In D. Reitter & F. E. Ritter (Eds.), Proceedings of the 14th International Conference on Cognitive Modeling (ICCM 2016). University Park, PA: Penn State.

An Account of Interference in Associative Memory: Learning the Fan Effect

Robert Thomson (robert.thomson@usma.edu)[1, 2]
[1]United States Military Academy, West Point, NY 10928 USA

Anthony M. Harrison[2] (anthony.harrison@nrl.navy.mil)
J. Gregory Trafton[2] (greg.trafton@nrl.navy.mil)
Laura M. Hiatt[2] (laura.hiatt@nrl.navy.mil)
[2]Naval Research Laboratory, Washington, DC 20375 USA

Abstract

Associative learning is an essential feature of human cognition, accounting for the influence of priming and interference effects on memory recall. Here, we extend our account of associative learning that learns asymmetric item-to-item associations over time via experience (Thomson, Pyke, Trafton, & Hiatt, 2015) by including link maturation to balance associations between longer-term stability while still accounting for short-term variability. This account, combined with an existing account of activation strengthening and decay, predicts both human response times and error rates for the fan effect (Anderson & Reder, 1999). This represents the highest fidelity replication of a human experiment for modeling that we are aware of.

Keywords: associative learning; interference; cognitive models; fan effects

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), similarity judgments (Hiatt & Trafton, 2013), and memory recall (Thomson, Pyke, Trafton, & Hiatt, 2015). Despite its ubiquity, it is difficult to model directly due to its entangled ties to other aspects of cognition (e.g., memory decay).

Perhaps associative learning's most studied effect is that of *priming* (and its converse *interference*). Priming occurs when the retrieval of one memory facilitates the retrieval of another. Conversely, interference occurs when a memory primes multiple other memories instead of just the ones that are useful or relevant to the current situation. Those other memories are said to *interfere* with the useful one. When there is high interference, recognition accuracies are relatively lower and recognition response times relatively longer when compared to situations where there is low interference, ostensibly due to having lower overall activation in memory. Assuming that the degree of interference is positively correlated with the number of competing associations, then having more competing associations (i.e., a higher *fan*) will lead to relatively higher error rates and latencies than memories having relatively fewer competing associations. This effect is most popularly known as the *fan effect* (Anderson, 1974).

In this paper we will extend our account of associative memory embodied in a cognitive architecture (Thomson, Bennati & Lebiere, 2014) to account for the fan effect experiment. This account of associative memory has already successfully predicted the complicated results of a multi-trial free and serial recall task, including asymmetric contiguity effects that strengthen over time (Thomson et al., 2015). Here, we extend our theory to include link maturation to balance associations between longer-term stability while still accounting for shorter-term variability. We then use the theory as part of a cognitive model that performs the fan effect experiment using the same stimuli and presentation times as the human participants.

By doing this, we become the first theory of associative memory to explain how associations are learned and updated throughout the fan effect experiment. Previous models considered only associations at the end of the experiment (Anderson & Reder, 1999; Schneider & Anderson, 2012; Anderson, 1974; Rutledge-Taylor and West, 2008); our model enhances their understanding of associative memory by describing the process of how these end-state associations are reached.

Associative Learning in Memory Recall

Our account of associative learning is situated in the cognitive architecture ACT-R/E (Adaptive Character of Thought-Rational / Embodied; Trafton 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). In ACT-R, recall and latency depend 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. The latency of the recall is also inversely correlated to the recalled item's activation.

Activation Strengthening

ACT-R's well-established theory of activation strengthening (also called *base-level* activation) has been shown to be a very good predictor of human declarative memory (Anderson et al., 1998; Anderson, 2007). Intuitively, activation strengthening depends on how frequently and recently a memory has been relevant in the past, and is calculated as:

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

where n is the number of times an element i has been accessed in the past, t_j is the time that has passed since the jth access, and d is the learning parameter, specifying an element's rate of decay. 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.

Associative Activation

In our account, associative strengths are learned, strengthened, and weakened over time as new elements are learned or prior elements re-experienced. These associations are learned between relevant working memory items within temporal proximity to one another, leading from earlier to later items (Thomson, Bennati, & Lebiere, 2014). The strength of the learned association (or how strongly an existing association is increased) is influenced 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 become 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. Additionally, associations are *asymmetric*; 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).

To balance the rate of associative learning between long-term stability and short-term variability, link maturation was included as an additional parameter. Associative link maturation slows the rate of strengthening and weakening based on the number of times the link has been used. This supports long-term stability of well-experienced associative links while allowing for rapid short-term learning of new associative elements. In neural networks, maturation is equivalent to the process of settling to reach a stable equilibrium (Wills et al., 2005; Eliasmith, 2005). Maturation is set using a logistic function:

$$M = 1 - \frac{1}{1 + e^{-(\ln(timesInContext) * MaturationRate)}} \quad (2)$$

The maturation rate controls the steepness of the curve, or, in other words, controls how quickly links will stabilize.

To compute associative strength from an item j to an item i, the learning mechanism computes an increment I_{ji}:

$$I_{ji} = lr * w * M \quad (4)$$

where lr is a learning rate parameter, w is the weight of the increment determined by the strength of the items in working memory (scales from 0 to 1), M is maturation, and R is refraction. This increment is used to update link strength as follows:

$$S_{ji} = S_{jiPrior} * (1 - I_{ji}) + (I_{ji} * mas) \quad (5)$$

where is the strength of the link from j to i, $S_{jiPrior}$ is the prior strength of S_{ji}, I_{ji} is the learning increment from above, and *mas* is a parameter controlling the maximum possible associative strength.

When a new link is learned or existing link updated that shares a source j with other existing links, then each of those other links are discounted proportionally to the *weight* that the original link is updated (e.g., S_{ji} is updated so S_{jk} is discounted):

$$S_{jk} = S_{jkPrior} * (1 - I_{jk}) \quad (6)$$

where I_{jk} is computed using the *weight* from the link from j to i, but using the maturation M from the link from j to k. Equation 6 normalizes the amount link j to k is discounted based on the degree to which it has settled. This allows for newer links to rapidly change while providing for long-term stability for more mature links.

This discounting function attenuates link strengths consistent with interference accounts of memory. As more concepts compete in memory, the amount of associative strength from each concept is reduced. In a balanced environment, this discounting will approximate the statistical likelihood $P(i/j)$, which is the odds of perceiving or retrieving i immediately prior to j.

Armed with an understanding of our modeling framework, we now turn to the fan effect experiment itself.

The Fan Effect Experiment

To understand the fan effect, we consider Anderson and Reder (1999) classical fan experiment They capture the fan effect in a recognition task where participants begin by learning 48 pairs of people and places. Persons and places could appear in multiple pairs, and each pair was shown for five seconds. Then, during testing, participants respond yes (*target*) or no (*foil*) to whether presented statements were previously studied: the *person* is in the *place* (e.g., 'the *hippie* is in the *park*'). In the testing phase, participants were provided a monetary reward based on their total score. The score was computed by providing 1 point for each correct response, plus an additional point for each 100 ms of response times faster than 1500 ms. This induced a speed-accuracy trade-off into the experiment.

The experiment proceeded according to three phases: a study phase, drop-out training, and then a testing phase. In the study phase, each stimulus pair was presented once on the screen for 5 seconds. In the drop-out training phase, participants were presented with questions 'Who is in the *location*?' and 'Where is the *person*?' Participants had to respond with all persons associated with the location (or vice versa). Participants had to correctly answer all these questions for person and location to complete the phase. Participants completed two of these drop-out training phases. Finally, in the testing phase, participants would respond *yes* or *no* to queries 'the *person* was in the *location*' with participants receiving feedback on their response.

The experiment manipulated the test stimuli in two different ways. The first was to manipulate the fan of the persons and places. In this experiment, fan is the number of persons associated with a place, and vice versa. Fan is controlled by varying the number of persons in each place, or the number of places with each person (e.g., 'the ***hippie*** is in the *bank*' or 'the *soldier* is in the ***park***'). Here, the fan of one term (person/place) was fixed at 2, while the fan of the other term (place/person respectively) was varied to be either 2 (low-fan) or 4 (high-fan).

The second manipulation was to control the composition of the set of test stimuli shown to participants by manipulating different target and foil conditions. There were four target conditions: facilitation, interference, suppression and control. In the facilitation condition, each target (e.g., 'the *biker* is the *tower*') was 'facilitated' by being repeated 5 times each in the stimuli set. In the interference condition each target was repeated only one time in the stimuli set, and was considered interference because the target's person or place overlapped with a target from the facilitation condition (i.e., 'the *biker* is in the *factory*', or 'the *doctor* is in the *tower*').

The other two conditions were the suppression and control conditions. In these conditions, each target appeared once in the stimuli set, and consisted of facts that were seen in the interference (but not facilitation) condition, such as *factory* and *doctor* in the above examples. Examples of suppression targets included: 'the *writer* is in the *factory*', and 'the *doctor* is in the *bank*'. Due to a particularity in the original study, there is limited difference between suppression and control stimuli, because the controls were designed such that they would functionally suppress stimuli from the suppression condition (e.g., 'the *monk* is in the *bank*'). They are different insofar as the suppression stimuli were effectively two steps removed from the facilitation condition, while the control trials were effectively three steps removed.

Foils were classified according to three conditions: high-frequency foils, which used person/place concepts from the facilitation condition but with novel pairings, and were repeated 4 times each in the stimuli set; low-frequency foils, which had novel pairings of person/place concepts from the interference, suppression, or control conditions and appeared once each in the stimuli set; and mixed foils, which created novel pairings using one high-frequency concept from the facilitation condition and one low-frequency concept from the interference, suppression, or control conditions and were repeated only once in the stimuli set. In total, there were 48 target sentences and 54 foil sentences in the stimuli set.

The test stimuli set was presented three times in successive blocks, and all stimuli were presented in each block. Feedback was provided for 1 second after participants' responses, with an additional 1 second inter-trial interval[1].

The results of this study were consistent with interference effects: there were longer latencies and more errors in the high-fan (i.e., fan of 4) conditions relative to the low-fan (i.e., fan of 2) conditions for both targets and foils, with both high-frequency (i.e., facilitation) targets and foils having relatively higher accuracy and quicker latencies than their corresponding low-frequency counterparts. They also predicted lower relative accuracy in the interference condition relative to the suppression and control conditions, and no difference between suppression and control.

Prior Modeling of the Fan Effect

There have been several attempts to mathematically model fan effects (Anderson & Reder, 1999). Most prominent is Anderson and Reder's (1999) model whose equations were grounded in the ACT-R cognitive architecture (Anderson and Lebiere, 1998). This model can be broken down into three related equations.

$$S_{ji} = S + \ln(1/fan_j\ ) \tag{7}$$

$$A_i = B_i + \sum_j W_j S_{ji} \tag{8}$$

$$T = I + Fe^{-A_i} \tag{9}$$

Equation 7 describes the spread of activation (S_{ji}) from element j to i as a function of associative strength intercept S attenuated by the fan of j, which is the number of concepts to which j is associated. In Equation 7, $1/fan_j$ is a simplification of $P(i/j)$ assuming equal frequencies of i and j. In Anderson and Reder, frequencies were not equal, and were instead set ahead of time according to the objective probabilities in the model. Equation 8 relates an activation function A_i to the base-level activation from Equation 1, and the sum of spreading activation from Equation 7 multiplied by an attentional weigh W_j. Prior efforts set B_i to 0 on the assumption that the drop-out testing would balance out base-level activation between stimuli.

Finally, Equation 9 computes retrieval time T based on an intercept I, time scale offset F, and the activation function A_i from Equation 3. The estimates for each parameter were as follows: I was 1197 ms, F was 773 ms, S was 2.5 ms, and W was .33 (reflecting an even weighting of '*person*' '*in*' '*place*'). Using these parameters, Anderson & Reder report a strong correlation with response times, $r = .956$. This model did not attempt to fit error patterns.

This model, while successful, does not focus on modeling both latency and error rates, which we believe is an important part of understanding priming and interference in associative learning. While describing nicely the final average performance of participants, they provide little intuition for *how* participants learn the associations via experiencing the task. For instance, Equation 7 uses a fixed value for each condition that does vary.

In contrast, our approach grounds our account of associative learning within the larger ACT-R/E architecture along with the constraints it places on cognition (Trafton et al., 2013) by using a production system simulating the time-course of perception, encoding, retrieval, and response. Once base-level activation is included as a factor, then the time-course of stimulus presentation and training becomes important in determining overall accuracy and response time. This added fidelity (and complexity) may test assumptions made in prior modeling efforts, and also may provide new insights or hypotheses about how participants learn the task. To that end, our model performs the experiment analogously to participants, and learns associations over time. This supports our theory of associative memory explaining how associations are learned and adapt over time.

Learning the Fan Effect

The model starts with only background knowledge of the words used in the experiment and is equipped with the procedural knowledge necessary to perform the experiment. It has no underlying knowledge of the concepts of person and place, and thus has no knowledge of targets or foils. As we have said, the model is presented with the same experimental paradigm as the human participants.

The model uses the same procedural knowledge at the start of each phase to perceive the person and place concepts:

[1] This 1 s ITI was not listed in the Anderson & Reder (1999) paper, however it was reported in subsequent research.

when it sees two concepts on the screen together, then they are linked into a common concept-pair, with each constituent person and place priming the concept-pair. This concept-pair is represented in memory as a single chunk of information containing three features: *person*, *place*, and *was-target*. Was-target is a binary *true/false* decision, where *true* indicates that the stimulus was a *target* and *false* indicates that the stimulus was a *foil*.

The external environment is a simulated computer screen. When perceiving stimuli, the model randomly encodes one symbol first and then the other. The model categorizes these symbols according to two features: *person-symbols* and *place-symbols*. These representations are functionally identical and are distinguished only for lexical purposes to simply categorize incoming words. Since stimuli are of the form 'the *person* is in the *place*' we present person-symbols on the left of the display and place-symbols on the right of the display.

When a concept-pair is learned or updated, then the associative strengths between the person/place concepts and the pair is strengthened while the strengths between the concepts and their other related concept-pairs are weakened. Since concepts are related to more concept-pairs on high-fan trials than on low-fan trials, high-fan concepts tend to have lower associative strengths to their related concept-pairs than low-fan concepts. This lower associative strength predicts that concept-pairs involving high-fan concepts will have slower response latencies and increased error rates. Also, since high-frequency stimuli have been seen more often, their strengths will be stronger than low-frequency stimuli, however maturation controls the degree to which these stimuli increase in strength (e.g., a link seen 4 times as often is not 4 times as strong).

In the study phase, the model automatically encodes all concept-pairs as targets. After encoding the stimuli and generating a concept-pair representation, the model then repeats this encoding until the stimuli are no longer presented on the display, averaging 2-3 rehearsals over the 5 second presentation time.

In the drop-out training phase, the model perceives either a *person* or *place* and attempts to retrieve all *places* where that *person* is (or all *persons* in that *place*). If the model correctly perceives all required elements then the model moves onto the next stimulus, otherwise it studies those stimuli again and returns to the drop-out training. Once the model has successfully retrieved all elements in both run-throughs of the query phase, the test phase begins.

The test phase is the critical phase where all response times and error rates were recorded. Similar to the study phase, the model begins to encode the concept-pair for *person* and *place* as an analogue to perceiving: was *person* in the *place*? The model then attempts to retrieve any decision containing said *person* and/or *place*. This decision is a prior concept-pair stored in memory including person, place, and the critical *was-target* decision. A response is then generated according to the following criteria: 1) if the model is unable to retrieve any concept-pair due to all pairs being below threshold, then it responds *foil*; 2) if the model correctly retrieves the matching *person*, *place*, and *was-target* decision then it responds with the respective decision: *target* for *true* and *foil* for *false*; or 3) if the model retrieves a mismatching concept-pair containing one *person* or *place* but not both, then it assumes that the response is a *target*. After the model responds it receives feedback, which it uses to encode the correct concept-pair. This includes encoding foils, which allows the model to be capable of correctly retrieving that an item was a foil seen in an earlier testing phase. This is a unique behavior of our model and reflects the fact that participants cannot 'ignore' decisions they've made and must encode feedback that they've seen.

Finally, if the model was incorrect or had retrieved a mismatched element, then for the feedback period it rehearses the correct response. This process is repeated across all trials through three test phases.

In the testing phase, response times are recorded from the stimulus onset time until the model has responded with the appropriate decision (target or foil). It is important to note that as a fully-implemented production system model, the complete time to respond include two relatively fixed durations: approximately 600 ms to encode the stimuli from the display, and approximately 350 ms to prime the motor command and press the response key. This is in a similar range to the structural offset I of 1197 ms that Anderson and Reder (1999) used in Equation 9. This means that fan effects in latencies occur mainly in the approximately 200 ms – 800 ms timeframe where the concept-pairs are retrieved. It is the retrieval of the concept-pair that determines fan effects in both latencies and accuracy.

One final difference between the present model and prior efforts is that our model incorporates base-level activation B_i (see Equation 8), but replaces the S_{ji} from Equation 7 with our learned S_{ji} from Equation 5. As previously mentioned, base-level activation reflects the recency and frequency of use of elements, and it is not a given that base-level would be equivalent for items across the different target and foil conditions, especially for the high-frequency vs. low-frequency elements where frequency is necessarily varied.

Results

The present model was run for 200 iterations with the parameters described in Table 1 below. As is apparent from Figures 1 and 2, the model qualitatively captured fan effects in both accuracy and latency, respectively, reflected by slower response times and higher error rates for high-fan concepts compared to low-fan concepts. We also predicted relatively higher accuracy for high-frequency conditions (facilitation for targets and high-frequency for foils) to the rest of the conditions; lower accuracy in the interference condition relative to the control and suppression conditions, and no difference between suppression and control target conditions. One difference is that we predict a smaller average fan (.03 s instead of .09 s) than did Anderson and Reder (1999).

Table 1. Parameters used in Fan Effect Experiment

PARAMETER	VALUE
Base-Level Learning (B_i)	.4
Learning Rate (*lr*)	1.1
Maturation Rate (*M*)	.5
Maximum Associative Strength (*mas*)	7.25
Mismatch Penalty	4

The model fits target accuracy with an $r = .83$. Interestingly, the source of errors is different between targets and foils. Target errors are due to failures in retrieving any concept-pairs, whereas foil errors are due to confusing foils with previously seen similar stimuli.

Table 2. List of Average Base-Level Activation (B_i) and Average Associative Strength per Link (S_i) per Condition across Drop-Out Training and Testing.

	Drop		*Test 1*		*Test 2*		*Test 3*	
	B_i	S_i	B_i	S_i	B_i	S_i	B_i	S_i
F2	.44	1.33	.54	.97	.73	.91	.86	.82
F4	.91	.69	.87	.55	.83	.48	.97	.44
I2	.40	1.34	.25	.99	.59	.94	.59	.86
I4	.70	.70	.11	.58	.40	.60	.60	.47
S2	.32	1.35	.11	1.03	.67	1.00	.34	.95
S4	.41	.72	.26	.61	.53	.60	.40	.49
C2	.47	1.40	.08	1.17	.25	1.05	.35	.98
C4	.44	.71	.37	.59	.60	.52	.60	.52
L2	N/A	N/A	.13	1.06	.57	1.00	.40	.92
L4	N/A	N/A	-.65	.57	.29	.60	.47	.48
M2	N/A	N/A	-.57	1.08	.42	1.01	.44	.92
M4	N/A	N/A	-.15	.61	.38	.52	.25	.51
H2	N/A	N/A	.63	3.21	.75	3.22	.78	3.20
H4	N/A	N/A	.47	2.65	.55	2.69	.55	2.71

Discussion

The present model describes the emergence of fan effects in both accuracy and latency using a theory of associative memory including an account of interference by discounting link strengths. As more stimuli (persons or places) are presented together (or within a short temporal window) they interfere with each other's associative strength, reducing overall activation. This has the effect of lowering overall accuracy and increasing response times. While prior explanations (Anderson & Reder, 1999; Schneider & Anderson, 2012) have presented good fits to static human performance, the present model learns both base-level activation (reflected recency and frequency of use) and associative weights throughout the entire experiment (including the testing phase) and predicts the presence of fan effects in both latency and accuracy across all target and foil conditions.

An advantage of modeling the fan effect experiment at this higher level of fidelity is that we are able to assess some if the assumptions made in prior modeling efforts. Most interesting is that, while the assumption that base-level would be similar between high-fan and low-fan stimuli, while this was valid (see Table 2) in aggregate, base-levels were highly variable between conditions and throughout the task. The present model was able to qualitatively match to human performance with associative activation strong enough to compensate for the differences.

While not obvious when examining end-state models, the average activation of concept-pairs between conditions (see Table 2) changes throughout the testing phase. Many models assume a fairly plastic study/learning phase and a fixed testing phase; however, our model learns throughout the experiment. For instance, the added interference from learning novel foils reduces the activation of targets, especially in the interference condition.

A potential concern that our model addresses that was foils were not encoded in Anderson and Reder (1999) when they were perceived. It seems odd that stimuli seen in training were encoded while stimuli seen in testing were not. For instance, when a novel foil is perceived it does not increment the fan of targets. For instance, if a studied fan-4 target 'the *biker* is in the *factory*' was tested after perceiving the novel foil 'the *hippie* is in the *factory*' then that fan-4 place term '*factory*' should in fact be incremented to be a fan-5 place term. In our model, the notion of fan-4 or fan-5 is solely for classification purposes of the various conditions. Link

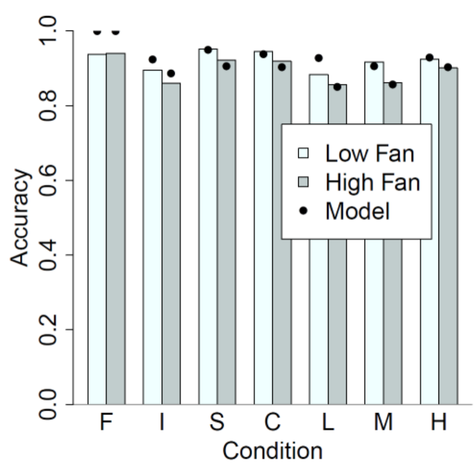

Figure 2. Error rates across target and foil conditions. Target conditions are Facilitation, Interference, Suppression, Control. Foil conditions are Low, Mixed, and High Foils.

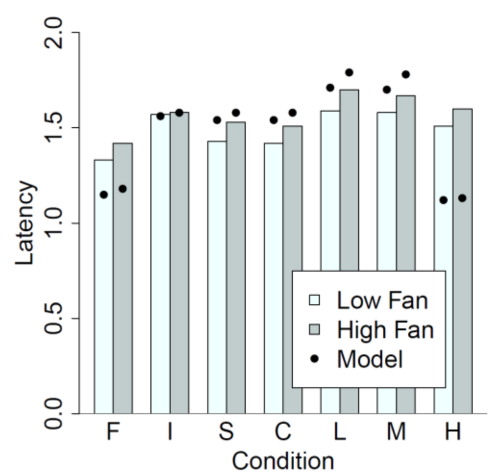

Figure 1. Latencies across target and foil conditions. Target conditions are Facilitation, Interference, Suppression, Control. Foil conditions are Low, Mixed, and High Foils.

strengths vary based on their use in the experiment. What is important is not whether an item is a fan-2 or fan-4 stimulus, but to what degree *biker* and *factory* prime the sentence (c.f., concept-pair) 'the *biker* in the *factory.*'

While the present model was able to predict fan effects in all conditions, it predicted a much faster response time for high-frequency targets and foils than humans exhibited. This was because the added frequency increased base-level activation too quickly. As seen in Table 2, associative strength was comparable between the high-frequency facilitation condition and the low-frequency interference/ suppression/control conditions. The difference was the higher base-level activation. The traditional activation equation (Equation 8) sums both base-level activation and associative strength, but it may be the case that the relative weighting of these factors changes over time based on some features of stimuli (such as relative strength, familiarity, or some other metacognitive feature). While the existing equation is well-justified in the literature, the inclusion of an adaptive frequency-based associative learning component replacing the fixed S_{ji} (Equation 7) may change the underlying balance between base-level and associative strength.

Another difference between the current model and human performance is that our model did not model speed-accuracy trade-offs reflecting the time-pressure based reward system of the original experiment. ACT-R does not have a mechanism to distinguish recognition from recall, and recall is an all-or-nothing event, thus it was not possible to have a meta-awareness of stimulus familiarity build-up throughout the retrieval process, something which could be leveraged to induce speed-accuracy tradeoffs. This speed-accuracy trade-off may result in relatively faster performance in the low-frequency conditions as participants' threshold to respond may be lower than the model's.

It is fair to argue that our model is substantially more complex than prior efforts, but we argue that this complexity is necessary to understand how fan effects arise from learning. By having our model perform the study equivalently to human participants and by having participants learn associative weights throughout the experiment, we present a model that supports our theory of associative memory and explains how associations are learned and adapt over time.

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. (1974). Retrieval of propositional information from long-term memory. *Cognitive Psychology*, 6, 451-474.

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., & Reder, L. M. (1999) The fan effect: New results and new theories. *Journal of Experimental Psychology: General*, 128, 186-197.

Anderson, J. R., Bothell, D., Lebiere, C. & Matessa, M. (1998). An integrated theory of list memory. *Journal of Memory and Language*, 38, 341-380.

Berry, M. J., & Meister, M. (1998). Refractoriness and neural precision. *The Journal of Neuroscience*, 18 (6), 2200-2211.

Eliasmith, C. (2005). A unified approach to building and controlling spiking attractor networks. *Neural Computation*, 17 (6), 1276-1314.

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.*

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.

Thomson, R., Pyke, A., Trafton, J. G., & Hiatt, L. M. (2015). An Account of Associative Learning in Memory Recall. *In Proceedings of the 37th Annual Conference of the Cognitive Science Society.* Austin, TX: Cognitive Science Society.

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),* 0-55.

Wills, C. L., Cacucci, F., Burgess, N., & O'Keefe, J. (2005). Attractor dynamics in the hippocampal representation of the local environment. *Science,* 308, 873 – 876.