Chaining models of serial recall can produce positional errors

Jeremy B. Caplan[1,2], Amirhossein Shafaghat Ardebili[2], and Yang S. Liu[3]

[1]Department of Psychology, University of Alberta, Edmonton, AB, T6G 2E9, Canada
[2]Neuroscience and Mental Health Institute, University of Alberta, Edmonton, AB, T6G 2E1, Canada
[3]Department of Psychiatry, University of Alberta, Edmonton, AB, T6G 2R3, Canada

Abstract

A major argument for positional-coding over associative chaining models of immediate serial recall has been the high probability that an error from a prior list will appear in its correct serial-position, so-called "protrusions." Here we show that a chaining model can produce protrusions if it includes three characteristics that have been incorporated into published chaining models: a) a "start-signal" item is associated with all first list-items, b) memory is not cleared following each list, and c) the retrieval cue for each item is always the full non-redintegrated retrieved information, regardless of the response. The model covertly recalls all studied lists in parallel (weighted by recency), such that when prior-list items intrude, they predominantly occur at the correct output position. In addition to fitting prior protrusion data, we report two new data sets that question the ubiquity of the simple protrusion-dominance characteristic. These findings show that protrusions cannot falsify an associative basis for serial-order memory and speak to the plausibility of mixture models.

Keywords: Associative chaining; positional coding; prior-list intrusions; serial-order memory; immediate serial recall; proactive interference

Jeremy B. Caplan ⓘ https://orcid.org/0000-0002-8542-9900 Yang S. Liu ⓘ https://orcid.org/0000-0003-0406-8056

Amirhossein Shafaghat Ardebili ⓘ https://orcid.org/0000-0001-9574-0623

Corresponding author: Jeremy B. Caplan. Department of Psychology, Biological Sciences Building, University of Alberta, Edmonton, Alberta T6G 2E9, Canada, E-mail: jcaplan@ualberta.ca, Tel: +1.780.492.5265, Fax: +1.780.492.1768.

Introduction

A major goal of verbal memory research is to determine how participants store list in serial-order in memory. Models can be broadly divided into two classes (Murdock, 1974). In associative chaining models, originally suggested by Ebbinghaus (1885/1913), it is assumed that participants store associations between pairs of list items. If only nearest-neighbour pairs are learned, this resembles a chain, A–B, B–C, ... Assuming the participant can retrieve the first list item (which we discuss below), A, item A can be used as a cue to retrieve B, B as a cue for C, and so on (e.g., Lewandowsky & Murdock, 1989; Logan, 2021; Murdock, 1993; Osth & Hurlstone, in press; Solway et al., 2012). Alternatively, positional-coding models, first proposed by Ladd and Woodworth (1911) and then revived by Conrad (1965), assume that some representation of absolute position, or relative position (the word "order" is often used to denote relative position, when coding is not by absolute position, but by a signal the offers relative earlier/later information), is explicitly associated with each list item. We can denote this symbolically as 1–A, 2–B, 3–C, ..., where the numbers stand in for one of the numerous ways in which position and order codes have been implemented. To recall the list, the model probes with each successive position code and retrieves the item that is most strongly linked to that position. Because these models have been developed as alternatives to associative chaining (e.g., Brown et al., 2007; Burgess & Hitch, 1999; Farrell, 2012; Hartley et al., 2016; Henson, 1998; Lewandowsky & Farrell, 2000, 2008), by design, they typically exclude direct associations between items (with some exceptions, such as Burgess and Hitch, 1992; Jensen and Lisman, 2005; Lisman and Idiart, 1995; Logan and Cox, 2021).

Some modellers continue to argue that inter-item associations cannot be part of the account of immediate serial recall, and pure positional-coding or order-coding models are better accounts of behavioural data (e.g., Hurlstone et al., 2014). However, the repertoire of support for associative chaining, sometimes at the expense of positional-coding, has been accumulating (Caplan, 2015; Caplan et al., 2015; Kahana et al., 2010; Lindsey & Logan, 2019; Serra & Nairne, 2000; Solway et al., 2012). Some have argued that such mixture findings suggest that a full account of serial recall behaviour may need to be a mixture of inter-item and item-position processes (Caplan, 2015; Osth & Dennis, 2015; Osth & Hurlstone, in press). The goal of research then becomes not to select between associative chaining and positional coding, but to identify which behavioural phenomena are produced by one or the other type of mechanism, to constrain such mixture or hybrid models.

Here we test the validity of one of the main arguments put forth in favour of position-coding and thought to rule out associative chaining (Henson, 1996; Kahana, 2012; Osth & Hurlstone, in press): the finding that prior-list intrusions, more often than expected by chance, appear at the correct response position (Conrad, 1960; Fischer-Baum & McCloskey, 2015; Henson et al., 1996), foreshadowed by Melton and von Lackum (1941) in serial anticipation, and termed 'protrusions' by Henson (1996). These authors have viewed the existence of above-chance rates of protrusions to be evidence of the presence of position–item cueing. If position (absolute or relative) is used to retrieve items during serial recall, it makes sense that items from prior lists might be retrieved (weighted down by forgetting), particularly if their position matches the cue. We question the reverse inference, that the finding protrusion-dominance of prior-list intrusions can only have been produced if position

were part of the retrieval cue. If it can be shown that a chaining model can produce pattern of errors, *without explicit position-coding or cueing,* this would suggest that the existence of protrusions is not diagnostic of chaining versus positional coding models. This would add to the body of evidence suggesting that chaining models should not be disregarded as accounts of serial-recall data, or if not full accounts, in conjunction with other mechanisms.

Dennis (2009) has done something similar to this. In his model, associative strengths are learned between pairs of items, but the model differs from prior chaining models, in that memory of the list is assumed to be retrieved as a whole, whereby the associative strengths act to constrain the retrieved sequence, and then the sequence is assumed to be read out after having been entirely retrieved. This model produced compelling patterns of protrusions, resembling those reported by Conrad (1960). Dennis' model produced intrusions because of a so-called "common component," which is supposed to result from inter-item similarity. Because of the way the model is designed, this common component accumulates from the edges of the list toward the middle, and does so in the same way for all lists. Thus, errors from prior lists, cued in part by this cumulative similarity value, will tend to be most strongly cued closest to their original serial position, appearing as protrusions. This already suggests that protrusions do not rule out all forms of associative-chaining models. Indeed, we do not dispute that Dennis' account could be the correct one. However, Dennis' model deviates significantly from other chaining models, particularly with respect to the holistic retrieval assumption (but this was also used by Shiffrin & Cook, 1978), which seems less tenable as the list length increases. Interestingly, the signal that emerges from the accumulation of the common component bears some resemblance to the start- and end-markers of the Start-End Model (Henson, 1998). The model still retrieves lists based on inter-item associations, and never uses the common component as an explicit positional retrieval cue. In this regard, one would classify it is as a chaining model and not a positional-coding model. Still, embedded within those associations is information about relative position. Our approach will differ, in that we will not have any positional information encoded in inter-item associations. In this regard as well, our model will stick closer to traditional chaining models.

Our goal here was to ask whether a more traditional chaining model (especially Lewandowsky & Murdock, 1989; Murdock, 1993; Solway et al., 2012), assuming sequential retrieval, could produce protrusions at a substantial rate, without introducing additional mechanisms. In fact, all the modifications we will introduce are either relaxing of assumptions or slight generalizations of assumptions that have been incorporated into chaining models. Three assumptions will offer the chaining model the ability to produce prior-list intrusions that are predominantly protrusions:

Assumption 1: A start-signal dedicated item. First, as has always been the case, a chaining model can start recall with the first item of the list (given A, recall B, etc.), but there needs to be a way in which the model retrieves that first item. In real behavioural data, although highly accurate, recall of the first list item is not strictly perfect. Borrowing from Lewandowsky and Murdock (1989), who were inspired by Shiffrin and Cook (1978), we assume the existence of some representation of the "start" of the list. This "start-signal" is assumed to be a vector with the same characteristics as item vectors, but is somehow already "known" to the model to have a special role in indicating the start of a list. Thus, when cued to begin serial recall, the model probes with the start-signal. Because the same start-signal will be used for every list, probing with the start-signal will retrieve a linear

combination of all previously studied lists (weighted by recency).

Assumption 2: Memory is not cleared between lists. This is necessary for any model whenever one is interested in prior-list intrusions. The consequence for the chaining model with the start signal explicitly implemented is that memory of prior lists will still be present (albeit weighted down by forgetting) to be retrieved in parallel.

Assumption 3: The retrieval cue for each successive recall attempt is the previously retrieved vector. This is arguably the most significant deviation from prior models, although it has a precedent. In the implementation by Lewandowsky and Murdock (1989), when an item is applied as a retrieval cue, the correct subsequent item might be retrieved, but in a linear superposition of other items along with noise. To produce a response, this noisy vector is redintegrated to produce a "cleaned-up" response that has a chance of being scored as correct— something like the entry for a word in the model's previously learned lexicon. In a typical chaining model, this redintegrated item is used as the cue for the next item (there is an interesting parallel to the classic method of serial anticipation, whenever the response is correct). Lewandowsky and Murdock (1989) proposed an elegant way to counter the obvious criticism of chaining, that if the chain is "broken" because a response is not produced, the model has no way to continue because the item-cue for the next item is not available. Rather than include remote associations (Logan, 2021; Solway et al., 2012), Lewandowsky and Murdock (1989) proposed that when an omission is made, the model should cue with the retrieved, non-redintegrated vector (essentially, the unrefined contents of memory). This enables the model to proceed with correct responses following an error of omission. For example, if the BC association is weak, leading the model to fail to produce C, the noisy retrieved vector might next retrieve D, if the CD association happens to be strong. Our modification of this rule is, in some sense, a simplification of the rule. We assume that the functional cue for the subsequent recall attempt is *always* the retrieved, non-redintegrated vector. Redintegration only serves the purpose of producing the response, itself. This type of implementation was recently used by Logan (2021) and has been noted to enable chaining models to explain some phenomena that may be partly immune to accuracy of the immediate-prior item, such as in phonological similarity paradigms, resembling accounts of such phenomena that have been implemented in positional-coding models (Caplan, 2015; Caplan et al., 2015; Logan, 2021; Osth & Hurlstone, in press).

Combined, these three assumptions have the following consequences. Denote the current list in uppercase, $ABCDEF$, and the immediate-previous list in lowercase, $abcdef$. Denote the start-signal as S. Let α_i denote an encoding strength selected from an independent, identical random distribution. Probing with S will retrieve $\alpha_A A + \alpha_a a$, where α are retrieval coefficients (strengths). Usually it will be the case that $\alpha_a < \alpha_A$ due to forgetting. The next cue will be $\alpha_A A + \alpha_a a$, which will tend to retrieve approximately $\alpha_A \alpha_B B + \alpha_a \alpha_b b$ (plus cross-terms). In general, this will ensure that the current-list items will be retrieved with greater strength than prior-list items, but in the very case for which recall of the current-list item fails, there is a higher probability that the prior-list term will be stronger than the current-list term, leading to the possibility of a prior-list item intruding. If this happens, the prior-list item will be most likely recalled in its correct output position, because the prior list has been retrieved in the background throughout the course of recall.

Next we describe the model formally and fit it to the data reported by Henson (1998),

first reported by Henson (1996), experiment 5. We then report new analyses of two data sets we had acquired for other purposes, to test the robustness of protrusion-dominance and then reconsider the generality and possible boundary conditions for the protrusion finding.

Model description

Overview. We describe a model that is designed like standard associative chaining models. Although described with vector representations of items, the simulations here are of a strength-model implementation (similar to Solway et al., 2012), for tractability and without loss of generality for the phenomena we investigate here (i.e., the large-dimensionality approximation, for which items will tend to be nearly orthogonal to one another). The model stores only nearest-neighbour associations, and we make explicit the standard assumption that the first item is, itself, linked to some representation of the start of the list. Making this concrete, we simply assume that the "start-signal" is simply one dedicated item that, in all other ways, is treated just the same as any other item. We build on existing ideas about associative chaining, with a few critical assumptions: 1) the same start-signal is associated to the first item of all lists, 2) memory is not cleared prior to each new list, but rather, prior lists are forgotten gradually in the same manner as forgetting occurs during list presentation and serial-recall; this is required to produce any non-trivial rate of prior-list intrusions, and 3) the cue for the next item is always the non-redintegrated retrieved information (a weighted sum of items), as in Lewandowsky and Murdock (1989). Unlike the latter, we assume this occurs regardless of the outcome of the current response (Logan, 2021).

Variables. We assume that, for a set of N items in the lexicon, every item, i, has an associative strength to every other item, j, which we denote X_{ij}. Thus, X is a $N \times N$ matrix. We also assume that there is a "special" item dedicated to mark the start of all lists, which we call the "start-signal." To keep the notation simple, and to emphasize that in all other ways, the start signal is treated equivalently to all other items, the start-signal is written as item 0. Thus, associations from the start-signal to all items are x_{0i}, $i = 1..N$ (and associations *to* the start-signal from all items are x_{i0}, $i = 1..N$):

$$X = \begin{pmatrix} \begin{array}{c|ccccc} x_{0,0} & x_{0,1} & x_{0,2} & \dots & \dots & x_{0,n} \\ \hline x_{1,0} & x_{1,1} & x_{1,2} & \dots & \dots & x_{1,n} \\ x_{2,0} & x_{2,1} & & & & \\ \dots & \dots & & & \dots & \\ \dots & \dots & & & & \\ x_{n,0} & x_{n,1} & \dots & & & x_{n,n} \end{array} \end{pmatrix}. \tag{1}$$

Study. Memory for a list of L items, $\{f_k\}$, $k = 1..L$, is learned by adding a strength drawn from a Gaussian distribution, $N(\mu, \sigma)$, where $\mu = 1$ (without loss of generality; $\mu = 1$ just sets the scale, and σ should be interpreted relative to μ), for each successive pair of items. The first item is always associated with the start-signal, and all list items are associated with their successors. Although associations are highly correlated in the forward and backward directions (Kahana, 2002), order can be well disambiguated when transitions within a serial list are probed (Caplan et al., 2006; Kahana & Caplan, 2002) and judgements of order are well above chance, though not perfect (Kato & Caplan, 2017).

To accommodate this, we include a parameter, η, to modulate the degree to which forward associative strengths exceed backward associative strengths, but the strengths are perfectly correlated (as also done by Solway et al., 2012), in line with preservation of a high forward–backward correlation of associations embedded within serial lists (Caplan, 2005; Caplan et al., 2006). Thus, $x_{ji} = \eta x_{ij}, \ \forall i < j$. In addition, in keeping with chaining models like TODAM, forgetting, parameterized by $\rho < 1$ (but the value of ρ is typically very close to 1), is implemented each time an item is encoded; thus:

$$X = \rho X \ \text{(incremental forgetting of all stored association strengths)} \tag{2}$$

$$X_{f_{k-1}f_k} = X_{f_{k-1}f_k} + N(\mu, \sigma)$$
$$X_{f_k f_{k-1}} = \eta X_{f_{k-1}f_k} \tag{3}$$
$$\forall \ k = 1..L$$

$$\text{where } f_0 = 0 \ \text{(the start-signal).} \tag{4}$$

Note that auto-associative strength (self strengths) are always fixed at zero, $S_{ii} = 0, \ \forall i$.

   Serial recall. Recall starts by probing memory with the start-signal. Such a single-item probe is represented as a Kronecker Delta vector, $\vec{\delta}(i)$, a column-vector (column vectors are denoted by $\vec{\ }$) with length $N + 1$, which has zeroes everywhere except at i, where the value is 1. To probe with the start-signal on its own, we multiply X by $\vec{\delta}(0)$, always preceded by one iteration of forgetting:

$$X = \rho_r X \ \text{(forgetting)} \tag{5}$$

$$\vec{r} = X\vec{\delta}(0) \equiv \sum_{i=0}^{N} X_{0i}. \tag{6}$$

We allow the amount of forgetting during recall (output interference) to vary, and to aid interpretation, express this forgetting parameter as a power of forgetting during study, $\rho_r = \rho^r$. Likewise, we parameterize forgetting between lists, $X = \rho_b X$, where $\rho_b = \rho^b$; r and b are free parameters to model the idea that forgetting proceeds at a different rate during recall than during study. We also expect that forgetting will be more pronounced between lists, otherwise the model would produce too many prior-list intrusions at the expense of accuracy, which is fairly high for the list lengths we consider. However, if between-list forgetting is excessive, the model will not commit any prior-list intrusions.

   The response is chosen using a winner-take-all rule, subject to a threshold, θ:

$$\left. \begin{array}{ll} k & r_k = \sup{(\vec{r})} \ \text{and} \ r_k > \theta \\ \text{Omission} & \text{otherwise} \end{array} \right\}, \tag{7}$$

where $\sup{()}$ indicates supremum (the largest, positive value). In a sense, θ has little theoretical importance. If θ is too high, omissions will dominate error-responses and very little information about prior-list intrusions will be available. Unsurprisingly, the fits produced fairly low values of θ, indicating that participants were somewhat likely to commit intrusions and not only omissions.

On each successive recall attempt, p, the retrieved vector from the previous step, $\vec{r}(p-1)$, is used to probe memory, regardless of the response made; thus:

$$X = \rho_r X \ \text{(forgetting)} \tag{8}$$

$$\vec{r}(p) = X\hat{r}(p-1) \equiv \sum_{j=0}^{N} \sum_{i=0}^{N} \hat{r}(p-1)_j X_{ji} \tag{9}$$

$$\text{Respond}:$$

$$\left.\begin{array}{cc} k & r(p)_k = \sup\left(\vec{r}(p)\right) \ \text{and} \ r(p)_k > \theta \\ \text{Omission} & \text{otherwise} \end{array}\right\}. \tag{10}$$

Note that we probe with the *normalized* retrieved strength vector from the previous recall, $\hat{r}(p-1) \equiv \vec{r}(p-1)/|\vec{r}(p=1)|$, where $|\cdot|$ denotes the magnitude of a vector. Without this normalization, retrieval strengths would increase without bound; among reasons this is problematic, it would render the threshold, θ, ineffective. Because associations are bidirectional, after the start signal (probably) retrieves item 1, item 1 will, in turn, be likely to retrieve the start signal. We assumed that participants have sufficient metaknowledge of the task to avoid producing the start signal as a response and to avoid continuing to use it as a retrieval cue. With this rationale in mind, we set the start signal strength to zero prior to response selection and prior to the cueing operation after the first recall (in fact, it is modelled as already suppressed in the repetition-suppression set, described in the next paragraph). Rather than implementing a stopping rule, we assume that the model continues to recall until $p = L$.

Because human subjects rarely repeated items within a recall sequence (Duncan & Lewandowsky, 2005), we implement response suppression by simply keeping track of recalled items (including the start-signal) and excluding them as response candidates—that is, they are not included in the $\sup()$ evaluation. Importantly, previously recalled items are not suppressed when it comes to cueing the next item. The response-suppression set is cleared after recall of a list terminates.

As the model tackles each successive list, the new first list-item is associated to the start-signal, and memory for the previous lists fades only gently due to forgetting (the ρ parameters). The forgetting parameter, ρ_b, allows for additional forgetting between lists, enabling the model to trade off current-list accuracy with prior-list intrusions.

Simulation

Target data set

We fit the model to the data from Experiment 5 of Henson (1996)[1], plotted by Henson (1998), in the lower panel of his Figure 10. Full methods can be found in the latter two sources. Briefly, the stimulus pool was a closed set of 14 words. Each list in the fixed-list-length condition of the experiment comprised 6 words, drawn at random from one half (seven words) of the stimulus pool. The lists alternated sampling from one half of the stimulus pool and the other half. In most closed-set experiments, the stimuli do not strictly

---

[1]Mean data generously provided by Rik Henson.

alternate, making it ambiguous whether an error is from the prior list or from the current list whenever an item appears on both lists. The strict alternation used by Henson (1996) makes identifying prior-list intrusions unambiguous. Each word was presented for 600 ms followed by a 200-ms blank screen and 1000-ms blank pause at the end of the list. Recall was written, with the number of response lines indicating the list length and a fresh page was given for each list. Although participants were instructed to recall in strict forward order, it is impossible to verify if and when they did so. For our purposes, we assume participants recalled in forward order and did not backtrack. In the upper panel of the figure plotting the data, Henson (1998) plotted a demonstration of his ordinal-coding model, the Start-End Model. In our replot of the data and model points on a single graph (Figure 2), one can see that the model followed the data in that protrusions were more prevalent than intrusions appearing at output positions other than their initial serial position. Although it should be kept in mind that the model was a demonstration, not a best-fit, the model over-predicted proportions of protrusions at several positions, and did not produce an additional peak at position 4, which Henson suggested reflected chunking, not incorporated into that version of the model.

Methods

Implementation. The associative chaining model was written in MATLAB (The Mathworks, Inc.). The word pool of size 14 plus the start-signal "item" resulted in a strength matrix, X, of size 15×15 which was initialized to zero. The main model function simulated 16 lists of length 6, simulating the fixed-list-length block of Henson's experiment. The main model was run 50,000 times for a given parameter set, simulating 50,000 "subjects," to produce model-behavioural data with sufficient resolution for parameter optimization. Such a large number of model-subjects was necessary in large part because prior-list intrusions are rare (particularly for large portions of the parameter space that fit the data poorly). To speed convergence, a parameter set was short-circuited if, after the first model run (16 lists), it produced either zero correct serial recalls or zero prior-list intrusions. If short-circuited, the parameter set was assigned a fitness value based on the maximum possible deviation from the data.

Associative strengths were drawn from a pseudorandom number generator (seeded with the clock time) using the function randn, from a Gaussian distribution with mean=1 and standard deviation σ (free parameter). Because negative associative strengths are hard to interpret, these strengths were then truncated at zero. Response suppression was implemented by accumulating recalled items, and removing the responded items as response options until the end of each recall phase. This leaves open the possibility for items that were not produced as responses nonetheless to be part of the functional cue for the subsequent item.

Parameter optimization. In early iterations of model-fitting, when best-fitting parameter values were bunched up at one end of the search range, the search range for that given parameter was extended. Best-fitting parameter values were found by running a Simplex search (Nelder & Mead, 1965) with a modified version of MATLAB's fminsearch function, minimizing root-mean-squared deviation (rmsd), and reporting the parameter set producing the best (lowest) rmsd. This rmsd was computed between model values and the empirical values for the serial-position curve (strict serial-position scoring; correct if a

Parameter	Meaning	Searched Range	Best-fit Value
σ	Strength CV	0.1–3	0.67
θ	Response threshold	0–0.5	0.0078
ρ (exponent, $z \ni \rho = 1 - 10^{-z}$)	Forgetting during study	1–4	2.12 ($\rho = 0.992$)
ρ^r (exponent, r)	Output interference	1–20	19.35
ρ^b (exponent, b)	Between-list forgetting	1–100	56.65
η	Symmetry	0.5–1	0.95

Table 1

Fits to Henson (1996) data from Experiment 5, fixed list length condition. Parameters are listed along with their searched ranges and best-fitting values. σ is equivalent to the coefficient of variation because the mean strength, μ, was fixed at 1. θ was fixed at zero. For ρ, the parameter searched was the exponent, z, such that $\rho = 1 - 10^z$. For forgetting during recall and between lists, the searched parameter was the power, r or b, respectively. This parameter set fit the Henson (1996) data with BIC=–96.34 (rmsd=0.19).

recalled item is both on the current list and in the correct position) and the full pattern of immediate-prior-list intrusions: probability of an intrusion from the immediately prior list at serial position a being recalled at current-list output position b, for a total of $L + L^2 = L(L+1)$ values. For prior-list intrusions, the deviation was set to its maximum possible value, 1, when the model produced no intrusions (within-list or prior-list, respectively), at the given output position (and thus, the probability is undefined, due to lack of evidence, rather than zero). Remote prior-list intrusions (from lists earlier than the immediately preceding list) were not evaluated because they are rare (Henson, 1996; Osth & Dennis, 2015) and because the alternating stimulus pools made it ambiguous whether a word was remembered from the current list or the second-last list.

In earlier attempts at parameter optimization, we noticed that the model was finding best-fits without protrusions dominating prior-list intrusions. This was because the model tended to over-predict the rates of protrusions, which produced a worse rmsd. Because the goal was to find out whether a chaining model could produce protrusions, when it was previously thought to be impossible, we did not want the model to be penalized for its own qualitative success. To urge the model to find a fit that captured the qualitative dominance of protrusions at all output positions, even if somewhat compromising the quantitative fit, we added a penalty of 1 to the rmsd, equivalent to the maximum deviation one could obtain at any one of the fit points, when the protrusion was not the most frequent source of prior-list intrusions at one or more output positions.

The parameter optimization was run 10 times from random starting points to reduce the chance of finding an idiosyncratic local minimum. The best-fitting parameter set across all ten fits was run once more when producing model-performance figures. Free parameters, their searched ranges and best-fitting values are reported in Table 1.

Results and discussion

The best-fitting model (Table 1 and Figure 1) fit the serial position curve fairly well. Importantly, at each output position, the model produced higher rates of protrusions, where the serial position of the prior-list intrusion matched the output position, than from any other serial position. It produced this characteristic at each output position, and even overestimated the rate of protrusions at all positions. Again, while the fit is not perfect, the high prevalence of protrusions is observed, a finding that was previously thought to be unachievable without position-cueing. Interestingly, Henson's own hand-fit of the Start-End Model to the same data also over-estimated protrusion rates (Figure 2).

To get a subjective feel for how the model produces prior-list intrusions, Figure 3 plots the strength matrix after study of the fourth list on a given run. First, note that earlier lists have weaker strengths (darker colour). Next, note that because symmetry is not perfect ($\eta < 1$), "backward" strengths, depicted on the lower-left half of the matrix, are slightly weaker than forward strengths. Recall that item 0 is the start-signal, itself. It is most strongly associated with the first item of the most recent list (item number 10), but also somewhat associated with the first item of the third list (item 3), albeit less so, due to forgetting (during study, ρ as well as during recall, ρ^r and between lists, ρ^b). This is sufficient to produce the hypothesized phenomenon: recall of the current list, initiated by the start-signal, will retrieve the first item of the first list, but also, generally to a far lower degree, the first item of the previous list. The weighted sum, dominated by these two items, will become the functional cue at output position 2; thus, the second item from the current list will be retrieved along with the second item of the previous list, etc. This covert recall of the previous list propagates until by chance (variability in encoding strength), an item from the previous list is retrieved with greater strength than the current item.

Osth and Dennis (2015) noted that a trivial way any model could produce a dominance of protrusions would be to occasionally retrieve the entire previous list. However, when we spot-checked, the model was not producing protrusions this way. As a quantitative way to verify this, given that a transition was from the immediate-prior list (excluding the final response), the proportions of the next responses being from the previous list, current list or an omission were 0.41, 0.56 and 0.03, respectively. In other words, the model more often returned to the current list than continuing with the previous list.

The covert retrieval that is responsible for the serial-position-preserving nature of the prior-list intrusions is made possible because the model probes with the weighted sum of retrieved items, rather than the redintegrated recalled item. When we switch to the latter (keeping all parameter values the same), Figure 4 shows that, actually the model still produces prior-list intrusions that are dominated by protrusions. However, it now does so, in fact, by often retrieving the entire previous list; the proportions of the next responses being from the previous list, current list or an omission were 0.95, 0.00 and 0.05, respectively.

Given that the chaining model over-estimated protrusions at all positions (as did Henson's fit of the Start-End Model, although less so at all positions apart from the final one; compare Figure 1b with Figure 2), we wondered if the model would be able to produce a tighter fit to the prior-list intrusions. In a single follow-up Simplex parameter optimization fitting only to the prior-list intrusions, the fit was indeed closer to the prior-list intrusion data (Figure 5b). The model was evidently able to do this by reducing current-list accuracy
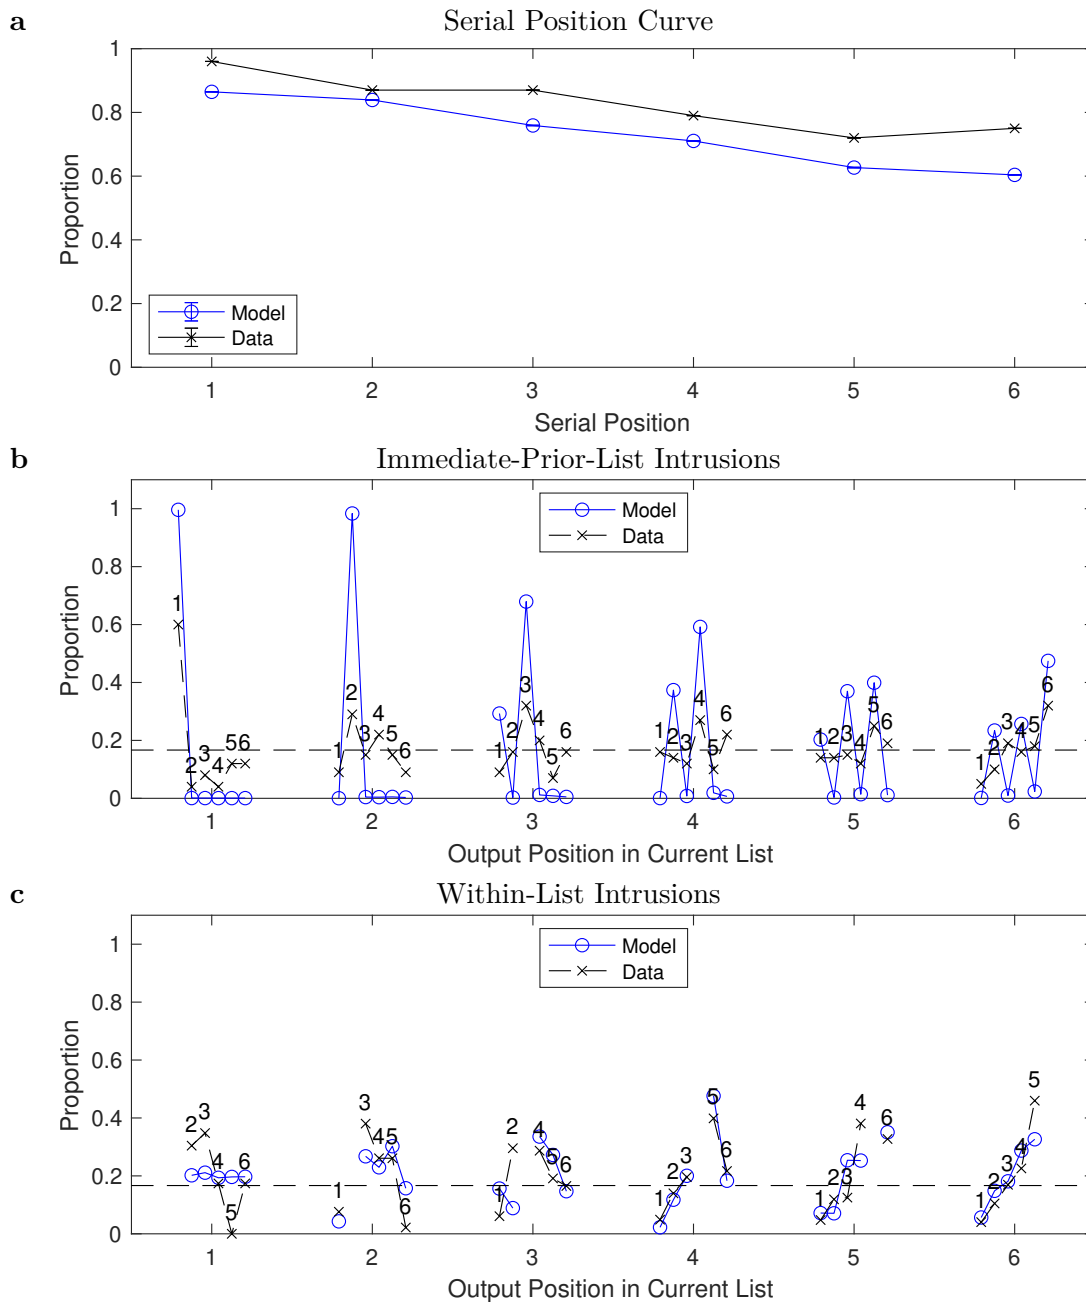
Figure 1

*Simulated LL=6 data plotted alongside behavioural data from Henson (1996), Experiment 5, fixed condition. Error bars plot 95% confidence intervals based standard error of the mean (model; although they are not visible because they are extremely small, due to the high number of model-subjects simulated) and omitted from the data (not available). **a**, Probability of recall as a function of serial position (strict-order scoring). **b**, Probability of an immediate-prior-list intrusion at current-list ouput position (x axis) recalled at each (previous) serial position denoted by numbers above the model values. **c**, Probability of within-list intrusion at current-list ouput position (x axis) recalled at each serial position denoted by numbers above the model values.*
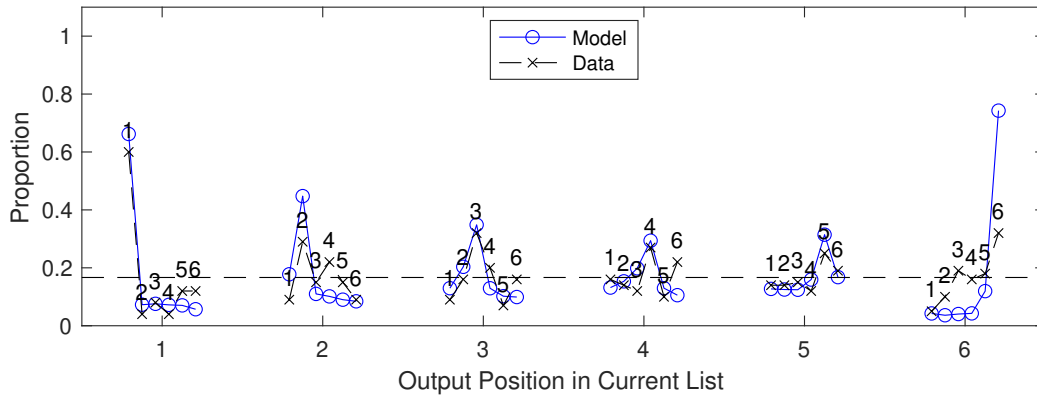
Figure 2

The fit (by hand-tuning to multiple data sets) of the Start-End Model reported by Henson (1998) to the same data we fit with the chaining model. Values were measured from Figure 10, top panel of Henson (1998). Compare with Figure 1b.
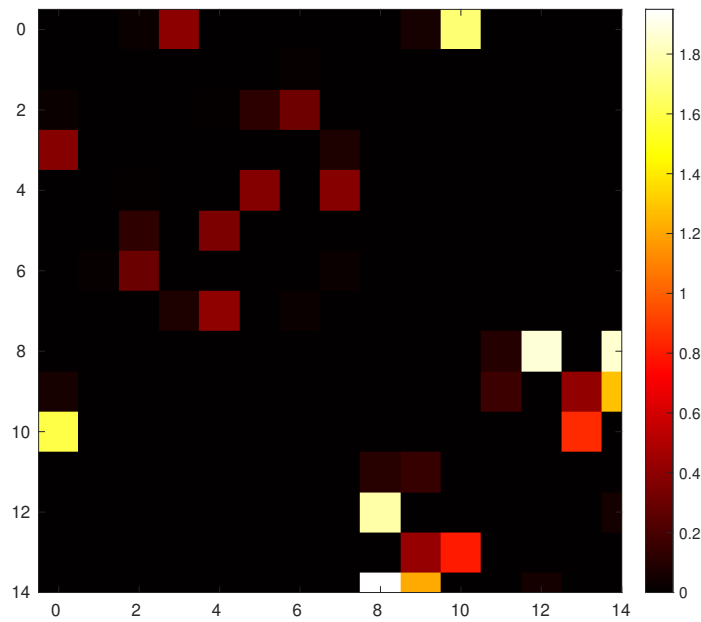

Figure 3

For the best-fitting parameter set: state of the strength matrix after the fourth list has been studied. Colour denotes the associative strength value (dimensionless) between the respective row and column item. Item "0" is the start signal. The session alternated between two pools, depicted here with indices 1–7 and 8–14, respectively.
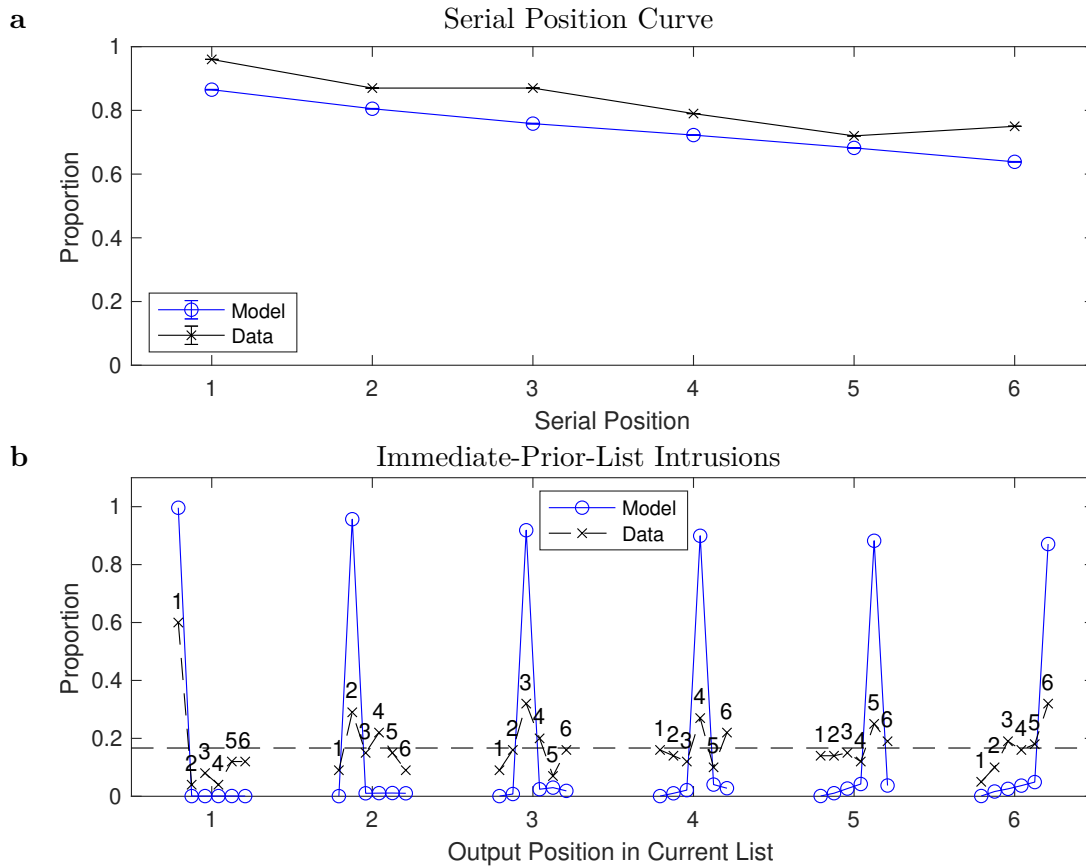
Figure 4

Simulated data using the best-fitting parameters that were fit to the Henson (1996) data, but where recall proceeds with the redintegrated item, when correct, rather than the weighted sum of retrieved items, plotted as in Figure 1, alongside the behavioural data. Error bars plot 95% confidence intervals based standard error of the mean. **a**, *Probability of recall as a function of serial position (relative-order scoring).* **b**, *Probability of an immediate-prior-list intrusion at current-list ouput position (x axis) recalled at each (previous) serial position denoted by numbers above the model values.*

by quite a lot (Figure 5a). In the main model fit, it seems that the model produced an excess of protrusions in order to preserve the high accuracy level. This may partly be an artifact of the fitting; because prior-list intrusions are relatively rare, a model that produces few such intrusions will produce a poor goodness of fit, because the estimate of prior-list intrusions will be inaccurate due to the small number of events produced by the model. A model with lower accuracy will produce more prior-list intrusions, thereby also increasing the precision with which those intrusions are measured. In any case, Figures 1b and 5b show that the model, despite its constraints and lack of positional cues, can fit a variety of patterns of prior-list intrusions with a protrusion-dominance.

Exploratory empirical replications of protrusion dominance

The predominance of protrusions in immediate serial recall has been reported a small numbers of times (e.g., Conrad, 1960; Fischer-Baum & McCloskey, 2015; Henson, 1996; Osth & Dennis, 2015). Osth and Dennis (2015) noted that the low report rate may be partly due to the fact that prior-list intrusions are rare. To test for the presence or absence of the protrusion-dominant pattern demands a very large data set. One can only speculate about the number of null findings that may have not been reported simply because the authors may have suspected they were underpowered to test it. We tested the predominance of protrusions in two experiments run recently in our lab. The experiments were selected based on the data sets being available, but were designed with other hypotheses in mind, not to test for protrusions.

Methods

The first data set was Experiment 1 from (Liu & Caplan, 2020)[2]. Liu and Caplan (2020) used the data to test whether effects of temporal grouping were a function of output position or serial position, confirming the former. Prior-list intrusions were not reported in that manuscript. The experiment presented participants with lists of 9 consonants at an average rate of 1 letter per 850 ms. Recall was typed, and participants were instructed to use the SPACE key to skip letters they could not recall. Between-subjects factors were direction of recall (forward or backward serial recall) and uniform versus temporally grouped presentation. We analyzed the data from both forward conditions (N=42 and 41, respectively). The experiment was thus a closed-set stimulus pool (consonants, repeated across lists). List length was somewhat longer than typical paradigms that have reported protrusions (9 items, compared to typically 5 or 6; e.g., Henson, 1996; Osth and Dennis, 2015). No participants were excluded from the analyses. Because the experiment used a closed set, letters that were on the current list could not be unambiguously scored as belonging to the current or previous list. Thus, only items from the immediate-prior list that were not also on the current list were included in the prior-list intrusion analyses. For each participant, the number of prior-list intrusions at each output position in the current list was tallied, then normalized for each output position. Thus, the final values are the proportion of immediate prior-list intrusions at a given output position, as a function of prior-list serial position. At each output position, a participant was included in the analysis if they had one or more prior-list intrusions at that output position.

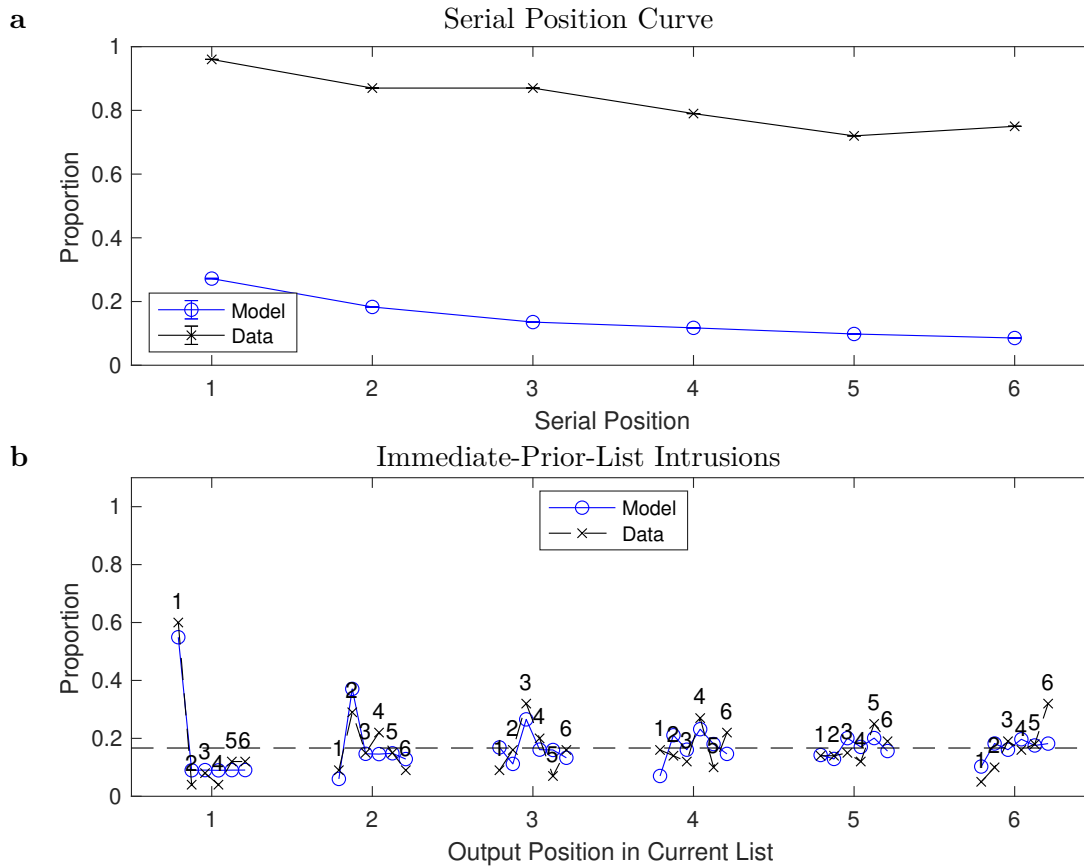[2]Data available from https://osf.io/evmct.

Figure 5

Simulated data using the best-fitting parameters that were fit to the Henson (1996) data, prior-list intrusions only, plotted as in Figure 1, alongside the behavioural data. Error bars plot 95% confidence intervals based standard error of the mean. **a**, *Probability of recall as a function of serial position (relative-order scoring).* **b**, *Probability of an immediate-prior-list intrusion at current-list ouput position (x axis) recalled at each (previous) serial position denoted by numbers above the model values. Parameter values were: $\sigma = 2.4243$, $\theta = 0.3238$, $\rho = 0.9923$ (exponent=3.48), $r = 11.99$, $b = 19.37$, $\eta = 0.50$.*

The second data set is not yet published (Shafaghat Ardebili, Liu and Caplan, in preparation)[3]. It was designed as a follow-up study to Liu and Caplan (2020) with words. The primary goals of this experiment were to test whether the output-position dependence generalized to word stimuli, open sets and to subjective chunking instructions. A secondary goal was to test whether subgrouped lists of words exhibit approximate all-or-none retrieval (Johnson, 1969, 1970). It used the same presentation rate, but the stimulus set was an open set, with no repetitions. List length was 6, very much in line with prior studies (Henson, 1996; Osth & Dennis, 2015) and recall was typed. In addition to a temporal grouping condition, this experiment had a condition in which participants were instructed to subdivide the list, but groups were not demarcated temporally. Only the forward-recall data are reported here. Due to relatively low rates of intrusions, all three forward-recall conditions are analyzed, collapsed together.

Results and discussion

The prior-list intrusions are plotted in Figure 6. We first took a broad-strokes look at the protrusion rates. Comparing all protrusion rates against that expected by chance (1/list length), and treating each output position of each participant as though it were an independent participant (the lack of independence indicates that this should be interpreted with caution), the t test was significant and well supported by the Bayes Factor (computed from the t values as suggested by Rouder et al., 2009) for each data set: $t(301) = 4.28$, $p < 0.0001$, $BF_{10} = 436.02$; $t(291) = 6.52$, $p < 0.0001$, $BF_{10} > 1000$; and $t(205) = 3.64$, $p = 0.0003$, $BF_{10} = 43.40$, respectively. However, protrusions were not dominant at all serial positions. Table 2 reports these tests conducted at each serial position, in turn. Robustly significant and supported by Bayes Factors, protrusions were dominant at serial position 1 in both Liu and Caplan data sets, and at position 2 for the ungrouped Liu and Caplan data. However, evidence against the predominance of protrusions (Bayes Factor favouring the null) was found at serial positions 3, 4, 5, 8 and 9 (Liu and Caplan, 2020, Forward, Ungrouped); 2 (Liu and Caplan, 2020, Forward, Grouped); and 2 and 5 (Shafaghat Ardebili, Liu and Caplan, in preparation). All other Bayes Factors were inconclusive. Nonetheless, statistical significance was also found at serial positions 4, 5 and 9 in the second data set. These statistical outcomes reinforce visual inspection. The Liu and Caplan grouped data show a nominal peak at the protrusion position at all positions apart from position 2, where the peak is instead at position 5, where many so-called interpositions (intrusions from one group to another, preserving within-group position) would arise from (Fischer-Baum et al., 2021; Henson, 1999; Henson, 1996). In the other two data sets, the nominal peaks are sometimes at the protrusion position and sometimes not, different than one would expect of a positional retrieval mechanism, or even the chaining-mediated mechanism we have proposed here.

In sum, the predominance of intrusions was nearly completely replicated for one condition of one data set, but not for the other condition, nor the other data set. This lack of robust replication suggests that protrusions, as a benchmark finding, should not inevitably fall out of a model, although they should be able to emerge under certain experimental conditions.

_____

[3]Data available at https://osf.io/z9sgy/?view__only=62b14f7b22c64aca8e33df29f88d54c3
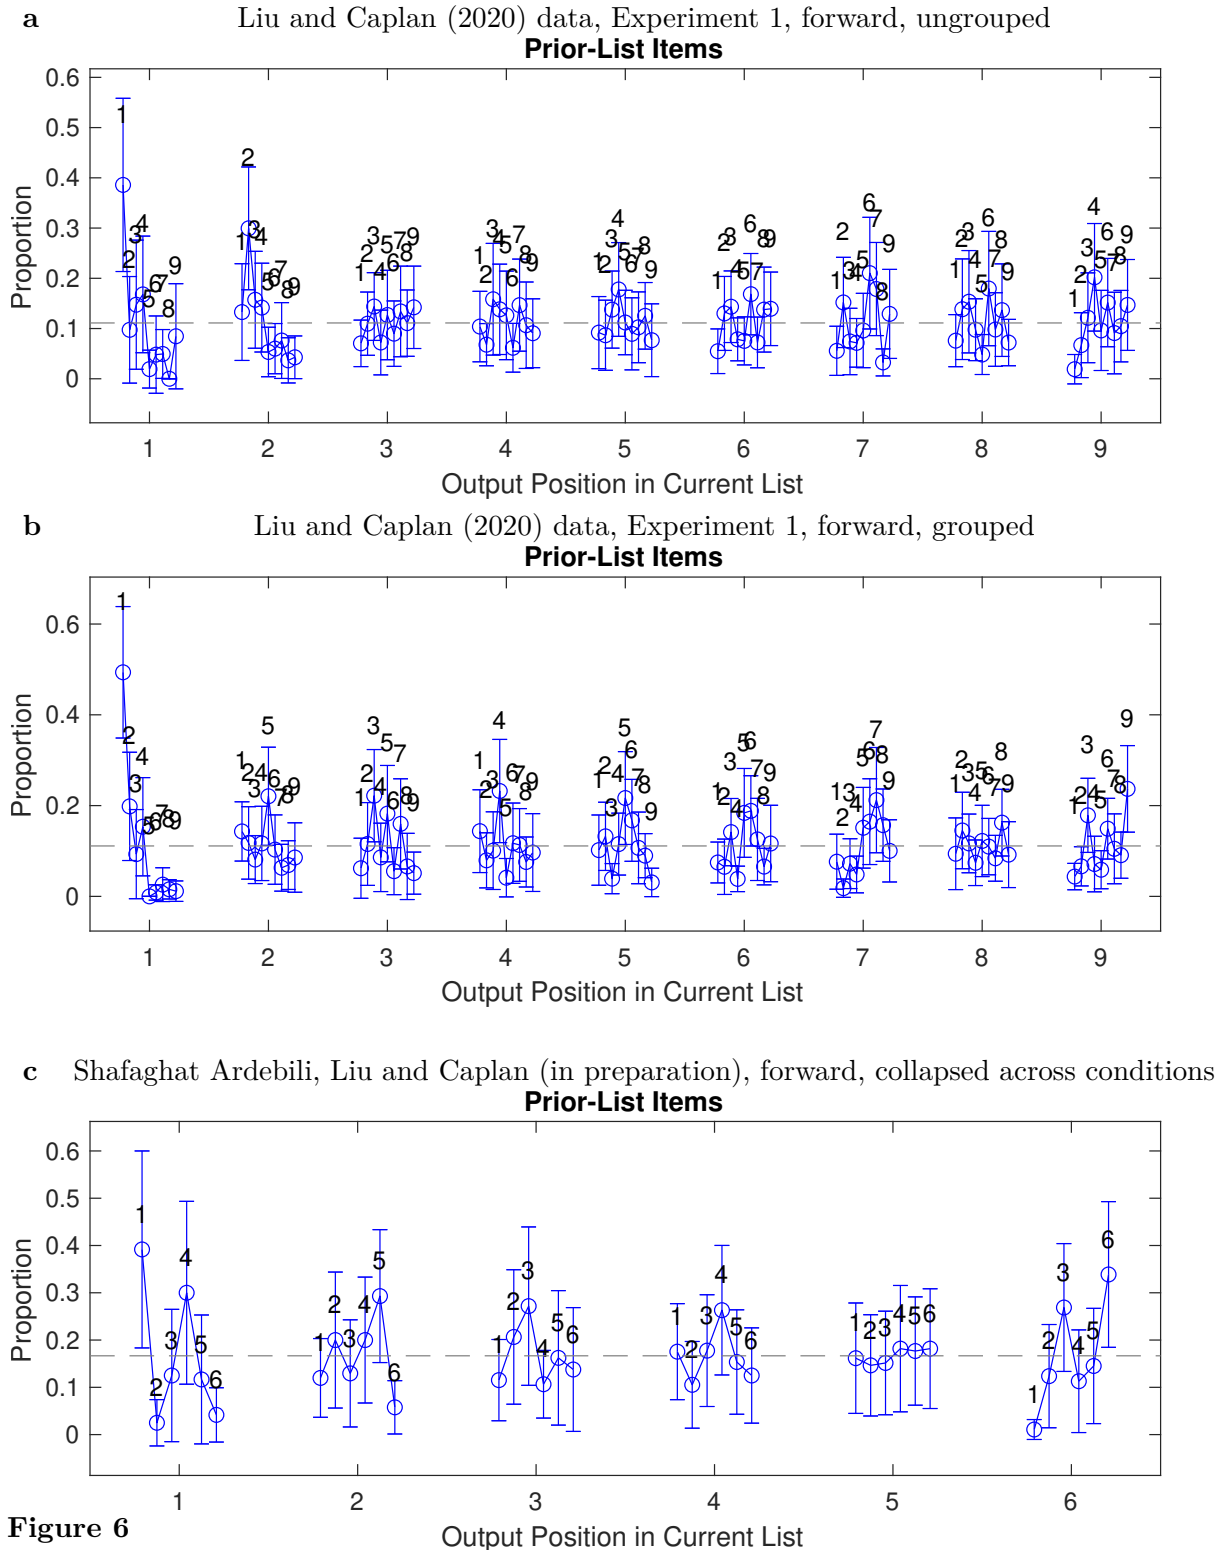
Figure 6

New empirical tests of the predominance of protrusions. a) Liu and Caplan (2020) data, Experiment 1, forward, ungrouped. b) Liu and Caplan (2020), Experiment 1, forward, grouped. c) Shafaghat Ardebili, Liu and Caplan (in preparation), forward, collapsed across ungrouped, temporally grouped and subjective-chunking instructions. Numbers over the data points designate the serial-position of the item on the immediate-prior list. Error bars plot 95% confidence intervals based on standard error of the mean. The dashed line denotes what would be expected if prior-list intrusions were equally probable from all prior serial positions (inverse of list length).

Liu and Caplan (2020), Forward, Ungrouped

SP	1	2	3	4	5	6	7	8	9
t	3.1239	3.0202	0.9502	0.5911	0.0269	1.3877	1.4280	0.5422	0.7762
df	25	29	38	33	32	37	33	32	34
p	0.0045	0.0052	0.3480	0.5585	0.9787	0.1735	0.1627	0.5915	0.4430
BF_{10}	9.3517	7.8304	0.2627	0.2161	0.1863	0.4223	0.4636	0.2135	0.2398

Liu and Caplan (2020), Forward, Grouped

SP	1	2	3	4	5	6	7	8	9
t	4.9934	0.2557	2.0063	2.1582	2.0892	1.9488	1.7159	1.4728	2.3991
df	27	28	30	32	37	35	27	31	36
p	< 0.0001	0.80	0.054	0.039	0.044	0.059	0.098	0.15	0.022
BF_{10}	759.67	0.20	1.11	1.42	1.22	0.97	0.73	0.50	2.18

Shafaghat Ardebili, Liu and Caplan (in preparation), Forward, all conditions

SP	1	2	3	4	5	6
t	1.90	0.57	1.92	1.61	0.66	2.03
df	21	29	34	43	37	36
p	0.071	0.57	0.063	0.12	0.51	0.050
BF_{10}	1.03	0.23	0.94	0.54	0.21	1.10

Table 2

At each serial position (SP), tests of whether the protrusion rate exceeds that expected if there were no dependence of prior-list intrusions on serial position (1/list length). Bayes Factors are computed from those t values as described in the main text.

Preservation of characteristics of previous chaining models

Our modifications to classic chaining models enabled the model to produce the protrusion-dominance effect, but it would be undesirable if this came with a cost of losing the ability of chaining models to fit other characteristics of serial recall. We checked two such characteristics. The first is the finding that within-list intrusions are more often from close than from distant serial positions, which Henson (1996) called the "locality constraint." Figure 1c plots within-list intrusions generated by the same run of the the model as in panels a and b. Without fitting directly to the within-list intrusion data, the model produced distance effects that hug the data quite closely. The obvious exception, again, is at the first output position. As with prior-list intrusions, because this simplified model only uses the start-signal to initiate recall, and accuracy at this position was quite high, error are rare, and when the model produces an item as an error, it comes from an arbitrary position. Although clearly a miss for the current model, a more fleshed-out chaining model would presumably have a way the model could skip position 1, for example, producing no response rather than an omission at position 1. In that case, a distance effect might appear at that position as well. Remote associations (Murdock, 1995; Solway et al., 2012), including from the start-signal, might be an alternative solution.

The second feature we tested the model on was the pattern of transitions from the first order-error committed in a list, to the next response, related to Henson's concept of "fill-in" errors (Henson, 1996). Henson's original argument was that if an item was recalled out of place, positional coding models would proceed to use the next position as a retrieval cue. Because of positional similarity, the skipped item would be very likely recalled, a so-called "fill-in" error. A chaining model would, instead, use the retrieved item as a cue and thereby proceed forward in the list, a so-called "in-fill" error. The high rate of fill-in errors was viewed as problematic for chaining models. It does seem to be problematic for chaining models that include forward-only associations (Logan, 2021; Osth & Hurlstone, in press). However, a chaining model that includes strong backward associations produces fill-in errors (Solway et al., 2012). Solway et al. (2012) introduced improvements to these transitional error analyses to make them more diagnostic of models. First, if the first order-error was item i, recalled at output position j, and the subsequently recalled item was from serial position k, they plotted probability of a transition as a function of lag of $j - k$, which we call position-aligned lag functions. Second, they plotted probability of a transition as a function of the lag of $k - i$, which we call response-aligned lag functions. Chaining models predict the response-aligned function should show steep distance effects, with transitions remaining quite close to the previously recalled item (as opposed to the previous output-position). Finally, we follow the advice of Solway et al. (2012) and normalize by availability of a particular lag. Thus, when a lag is not available as a response, either because it is beyond the limits of the list or because it has already been recalled (response suppression), the denominator of the ratio for that particular lag is not incremented. Data have largely shown contiguity effects for both lag functions, with asymmetries and relative tradeoffs of steepness of the functions varying across data sets (Caplan, 2015; Caplan et al., 2015; Farrell et al., 2013; Solway et al., 2012; Surprenant et al., 1999). Solway et al. (2012) found their chaining model produced both position- and response-aligned contiguity effects. Their implementation of a positional-coding model by Burgess and Hitch (2006) fit position-aligned effects well but predicted relatively flat response-aligned functions, deviating from the data. The hierarchical positional-coding model designed by Farrell (2012) was able to produce response-aligned distance effects (Farrell et al., 2013) by recalling entire chunks correctly out of sequence.

Figure 7 plots both the position- and response-aligned transition functions, computed from a fresh run of the model with the best-fitting parameter values. Apart from the peak at +4 (which is beyond the range these functions are typically plotted), the position-aligned transition function (Figure 7a) shows a contiguity effect and a moderate preference for the model to proceed later (lag=–1) than earlier (lag=+1), similar to the chaining model of Solway et al. (2012). This parameter set produced a very pronounced contiguity effect when response-aligned (Figure 7b), but differing from Solway et al.'s forward-transition advantage, this model produced a nearly symmetric contiguity effect with a slight tendency to backtrack, or fill-in.

Unlike the asymmetric associations in the Logan (2021) model, the near-symmetry of associations in our best-fitting model causes the model to backtrack quite frequently. This makes sense if one considers that if item i was skipped, the associative strength from item $i-1$ to item i was probably weak, but the association from item i to $i+1$ has a chance of being strong. Symmetry of associative strengths means that the backward association
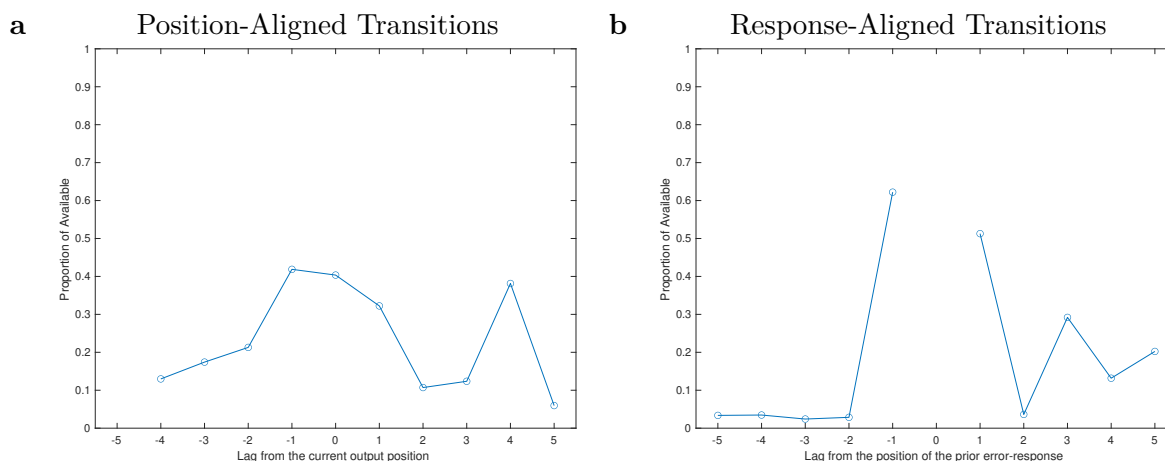
Figure 7

Transition functions produced by the chaining model using the best-fitting parameter values. Probability, following the first order-error, of a transition to a given serial position relative to the current output position (a) or the serial position of the just-recalled item (b). Proportions are relative to the number of opportunities to make the transition (availability), taking into account the limits of the list and response suppression. Following Solway et al. (2012), in (a), negative lags correspond to anticipations, recalling a later item too early, and positive lags correspond to postponements, recalling an item too late. In (b), negative lags correspond to recalling an earlier item than the out-of-order preceding response and positive lags correspond to recalling a later item. If the first order-error were an immediate transposition, recalling item $i + 1$ followed by item i (e.g., recalling ABCED), the position-aligned lag (for recall of item i) would be +1 (postponement) but the response-aligned lag would be –1 (transition to the immediately preceding item in the list).

from item $i + 1$ can retrieve item i. Note that the model has not, thus far, produced ratios of 2:1 or 3:1 fill-in:in-fill ratios as has often been observed, so this is not a clear-cut success of the model; however, it does not lose this functionality of previous chaining models like that of Solway et al. (2012). But note that the current model (like most models of serial recall) possesses very little metacognitive function. We left the transition choice to the passive retrieval process. We know that participants have some ability to discriminate the directions of associations (Kato & Caplan, 2017; Rehani & Caplan, 2011). It would seem plausible that upon skipping an item, participants may have a subjective preference for filling in via the backward association to maximize at least the number of items correctly recalled, even if slightly out of order. Having established that backtracking transitions occur often, both anticipations and postponements may be simultaneously available, and a metacognitive process might lead a model in some conditions (such as short list lengths) to prefer to respond with two well-retrieved candidate items in their relative-order of presentation. This may be an alternative approach to explaining what modulates the asymmetry of transitions without relying on entire chunks being displaced (Farrell et al., 2013), and could be implemented in any model, but remains to be tested.

In sum, contiguity effects in standard within-list intrusion distance functions, as

well as strong contiguity effects in transitions following the first order-error, were largely retained from earlier chaining models, suggesting that the modifications introduced here may not bring with them obvious new costs.

General Discussion

With a few additional assumptions, all of which involved relaxing assumptions or slightly expanding assumptions that have been designed into past associative chaining models, an associative chaining model was able to produce a substantial proportion of prior-list intrusions as protrusions. At the very least, this suggests that the existence of protrusions should no longer be considered damning evidence against associative chaining, adding to arguments by Dennis (2009) with respect to protrusions, as well as others who have argued that associative chaining models are not ruled out and positional-coding models are not as well supported as previously thought (e.g., Caplan, 2015; Fischer-Baum & McCloskey, 2015; Logan, 2021; Osth & Hurlstone, in press; Serra & Nairne, 2000; Solway et al., 2012).

Importantly, protrusions were created without any explicit encoding of position or order information into inter-item associations. One could argue that the start-signal functions like a position-code, similar to activation in the Primacy Model (Page & Norris, 1998) or the start-marker in the Start-End Model (Henson, 1998). However, our implementation is fully within the spirit of associative chaining for two reasons. First, start-signals, or some analogue, have always been necessary for chaining models to get off the ground, as one has to explain how the subject can retrieve the first list-item. Second, in our implementation, we treat the start-signal exactly the same as any other item; thus, it participates in inter-item associations, and those associations are still approximately symmetric. The only difference is that we assume the subject starts by encoding an association between the start-signal and the first list-item presented, and initiates recall by cueing with the start-signal. In this sense, the subject is not cueing with position, but with a special item that initiates the list. Unlike the Primacy and Start-End Models, the start-signal is not associated explicitly with any other list items (only indirectly, through the chain, or potentially accidentally, if noise were added to the model). This means that the subject does not have "random access," either by probing with a position to retrieve an item, or by probing with an item to retrieve its position. Such item–position and position–item probes would have to be answered by inference— by sequentially reading out the chain. In this sense, the basic distinction between associative-chaining and positional-coding still applies, with our model clearly belonging to the former class.

That said, it is quite plausible that a more accurate model of serial-recall behaviour will need to include a hybrid or a mixture of positional-coding and associative-chaining mechanisms (e.g., Burgess & Hitch, 1992; Caplan, 2015; Fischer-Baum & McCloskey, 2015; Jensen & Lisman, 2005; Lisman & Idiart, 1995; Logan & Cox, 2021; Osth & Hurlstone, in press). The important lesson is thus not that associative chaining models must be correct and positional-coding incorrect, but that associative-chaining should still generally be considered as a possible account of serial-recall data.

Boundary conditions for protrusion-dominance of prior-list intrusions. It is important to know whether the protrusion effect is ubiquitous or variable, perhaps sensitive to particular experimental conditions (points made by Osth and Dennis, 2015 and Osth and Hurlstone, in press). Reports of protrusions might suffer from a reporting

bias. That is, experiments designed to test for protrusions might be abandoned when if they fail to produce the expected result (a cause of the "file-drawer problem," e.g., Rotton et al., 1995) and experiments designed for other purposes might have either replicated or failed to replicate the protrusion effect but have not been reported if the authors did not see any obvious reason to do so. Our data sets were selected based on recency and the fact that we already were actively analyzing the data (and without having computed prior-list intrusion rates until needed for this manuscript), thus selected for convenience. Granted, these two experiments may still not be representative, but they do suggest that the ubiquity of the dominance of protrusions should at least not yet be assumed. Moreover, the presence of protrusions in the one condition also speaks against the opposite radical position: protrusion-dominance is also not exceedingly rare, and if there is publication bias, it is not absolute.

To our knowledge, the earliest report of protrusions is from Melton and von Lackum (1941), in a related paradigm, serial anticipation. However, they did not separate exact protrusions from ±1 serial position. They also found no such intrusions when the two lists were dissimilar, which suggested that serial-position as a cue was used primarily when inter-list similarity was high. Consistent with this, Henson (1996) found a predominance of protrusions when lists were drawn from a closed set (two pools of seven words that alternated from one list to the next). However, Henson (1996) also included a condition where list length varied from one list to the next (5, 6 or 7 words), still with the closed-set design, and found this condition to produce weaker evidence of protrusions. He suggested that when list length is less predictable, participants rely on position-cues less, identifying another possible boundary condition. Contradicting Melton and von Lackum (1941), Osth and Dennis (2015) did find protrusion-dominance with an open set. However, with very similar methods, our open-set word-list experiment did not clearly replicate this. Also, differing from Melton and von Lackum (1941), our closed-set (consonant-list) experiment did not result in a clear predominance of protrusions in the ungrouped condition, but fairly good support for predominance of protrusions for temporally grouped consonant lists. At the very least, the conditions in which protrusions are found are unclear, and the protrusion result appears not to be particularly robust. When present, protrusion-dominance might indicate that participants are making use of positional or order-cueing during serial recall. Alternatively, protrusions may result from covert retrieval of the prior list via chaining, as we have shown.

Another factor that might determine the presence or absence of protrusions is the presence of positional information during study or during recall. In particular, the data we fit, reported by Henson (1996), was written recall. Participants were given sheets with lines denoting position within the list. This might have biased participants to use position more as a cue. Of the new data sets we analyzed here, the one that produced some evidence of protrusion dominance (Liu & Caplan, 2020) used sequential presentation during study, but happened to use a recall procedure whereby each item (letter) remained on the screen, similar to Experiments 1 and 3 of Fischer-Baum and McCloskey (2015). This amount of position-cueing may have influenced the participant's strategy, either to use position as a retrieval cue or to an operational mode of our chaining model that increased the prevalence of protrusions. The other data set (Shafaghat Ardebili, Liu and Caplan, in preparation), which had words as stimuli had both sequential presentation and sequential recall; each word

disappeared before the subsequent word could be recalled. The complete absence of spatial information about position may be one reason why protrusions did not dominate in that data set. Complicating this, Osth and Dennis (2015) also used both sequential presentation and sequential (also typed) recall of words and did find high rates of protrusions, as did the vocal recall procedure of Experiment 2 of Fischer-Baum and McCloskey (2015), which also produced frequent protrusions.

To distinguish positional/ordinal cueing from chaining mechanisms of protrusions, one would require a more complex experimental design, in addition to replicating protrusion dominance. For example, consider a study in which, following study and serial recall of a list, sequential probes (Woodward & Murdock, 1968) were administered. In a sequential probe trial, the participant is presented with a list item and asked either for the item that followed or the item that preceded the probe item, for example, given item C, recall the item that followed it (D). A model that only had positional or ordinal cueing available (Caplan, 2005) would need to use the probe item to retrieve the associated positional code (use C to retrieve "3"), then shift along the position code and use the new position ("4") to retrieve the target item (here, D). The latter operation would be expected to produce protrusions, for the same reason as they are expected during serial recall. However, a chaining model presumes that items are associated directly to one another, so it would simply probe memory with the cue item in the desired direction and attempt to retrieve the target. Because sequential probes avoid probing the list with the start signal, and items are not assumed to be associated across lists, there is no reason to expect more protrusions than prior-list intrusions from other serial positions. Thus, the positional coding account of protrusions predicts that whenever protrusions are found during serial recall, sequential probes should also produce protrusions, whereas the associative chaining account of protrusions would make no such prediction.

Other models

Our goal here was to test a proof of principle that a model designed in the spirit of traditional associative chaining models could produce protrusion-dominance. This does not imply that our mechanism is the best or sole account or protrusions. Positional coding models quite naturally produce protrusions, as demonstrated as early as Henson (1996), with his Start-End Model. Although he did not report whether or not the model did so by retrieving entire lists or large chunks of lists, it seems plausible that such a model, when constrained also to fit high accuracy, would produce only one or two prior-list intrusions to a given list, satisfying this criterion as well. As we already noted, Henson (1996) commented that there must be an additional source of prior-list intrusions which might explain his failure to find protrusion dominance when list length varied (as well as the empirical findings we report here). Indeed, a model that can only retrieve via positional or ordinal cues might have trouble avoiding producing protrusion-dominance. In our casual investigations of the chaining model by hand-tuning parameters, there seems to be a large portion of parameter space that does not produce prior-list intrusions that are dominated by protrusions, but a strict positional-coding model might need additional components, whether inter-item associations or not, to expand the reach of the model across empirical data.

Two other models that are described as chaining models have addressed protrusions. First, as described in the introduction, Dennis (2009) designed a model that retrieves a list

entirely in one step, and produces protrusions as a side-effect of a similar similarity structure emerging over successive lists. Second, Logan (2021) designed a model in which memory of a list is built up from associations between items and weighted sums of previous items. As he described it, the model has some common characteristics with chaining models, and is arguably more similar to chaining than to positional coding models. Osth and Hurlstone (in press) were able to get this model to produce protrusions, but were disappointed that it seemed to do so by retrieving entire lists. It should be noted that they did not conduct a full search of the parameter space. However, some ways in which Logan's model diverges from conventional chaining models might, in fact, limit the ability of the model to fit some findings, although they may enable the model to fit others (Logan, 2021). Specifically, the model includes both remote associations and only forward-directed associations. Several limitations of the model identified by Osth and Hurlstone (in press) appeared to derive from these two properties of the model.

In sum, ours is not the only empirically supported model account of the dominance of protrusions. The mechanisms just summarized might explain some instances of protrusions, or be part of a mixture of mechanisms along with ours. To reiterate, our main argument is that conventional chaining mechanisms are able to produce a realistic pattern of prior-list intrusions and should not be ruled out.

Limitations of the current model

For logistical reasons, our chaining model was simplified, which apparently led to some unusual properties of the pattern of prior-list intrusions. Most notably, approximate, but imperfect symmetry was implemented by encoding a separate forward and backward association, fixing them to be perfectly correlated and ensuring that the forward association was stronger than the backward association. In fact, previous data suggests that associations are intrinsically symmetrical but include some information about order (Kato & Caplan, 2017). Because no currently developed model of association can satisfy all these constraints, the shortcuts were necessary, but also kept the model simple, and run-time short enough to fit to data. For correct recall of the current list, response suppression prevents the model from backtracking (unless an item is skipped; see so-called "fill-in" errors, Henson, 1996). However, the prior list that is retrieved simultaneously in the background gains no benefit from response suppression. Using the notation from the introduction, with upper-case letters for the current list and lower-case letters for the previous list, probe with the start-signal retrieve $A + a$ (weighted by encoding strengths and forgetting, omitted here for clarity). $A + a$ is then used as a cue to retrieve $B + b$ (and the start-signal, which is ignored). But then $B + b$ as a cue retrieves $A + C + a + c$. A is ruled out due to response suppression, leaving C as a response candidate but both c (the protrusion) and a (not a protrusion) as response candidates when a prior-list intrusion might be committed. This iterates, producing waves of forward and backward retrieval, with the result that weird peaks (such as prior item 2 intruding at output position 4; Figure 1b). Part of the reason this can occur is that the cued-recall steps are treated as discrete operations, and the model skips output positions perfectly, whereas it is plausible that real participants are not so vigilant about marking whenever they skip a position. Because associative strengths were initialized to zero (i.e., no prior knowledge), the only cause of prior-list intrusions is encoding of prior experimental lists, also an unrealistic assumption. Given that the current model presents a

new proof of principle that a chaining model is able to produce protrusions, and indeed, too many at early output positions, we suggest that adding realism through additional noise and making the successive cued-recall mechanism more fallible, the behaviour of the model would smooth out and be able to fit the data more tightly. Positional coding models would also need to be held to the same standard and would be given the same latitude.

Generalization of the idea

Although our focus was single-mindedly on the protrusion finding, the design of our model expresses a potentially farther-reaching concept. That is, as has been implied by prior chaining modellers, an associative model could gain functionality if one assumes the existence of specialized items with specific meaning or functions. The start-signal is thus only one of a potentially larger set of such items. The way in which verbal memory researchers think about experimental design leads us to presume that the concept of the start of the list is different than the elements of a list, such as digits, letters or words. However, there is no reason to assume that the way in which the concept of the start of the list operates fundamentally any differently than the elements of the list— apart from the participant's meta-knowledge that the start-signal is not a valid response option. In fact, it is possible that participants even use the word, "start," itself, in this way (as long as the word "start" is not part of the stimulus set). An obvious extension of the idea would be a special item standing in for the end of the list (an "end-signal," also with precedents in chaining models; Lewandowsky and Murdock, 1989; Shiffrin and Cook, 1978) and suggested by experimental data errors revealing aligned to the end of the list Fischer-Baum and McCloskey, 2015; Henson, 1999. Logan and Cox (2021) expressed a related idea with a different approach. They started with the chaining-like model of Logan (2021), which associates each item to a context, but that context is, in turn, a weighted sum of recent contexts, ultimately composed of the items, themselves. Logan and Cox (2021) showed how position codes resembling Henson's start- and end-markers (Henson, 1996) could be derived from the model, providing much of the functionality of positional-coding models without explicit representations of position, also reminiscent of Dennis (2009).

Another extension would be the concept of the start and/or end of a chunk (as a possible account of temporal grouping, Ryan, 1969, and subjective chunking, Johnson, 1969, manipulations), and the notion of an item being an exception within the list or a repetition, etc. Moreover, the way in which positional and ordinal coding models have been described, each list item is stored with positional or ordinal value or code, A–1, B–2, C–3, etc. These could arguably be stored as associations of the same type as inter-item associations. In fact, deliberately applied mnemonic strategies like the peg list method at least at face value resemble this. A peg list is a set of words or images that are pre-learned, and map directly onto a consecutive range of numbers (e.g., Bower, 1970; Roediger, 1980; Sahadevan et al., 2021; Wood, 1967). With the rhyming peg list, for example (1–BUN, 2–SHOE, 3–TREE, etc.), the peg words were deliberately chosen to be highly imageable words. Thus, learning the association between the first list-word— say, TABLE, and the peg, BUN, may be fundamentally no different than the inter-item association one would learn if TABLE and BUN were both list words and occurred immediately adjacently to one another. In the target data used to test chaining and positional-coding models, researchers try to avoid any such deliberate strategies. However, this line of thought shows that positional coding

and associative chaining may not be such distant cousins; they both build memory upon pairwise associations. This leads to an interesting compromise: One could start with a chaining model. In certain experimental conditions, specialized items that convey order might be additionally associated to list items (see the formulation by Logan and Cox, 2021, that is similar in spirit). The Start-End Model (Henson, 1998) could be integrated into a chaining model in this way, by representing the start and end markers in the same way as list items, but with specialized significance to the model. Unlike our start-signal, such a model could gain some of the desirable functionality of the Start-End Model by assuming that item–start associations decay exponentially with serial position and item–end associations decay from the end of the list. Other representations of order also seem amendable to coding as specialized items or sets of items Brown et al., 2007; Brown et al., 2000; Burgess and Hitch, 1999. And in more extreme conditions, peg lists could be integrated by extending a chaining model as just described, where every peg is an item for which the model has some meta-knowledge (e.g., position value). Thus, a mixture model built upon the basic unit of associations between items, could potentially span a large range of conditions and explain a large range of empirical findings. Unfortunately, such a model would have a large number of degrees of freedom and implies lots of heterogeneity across tasks and stimulus domains. It would not immediately be testable without first seeking ways to severely constrain it.

Conclusion

The predominance of protrusions has been used to argue against associative chaining accounts of immediate serial recall and in favour of position-cueing accounts. Newly analyzed data suggest that the predominance of protrusions is less robust than previously thought, and if anything, is not found at all serial positions. This makes the problem that chaining models need to solve far more circumscribed. On the other hand, positional coding models face the challenge of explaining why, if position is a dominant retrieval cue, do protrusions sometimes *not* predominate? The updated chaining model was able to produce the classic pattern of prior-list intrusions, with protrusions dominating such errors, by linking a start-signal vector to the first item of each list, resulting in covert retrieval of the prior list during recall of the most recent list. In short, the predominance of protrusions in some serial-recall data sets does not rule out associative chaining as an account of immediate serial recall, even for those data sets. The lack of predominance of protrusions in other data sets poses a challenge for positional-cueing models but is compatible with the updated chaining model presented here. With modifications that stick quite close to published chaining models, the associative chaining model developed here stands as a reasonable account of immediate serial recall, and ultimately may need to be combined in a mixture model with positional cueing to account for a more complete set of benchmark findings.

References

Bower, G. H. (1970). Organizational factors in memory. *Cognitive Psychology*, *1*, 18–46.

Brown, G. D. A., Neath, I., & Chater, N. (2007). A temporal ratio model of memory. *Psychological Review*, *114*(3), 539–576.

Brown, G. D. A., Preece, T., & Hulme, C. (2000). Oscillator-based memory for serial order. *Psychological Review*, *107*(1), 127–181.

Burgess, N., & Hitch, G. J. (1992). Toward a network model of the articulatory loop. *Journal of Memory and Language*, *31*, 429–460.

Burgess, N., & Hitch, G. J. (1999). Memory for serial order: A network model of the phonological loop and its timing. *Psychological Review*, *106*(3), 551–581.

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

Caplan, J. B. (2005). Associative isolation: Unifying associative and order paradigms. *Journal of Mathematical Psychology*, *49*(5), 383–402.

Caplan, J. B. (2015). Order-memory and association-memory. *Canadian Journal of Experimental Psychology*, *69*(3), 221–232.

Caplan, J. B., Glaholt, M., & McIntosh, A. R. (2006). Linking associative and list memory: Pairs versus triples. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *32*(6), 1244–1265.

Caplan, J. B., Madan, C. R., & Bedwell, D. J. (2015). Item-properties may influence item–item associations in serial recall. *Psychonomic Bulletin & Review*, *22*(2), 483–489.

Conrad, R. (1960). Serial order intrusions in immediate memory. *British Journal of Psychology*, *51*(1), 45–48.

Conrad, R. (1965). Order error in immediate recall of sequences. *Journal of Verbal Learning and Verbal Behavior*, *4*, 161–169.

Dennis, S. (2009). Can a chaining model account for serial recall? In N. Taatgen & H. van Rijn (Eds.), *Proceedings of the thirty-first annual meeting of the Cognitive Science Society* (pp. 2813–2818).

Duncan, M., & Lewandowsky, S. (2005). The time course of response suppression: No evidence for a gradual release from inhibition. *Memory*, *13*(3/4), 236–246.

Ebbinghaus, H. (1885/1913). *Memory: A contribution to experimental psychology*. New York: Teachers College, Columbia University.

Farrell, S. (2012). Temporal clustering and sequencing in short-term memory and episodic memory. *Psychological Review*, *119*(2), 223–271.

Farrell, S., Hurlstone, M. J., & Lewandowsky, S. (2013). Sequential dependencies in recall of sequences: Filling in the blanks. *Memory & Cognition*, *41*(6), 938–952.

Fischer-Baum, S., & McCloskey, M. (2015). Representation of item position in immediate serial recall: Evidence from intrusion errors. *Journal of Experimental Psychology: Learning, Memory, and Cognition*.

Fischer-Baum, S., Warker, J. A., & Holloway, C. (2021). Learning phonotactic-like regularities in immediate serial recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *47*(1), 129–146.

Hartley, T., Hurlstone, M. J., & Hitch, G. J. (2016). Effects of rhythm on memory for spoken sequences: A model and tests of its stimulus-driven mechanism. *Cognitive Psychology*, *87*, 135–178.

Henson, R. N. A. (1998). Short-term memory for serial order: The Start-End Model. *Cognitive Psychology*, *36*(2), 73–137.

Henson, R. N. A. (1999). Positional information in short-term memory: Relative or absolute? *Memory & Cognition*, *27*(5), 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*(1), 80–115.

Henson, R. N. A. (1996). *Short-term memory for serial order* (Doctoral dissertation). University of Cambridge. Start-End Model, positional coding models, order coding models, ordinal coding, associative chaining models, immediate serial recall, protrusions, prior-list intrusions, alternating lists, chunking grouping.

Hurlstone, M. J., Hitch, G. J., & Baddeley, A. D. (2014). Memory for serial order across domains: An overview of the literature and directions for future research. *Psychological Bulletin*, *140*(2), 339–373.

Jensen, O., & Lisman, J. E. (2005). Hippocampal sequence-encoding driven by a cortical multi-item working memory buffer. *Trends in Neurosciences*, *28*(2), 67–72.

Johnson, N. F. (1969). Chunking: Associative chaining versus coding. *Journal of Verbal Learning and Verbal Behavior*, *8*(6), 725–731.

Johnson, N. F. (1970). The role of chunking and organization in the process of recall. In G. H. Bower (Ed.), *The psychology of learning and motivation advances in research and theory* (pp. 172–247). Academic Press.

Kahana, M. J. (2002). Associative symmetry and memory theory. *Memory & Cognition*, *30*(6), 823–840.

Kahana, M. J. (2012). *Foundations of human memory.* Oxford University Press.

Kahana, M. J., & Caplan, J. B. (2002). Associative asymmetry in probed recall of serial lists. *Memory & Cognition*, *30*(6), 841–849.

Kahana, M. J., Mollison, M. V., & Addis, K. M. (2010). Positional cues in serial learning: The spin-list technique. *Memory & Cognition*, *38*(1), 92–101.

Kato, K., & Caplan, J. B. (2017). Order of items within associations. *Journal of Memory and Language*, *97*, 81–102.

Ladd, G. T., & Woodworth, R. S. (Eds.). (1911). *Elements of physiological psychology: A treatise of the activities and nature of the mind from the physical and experimental point of view.* Charles Scribner's Sons.

Lewandowsky, S., & Farrell, S. (2000). A redintegration account of the effects of speech rate, lexicality, and word frequency in immediate serial recall. *Psychological Research*, *63*, 163–173.

Lewandowsky, S., & Farrell, S. (2008). Phonological similarity in serial recall: Constraints on theories of memory. *Journal of Memory and Language*, *58*, 429–448.

Lewandowsky, S., & Murdock, B. B. (1989). Memory for serial order. *Psychological Review*, *96*(1), 25–57.

Lindsey, D. R. B., & Logan, G. D. (2019). Item-to-item associations in typing: Evidence from spin list sequence learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *45*(3), 397–416.

Lisman, J. E., & Idiart, M. A. P. (1995). Storage of 7 ± 2 short-term memories in oscillatory cycles. *Science*, *267*(5203), 1512–1515.

Liu, Y. S., & Caplan, J. B. (2020). Temporal grouping and direction of serial recall. *Memory & Cognition*, *48*(7), 1295–1315.

Logan, G. D. (2021). Serial order in perception, memory, and action. *Psychological Review*, *128*(1), 1–44.

Logan, G. D., & Cox, G. E. (2021). Serial memory: Putting chains and position codes in context. *Psychological Review*.

Melton, A. W., & von Lackum, W. J. (1941). Retroactive and proactive inhibition in retention: Evidence for a two-factor theory of retroactive inhibition. *American Journal of Psychology*, *54*(2), 157–173.

Murdock, B. B. (1974). *Human memory: Theory and data*. John Wiley & Sons.

Murdock, B. B. (1993). TODAM2: A model for the storage and retrieval of item, associative, and serial-order information. *Psychological Review*, *100*(2), 183–203.

Murdock, B. B. (1995). Developing TODAM: three models for serial-order information. *Memory & Cognition*, *23*(5), 631–645.

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

Osth, A. F., & Dennis, S. (2015). Prior-list intrusions in serial recall are positional. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *41*(6), 1893–1901.

Osth, A. F., & Hurlstone, M. J. (in press). Do item-dependent context representations underlie serial order in cognition? commentary on Logan (2021). *Psychological Review*.

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

Rehani, M., & Caplan, J. B. (2011). Interference and the representation of order within associations. *Quarterly Journal of Experimental Psychology*, *64*(7), 1409–1429.

Roediger, H. L. I. (1980). The effectiveness of four mnemonics on ordering recall. *Journal of Experimental Psychology: Human Learning and Memory*, *6*(5), 558–567.

Rotton, J., Foos, P. W., Van Meek, L., & Levitt, M. (1995). Publication practices and the file drawer problem: A survey of published authors. *Journal of Social Behavior and Personality*, *10*(1), 1–13.

Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian *t* tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review*, *16*(2), 225–237.

Ryan, J. (1969). Grouping and short-term memory: Different means and patterns of grouping. *Quarterly Journal of Experimental Psychology*, *21*(2), 137–147.

Sahadevan, S. S., Chen, Y. Y., & B., C. J. (2021). Imagery-based strategies for memory for associations. *Memory*, *29*(10), 1275–1295.

Serra, M., & Nairne, J. S. (2000). Part-set cueing of order information: Implications for associative theories of serial order memory. *Memory & Cognition*, *28*(5), 847–855.

Shiffrin, R. M., & Cook, J. R. (1978). Short-term forgetting of item and order information. *Journal of Verbal Learning and Verbal Behavior*, *17*, 189–218.

Solway, A., Murdock, B. B., & Kahana, M. J. (2012). Positional and temporal clustering in serial order memory. *Memory & Cognition*, *40*(2), 177–190.

Surprenant, A. M., Kelley, M. R., Farley, L. A., & Neath, I. (1999). Fill-in and infill errors in order memory. *Memory*, *13*(3/4), 267–273.

Wood, G. (1967). Mnemonic systems in recall. *Journal of Educational Psychology*, *58*(6, Part 2), 1–27.

Woodward, A. E., & Murdock, B. B. (1968). Positional and sequential probes in serial learning. *Canadian Journal of Psychology*, *22*(2), 131–138.