

proach is capable of addressing important questions about existing models of simple dynamic decisions, though it could undoubtedly shed light on an array of related problems.

Of course, there are limitations to this approach, many of which are computational. The agents we used had only 16 nodes, 4 of which were reserved for inputs and outputs, meaning that only 12 could be used for storing (memory) and processing information. Although more nodes could be added – and certainly an accurate model of even very simple nervous systems would have many times more – this would severely slow down the steps required for evolution. It might also lead to problems that are analogous to the over-fitting that occurs when more parameters are added to a model, though this is itself a question worth exploring.

Conclusions

In this paper, we presented a computational evolution framework that could be used to examine how environments lead to different behaviors. This framework allowed us to examine the strategies that might have arisen in organisms to address the problem of dynamic decision-making, where agents receive information over time and must somehow use this input to make decisions that affect their fitness.

We found that both the evolutionary trajectory and the strategies ultimately implemented by the agents are heavily influenced by the characteristics of the choice environment, with the difficulty of the task being a particularly notable influence. More difficult environments tended to encourage the evolution of complex information integration strategies, while simple environments actually caused agents to decrease in complexity, perhaps in order to maintain simpler and more robust decision architectures. They did so despite no explicit costs for complexity, indicating that mutation load may be sufficient to limit brain size.

Finally, we discussed these results in the context of existing models of human decision-making, suggesting that both non-compensatory strategies such as fast and frugal heuristics (Gigerenzer & Todd, 1999) and complex ones such as sequential sampling (Link & Heath, 1975) may provide valid descriptions – or at least serve as useful landmarks – of the strategies implemented by evolved agents. In doing so, we provided evidence that strategy use is environment-dependent, as different decision environments led to different patterns of information use. More generally, we have shown that a computational evolution approach integrating computer science, evolutionary biology, and psychology is able to provide insights into how, why, and when different decision-making strategies evolve.

Acknowledgments

This work was supported by Michigan State University's High Performance Computing Facility and the National Science Foundation under Cooperative Agreement No. DBI-0939454 and Grant No. DGE-1424871.

References

- Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, *113*(4), 700–765.
- Brandstätter, E., Gigerenzer, G., & Hertwig, R. (2006). The priority heuristic: Making choices without trade-offs. *Psychological Review*, *113*(2), 409–432.
- Briscoe, E., & Feldman, J. (2011). Conceptual complexity and the bias/variance tradeoff. *Cognition*, *118*(1), 2–16.
- Dougherty, M. R., Franco-Watkins, A. M., & Thomas, R. (2008). Psychological plausibility of the theory of probabilistic mental models and the fast and frugal heuristics. *Psychological Review*, *115*(1), 199–211.
- Edlund, J. A., Chaumont, N., Hintze, A., Koch, C., Tononi, G., & Adami, C. (2011). Integrated information increases with fitness in the evolution of animats. *PLoS Computational Biology*, *7*(10), e1002236.
- Gigerenzer, G., & Brighton, H. (2009). Homo heuristicus: Why biased minds make better inferences. *Topics in Cognitive Science*, *1*(1), 107–143.
- Gigerenzer, G., & Todd, P. M. (1999). *Simple heuristics that make us smart*. New York, NY: Oxford University Press.
- Griffiths, T. L., & Tenenbaum, J. B. (2006). Optimal predictions in everyday cognition. *Psychological Science*, *17*(9), 767–773.
- Isler, K., & Van Schaik, C. P. (2006). Metabolic costs of brain size evolution. *Biology Letters*, *2*(4), 557–560.
- Kuzawa, C. W., Chugani, H. T., Grossman, L. I., Lipovich, L., Muzik, O., Hof, P. R., . . . Lange, N. (2014). Metabolic costs and evolutionary implications of human brain development. *Proceedings of the National Academy of Sciences*, *111*(36), 13010–13015.
- Laughlin, S. B., van Steveninck, R. R. d. R., & Anderson, J. C. (1998). The metabolic cost of neural information. *Nature Neuroscience*, *1*(1), 36–41.
- Link, S., & Heath, R. (1975). A sequential theory of psychological discrimination. *Psychometrika*, *40*(1), 77–105.
- Marsteller, L., Hintze, A., & Adami, C. (2013). The evolution of representation in simple cognitive networks. *Neural Computation*, *25*(8), 2079–2107.
- Olson, R. S., Hintze, A., Dyer, F. C., Knoester, D. B., & Adami, C. (2013). Predator confusion is sufficient to evolve swarming behaviour. *Journal of The Royal Society Interface*, *10*(85), 20130305.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*(2), 59–108.
- Von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior*. Princeton, NJ: Princeton University Press.

The fan effect in overlapping data sets and logical inference

Kam Kwok (kamhung.kwok@carleton.ca),

Robert L. West (robert.west@carleton.ca),

Matthew Kelly (MatthewKelly2@email.carleton.ca)

Institute of Cognitive Science, Carleton University, 1125 Colonel By Drive
Ottawa, Ontario, K1S 5B6, Canada

Abstract

We examine the fan effect in overlapping data sets and logical inference. Three experiments are presented and modeled using the ACT-R cognitive architecture. The results raise issues over the scope of the memories that determine the fan effect and the use of search strategies to retrieve from memory.

Keywords: ACT-R; spreading activation; fan effect

Introduction

We apply the ACT-R model of the fan effect (Anderson & Reder, 1999) to a more complex experimental paradigm. We analyze how the ACT-R memory retrieval theory is linked to higher-level cognitive processes by examining the fan effect in overlapping data sets and in logical inference. Our focus is on the memory retrieval aspect of logical inference and not on the construction of situation models (Graesser, Singer & Trabasso, 1994).

The fan effect (Anderson, 1974) is the name for a collection of experimental results showing that people are slower to identify probes as facts that they have previously learned if the elements that make up the probe are also associated with other previously learned facts. For example, if asked about the fact (a probe) that the *hippie* is in the *park*, people will be slower to confirm they have previously learned this fact if they have also previously learned that the hippie was in other locations or that other people were in the park. Overall, the more associations there are, and the slower the retrieval of the correct matching fact from memory will be. The total number of associations for a fact is referred to as the fan of the fact.

The focus of this paper is on the question of how the fan effect plays out in more complex scenarios where there are overlapping components. In real life, facts have overlapping components. For example, consider the facts that the *apple* is in the *bucket* and the *bucket* is in the *yard*. A person who learned both these facts should be able to judge the truth of “*apple* is in the *yard*” - a fact they had not previously learned. So, overlapping components can be used for logical inference (a restricted type of inference).

To experiment with this type of inference, the *complex fan paradigm* was created. In this paradigm, subjects learn a series of related and overlapping relationships and are tested on the fan effect at various intervals. Subjects completed the first three experiments in one sitting and completed the

fourth experiment ten months later. Subjects first learned a list of *objects* in various *containers* and were then tested to confirm the basic fan effect (Experiment 1). Next, the same subjects learn a list of the same *containers* in various *locations*. Following this they were again tested to confirm the basic fan effect (Experiment 2). Then, subjects were presented with a previously learned *object* in a previously learned *location* and asked if it is a true fact, based on what they had previously learned (Experiment 3). Finally, after ten months, subjects returned to the lab to learn to associate colors with the objects from Experiment 1 and were tested on that (Experiment 4).

Fan Equations

The equations used in ACT-R for calculating reaction times for the fan effect (Schneider and Anderson 2012) are shown in Equations 1-4. In ACT-R, the retrieval of a proposition from declarative memory is based on its activation level. Activation (Equation 1) is central to the model:

$$A_i = B_i + \sum_j W_j S_{ji} \quad (1)$$

A_i is the *activation* for fact i , which is the sum of the base-level activation B_i of fact i and $\sum_j W_j S_{ji}$ which is the associative activation for fact i . B_i represents the influence of practice, time-based decay, and recent accesses to fact i . In the ACT-R analytical model of the fan effect (Anderson & Reder 1999), B_i is an estimated constant. The concepts in a fact are indexed by j . S_{ji} is the associative activation strength between fact i and constituent concept j . S_{ji} is a function of how many other facts (f) are associated with the concept j . If there are f facts connected to j , f is the fan of j . The conditional probability of retrieving fact i given concept j will be $1/f$. We calculate S_{ji} as:

$$S_{ji} = S - \ln(f) \quad (2)$$

where S is a scaling parameter. W_j is the attentional weight for concept j . The sum of $W_j S$ is equal to 1. It is generally assumed that attention is equally distributed over all the concepts in the fact; therefore W_j is equal to 1 divided by m , the number of concepts in a fact:

$$W_j = 1/m \quad (3)$$

Equation (4) is the ACT-R analytical model for fan effect:

$$R_i = I + Fe^{-A_i} \quad (4)$$

The latency of recognizing a fact i (R_i) is an exponential function of its activation level (A_i). F is a scaling parameter. The value of F depends on the scale of the units of time used by the model (seconds or milliseconds). I represents the estimated time for all other productions in the model, such as probe encoding, memory retrieval, and motor operations. Based on Anderson's 1974 experimental data, the values for I and F are estimated to be 845 and 613 milliseconds (ms.) respectively and S is set to 1.45 (Anderson & Reder, 1999).

Experiment 1: Replicating the fan effect

The purpose of Experiment 1 was to replicate Anderson's (1974) fan effect result and show that the original ACT-R fan model (Anderson & Reder, 1999) can be applied accurately. For this we used Anderson & Reder's (1999) parameter values, so this was a zero parameter model. That is, we did not fit the models to the data, all model predictions for all of the experiments were calculated before the experiments were run.

Experiment 1 was not an exact replication of Anderson's (1974) fan effect task. In Anderson's task, participants learned 26 propositions of the form "person in place". Experiment 1 used propositions of the form "objects in containers" and had fewer conditions. This was due to the need to link the propositions in Experiment 1 to the other, subsequent experiments.

Method

Participants

Six male and four female volunteers were tested. All were graduate students in cognitive science.

Materials

A single integrated system for learning and testing was developed (in Python) to support the execution and data collection for the experiments. We also developed software (in Python) to make sure all of the items produced the required fans and were correctly counterbalanced, including the foil (this is difficult to do by hand and it is easy for errors to occur). The test data consisted of 30 two-term propositions that each paired an object with a container. The propositions were designed with different fan combinations: 1-1 (the object and container occur uniquely in that one sentence), 2-1, (object occurs in one other sentence), 1-2 (the container occurs in one other sentence), 2-2, and 3-1. For the recognition test there were 30 target and 30 foil probes.

Procedure

Experiment 1 consisted of three-phases. In Phase I, participants studied 30 propositions. Propositions are

presented in sequences of three interspersed with fill-in-the-blank tests. Participants need to correctly fill in all the blanks to proceed to the next set of three propositions.

Phase II was the qualification test, where participants were tested for accuracy on a fill-in-the-blank test for the entire set of test data. Participants must achieve 90% accuracy before they can proceed to the recognition test.

Phase III, the recognition test, was similar in design to the recognition test in Anderson's 1974 experiment procedure. For the recognition test, 30 target and 30 foil probes were presented to each participant; participants had to respond as quickly as possible by pressing a key labeled in green ("L" key) if he or she recognized the probe from the study set, or pressed the red labeled key ("A" key) if the probe did not belong to the study set. The test begins with a screen of instructions. It is followed by a 2 seconds fixation cue. After each key-press there is a 2 seconds pause before the next probe appears. Reaction time is measured.

Model construction

All the ACT-R models in this paper are analytical models based on the previously described ACT-R equations and the analytical approach used in Anderson and Reder (1999). These models are fully ACT-R compliant. The models were constructed with MS Excel spreadsheets and the Spreading Activation Modeling Environment (Kwok, West, 2010).

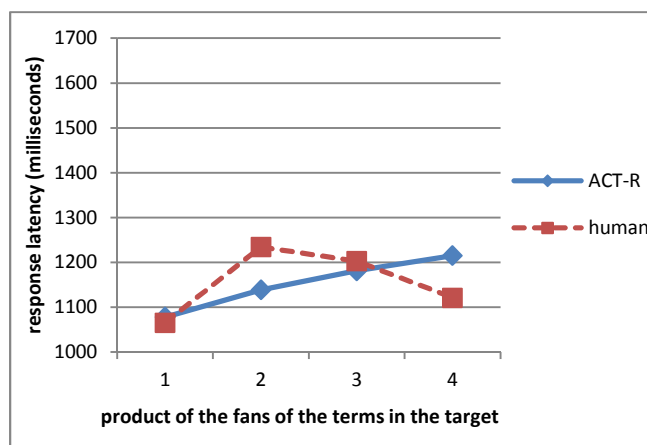


Figure 1: Model and human data for Experiment 1.

Results

Figure 1 shows the human results, with a small number of outliers removed, compared to the model predictions. The outliers were defined as RTs more than two standard deviations from the mean. All of the outliers were above the mean, indicating they were due to hesitation or second-guessing. All of the parameters in the model were taken from Anderson & Reder (1999).

Reaction time in the ACT-R model of the fan effect is a function of the product of the fans of each term that make up a proposition. For example, a two-term proposition with a term with a fan of 1 and a term with a fan of 4 is predicted to take the same amount of time to recall as a proposition

with two fans of 2 because $1 \times 4 = 2 \times 2 = 4$. Therefore, to maximize power we combine all fan combinations for which ACT-R makes the same prediction into the same condition. Participant reaction times are averaged together for all propositions where the fans of the terms that make up that proposition have the same product. In Figure 1, the conditions on the x-axis are the products of the fans of the two terms. The results show a good fit between the model and the human data demonstrating a successful replication of the fan effect. Given that this fit is obtained using Anderson and Reder's (1999) parameters without adjustment, this result demonstrates the accuracy of the ACT-R fan model.

Experiment 2: Overlapping data sets

In ACT-R, a *chunk* refers to a set of related elements, constituting a fact or proposition. According to the ACT-R theory of spreading activation, activation spreads from the constituent elements of a target chunk to the chunks stored in memory. The more spreading activation received, the higher the activation of the chunk in memory. During retrieval, the chunk with the highest activation is chosen and the retrieval time is proportional to its activation level. The amount of activation spread from an element in the target to a chunk in memory is theorized to be based on prior experience. Specifically, the amount of activation spread reflects the extent to which the element can be used to uniquely identify a chunk. For example, if prior exposure was equal (i.e. counterbalanced), a target element with a fan of 2 would have half the predictive power of a target element with a fan of 1. But, if prior exposure was not equal, e.g., the hippie was often seen in the park but only rarely in the bank, then the spreading activation cannot be determined from the fan of the target. The fan can only be used when prior experience with the co-occurrence of the elements has been counterbalanced, as in the current experiment and other fan experiments.

However, determining the size of the fan depends on the boundaries of the study set. This is because the probability that an element can uniquely identify a chunk depends on how many other chunks the element occurs in. Although we will not discuss the learning mechanism for this, any mechanism used would need to define the scope of the fan. Therefore, for Experiment 2, the scope could be all the facts in Experiment 2 (stand alone) or all the facts in Experiments 1 and 2 (combined). Note that the fan could still be used to calculate the effect because Experiment 2 alone and Experiments 1 and 2 combined were both designed to be counterbalanced. If the scope of the fan in Experiment 2 included the facts from Experiment 1 then it would raise the fan of the containers resulting in higher reaction times.

Method

In Experiment 2 the same subjects were asked to memorize a second set of 26 propositions that pair the containers from Experiment 1 with locations

Participants

The participants were the same six male and four female volunteers from Experiment 1.

Materials

The test data consisted of 26 two-term propositions. Each proposition paired a container from Experiment 1 to a location. The facts were designed with different fan combinations. For the recognition test phase there were 26 target and 24 foil probes.

Procedure and Model Construction

The procedures and model building process for Experiment 2 were the same as described in Experiment 1. Reaction time was recorded for each probe.

Results

Figures 2 and 3 respectively show the results for the stand-alone model and the combined model. A small number of outliers were removed in according to the same criterion as Experiment 1 the outliers were due to hesitation or second-guessing. As in Experiment 1, the conditions are defined by the model predictions. Specifically, the conditions are defined by the product of the fans of the two terms in a proposition. Here it is important to note that both the human data and the model predictions are different in the two graphs. The reason that the human data looks different is that boundaries for how it lines up with the two models are different. Since the two models predict different fans, the conditions differ for the two models. This can be seen in that the combined model divides the task into six conditions whereas the stand-alone model divides it into five.

The graphs clearly show a better fit for the combined model. To test this statistically, the RT scores for each subject were contrasted with the model predictions. This was done by subtracting the human RT scores from each model's predicted RT scores for each trial. This produced two *difference* scores for every trial, for every subject. The mean difference score for the stand-alone model was 358.61 and the mean difference score for the combined model was 276.20. Using a pairwise t-test we found that the difference scores for the combined model were significantly ($P < 0.001$) lower than the difference scores for the stand-alone model (note, we took a conservative approach and did not exclude the outliers from this analysis).

Experiment 3: Logical inference

Experiment 3 was conducted immediately after Experiment 2. In Experiment 3, participants were presented with probes that described an object in a location and they were asked to respond if it was true or not. Explicit instructions were given to the participants that the objects and locations were the same ones they learned about in Experiment 1 and 2 and that they can determine the answer by retrieving which container the object is in and then checking if the container is in the location.

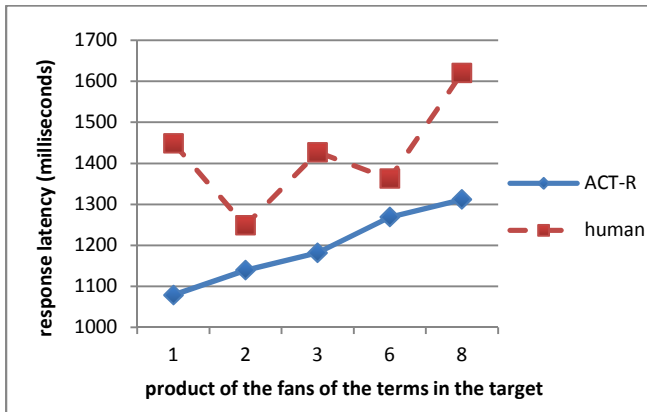


Figure 2. The stand-alone model predictions and the human data for Experiment 2.

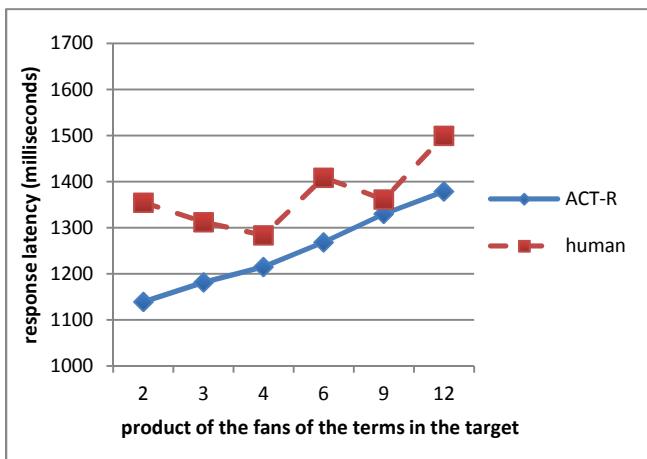


Figure 3. The combined model predictions and human data for Experiment 2.

Method

Participants

The participants were the same six male and four female volunteers from experiments 1 and 2.

Materials

The test data consisted of 24 targets and 25 foils for a total of 49 probes for the inference test.

Procedure

The procedure for Experiment 3 was the same as the recognition test procedure described in Experiment 1. Reaction time was measured for each probe.

Models

ACT-R has two ways to retrieve facts from memory. The first is to retrieve a match for a fact that has been presented (as in the models for Experiments 1 and 2). The second is to

construct a query by using a partial fact as a cue to retrieve a complete fact. For example, if the *apple* was in the *bucket*, the retrieval cue *apple ?* would retrieve *apple bucket*. In this case the fan is based on the elements present in the query (so the fan of *apple*).

One way to extend the ACT-R fan model to model Experiment 3 is to use a *query* to first retrieve a container for the object and then check to see if it's in the location. For example, if the target fact is *apple* in the *yard*, use *apple ?* to retrieve *apple bucket*, then use *bucket yard* to retrieve a match. The time for this would be the sum of the two retrieval times, as determined by the fan, plus the time for the associated productions (50 ms each) to direct the actions, plus the times for perceptual and motor actions. We will refer to this as the *dual retrieval model*.

However, although the facts were learned separately in Experiments 1 and 2, that does not necessarily mean they were stored separately. Subjects could have realized that the new information learned in Experiment 2 was related to the information learned in Experiment 1, leading them to store the new information as three element chunks (i.e., *object container location*). In this case, both the *object* and the *location* can be used as a cue to retrieve a *container*. This would require only one retrieval and, since activation would spread from both the object and the location, this single retrieval would be faster than the query retrieval in the dual retrieval model. We will refer to this as the *single retrieval model*.

The single retrieval model follows the logic of other ACT-R fan models so only one retrieval is needed to determine if the probe is a target or a foil. In the dual retrieval model, the situation is more complex. Anderson and Reder's (1999) foil identification strategy was meant for *matching* and the logic would not apply to *queries*. Therefore, in the dual retrieval model, there is no way to know if the first query retrieved the right container until the result of the subsequent match is evaluated. Under these conditions it makes sense to consider an ACT-R model that uses a search strategy. That is, if it fails to make a match it goes back and tries a different retrieval until either it correctly identifies the probe as a target or the search is exhausted.

The single retrieval model predicts a faster retrieval time than the dual retrieval model. We did not model the search strategy but instead modeled the dual retrieval process as if subjects always retrieved the correct chunk on the first query. This produced a baseline of the fastest possible time for responding. We expected that the search strategy would add time beyond this, with the exception of the lowest fan condition. This is because a fan of 1 guarantees that the initial query will produce the correct result.

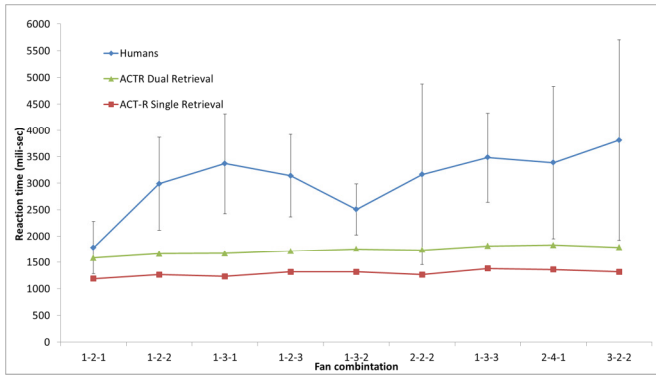


Figure 4. Human data (blue) from the logical inference task, compared to results from dual retrieval (green) and the single retrieval (red) model. The confidence intervals on the human data are set at 95%.

Results

As in Experiments 1 and 2, all of the models were constructed using the original perception/motor estimates from Anderson and Reder (1999), along with the times for the minimum necessary productions (Anderson and Reder's original model combined the times for perception, motor actions and production firing, however, it is possible to separate these out). Figure 4 shows the predictions of the single retrieval model and the dual retrieval model compared to the human data. The results indicate that the single retrieval model is too fast, whereas the dual retrieval model baseline prediction is plausible. As predicted, it is a close match on the lowest fan condition and slower on the other conditions that would involve additional search times if the first retrieval was incorrect. Since search is unnecessary for the single retrieval model we conclude that the dual retrieval model *plus search* is the more likely to be correct. Overall, this result is consistent with the hypothesis that people store the individual pairings they have learned (i.e., object-container and container-location) as separate chunks and do not construct integrated chunks (i.e., object-container-location)

Experiment 4: Long Term Memory

Experiment 4 was done 10 months after Experiment 3, using the same subjects. None of the subjects were aware that this would happen. In Experiment 4 the list of objects that the subjects had previously learned were paired with colors (e.g., the red pen). They were then tested for the fan effect in exactly the same way as in the other experiments. As in Experiment 2 we could construct two models, a stand-alone model, that assumed only what they studied and were tested on was relevant for calculating the fan effect, and a combined model, that assumed that the previously learned information about what containers the objects were in was still relevant for calculating the fan.

Results

Figures 5 and 6 display the results. Outliers were not an issue in this experiment so no trials were removed. Also, note that although the number of conditions was the same in each model, the internal boundaries between conditions were different across the models.

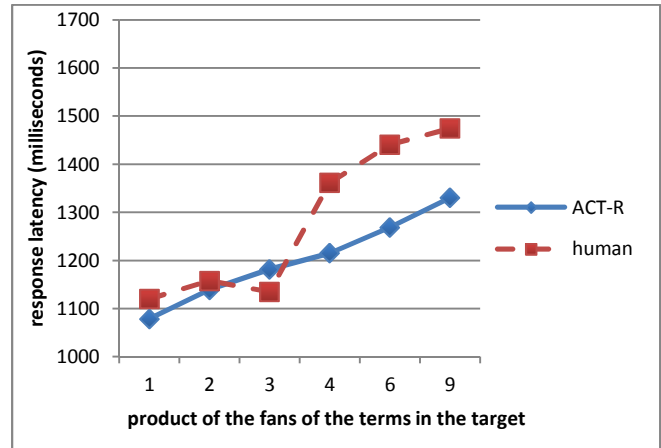


Figure 5. The stand-alone model predictions and the human data for Experiment 4.

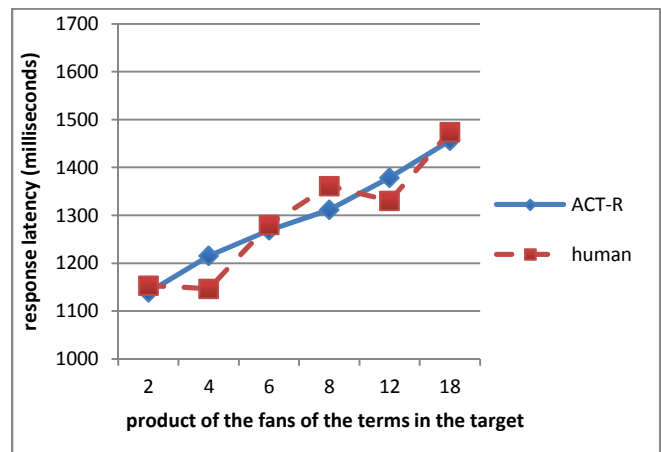


Figure 6. The combined model predictions and the human data for Experiment 4.

To test this difference we used the same method as in Experiment 2 to generate and test difference scores. The mean difference score for the stand-alone model was 72.69 and the mean difference score for the combined model was -11.65. These results were significantly different according to a pairwise t-test ($P=0.001$), and indicate that facts learned 10 months ago were still relevant in the retrieval process.

Discussion

The results of these experiments results confirm that a zero parameter ACT-R fan model can be used to accurately predict the results of different types of fan experiments (see

also West et al, 2010). However, the results of Experiments 2 and 4 show that recently learned facts and even facts learned 10 months earlier contribute to the fan effect as though they had been part of the learning set. In contrast, when ACT-R fan models assume that facts can be treated as counterbalanced because the material presented in the experiment is counterbalanced, it creates an implicit assumption that there is no effect of background experience on the fan. For example, the fact that you may have had more prior exposure to hippies in parks than in banks is not taken into consideration. Likewise, the fact that you may have seen more different types of people in parks than in banks (i.e., park has a higher real life fan than bank) is also not taken into consideration. Given that real life experience is much more extensive than in the experiments, real life experience should dominate and the assumption of counterbalanced exposure should not hold. One reason that it still works could be that the counter-balanced design of the test data corresponds roughly to the average human experience. Another possibility is that while the scope of the fan effect extends beyond individual experiments, it may still only apply to a limited set of data defined by the context of being in a set of related experiments. That is, there is a powerful effect of context that is undiminished with time. However, this would be problematic to model in the current version of the ACT-R fan model. More research is needed to explore these issues.

However, more generally, the results of Experiments 2 and 4 indicate that the strengths of association do not automatically decay with time. However, since it is unlikely that subjects thought about the facts between experiments it is still possible that in real life strengths of association can be eroded through interference. Therefore, these results are supportive of an interference-based account of memory as opposed to a decay-based account. This finding is consistent with other memory studies that found an effect of stimuli learned a year ago (e.g., Kolers, 1976; Salasoo, Shiffrin, & Feustel, 1985). The current ACT-R account of spreading activation is consistent with the interference view (although the ACT-R account of the effect of the passage of time in memory recall experiments is not consistent with this view).

The results of Experiment 3 show how the fan effect plays a role in logical inference. These results suggest that people do not combine logically related facts at the time of encoding but instead used sequential retrievals to do logical inference. This does not mean that people cannot combine logically related facts in memory. In fact, it seems clear that this does happen. However, these results indicate that it does not happen automatically. Experiment 3 also indicates that people use a search strategy to find the appropriate overlapping chunks to do logical inference. As we can see from the first point in Figure 4, when fan=1 there is no search. Arguably, the fan also plays an important role in determining the length of the search. In figure 4, the longest search times occur when the object and location both have a fan of 2 or more. Although we instructed subjects to first use the object to recall the container, this result suggests that

at least some subjects started by using the location to recall the container. The most efficient search strategy is to remember which objects and locations had a fan of 1 and to start the search with them. Some subjects may have used or partially used this strategy. Unfortunately, given the high variability in the results, it is possible that subjects were using different strategies from each other, in which case the averaged data is of limited use.

References

- Anderson, J. R. (1974). Retrieval of propositional information from long-term memory. *Cognitive Psychology*, 5, 451-474.
- Anderson, J. R. (1993). Rules of the mind. Hillsdale, NJ: Erlbaum.
- Anderson, J. R. (2007). How can the human mind occur in the physical universe. New York, NY, US: Oxford University Press.
- Anderson, J. R., Bothell, D., Lebiere C., & Matessa, M. (1998). An Integrated Theory of List Memory. *Journal of Memory And Language*, 38, 341-380
- 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., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review*, 111(4), 1036-60.
- Graesser, A. C., Singer, M., & Trabasso, T. (1994). Constructing inferences during narrative text comprehension. *Psychological Review*, 101(3), 371-395.
- Kolers, P. A. (1976). Reading a year later. *Journal of Experimental Psychology: Human Learning and Memory*, 2(5), 554-565.
- Kwok, K., & West, R. (2013). SAME: An ACT-R Spreading Activation Modeling Environment. *Iccm-Conference.org*, 366-367. Retrieved from <http://iccm-conference.org/2013-proceedings/papers/0065/paper0065.pdf>
- Salasoo, A., Shiffrin, R. M., & Feustel, T. C. (1985). Building permanent memory codes: Codification and repetition effects in word identification, *Journal of Experimental Psychology: General*, 114(1), 50-77.
- Schneider, D. W., & Anderson, J. R. (2012). Modeling fan effects on the time course of associative recognition. *Cognitive psychology*, 64(3), 127-160.
- West, R. L., Pyke, A. A., Rutledge-Taylor, M. F., & Lang, H. (2010). Interference and ACT-R: New evidence from the fan effect. In D. D. Salvucci & G. Gunzelmann (Eds.), *Proceedings of the 10th International Conference on Cognitive Modeling* (pp. 277-281). Philadelphia, PA: Drexel University.