

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

An Account of Associative Learning in Memory Recall

Permalink

<https://escholarship.org/uc/item/40b5v87q>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 37(0)

Authors

Thomson, Robert

Pyke, Aryn A

Trafton, J Gregory

et al.

Publication Date

2015

Peer reviewed

An Account of Associative Learning in Memory Recall

Robert Thomson (rob.thomson@knexusresearch.com)^{1,3}

¹Knexus Research Corporation, National Harbor, MD 20745 USA

Aryn A Pyke² (apyke@andrew.cmu.edu)

²Carnegie Mellon University, Pittsburgh, PA 15213 USA

J. Gregory Trafton³ (greg.trafton@nrl.navy.mil)

Laura M. Hiatt³ (laura.hiatt@nrl.navy.mil)

³Naval Research Laboratory, Washington, DC 20375 USA

Abstract

Associative learning is an important part of human cognition, and is thought to play key role in list learning. We present here an account of associative learning that learns asymmetric item-to-item associations, strengthening or weakening associations over time with repeated exposures. This account, combined with an existing account of activation strengthening and decay, predicts the complicated results of a multi-trial free and serial recall task, including asymmetric contiguity effects that strengthen over time (Klein, Addis, & Kahana, 2005).

Keywords: associative learning; priming; cognitive models; list memory

Introduction

Associative learning is an essential component of human cognition, thought to be part of many mental phenomena such as classical conditioning (Rescorla & Wagner, 1972), expectation-driven learning (Lukes, Thompson, & Werbos, 1990), similarity judgments (Hiatt & Trafton, 2013), and managing sequential tasks (Hiatt & Trafton, 2015). Despite its ubiquity, it is hard to model directly due to its entangled ties to other aspects of cognition (e.g., memory decay).

List learning is one task in which associative learning is increasingly thought to play a role and, because it involves fairly simple tasks, can be helpful in isolating and understanding any underlying associative mechanisms that may be at play (Howard & Kahana, 1999; Kahana, 1996). These tasks typically involve being shown a list of simple words or numbers, and being asked to recall them as accurately as possible.

One recent experiment studied list learning under both free recall (recalling list items in any order), and serial recall (recalling list items in the same order as they were presented), including an examination of how recall patterns change over several presentations of the list (i.e., multi-presentation recall) (Klein et al., 2005). In addition to serial position (SP), which shows each list item's recall accuracy, the study also considers conditional response probabilities as a function of lag (CRPs), which measures the distribution of successive recalls as a function of item distance from the current item. These two measures help to distinguish between effects arising from primacy and recency of items, and effects arising from the close temporal proximity of items to one another in the list (e.g., *contiguity effects*). The detailed results of this

study, which show how these measures change over multiple list presentations, present a challenge for other theories of memory (e.g., Henson, 1998; Brown, Neath, & Chater, 2007; Polyn, Norman, & Kahana, 2009), which generally match some, but not all, of the data.

We present here a theory of memory recall that heavily emphasizes the role of associative learning. This theory stems from the ACT-R architecture, which has been shown to perform very well on more limited list recall tasks in the past (Anderson, Bothell, Lebiere & Matessa, 1998). Since ACT-R is a general theory of cognition, and is not limited to memory, our use of ACT-R also connects this work with a plethora of literature across many different domains (Thomson, Lebiere, Anderson, & Staszewski, 2014; Hiatt & Trafton, 2013; Pyke, West, & Lefevre, 2007). While very promising, ACT-R's account of associative learning, however, is insufficient to capture data in list memory; specifically, it does not strengthen associations with repeated exposures in a manner that effectively accumulates over time, making it difficult for this account to predict the multi-presentation recall data we consider here.

In this paper, then, we heavily expand and improve the notion of associative learning in ACT-R. While we still keep many of its basic features – namely, spreading activation over asymmetric, item-to-item associations -- we developed a mechanism for associations to be learned, and strengthened, as new or repeated items are presented over time. We then combine activation via associative learning with ACT-R's second, and well supported, source of learning, activation strengthening (Anderson et al., 1998; Schneider & Anderson, 2011), which favors items that were recently or frequently in memory. Activation strengthening then serves to predict the exhibited primacy and recency effects, while associative learning predicts the shown contiguity effects. Overall, our account provides a good account for the data from Klein et al. (2005), showing primacy, recency, and asymmetric contiguity effects that strengthen over time.

The contributions of this paper are thus two-fold: a new, richer account of associative learning; and an overall theory of memory recall that combines this with an existing account of activation strengthening and decay. In the next section, we describe the list memory task we model in more detail. Then, we relate these results to other theories of list memory and

memory recall, and qualitatively distinguish our approach from these other theories. Then, we discuss our theory in more detail, present results, and end with a discussion of the implications of our theory.

Previous Experimental Results: Multiple-Presentation List Recall

To evaluate our theory of memory recall, we modeled the multi-presentation list recall task from Klein et al. (2005). A trial consisted of 5 separate presentations of a list of words. Each list consisted of 19 non-repeating words that were presented verbally, with a word presented every 1500ms. The words did not rhyme, and appear with similar frequencies in the English language. Each time the full list was presented, a tone and visual instructions then cued participants to recall the list by speaking the list items aloud.

The experiment included three conditions of presentation and recall. In the *free-varied* condition, list items were randomized between list presentations and participants were instructed to recall the list in any order (e.g., free recall). In the *free-constant* condition, list items were in the same position between list presentations and participants were instructed to recall the list in any order (e.g., free recall). Finally, in the *serial-constant* condition, list items were in the same position between list presentations and participants were instructed to recall the list in the order it was perceived (e.g., serial recall). Twelve participants completed 21 test trials each for each condition across several sessions. In each session, participants completed a set of trials from only one condition.

As mentioned earlier, participant responses were scored according to serial position (SP), and conditional response probability (CRP) as a function of lag. Serial position measures recall accuracy as a function of an item's position in the list; that is, for each list position (1st, 2nd, 3rd, etc.), the probability that participants report the corresponding item during recall. Typically recall is best for items in early positions (primacy) and late positions (recency). The conditional response probability as a function of lag measures the distribution of successive recalls as a function of item distance from the current item in the presentation of the list. Mathematically, it shows the probability of recalling item $i+lag$ after recalling item i (see Kahana, 1996 for more details). Typically, as per the contiguity effect, after recalling item i , learners are most likely to next recall the successive item ($i+1$) from the list presentation. Participants were also scored according to item score, which measures how many total list items were successfully recalled for each presentation regardless of order.

In accordance with prior findings (Howard & Kahana, 1999; Kahana, 1996), this study showed several noteworthy effects. First, the serial position measure (shown with our model results in Figure 2) indicated strong primacy and recency effects, where participants are biased towards recalling items at the beginning and end of the list. In both free recall conditions, the recency effect is generally favored, whereas in serial recall, the primacy effect dominates and the

recency effect is significantly lower. These effects are attenuated across learning, however, as subsequent presentations increase the accuracy rate overall.

Second, the conditional response probability measure (shown with our model results in Figure 3) indicated a clear contiguity effect, where participants are biased towards recalling neighbors of the item they just recalled. For all three conditions, this effect was also significantly asymmetric, where participants favored subsequent items as opposed to preceding items. Multiple significant interactions between presentation, transition direction and condition show that the asymmetry shows different characterizations for each condition over time. Specifically, the asymmetric effect was stronger in the serial-constant condition than the others, and for both it and the free-constant condition, the effect increased with the number of presentations. While the effect size in the free-varied condition was comparable to the free-constant condition after the first presentation, however, it decreased with further presentations until, after the fifth presentation, it was virtually absent.

The measure of item score showed a significant increase in accuracy over time. Although we correctly predict this increase, since this does not shed much additional light on distinguishing between the different theories of list recall, we do not focus on it much in this paper.

The study's authors interpret these results as being supportive of an associative account of list learning, as do we. To preview our approach, we explain the primacy effects via mental rehearsal, and we explain the recency effects via the decaying nature of activation strengthening, where more recent items in memory are more likely to be recalled. The asymmetric contiguity effect, and how it changes over time, is explained by asymmetries in associative learning. We go into this in more detail in the following section.

Associative Learning in Memory Recall

Our account of associative learning, as we have said, is situated in the cognitive architecture ACT-R/E (Adaptive Character of Thought-Rational / Embodied; Traflet et al., 2013), an embodied version of the cognitive architecture ACT-R (Anderson et al., 2004). ACT-R is an integrated theory of human cognition in which a "production system operates on a declarative memory" (Anderson et al., 1998). Key to this paper, in ACT-R, item recall depends on three main components: activation strengthening, activation noise, and associative activation. These three values are summed together to represent an item's total activation. When a recall is requested, the item with the highest total activation is retrieved, subject to a retrieval threshold; if no item's activation is above the threshold, the retrieval is said to *fail* and no item is recalled. We next discuss each of these components in turn, focusing on associative activation, which is the main contribution of this work.

Activation Strengthening

ACT-R's well-established theory of activation strengthening has been shown to be a very good predictor of human

declarative memory (Anderson et al., 1998; Anderson, 2007; Schneider & Anderson, 2011). Intuitively, activation strengthening depends on how frequently and recently a memory has been relevant in the past. It is designed to represent the activation of a memory over longer periods of time and, generally, is highest right after the memory has been accessed in *working memory*, slowly decaying as time passes. Working memory represents the items that are currently the model's focus of attention. Activation strengthening, A_s , is calculated as:

$$A_s(i) = \ln \left(\sum_{j=1}^n t_j^{-d} \right)$$

where n is the number of times an item i has been accessed in the past, t_j is the time that has passed since the j th access, and d is the strengthening learning parameter, specifying items' rate of decay, and which defaults to 0.5. Importantly, this equation predicts that items that have occurred recently, or have been rehearsed more, are more likely to be recalled than those that have not.

Activation Noise

The activation noise of a memory is drawn from a logistic distribution with mean 0 and standard deviation the parameter σ_n . It is a transient value that changes each time it is used, and models the neuronal noise found in the human brain. This parameter's default value was 0.25, a common value for this parameter across models.

Associative Activation

While associations are not new to the ACT-R framework (e.g., Anderson, 1983), we adopt a new account of associative learning as part of our approach (Thomson & Lebiere, 2013a). Like in the original version, a third contributor to the activation of items in memory is associative activation, which sources from the contents of working memory. Activation then spread along associations to items or memories related to those in working memory. Here, we describe this new account qualitatively, for the purposes of clarity; more technical details, formulations, and justifications of its mechanisms can be found in previous work (Thomson & Lebiere, 2013a; 2013b; Thomson, Bennati & Lebiere, 2014).

Important to this paper are that associative strengths are learned, strengthened, and weakened over time, as new or repeat items are encountered. Additionally, as in the original version, associations are *directional*; an association can be stronger from an item i to an item j , for example, than the association from item j to item i (or, there could be no association from item j to item i at all).

Associations are learned between items that are relevant in working memory in temporal proximity to one another, and lead from earlier items to later items. The strength of the association (or how strongly it is increased) is determined by the amount of time that passes between when the items were each in working memory. If one item is immediately followed by another in working memory, they will be very strongly associated; on the other hand, if an item has been out

of working memory for a while before another is added, they will be only weakly associated.

In this way, rich associations are formed that point forward in time, relating past items to current ones. Unlike explicit chaining models (e.g., Lee & Estes, 1977) that form only direct item-item chains between immediately adjacent neighboring items (i.e., between the last item and the current item entering working memory), we form multiple item-item associations between all items recently in working memory and newly added items.

There are two other substantial differences between ACT-R's original associative learning mechanisms and our new account's that are not relevant to this model, but that we mention here for completeness. First, our associative learning mechanism is based on Hebbian, not Bayesian learning; recently, we have argued that this is better suited to the types of large, complicated tasks that human memory is able to handle (Thomson & Lebiere, 2013a). Additionally, our mechanism includes buffer-specific associations that create a rich context for memory recall; again, however, that is outside the scope of this experiment.

Modeling Multi-Presentation List Recall

We wrote a model in ACT-R/E that completes both free and serial multi-presentation recall tasks, as were in Klein et al. (2005). We begin by assuming that, before a task begins, the model has a "start" concept in working memory that tells it to wait for the stimuli to start being presented; we also assume that the model has no *a priori* knowledge of these words (i.e., the words are not already associated with other items or concepts). Upon hearing a stimulus, the model initially encodes the stimulus as a word. The rapid pace of the experiment leaves little time for rehearsal; therefore, the model rehearses the first stimulus, but forgoes rehearsal after that due to the tight time constraints.

Once the full list has been presented, the model then attempts to recall each element of the list; at any given time, the item with the highest total activation is recalled. Retrievals proceed until the complete list has been recalled or until a recall request fails, at which point the presentation is considered complete.

The only difference between how the model performs the free and serial recall tasks is that, when beginning to recall a list in the serial recall task, the model first retrieves the "start" concept in an attempt to start at the beginning of the list. It forgoes this step in the free recall task.

When the model looks at a new item, the previous item immediately precedes the new item in working memory. Thus, a strong positive association is formed (or strengthened) from the preceding item to the new item. Additionally, associations from more distant items to the new item are also formed or strengthened, attenuated by their temporal distance to the new item. Figure 1 shows an example of what the associations look like after three items of a list have been presented.




Figure 1. A sample associative structure, including associative strengths, after three items of a list have been viewed. Of note is that association strengths weaken as items become farther removed in time, as well as the asymmetric structure of the associations. Note that, for clarity, we omit here associations not relevant to our discussion.

With respect to parameters, all ACT-R/E parameters were set at their default values. The three associative learning parameters (learning rate, interference rate, and residual activation decay rate; see Thomson, Lebiere, & Bennati, 2014) were set to represent a fairly moderate pace of associative learning (set at 1.5, .25, and .5, respectively). Note that these parameters were the same for both the serial and recall tasks and, thus, for all three conditions of the experiment we are modeling.

Model Explanations

The model explains the data according to both activation strengthening and associative activation. First, the decaying nature of activation strengthening implies that more recently presented stimuli will be more likely to be recalled, creating a recency effect among all conditions. Primacy is primarily explained by the rehearsal of the first few items. Primacy is relatively stronger in serial recall because the model makes the effort to retrieve the “start” concept before beginning list recall, which activates the beginning items of the list. On the other hand, the lower primacy effect in free recall implies that it will have a stronger recency effect. This is because the beginning items of the list will provide less competition to the items at the end of the list, leading to an increased bias towards those ending items.

The forward asymmetry of the associative structure created as the model learns the list clearly explains the forward asymmetry effects shown by the conditional response probability measure. When an item is in working memory, the subsequent item receives a strong amount of associative activation; the item after that, in turn, receives a much smaller, but still positive, amount. This boosts the probability that items in the forward direction will be recalled at any given time. The model also indicates that this asymmetry will only increase across multiple presentations of both free-constant and serial-constant conditions as the forward-leaning associations are strengthened. For the free-varied condition, the model explains why the asymmetry contiguity effect diminishes over multiple presentations: it is because, in this condition, associations are created and strengthened in various directions across various items.

Our model also explains why the serial-constant condition has a stronger contiguity effect than the free-constant condition. This is due to how the model learns during both the learning and recall phases of the experiment. In the serial-

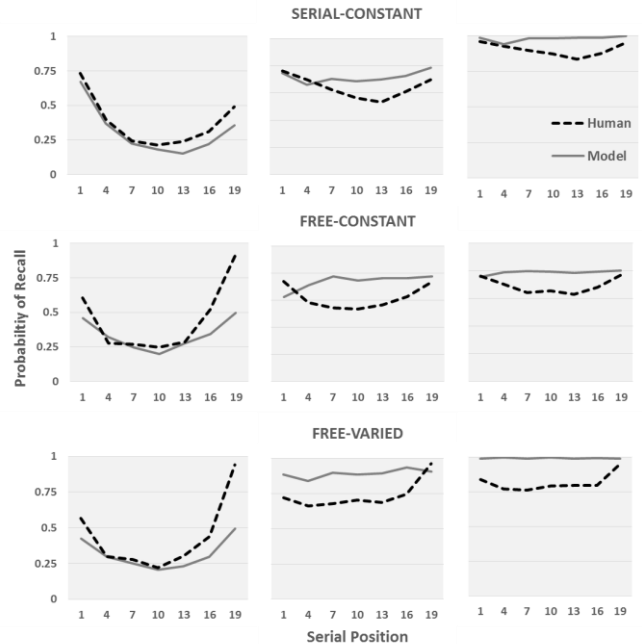


Figure 2. Serial Position curves, showing the overall recall probability for each list item, across serial-constant, free-constant, and free-varied conditions for both human and model. Panned from left to right are the results for presentations 1, 3, and 5, respectively. As seen, the model captures the broad primacy and recency effects in the first presentation, but not later ones; we believe this is due to a higher emphasis on rehearsal than we assume here.

constant condition, because the model attempts to report the items in order, the forward-facing associations are strengthened during the recall phase; in the free-constant condition, since items are not reported as serially, the forward-facing associations are strengthened to a lesser extent. This difference in associative strength ultimately predicts that the serial-constant condition will exhibit stronger continuity effects than the free-constant condition.

Model Results

To collect results, the model performed all three conditions of the Klein et al. (2005) experiment, performing the serial recall or free recall task as appropriate. All stimuli were presented at the same rate as they were to the human participants, and the same words were used as stimuli. The model was run for the same number of trials (252) per condition as all human participants (252 trials); we assume that the model begins each trial with no knowledge of any of the items.

As predicted, the model strongly predicts serial position curves in serial-constant condition ($r^2 = .92$; see Figure 2). The results of the serial position curve in the free-varied and free-constant, while acceptable, were not as strong ($r^2 = .71$ and 0.67 , respectively). An in-depth look at the data suggests that this lower-quality fit is due to us not accounting for primacy effects strongly enough in later presentations; we

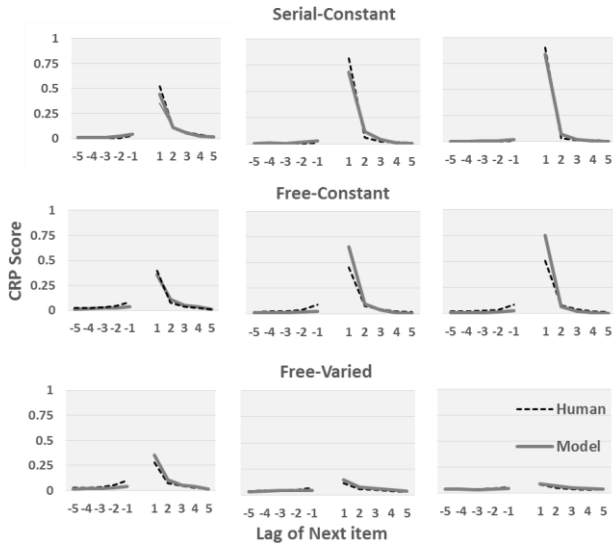


Figure 3. Conditional Response Probability curves, showing the probability of recalling item $i+lag$ after item i , across serial-constant, free-constant, and free-varied conditions for both human and model. Paneled from left to right are the results for presentations 1, 3, and 5, respectively. As seen, the model accurately captures not only the amount of asymmetric contiguity effect per condition, but also the change in the effect across multiple presentations.

believe this is due to participants putting more emphasis on rehearsal than we assume, and plan to investigate this further.

As seen in Figure 3, our model strongly matches the contiguity affects across all three conditions ($r^2 = .89$ for free-varied; $r^2 = .96$ for free-constant; and $r^2 = .99$ for serial-constant). As predicted, the asymmetric contiguity effect increases across presentations in the serial-constant condition and, to a lesser extent, in the free-constant condition, while it is reduced in the free-varied condition. The model slightly over-predicts contiguity in the free-constant condition while slightly under-predicting contiguity in the serial-constant condition. We argue that this is because the model learns no strategy while performing the task. Humans performed each condition in a block, and we argue, were able to adapt their encoding/recall strategies based on their task instructions. To avoid overfitting, all three of our models used the same encoding strategy. Our goal was to show the amount of variance that could be captured by a low-level, automatic, and stimulus-driven mechanism such as associative learning.

As a minor note, our model also correctly predicts increases in item score across presentations for all three conditions, with $r^2 = .96$. Our model predicts this due to increased associativity and activation strengthening over multiple presentations.

Alternate Accounts of List Learning

The detailed results from Klein et al. (2005) present a challenge for many of the current theories of memory that

explain serial and free recall of lists, which have modeled only a subset of its results. The temporal context model (TCM, also called the context maintenance and retrieval model, CMR) (Polyn, Norman, & Kahana, 2009) for example, associates items with contextual states; when an item is recalled, so is its contextual state, which drives the recall of other temporally similar items. They use this construct to account for both recency and asymmetric contiguity (Howard & Kahana, 2001). While they qualitatively describe how their model extends to serial recall, they do not explicitly model it, so it is unclear how good of a match it can achieve. More importantly, they also do not model how these curves change over multiple presentations. In contrast, the cornerstone of our theory of associative learning is explicitly modeling how associative strengths change with repeated exposure to items, allowing us to account for the multi-presentation recall data we discuss here.

The start-end model (SEM) (Henson, 1998) relies upon implicit start and end markers of the list sequence, as well as tokens for spatiotemporal markers for each item, to make its predictions. While these constructs allow it to successfully match data showing primacy and recency in single-trial serial recall, it does not explain serial recall's contiguity effects. It also does not model free recall, and the author also makes no predictions about how it would perform in a multi-trial setting.

SIMPLE (Brown, Neath, & Chater, 2007) models both serial and free recall tasks. Its predictions are generally based on the temporal distinctiveness of items in memory; it can also include other measures of distinctiveness (e.g., semantic distinctiveness). More importantly, it has been matched against only data showing primacy and recency effects, and it does not appear to correctly predict asymmetric contiguity effects, nor how these effects change across multiple list presentations. Like SIMPLE, we include a time-based component in the form of activation strengthening; our analog of their semantic distinctiveness, however, is our theory of associative learning, which more naturally explains the asymmetry that arises in conjunction with contiguity effects.

Anderson et al. (1998) models both free and serial recall tasks, as well. It also includes a simple conceptualization of item-item associations, and so it seems to predict contiguity effects after a single trial. It does not, however, seem to predict how contiguity effects would increase over time. This is because its associations, once learned, do not strengthen over time, they only potentially weaken as more and more items are encountered. As we indicated earlier, however, overall we view this approach as one of the most promising both because of its close capture of SP and CRP curves, and because of its strong foundation in general cognition; that is why we have expanded upon it in this paper by adding in a richer notion of associative learning.

Discussion

In this paper we presented a theory of memory recall that includes a rich account of how associations are learned and

strengthened over time. We described how a single model with fixed parameters presented an excellent fit to human data across both free and serial multi-presentation list recall tasks, including modeling asymmetric contiguity effects than change over time.

One criticism of other models of both free and serial recall has been that they do not well account for two notable effects that have been shown to differentiate between the two conditions (Murdock, 2008). First, similarity between list items seem to facilitate performance on free recall tasks, but hinder performance on serial recall tasks. Our model predicts this because of the nature of our associations, where similar items naturally become associated in memory; in fact, there is some evidence that similarity itself is based on associative learning (Hiatt & Trafton, 2013). This similarity would facilitate performance on a free recall task because remembering one item would activate similar items, boosting their recall probability. For the same reason, it would hinder serial recall accuracy since similar items that appear out of order would hinder the recall of items in the correct order.

Second, longer presentation rates have been shown to improve performance in free recall tasks, but do not affect performance on probe-digit experiments (a simplified version of serial recall). We predict this because longer presentation rates, as opposed to the rapid presentation rate in this experiment, promote rehearsal; rehearsal, in turn, increases activation strengthening for list items. While this intuitively helps recall performance for free recall tasks, the serial effects of the items' forward associations shield the probe-digit experiment from any negative (or positive) implications of the higher activation strengthening.

While associative learning account relies on item-item associations, these associations do not fall prey to the general criticisms against chaining models (Lee & Estes, 1977; see Henson, 1998 for critique). Specifically, since associations are formed between all items recently in working memory and a newly added item (i.e., what Henson (1998) refers to as *compound-chaining*) we avoid the brittleness of typical chaining theories, where a broken 'link' in the chain can cause cascading errors and leads to trouble matching behavioral data. Instead, our approach can recover from such problems due to its richer association structure.

Acknowledgments

This work was supported by the Office of the Secretary of Defense / Assistant Secretary of Defense for Research and Engineering (LH) and the Office of Naval Research (LH). The views and conclusions contained in this paper do not represent the official policies of the U.S. Navy.

References

Anderson, J. R. (2007) *How Can the Human Mind Occur in the Physical Universe*. Oxford University Press: Oxford.
 Anderson, J. R. (1983). A spreading activation theory of memory. *Journal of verbal learning and verbal behavior*, 22 (3), 261-295.
 Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., Qin, Y. (2004). An integrated theory of mind. *Psychological Review*, 111, 1036-1060.

Anderson, J. R., Bothell, D., Lebiere, C. & Matessa, M. (1998). An integrated theory of list memory. *Journal of Memory and Language*, 38, 341-380.
 Brown, G. D., Neath, I., & Chater, N. (2007). A temporal ratio model of memory. *Psychological review*, 114 (3), 539-547.
 Hiatt, L. M., & Trafton, J. G. (2015). An Activation-Based Model of Routine Sequence Errors. *Proceedings of the International Conference on Cognitive Modeling*.
 Hiatt, L. M., & Trafton, J. G. (2013). The Role of Familiarity, Priming and Perception in Similarity Judgments. *Proceedings of the Conference of the Cognitive Science Society*.
 Howard, M. W., & Kahana, M. J. (1999). Contextual variability and serial position effects in free recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 25(4), 923-940.
 Henson, R. N. (1998). Short-term memory for serial order: The start-end model. *Cognitive Psychology*, 36 (2), 73-137.
 Kahana, M. J. (1996). Associative retrieval processes in free recall. *Memory & Cognition*, 24, 103-109.
 Klein, K. A., Addis, K., & Kahana, M. J. (2005). A comparative analysis of serial and free recall. *Memory & Cognition*, 33(5), 833-839.
 Lee, C. L., & Estes, W. K. (1977). Order and position in primary memory for letter strings. *Journal of Verbal Learning and Verbal Behavior*, 16 (4), 395-418.
 Lukes, G., Thompson, G., & Werbos, P. (1990). Expectation driven learning with an associative memory. In *Proceedings of the IEEE International Joint Conference on Neural Networks*.
 Murdock, B. (2008). Issues With the SIMPLE Model: Comment on Brown, Neath, and Chater (2007). *Psychological review*, 115(3), 779-785.
 Polyn, S. M., Norman, K. A., & Kahana, M. J. (2009). A context maintenance and retrieval model of organizational processes in free recall. *Psychological Review*, 116 (1), 129-40.
 Pyke, A., West, R. L., & LeFevre, J. A. (2007). How Readers Retrieve Referents for Nouns in Real Time: A Memory-based Model of Context Effects on Referent Accessibility. In *Proceedings of the International Conference on Cognitive Modeling*.
 Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. *Classical Conditioning II: Current Research and Theory*, 2, 64-99.
 Schneider, D. W., & Anderson, J. R. (2011). A memory-based model of Hick's law. *Cognitive Psychology*, 62 (3), 193-222.
 Thomson, R., Bennati, S., & Lebiere, C. (2014). Extending the Influence of Contextual Information in ACT-R using Buffer Decay. In *Proceedings of the Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.
 Thomson, R., Lebiere, C., Anderson, J. R., & Staszewski, J. (2014). A general instance-based learning framework for studying intuitive decision-making in a cognitive architecture. *Journal of Applied Research in Memory and Cognition Special Issue on Intuitive Decision-Making*.
 Thomson, R. & Lebiere, C. (2013a). Constraining Bayesian Inference with Cognitive Architectures: An Updated Associative Learning Mechanism in ACT-R. In *Proceedings of the Conference of the Cognitive Science Society*.
 Thomson, R. & Lebiere, C. (2013b). A Balanced Hebbian Algorithm for Associative Learning in ACT-R. In *Proceedings of the International Conference on Cognitive Modeling*.
 Trafton, J. G., Hiatt, L. M., Harrison, A. M., Tamborello, F. II, Khemlani, S. S., & Schultz, A. C. (2013). ACT-R/E: An embodied cognitive architecture for human-robot interaction. *Journal of Human-Robot Interaction*, 2 (1), 30-55.