Kiyonaga, A., & Egner, T. (2014). Resource-sharing between internal maintenance and external selection modulates attentional capture by working memory content. *Frontiers in human neuroscience*, *8*, 670.

Lemaire, B., Pageot, A., Plancher, G., & Portrat, S. (2017). What is the time course of working memory attentional refreshing? *Psychonomic Bulletin & Review*.

Lewandowsky, S., Geiger, S., Morrell, D., & Oberauer, K. (2010). Turning simple span into complex span: Time for decay or interference from distractors? *Journal of experimental psychology. Learning, memory, and cognition*, *36*, 958-78.

Lewandowsky, S., & Oberauer, K. (2015). Rehearsal in serial recall: An unworkable solution to the nonexistent problem of decay. *Psychological Review*, *122*, 674-699.

Lucidi, A., Langerock, N., Hoareau, V., Lemaire, B., Camos, V., & Barrouillet, P. (2016). Working memory still needs verbal rehearsal. *Memory & Cognition*, *44*(2), 197–206.

Oberauer, K., & Lewandowsky, S. (2011). Modeling working memory: a computational implementation of the time-based resource-sharing theory. *Psychonomic Bulletin & Review*, *18*(1), 10–45.

Oberauer, K., Lewandowsky, S., Farrell, S., Jarrold, C., & Greaves, M. (2012). Modeling working memory: An interference model of complex span. *Psychonomic Bulletin & Review*, *19*(5), 779–819.

Tan, L., & Ward, G. (2008). Rehearsal in immediate serial recall. *Psychonomic Bulletin & Review*, *15*(3), 535–542.

Vergauwe, E., & Langerock, N. (2017). Attentional refreshing of information in working memory: Increased immediate accessibility of just-refreshed representations. *Journal of Memory and Language*, *96*, 23 - 35.