# Interference and ACT-R: New evidence from the fan effect

**Robert L. West (robert_west@carleton.ca)**
Institute of Cognitive Science, Department of Psychology, Carleton University, 1125 Colonel By Drive
Ottawa, Ontario, K1S 5B6, Canada

**Aryn A. Pyke (apyke@ccs.carleton.ca)**
Department of Psychology, Carleton University, 1125 Colonel By Drive
Ottawa, Ontario, K1S 5B6, Canada

**Matthew F. Rutledge-Taylor (mrtaylo2@connect.carleton.ca)**
Institute of Cognitive Science, Carleton University, 1125 Colonel By Drive
Ottawa, Ontario, K1S 5B6, Canada

**Hana Lang (hlang@connect.carleton.ca)**
Institute of Cognitive Science, Carleton University, 1125 Colonel By Drive
Ottawa, Ontario, K1S 5B6, Canada

**Abstract**

We present data demonstrating that interference plays a role in the fan effect. We also show that this cannot be accounted for using ACT-R. An ACT-R model is fit to the data and we discuss options for altering the model to account for the data.

**Keywords:** interference; fan; spreading activation; ACT-R; memory; cognitive model.

## Introduction

The fan effect refers to the fact that cues that are associated with more facts result in slower recall than cues that are associated with less facts. For example, in the study that established the fan effect, Anderson (1974) asked subjects to memorize facts about where various different characters had been seen. Subjects were then shown a cue with a character and a place and asked if it was true (i.e., if they occurred together in the set of facts subjects had memorized). Overall, the more places a character had been, the slower subjects were to confirm that it was either true or false. Also, subjects were slower to say false than they were to say true.

In the ACT-R architecture (Anderson & Lebiere 1998) the cue is held in a buffer as a chunk and each slot value of the cue spreads activation to chunks in declarative memory that have the same slot values. For example, if the cue chunk is person:hippy location:park, then hippy will spread activation to all chunks that have hippy as a slot value and park will spread activation to all chunks that have park as a slot value (note, the slot names do not play a role). The number of lines of activation leaving from a slot value in the cue is the fan of that slot value, and the fan of the cue is the sum of the fans of its slot values.

In ACT-R, the amount of activation spread from a cue to a chunk is theorized to be proportional to the number of past associations between slot values of the cue and the chunk. In the ACT-R architecture, the way of calculating this is based on an assumption that exposure to different facts has been counterbalanced, as in a psychology experiment (Anderson & Reder, 1999). If it is assumed that everything has been counterbalanced and the number of exposures per chunk is equal then the activation can just be divided evenly among the chunks. So, the higher the fan the lower the amount of activation delivered to each individual chunk (see Anderson & Reder, 1999, for how to use ACT-R when exposures have not been counterbalanced). Anderson and Reder (1999) modeled the fan effect in ACT-R by assuming that people retrieve the most active chunk and respond *true* (i.e., they have seen it before) if the retrieved chunk matches the cue, and *false* (i.e., they haven't seen it before) if it does not.

One consequence of this model is that only the spreading activation received by the chunk that is chosen affects the reaction time (RT). In other words, there is no interference from the activation of other chunks. However, as fan goes up so do the number of other chunks that receive activation. As part of a fan experiment we tested the effect of these "other" chunks to see if interference plays a role in the fan effect and how that might alter the ACT-R fan model.

### Experimental Design

In our experiment we created false cues by taking a true fact and replacing one element with a false element. For example, if subjects had studied the fact that the *red hat is in the kitchen*, we could create a false cue by replacing *hat* with *pen*. Under these conditions the ACT-R model predicts that the fan of the false element of the cue will have no effect on retrieval time, unless the original fact is not retrieved. However, we performed simulations with the ACT-R fan model and found that in our experimental design, the chunk representing the original version of the fact always received the most activation, and therefore was always chosen (as far as we can see this is also true for other fan experiment designs, but it is possible to create more extreme differences in fan where this would not be true). Related to this, the fan of the false element should also have no effect on the error rates. Although the ACT-R fan model does not explicitly model errors, errors would be due to noise and the retrieval threshold. This could conceivably interact with fan for the chunk that is being retrieved but the fan of the false element does not affect this chunk.

## Method

### Subjects

Twenty eight participants (11 males and 16 females: mean age 19.9 years, SD = 2.2) were recruited from introductory psychology courses a Carleton University to take part in the experiment. Participants received course credit as compensation for their time.

### Procedure

The experiment was divided into three main phases: A study phase, a recall phase and a recognition phase. In the study phase each participant was assigned one of three unique sets of study sentences and was instructed to memorize the sentences in the list. Once the participant indicated that they were prepared to proceed, the recall portion of the experiment began.

The study set consisted of sixteen sentences of the form, "The *color thing* is in the *place*". The color term was one of ten colors; the thing was one of ten house-hold items; and the place was one of ten locations in/around a typical home. Very typical item/locations combinations, such as 'comb'/'bathroom', were not used in generating the study set sentences. Each term could have a fan of either 1 or 4. Thus, the four possible sentence fans were: 3, 6 9, and 12.

In the recall phase each participant was tested three times. Each test began with the participant trading the study set with the experimenter for a new list of sentences identical to the study set, but with one term from each sentence replaced with a blank, and the order of the sentences randomized. The participant's task was to correctly fill-in each of the blanks with the missing word. The participant was given as much time as he or she needed. Once finished, the experimenter recorded the number of correct responses and for each error, provided the correct missing word to the participant. The participant was then given the opportunity to review the study set before being tested again. The three tests were balanced such that each term from each sentence in the study set was replaced with a blank exactly once. After the third iteration the recognition phase began.

The recognition phase of the experiment was conducted on a computer using Experiment Builder (by SR Research). The participant's task was to correctly judge whether sentences presented in the middle of a 17" CRT display were members of the study set, or not. Accuracy and reaction time data were recorded for each trial.

Each participant was presented with 96 test sentences, which consisted of three exposures to each of the study set sentences, and 48 sentences that were not from the study set. The participants were told that they should consider sentences from the study set to be *true*, while all others should be considered *false*. Each false sentence was generated by swapping one of the three terms from a true sentence with another term from the same category (e.g., color, thing, or place) and with the same fan. Each true sentence was used to generate three different false sentences. Thus, for each exposure to a true sentence there was a false sentence with the identical fan. Once the test sentence appeared the participant would indicate if the sentence was in the study set by hitting the 'z' key, or if it was not by hitting the '/' key.

## Results

The data from one of the participants was excluded from the results presented below. This was due to the fact that this participant's performance was significantly poorer than all other participants by a large margin ($P < 0.001$). The results below reflect the data collected from the remaining 27 participants. By the end of the third iteration of the recall phase the participants were correctly completing the sentences 91.4 percent of the time. The results of the recognition phase replicated the fan effect. These results are reported in Rutledge-Taylor, Pyke, West, & Lang (2010). However, in this paper we will focus on the results related to the predictions described above and fitting an ACT-R model to the data.

The hallmark of a good scientific theory is that it makes precise, falsifiable predictions. Many theories in Psychology fail to meet this criterion. However, because ACT-R is precisely specified, models built in ACT-R are more readily falsifiable. To test the predictions of the ACT-R fan model (Anderson & Reder, 1999) concerning the fan of the false items we ran an ANOVA testing for the effect of the fan of the false items on RT and error rate. RT was significantly higher when the fan of the false item was equal to 4 than when it was equal to 1 ($P<0.001$). The error rate was also significantly higher when the fan of the false item was equal to 4 than when it was equal to 1 ($P<0.001$). We also ran a correlation between the fan of the false items and RT, with the fans of the true items partialled out. We found a significant correlation of $r=0.156$ ($P<0.001$, one tailed). Similarly, we ran a correlation between the fan of the false items and % errors, with the fans of the true items partialled out. Here we also found a significant correlation of $r=0.193$ ($P<0.001$, one tailed). The size of these correlations was roughly similar to the same correlations done with the fan of the true items.

Contrary to the predictions of the ACT-R fan model, we found that the fan of the false items significantly affected RT such that a larger fan led to slower responses (see Figure 1). Consistent with this and also contrary to the predictions of the model, we found that the fan of the false element significantly affected the error rate such that a larger fan led to more errors (see Figure 2). These results indicate that interference from the false item plays a role in the fan effect.

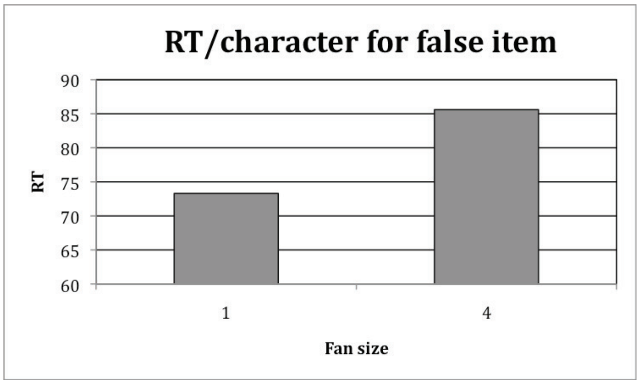


Figure 1: Reaction time in msec/character for responding false to a false cue as a function of the fan of the false item in the cue.

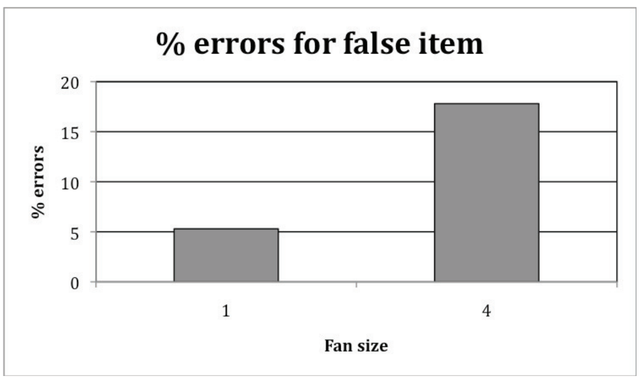


Figure 2: Percent errors for responding false to a false cue as a function of the fan of the false item in the cue.

.

## Model Evaluation

Although falsification of a model is sometimes viewed as a bad thing, falsification actually shows that a model was well specified. Falsification also creates an opportunity to move toward a better model. To this end we fit the ACT-R fan model to our data. Anderson and Reder (1999) used the following function to calculate RT:

$$RT=I+Fe^{-Ai}$$

Where $F$ is a scaling constant for time, $I$ is a constant representing how long it takes subjects to make their response after they know it, $e$ is the base for natural logarithms and $A$ is the activation of the chunk (which includes spreading activation). Activation was calculated as:

$$A=B+S$$

Where $B$ is base level activation and $S$ is spreading activation. We fitted the Anderson and Reder (1999) ACT-R fan model to our data using identical parameter values, except that we had to increase $S$ slightly from 1.45 to 2 to avoid getting negative activation values. As in Anderson and Reder (1999), $B$ was set to zero because it trades off with $S$.

We eventually figured out that the slight difference in $S$ was because we used the current method of calculating fan size in ACT-R 6, which is to add 1 to the fan of each item in the cue to represent the match between that item and a chunk in memory representing that item. For example, 1 would be added to the fan of cup because it is assumed that we all have a chunk in declarative memory representing cup. In contrast, Anderson and Reder (1999) calculated the results without adding 1 to fans of the items in the cue. Whether or not to do this is an interesting issue. However, we recalculated our results without adding 1 and found it made very little difference to our results or conclusions.

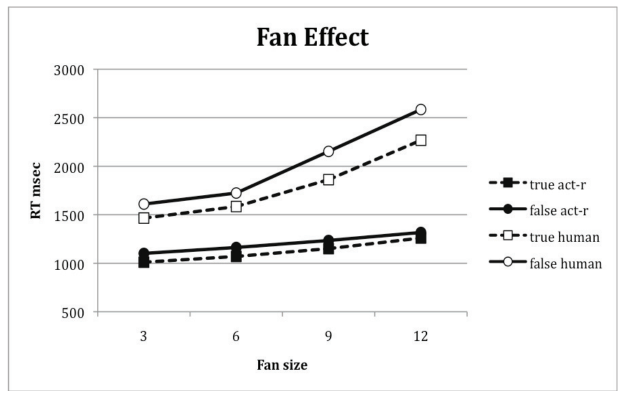


Figure 3: The original Anderson & Reder (1999) ACT-R fan model fit to our data.

Figure 3 shows the fit of the original ACT-R fan model to our data. The fact that the model predicts an overall lower RT is not significant as it can be accounted for by assuming our subjects took longer to press the true/false keys, which can be modeled by increasing the $I$ parameter. However, the shape of the functions and the relationship between the functions is clearly different. The human data shows a clear upward curve that the model does not and the model RTs converge as fan increases wile the human data diverges.

Figure 4 shows the original fan effect data from Anderson and Reder (1999) re-plotted. Note that it shows the same divergence and upward curve. In fact, the original ACT-R model for this data (faithfully recreated and shown in Figure 5) also shows a slight upward curve for the *true* cues, but not for the *false* cues. Also, as with our data, the *false* function diverges from the *true* function as fan goes up. However, it is important to keep in mind the scale of the graphs and realize that these effects are much smaller in the Anderson and Reder (1999) data and may not even be real, although, the consistency of this result across conditions and studies indicates that we should take it seriously.

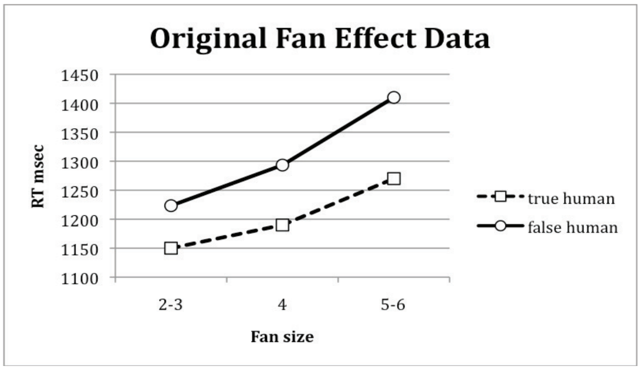


Figure 4. Re-plotted data from Anderson and Reder (1999)

## An Alternative Model

Next we addressed the issue of the parameter values. Specifically, we wanted to know if the ACT-R model could be made to fit the data. The only way that we could find to fit the data was to use the latency exponent parameter (*f*) that is available in ACT-R 6. This parameter, which has rarely been used, changes the RT function to:

$$RT=I+Fe^{-(f*A)}$$

By setting *f=3* and increasing *F* from *613* to *2000* we obtained a good fit to the data (see Figure 6 - note, that the *I* parameter could be increased to overlap the functions but it is easier to see this way). Increasing *f* lowers overall RT, so increasing *F* can be viewed as a way of compensating for this. The other effect of raising *f* was to increase the acceleration of the rate at which lowering activation raised RT. We will refer to this as the ACT-R(*f*) model (see Figure 6). However, please note that this model violates the ACT-R modeling convention of using establish parameter values unless you have a justification (Anderson & Lebiere, 1998).

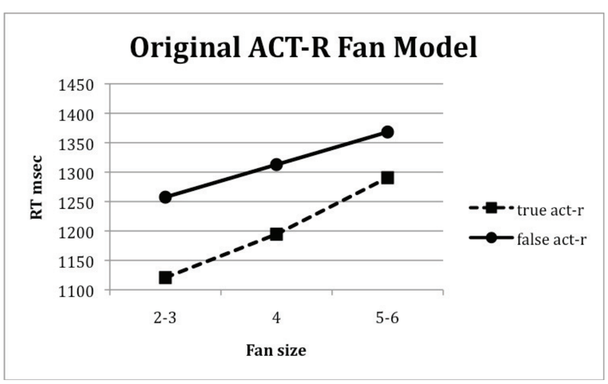


Figure 5. A re-creation of the ACT-R fan model from Anderson and Reder (1999)

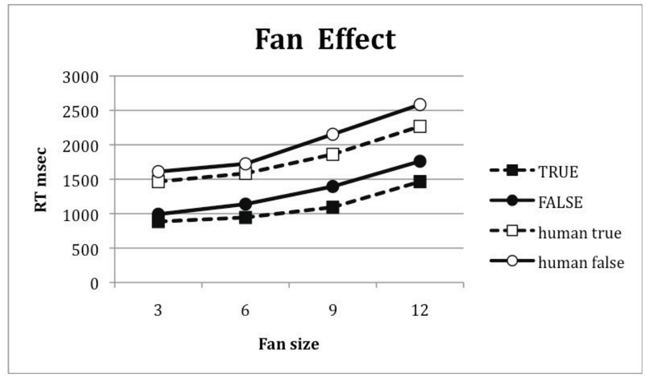


Figure 6: The ACT-R(f) model fit to our data (note, that the *I* parameter could be increased to overlap the functions but it is easier to see this way).

**Rationalizing the alternative model**

There are three ways we can interpret the ACT-R(*f*) fan model. We know that it cannot account for our finding that the fan of the false item in the cue affects RT and % error any better than the normal ACT-R fan model. However, it is possible to interpret the manipulation of *f* as representing the aggregate effect of interference. In this case, *f* would be related to the total effect of interference in the task. If we assume that our use of more cue items and higher fans produced greater overall levels of interference, then the fact that our results show a more pronounced nonlinear effect than the Anderson and Reder (1999) results could be modeled by increasing *f* to represent higher levels of interference. In this sense, ACT-R could be adjusted to account for the presence of interference but could not be said to include a (process) model of interference. More studies would be required to see if *f* actually does function this way.

A less charitable approach to understanding the ACT-R(*f*) model would be to point out that adding *f* to a model that already has a lot of parameters creates a system capable of fitting a lot of different functions. We had no principled reason to adjust *f* and found that it worked as part of a parameter tweaking process that involved all of the available parameters. So possibly the fit of this model is merely fortuitous.

A third, more constructive way of viewing it is to see the manipulation of *f* as a proxy for an additional mechanism or process - in this case, interference. Although ACT-R does not include an interference mechanism, modifications have been introduced to do this. For example, the spacing effect modification of Pavlik and Anderson (2005) assumes that interference plays a role in order to account for the spacing effect in memory. Similarly, the semantic interference modification proposed by Van Maanen & Van Rijn (2007) assumes a form of interference to account for the Stroop effect. Likewise, our findings indicate the need for an explicit model of interference in ACT-R. A simple way of doing this that is consistent with our manipulation of *f* is to introduce a penalty that reduces activation based on the total fan of the information in the cue – the higher the overall fan, the greater the penalty for all chunks receiving spreading activation. We could create such a function but it would not be meaningful at this point since it would be custom made to fit our data. Essentially, this would have the same effect as raising *f*, but the effect would be tied to the overall fan and therefore would account for our finding that the fan of the false item in the cue affects RT and % error.

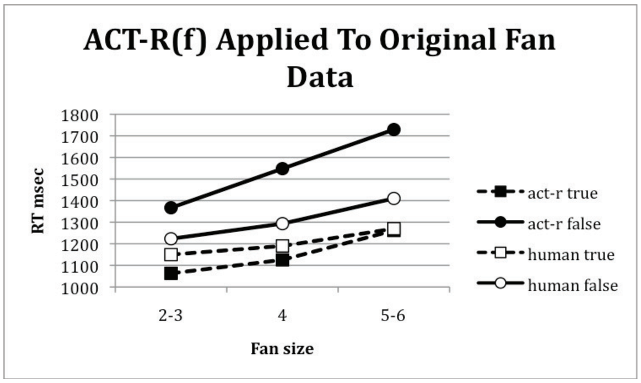


Figure 7. The ACT-R(f) model applied to the data from Anderson and Reder (1999).

### Model Re-Evaluation

To gain further insight into the ACT-R(*f*) model we applied those parameter settings to our recreation of the Anderson and Reder (1999) fan model. The results are illustrated in Figure 7. The *true* results actually produce a reasonable fit to the data but the *false* results clearly do not fit. This could be because the fit of the ACT-R(*f*) model to our data was merely fortuitous, or it could be because higher *f* values are only appropriate when interference is higher, as suggested above.

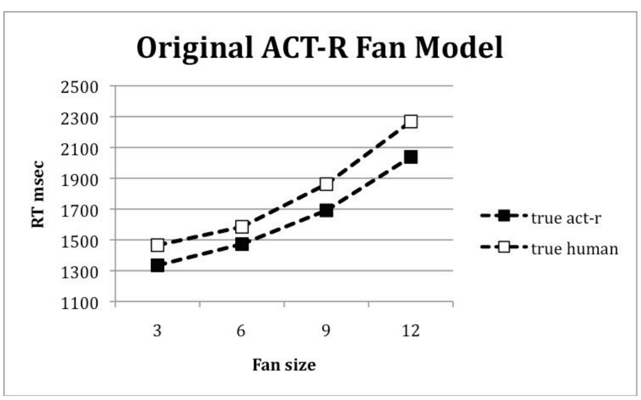


Figure 8. The Anderson and Reder (1999) ACT-R fan model fit to our data for correctly identifying true cues only.

Based on our experimental findings showing that the ACT-R fan model for correctly identifying *false* cues cannot be correct, we also tried fitting Anderson and Reder's (1999) fan model to our data for the *true* results only (see Figure 8). Without having to fit the *false* data we were able to get a good fit by adjusting only *F* and *I* (*F*=1000; *I*=1100; *S*=1.45; similar to Anderson and Reder we did not add 1 when calculating the fan). This is much less problematic because it avoids adjusting *f*, which is almost never altered in ACT-R modeling. Also, it is important to remember that there is variability associated with the human data so it is likely that a single intermediate value of *F* could be used to obtain a reasonable fit to our data and Anderson and Reder's (1999) data.

## Conclusions

Our results show that the ACT-R fan model for correctly identifying false cues cannot be completely correct. Also, although fitting ACT-R to our data was possible, it was also problematic because it required unprecedented alterations to the parameter values as well as assumptions about the meaning of those alterations that remain untested. However, when we did *not* try to fit the ACT-R fan model for correctly identifying false cues, the ACT-R fan model for correctly identifying true cues fit our data well, without any problematic parameter alterations. Based on this, it appears most likely that the problem lies with the assumptions and processes behind the ACT-R fan model for correctly identifying false cues.

## References


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

Anderson, J. R. & Lebiere, C. (1998). *The atomic components of thought.* Mahwah, NJ: Erlbaum.

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

Pavlik, P. I., Jr., & Anderson, J. R. ( 2005). Practice and forgetting effects on vocabulary memory: An activation-based model of the spacing effect. *Cognitive Science, 29*, 559-586.

Rutledge-Taylor, Pyke, West, & Lang (2010). Modeling a three term fan effect. *Proceedings of the International Conference on Cognitive Modeling*.

Van Maanen, L., & Van Rijn, H. (2007). An accumulator model of semantic interference. Cognitive Systems Research, 8(3), 174-181.

# An Online Database of ACT-R Parameters: Towards a Transparent Community-based Approach to Model Development

**Tsunhin John Wong**[1] (jwong@mpib-berlin.mpg.de)

**Edward T. Cokely**[1,2] (cokely@mpib-berlin.mpg.de)

**Lael J. Schooler**[1] (schooler@mpib-berlin.mpg.de)

Max Planck Institute for Human Development, Lentzeallee 94 14195 Berlin, Germany[1]
Michigan Technological University, Cognitive & Learning Sciences, Townsend Drive, Houghton, MI 49931[2]

**Abstract**

We present a database that provides an interface for the ACT-R modeling community to interact with each other (http://www-abc.mpib-berlin.mpg.de/actrdb/). The database includes estimated values of ACT-R parameters from a wide range of ACT-R modeling studies, selected from the studies available on the ACT-R website. It serves as a tool to query studies and estimated values for ACT-R parameters, providing the exact range of values for each of the available free numerical parameters. In short, the database supports an alternative community-based approach to manage the challenges associated with parameter estimation for complex cognitive architectures like ACT-R.

**Keywords:** ACT-R; modeling; parameter.

## Managing Parameters for ACT-R Models

Unified theories of cognition allow us to approach mechanisms of human cognition in a holistic, cumulative manner (Simon & Newell, 1973). Among the existing unified theories of cognition, ACT-R is one of the most widely used architectures, producing the largest body of sustained research and application. In order to study a wide range of cognitive mechanisms, ACT-R includes a variety of modifiable parameters. While these parameters enable flexibility they also result in fundamental challenges.

Wexler (1978) criticized the early framework of the ACT research program (Anderson, 1976), stating that "There is no explanatory power in ACT because there are no restrictions on human abilities". He also posited that "the general problem with ACT is (its flexibility), it is simply so weak that there is no way to find evidence for or against it". About twenty years later, Pashler and Roberts (2000, 2002) again brought these concerns to the fore, arguing that the practice of using good fits as major evidence for complex theories is "rotten to core". Indeed, goodness-of-fit metrics remain a very common means of model validation. These concerns not only hold when criticizing ACT-R and some other unified models, but also address a wide-spread misuse of goodness-of-fits as key evidence in psychology. Sound scientific theory requires that models not only fit but also predict data (Gigerenzer, 1998; Gigerenzer & Brighton, 2009). How can modelers of the ACT-R architecture deal with these concerns about parameter estimation and model fitting?

There have been some attempts to understand the relation among ACT-R parameters that result from parameter fitting. For example, Anderson, Bothell, Lebiere, and Matessa (1998) suggested that there is a systematic linear relationship between the estimated values of activation thresholds and the logarithm of estimated latency factors. Their data also implied that estimated values of these parameters are exceedingly regular. To date, however, there has been no meta-analytic assessment to evaluate whether there is any sustained regularity of these estimated parameters for ACT-R models across other published studies.

Computational cognitive models are often evaluated by their fit and generalizability. These properties of a model are related to two aspects of model complexity: (1) number of parameters and (2) the functional forms of computation. In part, such evaluations seek to evaluate the extent to which noise is unnecessarily captured (Pitt, Myung, & Zhang, 2002; Oaksford, 2002). Using cross-validation, Taatgen, van Rijn, and Anderson (2007) estimated parameters of a base-model once and then made use of these estimated values throughout subsequent models. This study exemplifies a strict practice that allows minimal parameter estimation; however, like many ACT-R studies, the work of Taatgen et al. still relied on superior goodness-of-fits as the major support for their proposed models.

The latest ACT-R architecture version 6.0 has 62 free parameters with numerical values, together with the flexibility of mapping these parameters to tailor-made handlers and tens of other non-numerical parameters. Different instantiations of specific ACT-R models do not typically require setting and optimizing all these numerical parameters, as default values are provided. However, our analyses of a large and representative sample of ACT-R studies indicates that on average each ACT-R model modifies nearly six free numerical parameters for better model fitting. Moreover, many of these studies added task-specific parameters.