Modeling Statistical Learning and Response Inhibition with the Change Signal Task

L. Richard Moore Jr. (larry.moore@mesa.afmc.af.mil)
Lockheed Martin
Air Force Research Laboratory
Mesa, AZ 85212 USA

Glenn Gunzelmann (glenn.gunzelmann@mesa.afmc.af.mil)
Air Force Research Laboratory
Mesa, AZ 85212 USA

Joshua W. Brown (jwmbrown@indiana.edu)
Dept. of Psychological & Brain Sciences, Indiana University
Bloomington, IN 47405 USA

Abstract

The change signal task is a two alternative forced choice task with the addition of a *change signal* presented on 1/3 of the trials at some delay relative to the initial stimulus. The change signal indicates to participants that they should inhibit their initial choice and respond with the other alternative. It provides an opportunity to examine the cognitive mechanisms involved in statistical learning and response inhibition. Within the task, change signal delays are associated with stimulus color, and are adjusted independently with a step function to produce high and low error conditions. Observed data show a significant difference in reaction times between these two conditions. In this paper we propose a model for the change signal task that leverages existing declarative memory mechanisms in ACT-R as a surrogate for the implicit contextual learning observed in human trials. We compare the mechanisms in this model briefly to an existing neural account, and use the model to predict the consequences of cue-conditional reversal.

Keywords: response inhibition; statistical learning; declarative memory; ACT-R.

Introduction

Executive control of behavior is a defining component of high-level cognition. One aspect of executive control, response inhibition, has been explored extensively using the stop signal paradigm. The classic task from Logan and Cowan (1984) visually presented subjects with one of four letters, which the subjects then classified into groups by pressing one of two buttons. On 25% of the trials an audible tone signaled that they should inhibit their response. The probability of responding was related to the timing of the *stop signal* (with a greater chance of inhibition with shorter delays) and so the authors proposed a "horse race" model for resolving executive conflict.

Brown and Braver (2005) extended the stop signal paradigm to assess error-likelihood effects. In their *change signal* task, a two alternative forced choice task is presented. On one third of the trials, however, a second stimulus is presented at some delay following the original stimulus. The second stimulus – the change signal – indicates to subjects that they should inhibit their response to the original stimulus and respond with the other alternative instead. To ensure a particular error rate in the task, the delay between the initial stimulus and the change signal is manipulated.

In Brown and Braver (2005), two colors were used for the stimuli, each of which was associated with a different target error rate. They collected fMRI data from participants across the four stimulus conditions (i.e., Change versus No Change trials crossed with High versus Low error probability) to evaluate two alternative models of anterior cingulate cortex (ACC) function. The successful model, known as the error-likelihood model, correctly predicted a learned response in the ACC that was sensitive to the stimulus color (error rate condition), for both the "go" and "change" trials.

The model presented in Brown and Braver (2005) was focused on understanding the role of the ACC in learning to recognize situations in which the risks of errors are high. Previous work suggested that the ACC detected actual errors (Dehaene et al., 1994) as well as conditions of response conflict (Botvinick et al., 2001). The error likelihood model further suggested that the ACC activity warns of an impending error as a basis for implementing proactive control.

There are other interesting aspects to the empirical data that are not addressed directly by Brown and Braver (2005). For instance, the model does not address the sequential behavior of participants in terms of their reaction times. In addition, the model does not explicitly account for differences in reaction times for the two different error conditions. These effects in the empirical data provide further evidence regarding the cognitive mechanisms involved in human performance on this task that will be explored in the current research.

To better understand the mechanisms influencing human performance on the change signal task, we have created a complementary model that focuses on the detailed behavior

of participants. For instance, the data illustrate that the conditional learning predicted by the error-likelihood model (i.e., differences in ACC activation for High versus Low error conditions) has an impact on reactions times that unfolds over the course of many trials. We used the ACT-R (Anderson, 2007) computational cognitive architecture to model these results from Brown and Braver (2005) study. After we describe the model and results in detail, we discuss the distinct and complementary insights afforded by the modeling approach used here versus Brown and Braver (2005).

The Task

We reimplemented the original Brown and Braver change signal task in Lisp to accommodate integration with ACT-R. The only known differences include color choices, symbols presented, and response keys. Although these items were altered for implementation convenience, they have no bearing on model behavior or performance. The remaining description will focus on the task as presented to human subjects. Additional details regarding the task and instructions can be found in the supplementary materials from Brown and Braver (2005).

A schematic of the change signal task is shown in Figure 1. After a .5s blank inter-trial delay, subjects were presented with a cue stimulus in one of two colors. Unbeknownst to the subjects, the cue color represented either a high error rate condition or a low error rate condition. After one second, the cue was replaced with a similarly colored *go signal*—an arrow pointing either right or left. The subjects were instructed to respond to the go signal by pressing the corresponding right or left arrow key on the keyboard.

On one third of the trials, a larger arrow pointing in the opposite direction of the go signal appeared after a change stimulus delay (CSD). (Again, the coloring was consistent with the error rate condition.) In this case, subjects were instructed to inhibit their initial response to the go signal, and instead respond to the "change signal." A response ended the trial, which progressed directly to a blank screen and the inter-trial delay. No feedback was presented. If the subject failed to respond within one second after the go signal appeared, the trial timed out.

The high and low error rate conditions were bound to unique CSDs, which were adjusted independently using a step function to maintain a consistent error rate defined for each condition. In both error rate conditions, CSDs were constrained to a range of 20 to 800ms and incorrect responses reduced the CSD by 50ms. In the low error rate condition, correct responses led to a 2ms increase in the CSD, while in the high error rate condition the CSD increased by 50ms when a correct response was made. These adjustments were made to motivate a 4% error rate, and a 50% error rate, respectively. Responses made prior to the presentation of the change stimulus were considered errors.

The original experiment used five blocks with approximately 107 trials each, although the trial count varied slightly across subjects. Our task fixed this number to 107, and the direction of the go signals and error rate conditions was randomized and counterbalanced within each block as best as possible. Stimulus colors were consistent with the error rate condition in all blocks except the last. For the final block, the relationship between stimuli colors and error rate conditions was reversed.

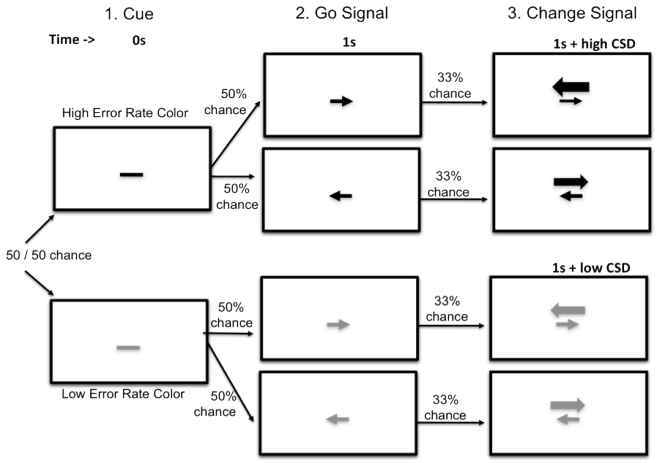

Figure 1: Task schematic. A cue signal is presented in one of two colors, followed by a go signal 1 second later. There is a 33% chance that a subsequent change signal will be presented, the timing of which is determined by a change stimulus delay bound to the signal color.

Human Performance

Figure 2 shows aggregate reaction times across trials collapsed across subjects and conditions. The solid line represents the central tendency as predicted by a simple linear regression of a logarithmic model, although the regression is intentionally discontinuous at the start of the reversal block, indicated by the grey area. The subjects performed more slowly across trials until they reach an asymptote. The regression model coefficient affecting the rise and asymptote of the curve is significantly greater than zero for the normal trials ($p < .001$), and not significant for the reversal block. This suggests that there are not enough reversal block trials to reveal an effect, if there is one.

Time on task effects may account for some of the performance decline (e.g., Gunzelmann, G., Moore, L. R., Gluck, K. A., Van Dongen, H. P. A., & Dinges, D. F., 2010), but we believe that the more influential factor is that subjects were strategically hedging their responses to improve their odds of successfully responding to change signals. (Of course, such a strategy is futile in this experiment because the CSDs were adjusted to encourage a consistent error rate.) Evidence for strategic hedging becomes apparent when we examine reaction times for each condition, also shown in Figure 2. The dashed line shows the central tendency for the high error rate condition, and the dotted line shows the central tendency for the low error

rate condition. Again, the regression is intentionally discontinuous at the start of the reversal block.

Not surprisingly, the statistics for the two error rate conditions match those of the collapsed data, with highly significant coefficients for the normal blocks ($p < .001$) and insignificant coefficients for the reversal block. The confidence intervals for the normal block coefficients, however, are more interesting because they do not overlap. ($17.8 < A_{high} < 27.7$, and $3.0 < A_{low} < 11.2$) The significant difference between these coefficients suggests a relationship between stimulus color and reaction time. In other words, over the duration of the experiment, subjects learn to delay their response more for the high error rate condition than for the low error rate condition. A simple time on task effect would not produce a disparate hedge times across error rate conditions.

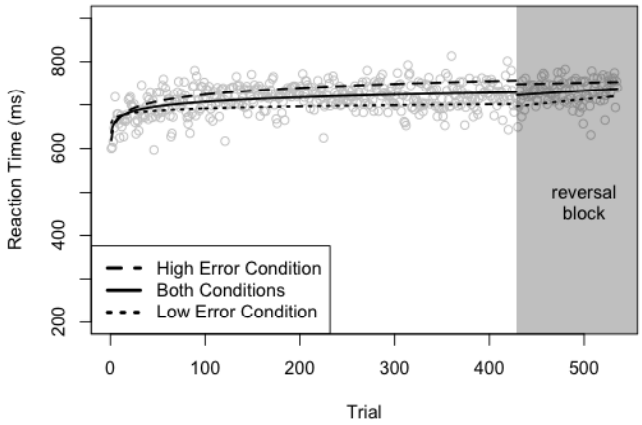

Figure 2: Reaction times collapsed across conditions are shown in the grey scatterplot, with the central tendency shown as a solid black line. Central tendencies for the high and low error conditions are shown as dashed and dotted lines, respectively. The central tendencies, generated through regressions, are discontinuous at the start of the reversal block, shown in grey.

The Model

The ACT-R 6 (Anderson, 2007) cognitive architecture provides the computational framework for our model. It integrates perceptual, cognitive, and motor processing mechanisms from the psychological literature. At its core, it is a symbolic production system with a semantic network memory and simulated subsymbolic effects. Specifically, our model leverages the procedural and declarative capabilities, the intentional module, and a timing capability derived from a temporal module (Taatgen, Van Rijn, & Anderson, 2007).

The empirical data from Brown and Braver (2005) demonstrate that subjects implicitly learned the association of stimulus color to error rate condition. In this paper, we show that this learning measurably influenced subject performance—their response times were strategically mediated by stimulus color. Out of several possible approaches to model this in ACT-R, we chose to use the declarative module to emulate the statistical learning attributed to the ACC.

From the perspective of the ACT-R theory, the declarative module is not intended to represent the functional properties of the ACC (see Anderson, 2007), but it does provide the appropriate Bayesian dynamics to represent the learning we hypothesize may be involved. Thus, we treat the declarative module as a surrogate for the ACC functionality that is not represented by existing mechanisms in ACT-R. This absent functionality would appear to appropriately reside within ACT-R's intentional module, which is associated, in part, with ACC function (Anderson, 2007).

The model employs a simple hedging strategy to accomplish the task. Upon attending to a cue, it attempts to retrieve a similar trial from declarative memory based on the cue color. When the subsequent go signal is attended, the model does not respond immediately. Instead, it waits according to a remembered "hedge time" from the trial that was retrieved from declarative memory. If no similar trial exists (i.e., nothing was retrieved), a default initial hedge time is used, which is a free parameter discussed below. If a change signal is seen prior to the expiration of the hedge time, a response is made accordingly. If no change signal is seen and the hedge time expires, the model responds to the go signal.

Even when the model responds to the go signal, the key press does not occur immediately. Instead, the ACT-R motor module initiates a 3-phase motor movement that can take well over 100 milliseconds before the actual key press is registered by the task (Byrne & Anderson, 2001). During this time, the model can detect a change signal, although it is too late to cancel the requested motor action. The model learns from its failure by associating the CSD with the color for that trial in its goal buffer of the intentional module. This timing information is based upon estimates from the temporal module (Taatgen et al., 2007).

At the start of the next trial, the contents of the goal buffer, which includes the association between the stimulus color and hedge time, is stored in declarative memory to serve as an exemplar for future trials. Because detected errors typically associated a longer hedge time than what was originally retrieved, they have the effect of increasing future hedge times (Rabbitt, 1966). As currently written, the model has no specific mechanism to reduce hedge times.

Without a mechanism to reduce hedge times, it might seem that model response times would always increase and never asymptote. Indeed, sharp increases in hedge times do occur in early trials. However, because each stimulus color / hedge time pairing is stored as an independent chunk (i.e. there is no merging) the likelihood of retrieving a particular hedge time increases the more often it is used, in part due to the influence of stochasticity in declarative memory. After a large number of trials, the declarative memory becomes so saturated with hedge times associated with each stimulus

color, that the model's hedging essentially reaches a steady state.

Three parameters were involved with fitting the model to observed data. The first is the initial hedge time, which we believe was established either through practice trials or as a side effect of instructions that informed subjects of delayed change signals. This has the simple effect of moving the y-intercept in Figure 3.

The second free parameter was activation noise, which reflects the effect of subsymbolic processes in the declarative memory system. Noise influences the likelihood that recent and correct declarative information will be retrieved. In terms of the curve in Figure 2, noise affects the overall shape—higher noise flattens it out. In ACT-R, activation noise is set with the ans parameter, for which we settled on a value of .53. This produces a standard deviation of .96 in the distribution of noise that is sampled to add stochasticity to the activation of declarative memories.

Lastly, the ACT-R declarative memory system allows for errors of commition through a mechanism called partial matching. We used this mechanism so that the model would be indifferent to stimulus colors in early trials and develop a differentiation in later trials. The mechanism requires us to specify a degree of similarity between stimulus colors, which we set to 50%. We did not use this as a free parameter in the fitting processes because the other parameters provided the necessary degrees of freedom.

Results

Using the parameter values described above, we aggregated the results from 100 model runs to obtain reliable measures of central tendency. A comparison of reaction times between model and human data are shown in Figure 3. Because a large amount of stochasticity still remains even after aggregation, the model results are represented using linear regressions of a logarithmic model in the same way the human data is shown. (The standard deviation is considered as a separate measure of fitness below.)

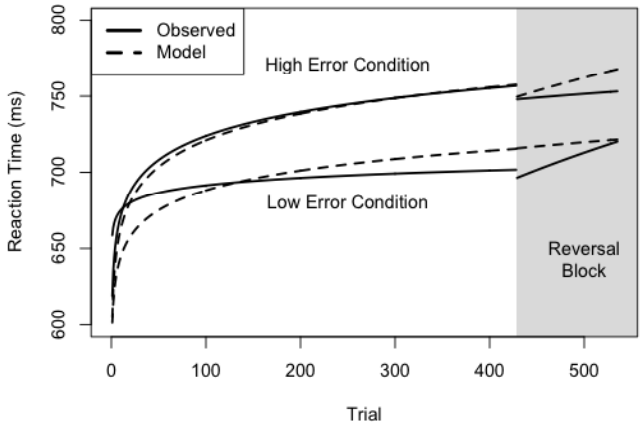

Figure 3: ACT-R model results are shown as dashed lines on top of the human data shown as black lines.

The RMSD values calculated from the model and human reaction time data are shown in Table 1. The overall mean RMSD was 58.5ms, which seems reasonable given that some of the deviation is a result of remaining stochasticity in the model and human data.

Table 1: RMSD values between model and human data.

Condition / Block	RMSD (ms)
High Error / Normal	51.6
High Error / Reverse	48.8
Low Error / Normal	74.3
Low Error /Reverse	59.1

The high stochasticity suggests that the standard deviation of the reaction time is another important measure of fitness (non-responses were removed for this analysis). Figure 4 overlays model performance on top of a box plot of the subject data. The model's standard deviation was in the middle of the 1st quartile for the subject data. This could be improved by increasing noise in other areas of ACT-R, but we opted against doing so in the interest of parsimony.

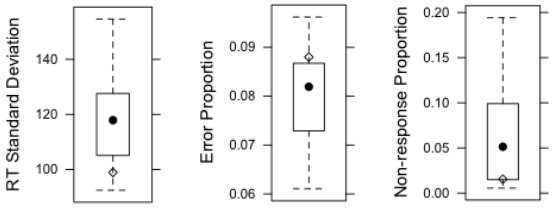

Figure 4: ACT-R model standard deviation, error proportion, and non-responses overlaid on subject data. The hollow diamonds indicate ACT-R values.

The proportion of incorrect responses made was also a consideration. For purposes of this analysis, an incorrect response occurs when the subject presses the wrong arrow key, regardless of condition. Since a response is actually made, this does not included non-responses, which are analyzed separately below. Also shown in Figure 4, the results were within the range of humans, although on the high side.

The remaining measure of fitness is the proportion of non-responses. A non-response occurs when the model fails to respond to a go signal within 1 second. The temporal module in ACT-R adds some stochasticity to the timing so this can occur even if the intended hedge time is within the trial period. Again, the non-responses were well within the human range (see Figure 4), but on the low side of the second quartile. As was the case with standard deviation, this could be improved if we allowed the model another degree of freedom.

Finally, fMRI studies, including the Brown and Braver (2005) work, often use a blood oxygenation level-dependent (BOLD) contrast mechanism. With this technique, regions

of the brain with higher blood oxygenation appear more intensely on images, which indicates greater neural activity. ACT-R uses buffer activity to make BOLD predictions (Anderson, Bothell, Byrne, Douglass, Lebiere, & Qin, 2004), as shown in Figure 5. In this figure, ACT-R makes BOLD predictions for the ACC region based on activity in the goal buffer of the intentional module. To produce this graph, the inter-trial delay was extended to 10 seconds to isolate responses. Data was aggregated from 12 normal blocks of 107 trials.

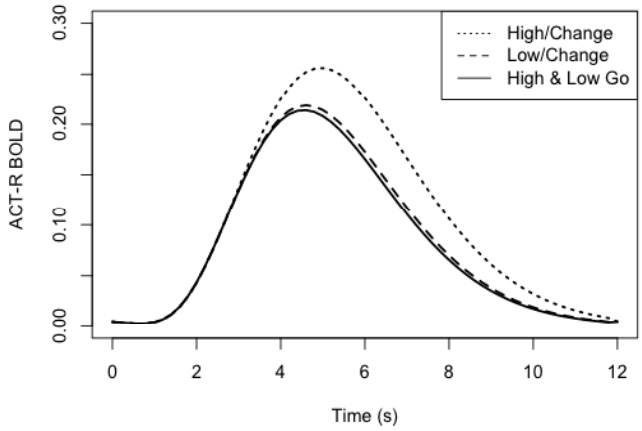

Figure 5: ACT-R BOLD predictions for the ACC region in each of the four conditions.

Discussion

As modelers, we often confront (and perhaps carry our own) biases related to specific modeling approaches, whether it be production level architectures like ACT-R, connectionist approaches like the error likelihood model, diffusion models, dynamic systems, or others (Anderson & Lebiere, 2003). This is unfortunate, because as this research demonstrates, each methodology maintains distinct advantages as well as disadvantages that may be overcome using a variety of techniques. Specifically, the error likelihood model makes detailed predictions about neurological processes in the ACC beyond the current scope of ACT-R. However, ACT-R brings to the table a generalized account of end-to-end perceptual-cognitive activity, which can reproduce observed behavior.

If we accept that both models contain elements of truth, there must be some functional overlap despite the differing levels of abstraction. Recent work on the theory of ACT-R has focused on mapping functionality to specific brain regions (e.g., Anderson, Bothell, Byrne, Douglass, Lebiere, & Qin, 2004). Specifically, the ACC is attributed to the ACT-R intentional module, which includes the goal buffer (Anderson, 2007). The goal buffer typically maintains the internal and relevant external information required to make decisions. This is intended to include the conflict resolution typically attributed to the ACC, but it is a functionally broader interpretation.

In our change signal model, the goal buffer contains the stimulus color and hedge time, among other state information. The current implementation of ACT-R provides no functional computation in the intentional module, so the statistical learning demonstrated by the error likelihood model involves knowledge maintained in the declarative module, which acts as a surrogate. Our position that the declarative memory acts as a surrogate is largely based on that fact that many subjects were unable to explicitly distinguish the difference between stimulus colors in terms of their pairing with error likelihood even after the experiment.

This is not a firm position, and we are planning a follow-up study to guide our modeling direction. A more detailed participant debriefing will help determine the degree of declarative learning and influence on behavior. The results may suggest that the declarative component is more than just a surrogate—perhaps the ACC activity is epiphenomenal to declarative function. On the other hand, it may be confirmed that there is little relation between declarative knowledge and subject behavior with respect to high and low error conditions. In this case, the model may evolve towards a bottom-up learning approach, perhaps though augmenting the intentional module in ACT-R or focusing on a procedural learning approach.

In the mean time, the declarative module provides a reasonable proxy for ACC function because it employs a similar statistical learning process. Because the information managed in declarative memory relates stimulus color and hedge times, greater activity occurs when change signal errors are detected. This is reflected in the goal buffer, which results in higher predicted BOLD responses in ACT-R. Furthermore, because errors are detected 3x more often in the high error rate change condition, its mean BOLD response will rise above all other conditions. This is supported in Figure 5.

The ACC BOLD responses recorded in the Brown and Braver (2005) study aligns with some, but not all, of the ACT-R predictions. Specifically, the high error change condition shows the highest activation, followed by low error change and high error go conditions which are essentially tied.

The low error rate go condition is a significant divergence, as the BOLD response show that the activation is clearly lower than the other conditions in that region. Unfortunately this was one of the key findings that distinguished the error likelihood model from the alternative "conflict" model. The current ACT-R model does not produce a similar prediction because extra goal manipulation only occurs when errors are detected in change conditions. One could argue that this is a response to the statistical learning that was delegated to the declarative memory system in our model. In this regard, the ACT-R model stands in contrast with the Brown and Braver (2005) model, which predicted greater fMRI activity in ACC for

high vs. low error likelihood trials, even when restricted to correct trials with no change signal. Nevertheless, if the hedge time in the declarative memory were to increase the simulated fMRI activity, then our model might be able to simulate an error likelihood effect in ACC activity.

Finally, with an ACT-R model of the change signal task performing reasonably well, we have an opportunity to make a prediction. The reversal block in the observed human data had surprisingly little effect, and the ACT-R model produced similar results. By extending the number of reversal blocks, we can predict how many trials will be required to see an effect, and what that effect might be.

The predicted results of 24 reversal blocks are shown in Figure 6. As mentioned previously, the model does not currently have a mechanism to reduce hedge times. Both conditions achieve a steady state at their asymptotes through a combination of accumulated statistical evidence and retrieval noise. Even when failures to respond to change signals are detected and increased hedge times remembered, noise in the declarative retrieval process makes it unlikely that the latest trial information will be retrieved over the large number of older, lower trial hedge times available.

Without this statistical influence, the low error rate condition would never achieve an asymptote below the high error condition without a mechanism to hedge downward. This also provides an explanation for the predictions in Figure 6, which continue on the same trajectory as the normal block. In contrast, the error likelihood model of Brown and Braver (2005) would predict that over time, the ACC will learn the reversed error likelihood pairings, leading to a reversal of error likelihood effects on reaction time. Although our current data is insufficient to make concrete statements about which prediction is correct, our follow-up study will extend the number of reversal blocks with hopes to allow such a test. Once again, this will help inform future model development.

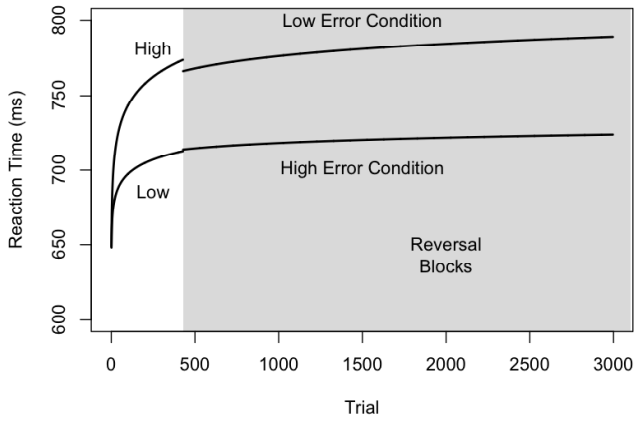

Figure 6: ACT-R model prediction of color reversal over 24 blocks, shown in the grey region.

Acknowledgments

The views expressed in this paper are those of the author and do not reflect the official policy or position of the Department of Defense, the U.S. Government, or Lockheed Martin Corporation. This research was sponsored by grants 07HE01COR and 10RH04COR from the Air Force Office of Scientific Research. We would like to thank Tim Halverson for his contributions to this research.

References

Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* Oxford, UK: Oxford University Press.

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). 1036-1060.

Anderson, J. R. & Lebiere, C. L. (2003). The Newell test for a theory of cognition. *Behavioral & Brain Science* 26, 587-637.

Botvinick, M. M., Braver, T. S. , Barch, D. M., Carter, C. S., Cohen, J. C. (2001). Conflict monitoring and cognitive control. *Psychological Review,* 108, 624-652.

Brown J. W, & Braver T. S. (2005). Learned predictions of error likelihood in the anterior cingulate cortex. *Science* 307, 1118–1121 doi: 10.1126/science.1105783.

Byrne, M. D., & Anderson, J. R. (2001). Serial modules in parallel: The psychological refractory period and perfect time-sharing. *Psychological Review*, 108, 847-869.

Dehaene, S., Posner, M.I., & Tucker, D. M. (1994). Localization of a neural system for error detection and compensation. *Psychological Science,* 5, 303-306.

Gunzelmann, G., Moore, L. R., Gluck, K. A., Van Dongen, H. P. A., & Dinges, D. F. (in press - 2010). Fatigue in sustained attention: Generalizing mechanisms for time awake to time on task. In P. L. Ackerman (Ed.), *Cognitive Fatigue: Multidisciplinary Perspectives on Current Research and Future Application.* Washington, DC: American Psychological Association.

Logan, G. D. & Cowan, W. B. (1984). On the ability to inhibit thought and action: A theory of an act of control. *Psychological Review*, 91, 295-327.

Rabbitt, P. M. A. (1966). Errors and Error Correction in Choice-Response Tasks. *Journal of Experimental Psychology*, 71, 2, 264-272.

Taatgen, N., Van Rijn, H., Anderson, J. R. (2007). An integrated theory of prospective time interval estimation: The role of cognition, attention, and learning. *Psychological Review*, 114(3), 577-598.