Halverson, T., Gunzelmann, G., Moore Jr., L. R., & Van Dongen, H. P. A. (2010). Modeling the effects of work shift on learning in a mental orientation and rotation task. In D. D. Salvucci & G. Gunzelmann (Eds.), Proceedings of the 10th International Conference on Cognitive Modeling. (pp. 79-84). Philadelphia, PA: Drexel University.

Modeling the Effects of Work Shift on Learning in a Mental Orientation and Rotation Task

Tim Halverson (thalverson@gmail.com)
Oak Ridge Institute for Science and Education
Air Force Research Laboratory
Mesa, AZ 85212 USA

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

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

Hans P.A. Van Dongen (hvd@wsu.edu)
Sleep and Performance Research Center, Washington State University
Spokane, WA 99210 USA

Abstract

Circadian rhythms cause alertness declines at night, producing performance decrements across cognitive domains and tasks. Building on the learning mechanisms for declarative knowledge instantiated in the ACT-R cognitive architecture, this research seeks to explain the effects of circadian rhythms on performance of an orientation task performed repeatedly across two weeks by participants working either day or night shifts. The differences in performance between the two groups are best explained by varying the decay rate in declarative knowledge as a function of the time of day the task was performed. The model accounts well for task learning reflected in decreases in response times across days, as well as differences in learning between the day and night shift conditions.

Keywords: sleep; circadian rhythm; fatigue; learning; shift work; declarative memory; spatial; ACT-R

Introduction

Variations in alertness due to circadian rhythms and sleep loss have been shown to affect various components of cognitive functioning (e.g. Jackson & Van Dongen, in press). For example, vigilant attention (Lim & Dinges, 2008), perceptual learning (Mednick, Nakayama & Stickgold, 2003), and motor learning (Walker, Brakefield, Morgan, Hobson & Stickgold, 2003) are all affected by fluctuations in alertness associated with time awake and circadian rhythms.

Despite well-documented behavioral changes, it is not well understood how nighttime operations affect learning in different contexts. Most research on night and shift work has focused on how shift differences affect sleep and frequency of accidents (e.g. Åkerstedt, 1988). The affect of changes in alertness (e.g., as associated with work shift differences) on learning is one area of research where a better understanding of the mechanisms involved is needed. More detailed explanations hold the promise of enabling predictions about how learning experiences at different times of the day may differ, and how this may impact eventual task performance.

Previous cognitive modeling efforts have explored some effects of moderators on cognitive processes. In fact, several studies have examined such effects in the context of declarative knowledge. For instance, the effects of caffeine on memory retrieval have been modeled by increasing the activation of declarative knowledge (Kase, Ritter & Schoelles, 2009). Conversely, the effects of sleep loss on memory retrieval have been explained using decreases in declarative activation (Gunzelmann, Gluck, Kershner, Van Dongen & Dinges, 2007). The negative effect of time-on-task on response accuracy has been explained by increasing noise, making misretrievals more common (Fu, Gonzalez, Healy, Kole & Bourne Jr, 2006).

These research efforts focused on processes associated with retrieving declarative knowledge by impacting the availability or confusability of chunks when they are requested. In contrast, the effects of alertness on the learning and retention of declarative knowledge have not been addressed.

In the research presented here, we investigate how long-term learning may be affected by fluctuations in alertness resulting from circadian rhythms during laboratory-simulated shift work. This is accomplished within the context of a spatial direction task based on Gunzelmann, Anderson, and Douglass (2004), which was performed repeatedly by participants over two weeks. A computational cognitive model is presented that accounts for changes in observed response times across successive days of the study, including differences in learning rates as a function of simulated work shift. Differences in performance between shift conditions are explained by manipulating the decay rate parameter in ACT-R's declarative knowledge activation function. Increased decay (reduced learning) in the night shift condition leads to performance decrements that match the human data. The details of the task, the human performance data, and the model are described in the following sections.

Orientation Task

This experiment was conducted as part of a larger study to understand how circadian rhythms and sleep disruption affect performance in a variety of domains. The participants were screened to be healthy and without sleep disorders, with no evidence of brain damage or learning disabilities, and free of drugs of abuse. Participants gave written informed consent, and were paid for their participation.

Figure 1 shows the orientation task used in this study. There are 8 possible target locations (left) and 8 possible misalignments (right; 45 degree intervals). Twenty-five randomly ordered trials were presented per session — 5 target locations (bottom, near, middle, far, and top) crossed with 5 misalignments (0, 45, 90, 135, and 180 degrees). Because performance is roughly equivalent for right-left mirrored stimuli for both target location and misalignment (see Gunzelmann, Anderson & Douglass, 2004), the location was selected at random from the left or right positions.

Participants received instructions that encouraged them to mentally rotate the relative positions of the viewpoint (indicated by the "You" arrow) and the target on the overhead view (left side filled circle) to align them with the viewpoint indicated on the map view (right side arrow). Specifically, they were taught to imagine an angle that connects the viewpoint to the target on the overhead view, with the vertex at the center of the field (a 90 degree angle in Figure 1). They were then told to mentally shift to the map view, and to rotate the angle so that the arrow in the overhead view was aligned with the arrow in the map view (a rotation of 90 degrees clockwise in the trial shown in Figure 1). At this point, the answer could be determined by finding the target end of the angle.

Participants responded using the numeric keypad portion of a computer keyboard, which was spatially mapped to the possible response locations on the map view. So, if the target was in the bottom position on the map (as it is in the sample trial shown in Figure 1), participants responded by pressing the "2" on the numeric keypad.

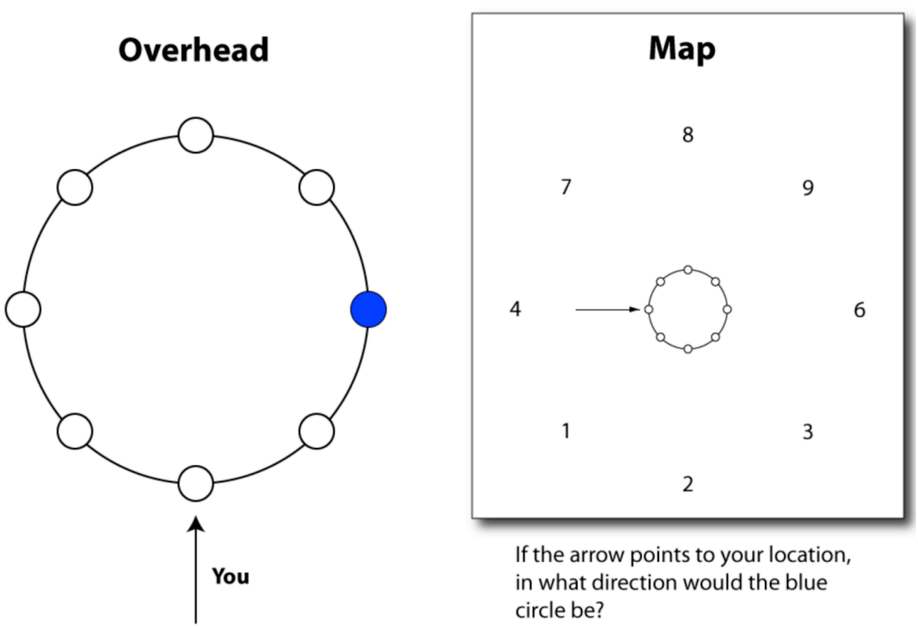

Figure 1: An example trial. The target on the overhead ego-oriented view (left side), indicated by the filled circle, is at middle distance to the right of center. The perspective on the map view (right side), indicated by the arrow, is misaligned by 90° clockwise. The correct response in this example trial is "2."

Method Thirteen participants, ranging in age from 22 to 39 years old (mean = 28), were in the laboratory for fourteen consecutive days. The first day was a baseline day with 10 hours in bed for sleep (22:00–08:00). Subsequently, some of the participants (n = 6) changed to a simulated night shift. Night shift participants were given five hours in bed (15:00–20:00) on the second baseline day, before starting five consecutive work days with 10 hours in bed during the daytime (10:00–20:00) on each day. On the seventh and eighth day, night shift participants had a simulated "day off" during which they had 5 hours in bed (10:00–15:00), 7 hours awake, 10 hours in bed during the night (22:00–08:00), 7 hours awake, and then 5 hours in bed (15:00–20:00), before resuming their night shift schedule for the next 5 days. This schedule represented a common schedule for individuals working a night shift, who frequently switch back to a nighttime sleep schedule during weekends. After the last night shift day, night shift participants received 5 hours in bed (10:00–15:00), 7 hours awake, and then, on the final day of the study, were given 10 hours in bed at night (22:00–08:00) for recovery.

Participants on the day shift (n = 7) were subjected to the same pattern of two baseline days, five consecutive work days, a "day off", another five consecutive work days, and a recovery day. They maintained the same sleep schedule throughout the study, however, with 10 hours in bed (22:00–08:00) each night. Note that participants on the day shift and night shift schedules were given the same amount of time in bed over the course of the experiment; it was merely distributed differently.

Over the course of the study, participants completed fifty-one test sessions of the spatial direction task, with 2 to 4 sessions per day. Before the first session, participants were presented with instructions for the task.

Eight to sixteen days prior to the first session (mean = 14 days), participants were given baseline training on the spatial direction task. This included a set of instructions for the task and four training sessions (these data are not modeled here).

Observed Data

Average response times for each day of the study are presented in Figure 2 for both the day and night shift conditions. Performance during the baseline days of the study (days 1 and 2) was similar for the two groups, and there was no significant difference in mean RT at that point. However, when the conditions diverged, so did performance on the spatial direction task. The performance of the night shift group did not recover during the simulated “day off”, and differences in mean response time remained at the end of the experiment.

To evaluate the differences between shift conditions, we compared response times on the days when they were awake for different shifts (ten days; excluding the baseline, day off, and recovery day) using a linear mixed-effect model with subject as a repeated-measure grouping factor. This was planned *a priori* to most effectively evaluate the impact of shift on performance. However, for the model comparisons later in the paper, all of the observed data was used. See Halverson, Gunzelmann, Moore, and Van Dongen (in press) for more complete analyses of the human data.

Figure 2 shows the mean participant response times (solid lines) as a function of day in study and simulated work shift. There was a steady decrease of response time between days 1 and 14, as corroborated by a main effect of day, $F(9, 7769) = 112.2$, $p < .001$. While there was no evidence of an overall effect of shift, $F(1, 11) = 0.8$, $p = .37$, there was an interaction between shift and day, $F(9, 7769) = 2.1$, $p = .03$. Response times did not improve as quickly when a participant was on the night shift. Observed error rates were low (M = 4%, SD = 3%) and are not addressed in this work.

Mental Rotation Model

A computational cognitive model of the orientation task was developed using the ACT-R 6.0 cognitive architecture (Anderson et al., 2004). The model behavior is primarily driven by mental rotations and learning. The mental rotation operation is implemented using ACT-R’s imaginal module and the imaginal-action buffer. Learning in the model occurs both in the declarative module and through the compilation mechanisms in procedural knowledge. The task procedure implemented in the model was based on the instructions given to the participants in the empirical study.

Model Overview

The model executes the task as follows: In the overhead view, the model encodes the angle defined by the target (blue circle), the center of the overhead view, and the viewpoint (circle nearest the “You” arrow) by visually attending those locations and encoding their coordinates in the imaginal buffer. The model then switches to the map view, encoding the vector defined by the viewpoint (circle

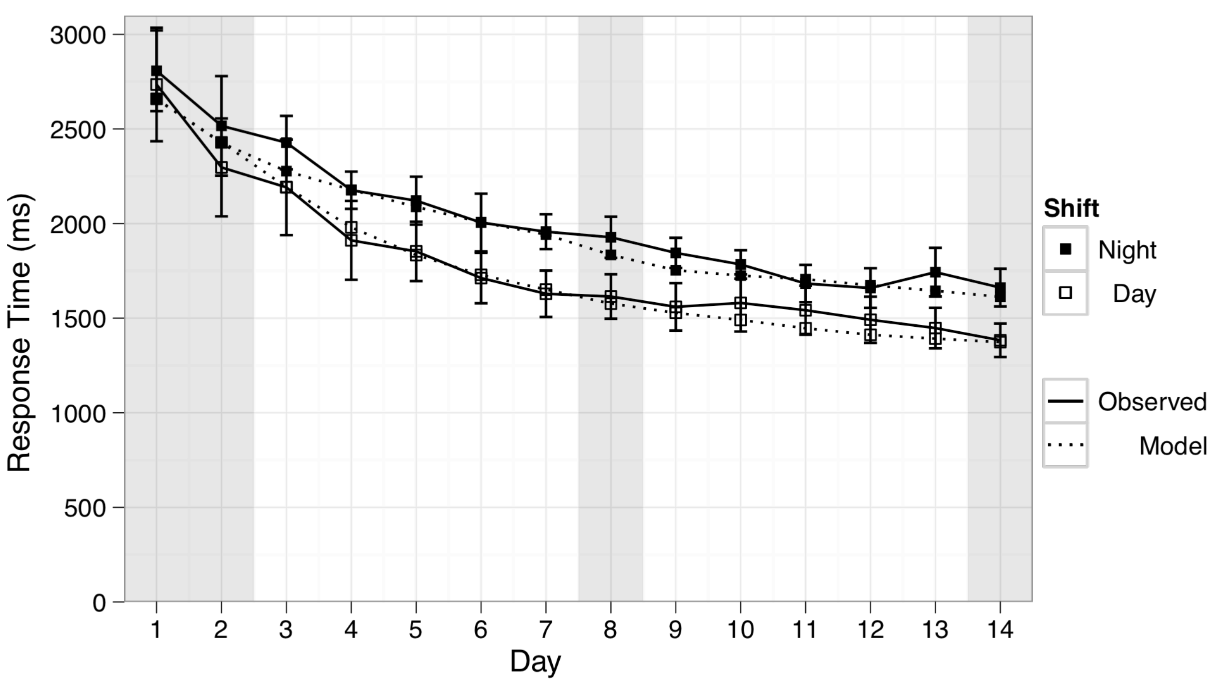

Figure 2: Observed and predicted mean response times as a function of day and simulated work shift (night or day). The shaded regions indicate simulated “days off” in which night shift participants (and the model) performed the task during the day at the same time as day shift participants. Shaded days are not included in the human data analysis. Error bars indicate ±1 standard error.

nearest the arrow) and center of the map view by attending those locations and encoding their coordinates.

The angle that was encoded in the overhead view is then translated to center it on the map view (an imaginal action; 200 ms) and rotated to align the viewpoints of the overhead and map views. The model visually attends the response location closest to the transformed location of the target, encodes the response digit, and presses the corresponding keyboard key.

Mental rotations were implemented using the ACT-R imaginal module. The time to perform the rotation was based on previous mental rotation research (e.g. Bethell-Fox & Shepard, 1988) and was a linear function of the angle of rotation. The slope of the linear function was a free parameter, as the slope can vary by task depending on the relative complexity of the object to be rotated.

Learning

The model's performance improves over time by learning in three ways. First, the angle from the overhead view is encoded in declarative memory when the first subtask is completed. In subsequent trials, the model attempts to retrieve an existing chunk based on the target's location. If a chunk exists and gets retrieved before the model completes the process of visually encoding the angle, then the information from the chunk that was retrieved from declarative knowledge is used. Over time, retrievals become more likely and faster than explicitly encoding the angle using perceptual and imaginal actions. This leads to a speed-up in the model's execution of the task.

In addition to an increasing reliance on declarative representations for target location information, the second step of the solution process is also stored in declarative knowledge once the response is made. These chunks contain information about the target location from the overhead view as well as the perspective on the map view (i.e., the misalignment). Consequently, with experience the model can attempt to retrieve the response based on the target location and map view perspective location. Like encoding the target location on the overhead view, if a chunk is retrieved before the model completes the mental transformations on the map view, the response is based upon the chunk retrieved from declarative knowledge.

The final learning process in the model involves ACT-R's production compilation (i.e. proceduralization). Production compilation is a process by which new productions are created dynamically to represent in one step the consequences of two productions that execute consecutively. With experience, it becomes increasingly likely that the new production will be used, as the model learns that the utility of the new production is greater than the utility of the original, separate productions. However, due to the many constraints imposed on production compilation by the architecture and the structure of this model, the only compilation that occurs in the current model involves encoding the mental rotation into productions specific to each pair of overhead target and map view perspective locations. Therefore, the only savings introduced by production compilation were the infrequent, but substantial, time savings from the mental rotation of trial layouts that were only seen once per session.

Explaining Night Shift Performance Decrements

Several alternatives were explored to explain the decrement in performance observed for participants on the night shift. The solution that resulted in the best explanation of the data was a variation of the decay rate of declarative chunks activation as a function of simulated work shift. Alternative solutions that did not explain the observed trends as well are described in the Results and Discussion section.

By default, the decay rate parameter is not allowed to vary in the implementation of ACT-R. That is, the decay rate can be set, but it assumes the same value for the duration of a model run. There have been various efforts to implement more dynamic mechanisms for decay in ACT-R. Most of these have been related to accounting for the spacing effect (Anderson, Fincham & Douglass, 1999; Jastrzembski & Gluck, 2009; Pavlik & Anderson, 2005).

In our case, we utilize the decay rate to represent differences in the effectiveness of learning as a function of when during the day the task was performed. To implement the mechanisms, the equation to calculate the base-level activation of declarative chunks was modified (Equation 1). The only change to the standard ACT-R base-level learning equation is that the value of the decay rate parameter can vary according to the time when a chunk was added to declarative memory or when the chunk was rehearsed (d_j), as opposed to a constant decay rate across all rehearsals (d) in the original equation. This modification does not change the effect of decay for current ACT-R models.

$$B_i = \ln\left(\sum_{j=1}^{n} t_j^{-d_j}\right) + \beta_i \qquad (1)$$

The current model was implemented with the simplifying assumption that the level of alertness, and thus the value of d_j, is constant across all hours of a work shift (day or night). It is well known that alertness due to circadian rhythms varies throughout the day and night (Van Dongen & Dinges, 2005). However, while the model executed the task the same number of times as the participants did through a simulated workday, we aggregated the data across each day to reduce noise. We have not yet evaluated the capacity of the mechanism to account for finer grained circadian rhythm fluctuations or varying inter-session intervals.

The model was fit to the day shift data using the retrieval threshold (best fit = 1.2), retrieval latency factor (8.0), and rotation slope (0.009 sec/degree) parameters. The rotation slope is similar to the slope found in previous research for simple rotations (Bethell-Fox & Shepard, 1988). The base level learning, which controls the rate of activation decay (d_j), was left at the ACT-R default (0.5) during sessions when participants were on the day shift. For predicting the night shift data, the declarative chunk decay rate was allowed to vary. The best fitting decay parameter for the night shift sessions was 0.6.

Results and Discussion

Figure 2 shows observed (solid lines) and best fitting model (dashed lines) mean reaction times as a function of day in the study and simulated work shift (night or day). For both shifts, the observed behavior is well predicted (RMSD = 65 ms, $r^2 = .98$ for day shift; RMSD = 79 ms, $r^2 = .98$ for night shift). The night shift predictions are particularly noteworthy, as only one parameter was varied relative to the day shift model.

The model is able to predict the observed response times well across fourteen days, including differences across work shifts (i.e. the interaction of day and shift). The model is able to predict the effects of work shift changes well with variations in declarative memory decay rates based on the time at which the tasks are performed. While the declarative decay mechanism explains the observed decrements well, several alternative mechanisms for explaining the trends were considered.

One alternative mechanism involves manipulating overall declarative chunk activation at the time of retrieval, as was done in Gunzelmann et al. (2007). This model did fit the observed data on most days, but did not correctly predict the effect on the overall learning rate when the participants in the night shift condition temporarily switched to the day shift on days 8 and 14. On these days, the model predicts that the performance of participants in the night shift group is nearly equivalent to that of participants in the day shift group. This is because the model assumes that the participants' alertness recovers when performing the task during the day. There is some evidence in associated data (not reported here) to support this, although we do not have conclusive evidence. Regardless, if the impact of degraded alertness were only on activation levels, then the knowledge should be more available during the day. As the human data illustrate, however, the deficits associated with performing the task on the night shift persisted.

Another alternative mechanism for explaining the decrements of alertness is a decrement to utility values associated with production selection and execution. This mechanism has been used to predict performance decrements due to decreased alertness in vigilance tasks (e.g. Gunzelmann, Moore Jr, Salvucci & Gluck, 2009). However, such a mechanism in the model presented here does not explain the observed data for the current task. The same issue is encountered as with the previous alternative — the model recovers to day shift levels of performance on the "day off" and "recovery" days. This is likely a result of the current task requiring constant engagement, over short periods, and thus mechanisms employed for sustaining attention throughout the task would not be stressed.

A third alternative mechanism that was explored is a variation in *procedural* learning as a function of shift. The model presented in this paper has both procedural and declarative learning enabled. It may be that the observed night shift decrement resulted from a slowing of procedural learning rather than a slowing of declarative learning. To test this, the rate of learning for productions rule utilities was varied. This made little difference in the predicted results. This lack of predictive power may result from either the way in which the model was constructed, with an emphasis on declarative knowledge, or a result of the study design, with most of the procedural learning occurring early in the protocol when all participants performed the task during the day.

Thus, the model presented here provides support for the hypothesis that variations in alertness have an impact on learning that may persist beyond immediate task performance. This is consistent with previous research that has indicated that sleep loss causes deficits in encoding declarative knowledge (see Jackson & Van Dongen, in press, for a review). In the ACT-R theory of memory, decay rate is arguably the parameter that most closely corresponds to encoding and rehearsal, as this parameter determines how much the previous exposures to knowledge will affect future retrievals. While there is no conclusive evidence in the literature to attribute either encoding or retrieval deficits to the observations, the current modeling helps support the claim that decreased alertness affects encoding.

A useful future extension to the proposed mechanism for predicting the effects of alertness on learning would be to account for the inter-session intervals. Currently the model does not specifically take into account the 2 to 26 hour intervals between consecutive sessions, which is problematic if we want to generalize the model to tasks in which the time between sessions varies. Incorporating mechanisms proposed in previous modeling to account for inter-session intervals (Anderson, et al., 1999) or practice spacing effects (Jastrzembski & Gluck, 2009) may allow the current model to predict these inter-session intervals.

Conclusion

Performance variations based on alertness have both theoretical and real-world importance. The present results illustrate how specific cognitive processes may be affected by circadian rhythms, and have implications for task training and performance in real-world contexts.

The cognitive modeling presented here illustrates how learning rates may be impaired at night, during the nadir of circadian rhythms. Because degraded learning has potential consequences that extend beyond the immediate situation, brief transitions to day shift may not result in immediate recovery. While the benefit in response time was fairly small in this study (300 ms), the modeling suggests that the effects of learning under conditions of lower alertness may accumulate over time and thus the benefit of training during the day will grow. Moreover, tasks in which exposures to declarative facts are less frequent, as seen in many real world tasks, are expected to encounter an even greater effect of decreased alertness due to a greater time between rehearsals *and* a greater (exponential) decay rate.

Several mechanisms were explored to explain the observed night shift response time decrement. Some mechanisms that have been used previously to explain observed decrements of alertness could not explain the results found in this research. We do not find this outcome particularly troublesome, or even surprising. Rather, in the current study and others, the tasks were specifically selected to ascertain the various ways in which reduced alertness may affect performance on particular mechanisms within the ACT-R architecture.

Our goal is to identify a general set of mechanisms to account for the ways in which variations in alertness impact various components of cognitive functioning. Focusing on laboratory tasks allows us to better isolate various components and evaluate particular computational mechanisms. Such an understanding is necessary in order to predict performance in more complex tasks where various cognitive functions, and mechanisms, interact in complex ways. This represents the focus of this research in the long term (e.g. Gunzelmann & Gluck, 2009; Gunzelmann, Moore, Salvucci, & Gluck, 2009; Tucker et al., 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 or the U.S. Government. The research was supported in part by the Air Force Research Laboratory's Warfighter Readiness Research Division and grants 07HE01COR, 09RH06COR, and 10RH04COR from the Air Force Office of Scientific Research (AFOSR). The first author was supported by an appointment to the Postgraduate Research Participation Program at the U.S. Air Force Research Laboratory administered by the Oak Ridge Institute for Science and Education through an interagency agreement between the U.S. Department of Energy and USAFRL. The experimental research was supported by FMCSA grant DMC75-07-D-0006. The fourth author was supported by AFOSR grant FA9550-09-1-0136.

References

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., Fincham, J. M., & Douglass, S. (1999). Practice and retention: A unifying analysis. *Journal of Experimental Psychology. Learning, Memory, and Cognition, 25*(5), 1120-1136.

Åkerstedt, T. (1988). Sleepiness as a consequence of shift work. *Sleep, 11*(1), 17-34.

Bethell-Fox, C. E. & Shepard, R. N. (1988). Mental rotation: Effects of stimulus complexity and familiarity. *Journal of Experimental Psychology: Human Perception & Performance, 14*(1), 12-23.

Fu, W. T., Gonzalez, C., Healy, A. F., Kole, J. A., & Bourne Jr, L. E. (2006). Building predictive human performance models of skill acquisition in a data entry task. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting,* 1122-1126.

Gunzelmann, G., Anderson, J. R., & Douglass, S. (2004). Orientation tasks with multiple views of space: Strategies and performance. *Spatial Cognition and Computation*, *4*(3), 207-253.

Gunzelmann, G. & Gluck, K. A. (2009). An integrative approach to understanding and predicting the consequences of fatigue on cognitive performance. *Cognitive Technology*, *14*(1), 14-25.

Gunzelmann, G., Gluck, K. A., Kershner, J., Van Dongen, H. P. A., & Dinges, D. F. (2007). Understanding decrements in knowledge access resulting from increased fatigue. *Proceedings of the Annual Meeting of the Cognitive Science Society,* Austin, TX, 329-334.

Gunzelmann, G., Moore Jr, L. R., Salvucci, D. D., & Gluck, K. A. (2009). Fluctuations in alertness and sustained attention: Prediction driver performance. *Proceedings of the International Conference of Cognitive Modeling,* Manchester, UK.

Halverson, T., Gunzelmann, G., Moore Jr, L. R., & Van Dongen, H. P. A. (in press). The effects of work shift and strategy on an orientation task. *Proceedings of the Annual Meeting of the Cognitive Science Society.*

Jackson, M. L. & Van Dongen, H. P. A. (in press). Cognitive effects of sleepiness. In M. Thorpy & M. Billiard (Eds.), *Sleepiness.* Cambridge University Press.

Jastrzembski, T. S. & Gluck, K. A. (2009). A formal comparison of model variants for performance prediction. *Proceedings of the International Conference of Cognitive Modeling,* Manchester, England.

Kase, S. E., Ritter, F. E., & Schoelles, M. (2009). Caffeine's effect on appraisal and mental arithmetic performance: A cognitive modeling approach tells us more. *Proceedings of the International Conference on Cognitive Modeling,* Manchester, England, 174-179.

Lim, J. & Dinges, D. F. (2008). Sleep deprivation and vigilant attention. *Annals of the New York Academy of Science*, *1129*, 305-322.

Mednick, S., Nakayama, K., & Stickgold, R. (2003). Sleep-Dependent learning: A nap is as good as a night. *Nature Neuroscience*, *6*(7), 697-698.

Pavlik, P. I. & Anderson, J. R. (2005). Practice and forgetting effects on vocabulary memory: An activation-based model of the spacing effect. *Cognitive Science*, *29*(4), 559-586.

Tucker, A. M., Whitney, P., Belenky, G., Hinson, J. M., and van Dongen, H. P. A. (2010). Effects of sleep deprivation on dissociated components of executive functioning. *Sleep, 33*(1), 47-57.

Van Dongen, H. P. A. & Dinges, D. F. (2005). Sleep, circadian rhythms, and psychomotor vigilance. *Clinics in Sports Medicine,* 24(2), 237-249.

Walker, M. P., Brakefield, T., Morgan, A., Hobson, J. A., & Stickgold, R. (2003). Practice with sleep makes perfect: Sleep-Dependent motor skill learning. *Neuron*, *35*(1), 205-211.