One Model, Two Tasks: Decomposing the Change Signal Task

L. Richard Moore Jr. (larry.moore@L-3com.com)
L3 Communications
5950 East Sossaman Road, Suite 102
Mesa, AZ 85212 USA

Glenn Gunzelmann (glenn.gunzelmann@us.af.mil)
Cognitive Models and Agents Branch, Air Force Research Laboratory
Wright Patterson Air Force Base, OH 45434 USA

Marissa Daigle (mdaigle@g.clemson.edu)
Clemson University
Department of Mathematical Sciences
Clemson, SC 29634 USA

Abstract

The change signal task is a variant of a two-alternative forced-choice (2AFC) task where the initial stimulus is superseded with the alternative stimulus (the change signal) at a delay on a proportion of trials. Taking advantage of the overlap in task requirements, we present a single model that can perform both tasks. We validate the model using the empirical data from participants who performed both tasks sequentially. The results confirmed the existence of a dynamic hedging strategy, and showed that cognitive fatigue had little, if any, role in slower response times with increased time on task. When fitting the 2AFC task, the model required adjustment to one architectural parameter while the rest were left to defaults. That parameter is then constrained while fitting the remaining three task-specific parameters of the change signal task. This effectively reduces a degree of freedom in the larger task, and increases confidence in the model as it closely matches human performance in multiple tasks.

Keywords: ACT-R; change signal; two-alternative forced-choice; cognitive model.

Introduction

Roberts and Pashler (2000) point out that modelers should consider criteria other than fit when evaluating the credibility of a model. Specifically, they propose examining the data that the model is unable to fit: how much data disagree with the theory, and how strongly does it do so? And could the model fit any data?

Within the computational cognitive modeling community, other approaches to bolster model confidence exist. For example, an established cognitive architecture (Anderson, 2007; Rosenbloom, Laird & Newell, 1993; Meyer & Kieras, 1997) provides the software framework to constrain a model to specific theories of cognitive processes that have been validated independently in the literature. With an active cognitive architecture user community, models can be tested against new empirical data and different experimental conditions (Gobet & Ritter, 2000; Gunzelmann, Byrne, Gluck & Moore, 2009; Gunzelmann, Moore, Salvucci & Gluck, 2011). Even when parameters are adjusted to account for individual and group differences, successful model reuse instills confidence in the theory behind the model.

More complex models provide another validation opportunity through task decomposition. For example, Myers (2009) reports on a composite model that integrates two previously published models to perform a more complex task. Independent validation of each sub-task instills greater confidence in the composite model (cf., Halverson & Gunzelmann, 2011).

In this paper, we present a single model that can perform two tasks with the same set of knowledge. Model performance on the simpler task relies entirely on a subset of knowledge from the more complex task. Because of this relationship, the simpler model can be fit to empirical data, and we would expect any relevant architectural parameters to be identical for the complex task when fitting performance of the same participants on both tasks. Even though the complex task introduces several task-specific parameters that require fitting, all architectural parameters are constrained by the simpler task fit.

The proposed fitting strategy requires a suitably designed empirical study. Specifically, a repeated measures design where each participant performs both tasks in sequence is a necessity. Therefore, we also report on a study that allows for the independent task fitting approach described above.

Change Signal Task

The *change signal task* (Brown & Braver, 2005) and the *two-alternative forced choice* (2AFC) task provide the context for the empirical study as well as the model discussed in this paper. The 2AFC task is the simpler of the two, yet it provides all the necessary fundamentals to perform the more complex change signal task. (There is, however, a parameterized difference in strategy between the two tasks that is discussed in the model fitting section below.)

The change signal task was originally devised by Brown and Braver (2005) as a variation of the classic Logan and Cowan (1984) *stop signal task*. Whereas the stop signal task focused on response inhibition, the Bown and Braver variant focused on changing responses.

At its core, the change signal task is a modified 2AFC task. In the basic 2AFC design, participants respond to arrows pointing right or left by pressing the associated arrow key on the keyboard. The modification for the change signal task is that on 1/3 of the trials, a larger arrow appears after the initial stimulus, critically timed to interrupt their normal response. The larger arrow always points in the opposite direction of their initial response, and signals participants to inhibit their initial response and instead respond to the change signal.

In Brown and Braver (2005), the timing of the change signal was dynamically adjusted to induce consistent error rates. In fact, the task implemented two change stimulus delays to produce different error rates: a 50% high error rate condition and a 4% low error rate condition. These two conditions were not explained to the participants prior to the experiment, but they were differentiated by stimulus colors (color cue conditions).

The experimental tasks are described in more detail below, as well as the modeling insights gleaned.

Empirical Work

In Moore, Gunzelmann, and Brown (2010), we examined the Brown and Braver (2005) data and found that participants responded more slowly over time. We proposed that subjects were "hedging" their responses in anticipation of a possible change signal, and that participants hedged for longer periods of time with increased task experience. We also raised the possibility that some slowing may be the result of time on task effects, but the data from the change signal task alone did not allow us to assess that possibility in detail. One motivation for the study described in this paper was to evaluate the role of within task fatigue more thoroughly.

Another important result in the Brown and Braver study was that participants responded to the high error rate condition more slowly (allowing more time for the change signal to appear) compared to the low error rate condition. One interpretation of these data is that participants were forming an implicit association of stimulus color to error condition, which was implemented in our earlier model (Moore et al., 2010). However, in this paper we will demonstrate that this is not necessarily the case, and present a model that embodies a more parsimonious explanation for the data. This will be discussed in the results section.

Experiment

The experiment included 33 participants between the ages of 18 and 50, with 18 females and 15 males. Participants were asked to perform two tasks during the hour-long experiment. One was the change signal task, and the second was a 2AFC task. The order of the two tasks was counterbalanced across the participants. All participants completed 642 trials for each task, except one who mistakenly only completed the change signal task. Data from that subject is excluded from this paper.

At the start of the experiment, participants were shown instructions and allowed to perform six sample trials for each task. Instructions for each task were redisplayed before the participants performed it, and there was an optional break between them.

The change signal task consisted of 6 blocks of 107 trials each. (The trial count was selected for consistency with the Brown and Braver (2005) experiment.) After each block, subjects were allowed to take a brief break. A diagram of the possible sequences of events during a trial and their probabilities is shown in Figure 1. At the start of each trial, a cue was presented in one of two colors, which was associated with either a high error condition or a low error condition. The significance of the colors in the task was not explained to the participants, but they were made aware that there would be two colors, just as in Brown and Braver (2005).

After 1 second, the cue was replaced with an arrow in the same color that pointed right or left. The participant was instructed to respond to this "go signal" with the appropriate arrow on the keyboard. On 1/3 of the trials, a larger arrow pointing in the opposite direction appeared after a brief delay (the "change signal delay," or CSD). The participant was instructed that, in these circumstances, they should inhibit their initial response and instead respond to the larger arrow. The larger arrow, or "change signal," always pointed in the opposite direction as the go signal.

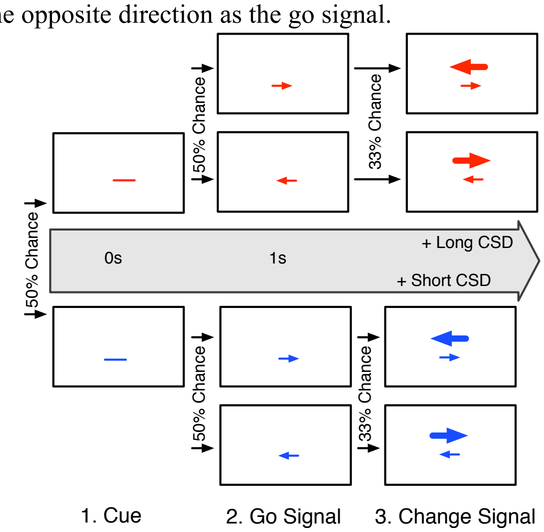

Figure 1: State diagram of a change signal trial. Error condition and arrow direction appear with equal probability. Change signals only appear on 33% of the trials.

The change signal delay was dynamic, and varied according to the error rate condition. At the start of the task, the CSDs for both conditions were set to 250ms. For the high error rate condition, a correct response increased the CSD by 50ms, while an incorrect response decreased it by 50ms. The low error rate CSD behaved similarly, except that it only increased by 2ms when a correct response was made. These manipulations replicate Brown and Braver (2005), and were designed to produce different error rates in the two conditions. For both conditions, the CSD was constrained between 20 and 800ms. Each trial allowed responses up to a full second after the go signal was displayed. If no response was provided within that time, the trial was recorded as a non-response and another trial was automatically initiated.

Halfway through the trials for the change signal task, the mapping between color cue and error condition was reversed, although the CSDs for each condition were not reset. After the participant completed the experiment, he/she was asked whether they noticed the role of the colors in the experiment. In doing so, the experimenter was testing whether the participant acquired any explicit knowledge. We considered it adequate evidence if the respondent indicated that one color was faster or more difficult than the other.

The 2AFC task was identical to the change signal task in every respect, except that no change signals were presented.

Results

Generally speaking, the results from our study were consistent with the Brown and Braver (2005) study. The aggregate data for the change signal task shows the expected slowing in reaction time as the experiment progresses. Conversely, the 2AFC task shows slightly improved reaction times over the duration of the experiment when reaction time is regressed against trial index (b=-.032, R^2=.00040, F(1,20342)=82.49, p<.001; see Figure 2). This result argues against time-on-task based declines in cognitive performance as the source of slowing in the change signal task, and supports the hypothesis that participants were strategically “hedging” their response times. Furthermore, an ANOVA with factors of block and error likelihood confirms that the response times between the two error conditions are significantly different, (F(5,17552)=62.48, p<.001), as found in our previous research (Moore et al., 2010). Overall, participants made errors on 34% of the trials in the high error rate condition, and on 5% of the trials in the low error rate condition.

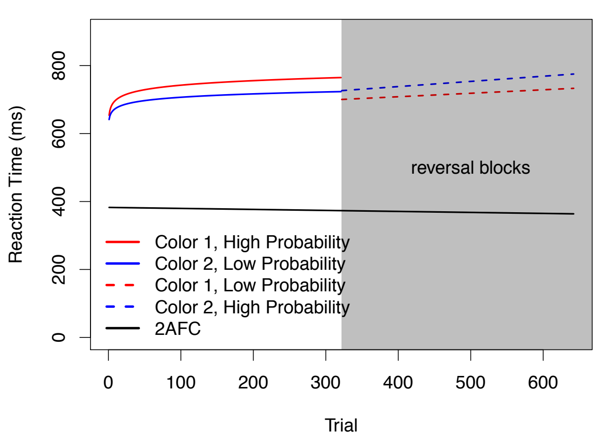

Figure 2: Regression lines for aggregate participant data, showing the strategic hedging over time, as well as the significant disparity across the two conditions.

A more revealing perspective on these results can be observed in Figure 3, which breaks apart trials where a change signal was presented (change condition) versus trials where no change signal was presented (go condition). The experimental condition permutations then become:

1. go-low: go signal only (no change signal) presented in the low error condition color,
2. go-high: go signal only (no change signal) presented in the high error condition color,
3. change-low: change signal present in the low error condition
4. change-high: change signal present in the high error condition.

The figure also includes the relative reactions times in the 2AFC task. All reaction times are measured from the onset of the go signal, and results are aggregated across blocks.

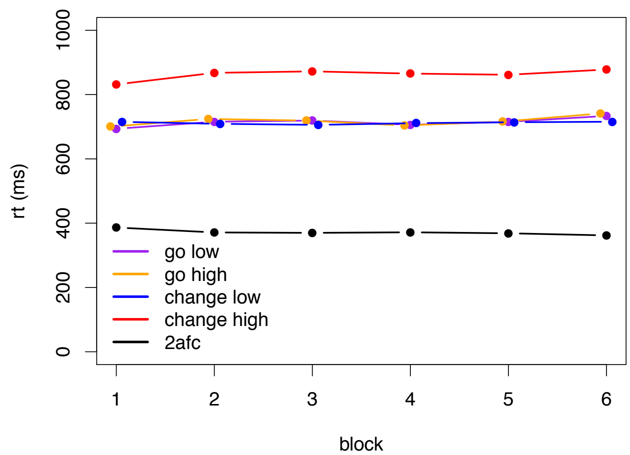

Figure 3: Mean reaction times aggregated by block for correct responses in each of the four change signal task conditions as well as the 2AFC task. Reaction times are measured from the onset of the go signal.

Notice that the 2AFC response times are substantially faster than the response times in the change signal task. As mentioned previously, our theory proposes that this is strategic; participants are hedging their responses to go signals in order to allow for the possibility of a change signal being presented. Under this account, the go conditions represent the situation where individuals exhaust their hedge time and produce a response. Therefore, the difference between the 2AFC and go condition reaction times would represent the mean participant hedge time, which is about 300ms.

Also notice the disparity between the change-high and change-low conditions in Figure 3. In Moore and Gunzelmann (2010) we suggested that implicit association of stimulus color might explain the disparity. This was also supported by the original Brown and Braver (2005) work, which focused on learned responses to error conditions in the anterior cingulate cortex (ACC). However, if there is learning in the ACC, it is not reflected in the empirical data because the stimulus color only impacts response time on trials where a change signal is presented. There is no difference in RT between the two "go" conditions (F(1,12607)=.12, p=.73). If participants were learning the association between error-likelihood and stimulus color, we would expect to see an identical disparity between the go-low and go-high conditions.

Rather than implicit learning, our finding suggests that the emergent difference in reaction times between the change-high and change-low conditions can be explained as an artifact of the task itself. Recall that correct responses to change signals increase the change signal delay by 2ms in the low error condition and 10ms in the high error condition. The different step functions result in change signal delays that tend to be shorter in the low error condition than in the high error condition. Regardless of the error condition, however, participants respond immediately to change signals that appear within the hedge period. Thus, reaction times for the low error condition are faster than the high error condition because of the shorter change signal delay.

To demonstrate that change signal delays are driving the difference in reaction times across the two change conditions, Figure 4 removes the change signal delay from change condition reaction times. (i.e. Reaction times are now measured from the onset of the change signal for the trials with a change signal). The disparity between the high and low change conditions in Figure 3 is greatly reduced, which reinforces the position that it is an artifact of the task itself. In fact, in this analysis response times are slightly *faster* in the high error likelihood condition.

A descriptive analysis of the response distributions, as well as quantitative evidence of a disparity in lapses across the two conditions (14% for the change-high condition versus 1% for the change-low condition), both suggest that the remaining discrepancy in reaction times may be accounted for by a truncation of the response distribution in the change-high condition. This truncation occurs because some of the response times are very close to the 1-second trial time limit.

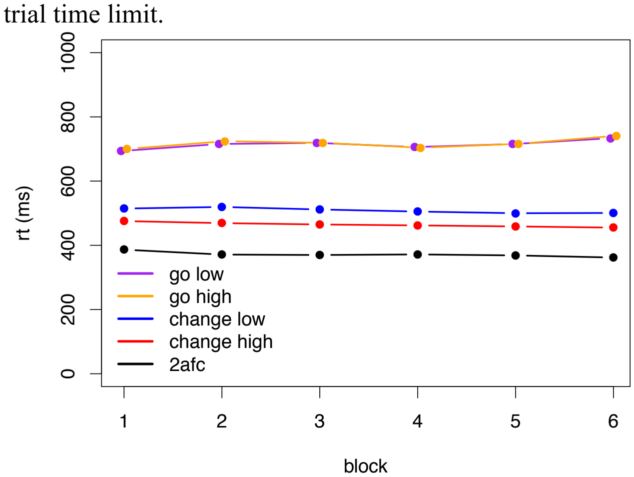

Figure 4: Mean reaction times aggregated by block for correct responses in each of the four change signal task conditions as well as the 2AFC task. Reaction times are measured from the onset of the go signal for go conditions, and the onset of the change signal for change conditions.

Another item in Figure 4 that warrants attention is the fact that both change condition reaction times are higher than the 2AFC reaction times. In all three of these conditions participants can respond immediately to the stimulus, which suggests that the change conditions impose an extra cognitive penalty. This is an important consideration for the model, which will be discussed in the following section.

The Change Signal Model

The change signal model was developed within the Adaptive Control and Thought – Rational (ACT-R) cognitive architecture (Anderson, 2007). ACT-R is a symbolic production system coupled with mathematically grounded mechanisms that reflect sub-symbolic influences. The change signal model is instantiated within the architecture by supplying knowledge in the form of production rules and declarative chunks. Our model is relatively simple, consisting of 14 productions and no initial declarative knowledge. We characterize it as a "procedural" model because it does not rely upon declarative retrievals to function.

At a high level, the critical feature of the model is a strategic delay of its response to the stimulus to accommodate the possibility of a change signal. As described in the task section, we refer to this as the hedge time. If a change signal does occur during the hedge time, the model generates a response to the larger arrow as soon as it appears (change signal). If no change signal occurs during the hedge time, the model responds to the original arrow (go signal). Time estimates are noisy, and are derived using a mechanism proposed by Taatgen, van Rijn, and Anderson (2007).

The model will also adjust its hedge time (which is maintained as a slot in the goal chunk) dynamically based on responses to change signals. When a change signal is detected after it has already responded to a go signal, the hedge time is increased in hopes that it will correctly catch the change signal in the future. When the model does correctly respond to a change signal, or when the model sees an unexpected cue because it failed to respond before the trial expired, the hedge time is decreased.

The 2AFC task is identical to the change signal task except no change signals are ever presented. As a result the model can perform the 2AFC task unaltered using a subset of the full procedural knowledge; those productions involved with responding to change signals never fire.

As discussed in the previous section, Figure 4 shows that change signal responses incur an extra penalty in response time. There are several plausible theories to account for the increased response time when a change signal is encountered. In our model, the delay is attributed to motor control. The model prepares its response to the "go" signal when it is presented. Thus, when a change signal is observed, the motor system is reset, which negates the preparation that was done. To respond, the model must first plan the motor movements and then execute them. The extra planning adds time to the response process.

Model Fitting

The initial hedge time, hedge increase, and hedge decrease are tunable parameters in the model. We also chose to adjust the ACT-R default action time, which describes the speed of the average production cycle.

Because the model shares knowledge between the 2AFC and change signal tasks, it was fit in two stages. The 2AFC task was fit first because it requires fewer degrees of freedom. The empirical data supports the conclusion that participants did not engage in strategic hedging while performing the 2AFC, so initial hedge time, hedge increase, and hedge decrease parameters are all set to 0.

The only remaining parameter to fit in the 2AFC task is default action time (DAT). It defaults to 50ms, but Stewart, Choo and Eliasmith (2010) have shown that tasks composed of simple productions (such as those implemented to perform the 2AFC and change signal tasks) will have shorter cycle times. This is also consistent with our work with the psychomotor vigilance task (PVT), where we typically find default action time values of approximately 40ms (Gunzelmann, Moore, Gluck, Van Dongen & Dinges, 2010). Rather than refitting the 2AFC task independently, we chose to constrain DAT to the value obtained from the PVT, resulting in the dashed black line in Figure 5. This 40ms value was held constant while fitting the model with the change signal task, as well.

The remaining parameters include the initial hedge time, the increase in hedge time when the model detects an error, and the decrease in hedge time when the model correctly responds to a change signal. All three parameters are specific to the hedging strategy, and are not general parameters of cognition. (We also enabled a mechanism to provide some stochasticity to production cycle times, but left that parameter at its default value.) To resolve the three dimensional parameter space, we used large scale computational resources running our in-house search software (Moore, 2011). The model was then rerun using the predicted optimal values to produce the change signal model results in Figure 5.

Results

The overall RMSD for the block-aggregated data across all 5 conditions in Figure 5 was 33.4ms. The simpler 2AFC task fit the best at 15.9ms RMSD, while the go-high condition fit the worst, at 50.1ms.

The model was not very sensitive to the initial hedge time parameter. Its behavior was primarily driven by the hedge-up and hedge-down values, as they are critical to establishing and maintaining the equilibrium between the model and task. The optimal values (27% increase when hedging up, and 6.5% decrease when hedging down) suggest that participants were more liberal with hedging up (waiting longer when an error is detected) than they were hedging down (responding sooner when a change signal is correctly detected). Furthermore, there was clearly a relationship between the two variables: larger upward hedges could be compensated by larger downward hedges to maintain a reasonable fit. A degree of freedom could potentially be reduced if one of the two parameters could be experimentally isolated.

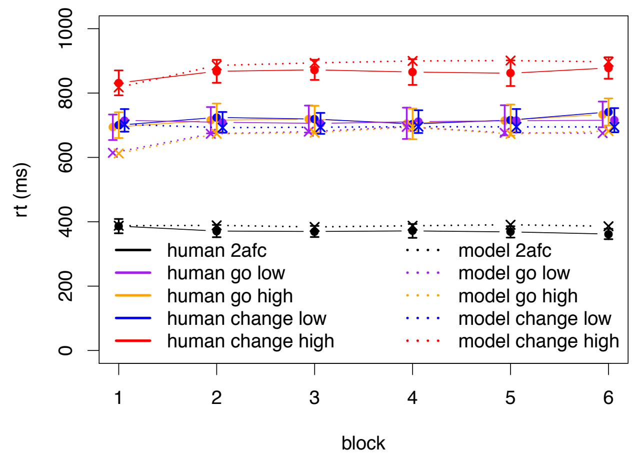

Figure 5: Model fits to empirical data for correct responses.

In addition to the reaction time across blocks, there are several other statistics that can be examined to evaluate the model's performance relative to human participants. The percent of correct responses, the standard deviation in reaction time, and the number of non-responses are three measures shown in Figure 6. The model's performance was within the inter-quartile range on all three measures, and it performed particularly well with percent correct and proportion of non-responses.

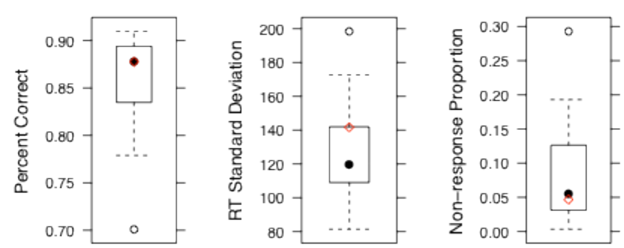

Figure 6: Secondary measures of fitness in the change signal task. The box and whiskers demonstrate the variation across subjects, while the red diamond indicates the mean for the model.

Conclusion

In Moore et al. (2010) a cognitive model of the Change Signal task raised some important questions that could not be fully addressed with the available data. The follow-up study reported in this paper provided an opportunity to further solidify our understanding of the cognitive mechanisms associated with the task.

A primary benefit was having performance data from participants for both the 2AFC and change signal tasks, for several reasons. First, participants performed the 2AFC consistently and much more quickly than they did in the change signal task. This supports our theory of a hedging strategy to manage change signal responses. The lack of performance degradation in the 2AFC task shows that within task fatigue plays a very small role, if any. Furthermore, participant hedging appears to be dynamic, as is evidenced by the slowing reaction time throughout the session. This informs the model, and justifies the model's hedging parameters.

The dual task / repeated measures design of the experiment also allowed us to isolate one of the parameters in the change signal model (default action time) and fit it independently within the context of the 2AFC. It was an unexpected additional benefit that we were able to use a value for that parameter derived from previous research using a different task. The remaining parameters, which were all related to hedging, could then be fit within the context of the change signal task.

The change signal model has demonstrated that it can perform two tasks with overlapping knowledge using a single set of model parameters. In doing so, it inspires confidence in the theory that the model represents. We believe this is well within the spirit of looking beyond a simple model fit for validation (Salvucci, 2010).

Acknowledgments

The views expressed in this paper are those of the authors and do not reflect the official policy or position of the Department of Defense, the U.S. Government, or L3 Communications. This research was sponsored by grants 07HE01COR and 10RH04COR from the Air Force Office of Scientific Research. It was also was supported through the Repperger Internship Program sponsored by the 711 HPW/RH Chief Scientist's Office.

References

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

Brown, J. W., & Braver, T. S. (2005). Learned predictions of error likelihood in the anterior cingulate cortex. *Science*, *307*(5712), 1118. doi:10.1126/science.1106642

Gobet, F., & Ritter, F. E. (2000). Individual data analysis and Unified Theories of Cognition: A methodological proposal. In *Proceedings of the 3rd International Conference on Cognitive Modeling* (pp. 150-157). Veenendall, The Netherlands: Universal Press.

Gunzelmann, G., Byrne, M. D., Gluck, K. A., & Moore, L. R., Jr. (2009). Using computational cognitive modeling to predict dual-task performance with sleep deprivation. *Human Factors*, *51*(2), 251-260. From http://handle.dtic.mil/100.2/ADA504480.

Gunzelmann, G., Moore, L. R. Jr., Salvucci, D. D., & Gluck, K. A. (2011). Sleep loss and driver performance: Quantitative predictions with zero free parameters. *Cognitive Systems Research*, *12*, 154–163.

Gunzelmann, G., Moore, L. R., Jr., Gluck, K. A., Van Dongen, H. P. A., & Dinges, D. F. (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* (pp. 83-96). Washington, DC: American Psychological Association.

Halverson, T., & Gunzelmann, G. (2011). Visual search versus memory in a paired associate task. In *Proceedings of the Human Factors and Ergonomics Society 55th Annual Meeting* (pp. 875-879). Thousand Oaks, CA: SAGE. (DOI: 10.1177/1071181311551182)

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.

Meyer, D. E., & Kieras, D. E. (1997). A computational theory of executive cognitive processes and multiple-task performance. Part 1: Basic mechanisms. *Psychological Review*, *104*(1), 3-65.

Myers, C. W. (2009). An account of model inspiration, integration, and sub-task validation. In *Proceedings of the 9th international conference on cognitive modeling* (pp. 333-335). Manchester, United Kingdom.

Moore, L. R., Jr. (2011). Cognitive model exploration and optimization: A new challenge for computational science. Computational. *Computational & Mathematical Organization Theory, 17,* 296-313.

Moore, L. R., Jr., Gunzelmann, G., & Brown, J. W. (2010). Modeling statistical learning and response inhibition with the change signal task. In D. D. Salvucci & G. Gunzelmann (Eds.), *Proceedings of the 10th international conference on cognitive modeling* (pp. 169-174). doi:10.1126/science.1105783

Roberts, S., & Pashler, H. (2000). How persuasive is a good fit? A comment on theory testing. *Psychological Review, 107*, 358-367.

Rosenbloom, P. S., Laird, J. E., and Newell, A. (1993). *The SOAR Papers*. Cambridge, Massachusetts: MIT Press.

Salvucci, D. D. (2010, April). *Cognitive supermodels*. Invited talk, European ACT-R Workshop, Groningen, The Netherlands.

Stewart, T. C., Choo, X., & Eliasmith, C. (2010). Dynamic behavior of a spiking model of action selection in the basal ganglia. In D. D. Salvucci & G. Gunzelmann (Eds.), *Proceedings of the 10th international conference on cognitive modeling* (pp. 235–240). Philadelphia, PA: Drexel University.

Taatgen, N. A., Rijn, H. v., 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.