

Comparing Accounts of Psychomotor Vigilance Impairment Due to Sleep Loss

Matthew M. Walsh (mmw188@gmail.com)

Glenn Gunzelmann (glenn.gunzelmann@us.af.mil)

Air Force Research Laboratory, 2620 Q Street, Building 852
Wright-Patterson Air Force Base, OH 45433

Hans P. A. Van Dongen (hvd@wsu.edu)

Washington State University Spokane, Sleep and Performance Research Center, P.O. Box 1495
Spokane, WA 99210

Abstract

The effects of fatigue on cognitive processing are not fully understood. Computational modeling research has led to two distinct accounts of fatigue, and specifically its effects on psychomotor vigilance performance, which are both supported by empirical findings. The first account is based on ACT-R and posits that fatigue increases the probability of microlapses. A biomathematical model of fatigue modulates this probability. The second account is based on a diffusion model and posits that fatigue decreases the drift rate of the diffusion process. The same biomathematical model of fatigue is used to control the drift rate parameter. We compare the models' predicted reaction time distributions to one another and to human data in a psychomotor vigilance performance task. Though they embody entirely different modeling approaches and different levels of abstraction, the accounts generate equivalent predictions and capture the detrimental effects of fatigue through mechanisms that have similar theoretical interpretations. In both accounts, fatigue effectively increases the contribution of noise to cognitive processing and decreases neural inhibition. This unexpected convergence supports a more general account of how sleep deprivation impairs psychomotor vigilance performance through degradation of the quality of cognitive processing.

Keywords: ACT-R; biomathematical model of fatigue; cognitive processing; diffusion model; leaky accumulator; local sleep.

Introduction

In this paper, we use computational models to study the effects of fatigue from sleep loss on cognition. More than 20% of the population systematically obtains insufficient sleep (Hublin et al., 2001). The societal cost of fatigue from sleep loss, in terms of property damage, lost human capital, injury, and death runs in the billions of dollars annually. Fatigue increases the risk of transportation accidents, physicians make more mistakes when tired, and fatigue has contributed to major industrial accidents (Dinges, 1995).

This topic is relevant and important from a cognitive science perspective as well. Computational models play a central role in cognitive science. Yet most models assume that the cognitive system and its various components operate effectively. This assumption is not valid when individuals are fatigued (Jackson et al., 2013). Representing the effects of behavioral moderators such as fatigue in cognitive models remains a critical challenge for basic and applied research.

Psychomotor Vigilance Deficits from Fatigue

We focus our attention on one of the most extensively studied aspects of fatigue due to sleep loss, that is, deficits in psychomotor vigilance performance. This is well captured by the psychomotor vigilance test (PVT), a 10-minute simple reaction time task in which stimuli are presented with random inter-trial intervals ranging from 2 to 10 seconds (Lim & Dinges, 2008). Participants monitor a display for the presentation of a stimulus, after which they respond as fast as possible. Performance impairment on the PVT demonstrates a hallmark effect of fatigue on cognition, namely increased performance instability. That is, the distribution of reaction times on the PVT, which has a long right tail even at baseline, becomes increasingly skewed to the right with greater fatigue.

When PVT results are reported, responses are typically divided among three categories: *false starts* occur before or within 150 ms of stimulus presentation, *alert responses* occur from 150 to 500 ms of the stimulus, and *lapses* occur after 500 ms of the stimulus. Fatigue due to sleep loss leads to three general behavioral changes in PVT performance. First, alert responses slow. Second, lapses occur more often. Third, participants commit more false starts. These outcomes of the PVT are highly replicable and predictable (Lim & Dinges, 2008).

Computational Models of Fatigue

Several classes of models have been proposed to account for the cognitive effects of fatigue. First, biomathematical models of fatigue predict cognitive performance based on changes in alertness over the course of hours and days (Fig. 1). Recent instantiations of these models are based on the two-process model of sleep regulation (Borbély & Achermann, 1999), in which fatigue increases with time awake (homeostatic process) and is modulated by time of day (circadian process). Model outputs are typically scaled to human performance data to predict response speed and/or accuracy. Although these models can accurately predict performance impairments, they are silent with respect to the effects of fatigue on specific cognitive components.

The second class of models uses cognitive architectures to study how fatigue affects components of cognition. Cognitive architectures specify foundational information processing mechanisms. Gunzelmann et al. (2009b)

provided an account of how the temporal dynamics of fatigue, as controlled by a biomathematical model, may impact specific information processing mechanisms in ACT-R. The integrated approach captured PVT reaction time distributions in sleep deprivation experiments and predicted performance in other tasks, too (e.g., Gunzelmann, Byrne, Gluck, & Moore, 2009; Gunzelmann Moore, Salvucci, & Gluck, 2011).

The final class of models is based on a diffusion process representing the internal flow of information during cognitive task performance. The diffusion model accounts for performance on simple and two-alternative forced choice reaction time tasks (Ratcliff & McKoon, 2008; Ratcliff & Van Dongen, 2011). By varying parameter values in the diffusion model, researchers have reproduced the effects of various experiment manipulations on choice accuracy and response time distributions (e.g., Ratcliff & McKoon, 2008). Ratcliff and Van Dongen (2011) showed that the temporal dynamics of fatigue may be seen as affecting a composite diffusion model parameter, drift rate divided by drift rate variability. This approach captured PVT reaction time distributions in sleep deprivation experiments and predicted performance in other decision tasks, too.

Research Approach

Our primary goal is to develop and compare an ACT-R model and a diffusion model that simulate the processes involved in performing the PVT. We compare these models because they are the only published models predicting the complete distribution of reaction times in this task, yet they represent different levels of abstraction and connect with different aspects of the underlying neurobiology (Van Dongen et al., 2011). By comparing these models, we seek to develop a general account of fatigue not tied to a single modeling formalism.

The ACT-R model is based on Gunzelmann et al. (2009b). We describe a version that has been updated for the latest release of ACT-R. Dynamic parameters are constrained using the output of a biomathematical model of fatigue (McCauley et al., 2013). The diffusion model is based on Ratcliff and Van Dongen (2011). We augment their model with a leaky accumulator to capture data that fall beyond its original scope. Additionally, we constrain dynamic parameters of the diffusion model with the same biomathematical model of fatigue. We use the empirical

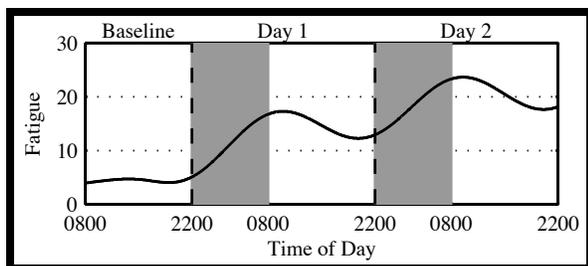


Figure 1. Biomathematical model predictions of fatigue over 62 hours of sleep deprivation (McCauley et al., 2009).

data from a 62-hour total sleep deprivation experiment described by Van Dongen et al. (2013) as the basis for model comparison.

Integrated Models

ACT-R

ACT-R is an integrated cognitive architecture that contains specialized information processing modules (Anderson, 2007). Buffers connect these modules with a central procedural module. Procedural knowledge is represented in the form of production rules. Each rule has conditions that must be met for it to be selected, and a set of actions that modify the external state of the world and the internal state of the architecture when it is performed. The temporal dynamics of cognition unfold across a sequence of production cycles. During each cycle, conditions for different productions are compared against the contents of the buffers, a production is selected, and the production fires. The resulting state of the world and architecture serve as the starting point for the next production cycle.

Task Model Gunzelmann et al. (2009b) built an ACT-R model of the PVT that contained three core productions: *wait* for the stimulus to appear, *attend* to the stimulus, and *respond*. Different conditions favored different production: *wait* fired when the screen was blank, *attend* fired when the stimulus was present but had not yet been attended to, and *respond* fired after the stimulus had been attended to. A fourth production, *respond randomly*, could fire at any time. This production, though rarely chosen because of its low utility, produced false starts.

Subsequent to that model, ACT-R underwent several major changes, including the addition of production partial matching. With production partial matching, productions whose conditions are not perfectly met remain eligible for selection, but their utility values are penalized,

$$U'_i = U_i - MMP_i + \varepsilon_i \quad (1)$$

U_i is the stored utility for production i , MMP_i is the mismatch penalty, and ε_i is logistically distributed noise. The production with highest utility is enacted if its utility exceeds the threshold,

$$\text{Production} = \max(U'_i) \text{ if } \max(U'_i) > \text{threshold} \quad (2)$$

If no production's utility exceeds the threshold, a microlapse occurs; the model becomes inactive for the duration of a production cycle (*cycle time*) before initiating a new cycle.

Because of production partial matching, *respond* can be selected at any time. This obviates the need for the *respond randomly* production: when *respond* is selected before the stimulus appears, a false start occurs. But because *respond* is subject to the mismatch penalty when its conditions are not perfectly met, this happens infrequently.

Integration Fatigue affects three components of the ACT-R

model. First, fatigue reduces productions' utility values,

$$U' = FP \cdot (U_i - MMP_i) + \varepsilon_i \quad (3)$$

This differs from Gunzelmann et al. (2009b) because the function that defines U_i in ACT-R has changed. The FP parameter is derived from a biomathematical model of fatigue (McCauley et al., 2013; Fig. 1) by the equation, $FP = 1 + FP_{slope} \cdot Fatigue$. When FP_{slope} is negative, as we found by fitting the model, production utilities decrease with fatigue. Consequently, selections are increasingly driven by noise. Also, productions' utilities increasingly fall below the utility threshold, causing more microlapses.

Second, fatigue lowers the utility threshold,

$$\text{Production} = \max(U'_i) \text{ if } \max(U'_i) > FT \cdot \text{Threshold} \quad (4)$$

The FT parameter is derived from a biomathematical model of fatigue (McCauley et al., 2013) by the equation, $FT = 1 + FT_{slope} \cdot Fatigue$. When FT_{slope} is negative, as we found upon fitting the model, the utility threshold decreases with fatigue. This partially offsets the effect of fatigue on utility values, reflecting compensation. But this also offsets the effect of the mismatch penalty on the *respond* production, allowing more false starts.

Third, when no production has sufficiently high utility to enact, a microlapse occurs. Concurrent with the microlapse, FP decreases by a small amount ($0 < FP_{dec} < 1$),

$$FP \leftarrow FP \cdot FP_{dec} \quad (5)$$

The small drop in FP increases the likelihood of a microlapse in the subsequent cycle. Across such a series of cycles, the probability of responding decreases progressively, producing behavioral lapses.

Diffusion Model

The diffusion model belongs to a class of sequential sampling models for simple reaction time tasks (Ratcliff & McKoon, 2008). In these models, information is accumulated sequentially over time, and a response occurs when accumulated information reaches a criterion. Response time is the sum of *decision time* (i.e., the duration of the diffusion process) and *non-decision time* (i.e., the duration of perceptual and motor processes).

Task Model Ratcliff and Van Dongen (2011) built a diffusion model of the PVT. Because the PVT is a one-choice reaction time task, their model included a single decision criterion, A . Evidence accumulated from zero to the criterion. Drift rate varied across trials according to a normal distribution with mean V and standard deviation η . Non-decision time was uniformly distributed with mean T_{ND} and range SD . Ratcliff and Van Dongen (2011) found that as fatigue increased, drift rate (V) decreased, and the standard deviation of drift rate (η) increased. These dynamics were captured by decreasing the ratio V/η .

Extension Ratcliff and Van Dongen's model (2011)

pertains to processing *after* the stimulus appears, and therefore makes no predictions about false starts. To produce false starts, we implemented a diffusion process that began during the pre-stimulus interval and persisted beyond stimulus presentation. Conceptually, the rate of information accumulation should increase once the stimulus appears. Thus, we set drift rate to zero during the pre-stimulus interval, and allowed drift rate to become positive after stimulus presentation.

Neurobiological models of simple decisions further informed our implementation. Passive decay is an intrinsic property of neural activity. Excitatory feedback connections offset decay. When excitatory feedback dominates, the neural circuit moves away from baseline and is said to be unstable. When decay dominates, the circuit returns to baseline and is said to be stable. One function of decay is to dampen noise from excitatory feedback connections. This ensures that representations remain near baseline in the absence of input signals. Usher and McClelland (2001) incorporated decay into a sequential sampling model of perceptual choice. Likewise, we incorporated decay into the PVT model, which reduces the likelihood of responding during the pre-stimulus interval. Preliminary simulations showed that the diffusion model with decay committed false starts about as frequently as participants did.

Integration Ratcliff and Van Dongen (2011) found that the ratio V/η , as estimated repeatedly over a period of extended wakefulness, was correlated with the output of a biomathematical model of fatigue (McCauley et al., 2013). In our simulations, we allowed η to take one value across all sessions, and V to take different values for each session (cf. supplemental material in Ratcliff & Van Dongen, 2011).

Drift rate was set to zero during the pre-stimulus interval. As such, changes in drift rate with time awake affect post-stimulus processing but not the frequency of false starts. Two other model parameters could affect false starts, however. First, if decay decreased with time awake, reduced drag on evidence accumulation would allow more decision processes to reach the criterion during the pre-stimulus interval. Second, if the decision criterion decreased with time awake, more decision processes would terminate during the pre-stimulus interval. Because Ratcliff and Van Dongen (2011) reported that the decision criterion did not vary with time awake, we allowed decay, but not the decision criterion, to take different values for each session.

We implemented the model as a random walk approximation of a diffusion process. At each time step, a displacement of Δ (toward the decision criterion) occurs with probability p , and a displacement of $-\Delta$ (away from the decision criterion) occurs with probability $1-p$. The size of the displacement depends on the granularity of the time step ($\tau = .005$ sec) and the within-trial stochastic component of the accumulation process ($s = 0.1$; Tuerlinckx et al., 2001),

$$\Delta = s\sqrt{\tau} \quad (6)$$

The probability of positive displacement depends on these factors and drift rate,

$$p = 0.5 \cdot \left(1 + \frac{V\sqrt{\tau}}{s} \right) \quad (7)$$

To implement decay, we modified the displacement term,

$$\Delta = s\sqrt{\tau} - \lambda \cdot Evidence \quad (8)$$

Evidence is the accumulated information to that point in the trial, and λ is the decay term. Decay acts in the opposite direction of accumulated information because of the negative sign, and the strength of decay scales with the amount of accumulated information.

The V parameter is derived from the biomathematical model of fatigue (McCauley et al., 2013; Fig 1) by the equation, $V = V_{intercept} + V_{slope} \cdot Fatigue$. When V_{slope} is negative, as we found by fitting the model, drift rate decreases with fatigue. Consequently, information accumulation occurs more gradually. The λ parameter is also derived from the biomathematical model by the equation, $\lambda = \lambda_{intercept} + \lambda_{slope} \cdot Fatigue$. When λ_{slope} is negative, as we found upon fitting the model, decay decreases with fatigue. This partially offsets the effect of fatigue on drift rate. But this also offsets the stabilizing effect of decay during the pre-stimulus interval, allowing more false starts.

Statistical Methods

Van Dongen et al. (2013) administered PVT sessions to participants approximately once every 2 hours over 62 hours. Each session contained relatively few observations. Therefore, we collapsed data over the circadian cycle, that is, across sessions that occurred in the same day of the experiment for each participant (Baseline: 0 to 14 hours awake; Day 1: 14 to 38 hours; Day 2: 38 to 62 hours).

Responses in the PVT occur before (false starts) or after (alert responses and lapses) the stimulus appears. We fit the models to the complete reaction time distributions. To do so, we binned reaction times corresponding to the 5% quantiles of responses from 150 ms to 500 ms after stimulus presentation for each participant and on each day of the experiment. We created cumulative distribution functions (CDFs) for each participant by calculating the proportion of trials with responses before 150 ms (i.e., false starts), and the proportions of trials with responses by each of the 20 quantiles of reaction times. We used a simplex search algorithm to find parameter values that minimized mean squared error (MSE) between the predicted and observed CDFs. Thus, for each participant, we estimated one set of parameter values across all days of the experiment, and parameters that changed across sessions did so according to the biomathematical model. Simulations leveraged large-scale computational resources (Harris, 2008).

For ease of interpretation, we display group-level data, but reported fits are at the level of individuals. We present probability density functions (PDFs) showing the proportion

of responses in 10 ms bins from 150 to 500 ms (i.e., the distribution of alert responses), as well as the proportion of trials with false starts and lapses.

Results

ACT-R The ACT-R model closely matched the observed PDFs (Fig. 2). MSE for each participant ranged from $.08 \cdot 10^{-3}$ to $.16 \cdot 10^{-3}$ with a mean of $.11 \cdot 10^{-3}$. Correlation for each participant ranged from .89 to .96 with a mean of .93.

The ACT-R model captures the three previously described effects of fatigue on PVT performance: (1) slowing of alert responses, (2) more lapses, and (3) more false starts. In the ACT-R model, the first and second effects occur because of the greater frequency of microlapses with time awake. FP_{slope} is more negative than FT_{slope} , $t(12) = 12.2$, $p < .001$ (Table 1). Consequently, production utilities drop below the utility threshold after extended wakefulness, resulting in more microlapses. Microlapses slow alert responses, and sometimes delay responses beyond 500 ms. This is accentuated by the fact that with each microlapse, the probability of another microlapse occurring increases (FP_{dec} , Eq. 5). The third effect occurs because of the negative value of FT_{slope} , $t(12) = 6.7$, $p < .001$. Lowering the utility threshold partially offsets the effect of fatigue on production utilities. But lowering the threshold also reduces the inhibitory influence of the mismatch penalty on the *respond* production, leading to more false starts.

Table 1: Mean and Standard Deviation over Subjects of Parameters for the ACT-R and Diffusion Models

| ACT-R Model | | Diffusion Model | |
|-------------------|-----------------|-----------------------|--|
| Parameter | Value | Parameter | Value |
| U_i | $4.96 \pm .50$ | T_{ND} | $.12 \pm .03$ |
| FP_{slope} | $-.01 \pm .01$ | SD | $.09 \pm .03$ |
| <i>Threshold</i> | $4.35 \pm .46$ | A | $.10 \pm .01$ |
| FT_{slope} | $-.005 \pm .01$ | V_{slope} | $-.02 \pm .01$ |
| FP_{dec} | $.98 \pm .01$ | $V_{intercept}$ | $1.30 \pm .11$ |
| <i>Cycle time</i> | $.04 \pm .01$ | λ_{slope} | $-3 \cdot 10^{-4} \pm 2 \cdot 10^{-4}$ |
| | | $\lambda_{intercept}$ | $.04 \pm .01$ |
| | | η | $.16 \pm .07$ |

Diffusion Model The diffusion model also closely matched the observed PDFs (Fig. 2). MSE for each participant ranged from $.07 \cdot 10^{-3}$ to $.12 \cdot 10^{-3}$ with a mean value of $.10 \cdot 10^{-3}$. Correlation for each participant ranged from .91 to .96 with a mean value of .94.

Like the ACT-R model, the diffusion model captures the three previously described effects of fatigue on PVT performance. The first and second effects occur because of the more gradual accumulation of information with time awake. V_{slope} is negative, $t(12) = 9.0$, $p < .001$ (Table 1), so that although drift rate is initially high, it decreases with time awake. Consequently, the reaction time distribution becomes more skewed to the right. The third effect, increased false starts, occurs because of the negative value of λ_{slope} , $t(12) = 3.7$, $p < .01$ (Table 1). Reducing decay partially offsets the effect of fatigue on drift rate, but this

also creates greater instability. Noise thus begins to drive the diffusion process beyond the decision criterion during the pre-stimulus interval, leading to more false starts.

Comparison Better than the fits of the models to the data were their fits to each other. MSE between the models' PDFs ($.05 \cdot 10^{-3}$) was lower than MSE between either model and the data. Likewise, the correlation between the models' PDFs (.97) exceeded the correlation between either model and the data. Moreover, the models' residual errors were highly correlated, $r = 0.85$. Thus, there was substantial overlap between the models in terms of what they succeeded and failed to account for in the data (Fig. 2).

Discussion

In this paper, we compared two computational models of the effects of fatigue from sleep loss on psychomotor vigilance performance. The first uses ACT-R; the second uses a diffusion process. In both accounts, a biomathematical model controls dynamic parameters affected by sleep loss.

On the surface, the accounts appear dissimilar. They differ in whether they treat the decision to respond as a repeated or unitary event. In ACT-R, a production only has effect when its value exceeds the utility threshold. Slow responses occur because the productions *attend* and *respond*, though repeatedly selected, have insufficient utility to enact. The diffusion model represents decision time as the duration of a single diffusion process. Slow responses occur because the diffusion process moves inconsistently toward the decision boundary when drift rate is low.

The accounts also differ in whether they treat evidence accumulation as discrete or continuous. In ACT-R, the *respond* production fires once its value exceeds all other productions and the utility threshold. This depends largely on moment-by-moment fluctuations in noise (Eq. 1). The diffusion model initiates a response when accumulated information exceeds the decision criterion. Accumulation occurs gradually in the moments leading up to the response.

Notwithstanding these differences, the accounts capture the negative effects of fatigue through the same basic underlying mechanisms. First, fatigue increases the relative contribution of noise to the decision process. In ACT-R, this is achieved by decreasing production utilities. In the diffusion model, this is achieved by decreasing drift rate. Second, fatigue reduces neural inhibition. In ACT-R, this is accomplished by reducing the utility threshold. In the diffusion model as implemented here, this is accomplished by reducing decay. The unexpected convergence present across modeling formalisms supports a more general account of how fatigue from sleep loss impairs psychomotor vigilance performance.

Connection with Local Sleep Theory

According to a recent theoretical proposal, cognitive deficits from fatigue due to sleep loss are caused, at least in part, by use-dependent, local sleep interfering with cognitive processing (Van Dongen, Belenky, & Krueger, 2011). Sleep occurs locally at the level of cortical columns in response to neuronal use, while overall the brain maintains wakefulness. This degrades the processing capacity of neuronal circuits that subservise the cognitive processes associated with the performance task at hand (Chee & Asplund, 2013). The ACT-R and diffusion modeling accounts described in this paper are both consistent with this theory.

The local sleep theory also posits that use-dependent local sleep underlies the time-on-task effect (i.e., the increase of performance impairment and instability across task duration) and its interaction with sleep deprivation, as has been documented for the PVT (Van Dongen, Belenky, & Krueger, 2011). The time-on-task effect was not adequately explained by earlier theories of the effects of fatigue on cognition, and is also not accounted for in the ACT-R and diffusion models described here because neither embodies use-dependence. However, it should be straightforward to incorporate use-dependence in future versions of these models. Indeed, a proposal already exists for capturing time-

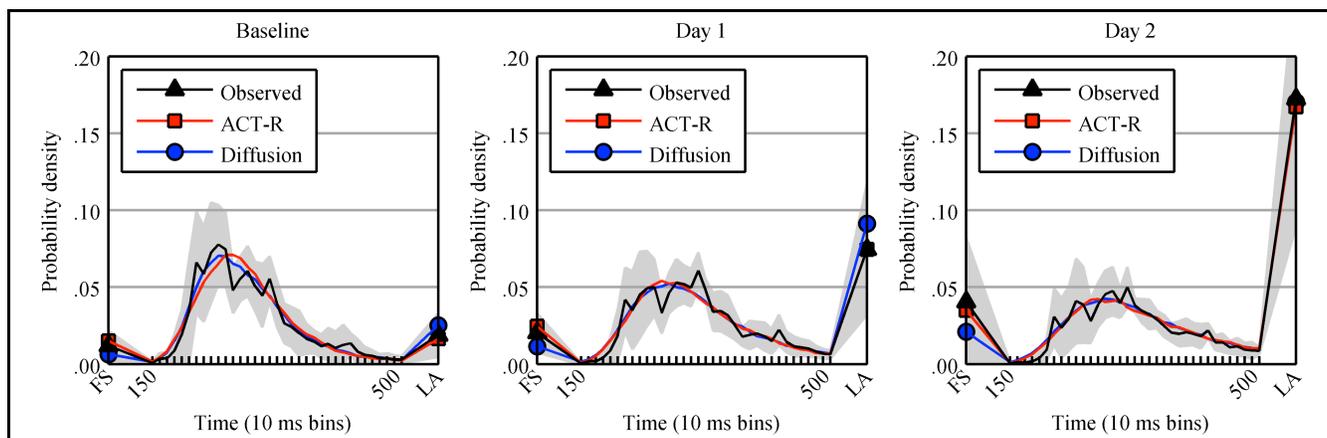


Figure 2. PVT response time distributions across 62 hours of total sleep deprivation (± 1 STD), averaged over each day of the experiment. The first bin shows the proportion of false starts (FS), and the final bin show the proportion of lapses (LA). The middle bins show the proportion of responses occurring in 10 ms intervals from 150 to 500 ms.

on-task effects on psychomotor vigilance in ACT-R using the mechanisms described here (Gunzelmann et al., 2010).

Future Research

The present research is based on a single data set that was used to fit parameters of the integrated computational model accounts. Work currently in progress will employ independent data sets for validation. In this further research we will also pursue a formal mathematical comparison of the reaction time distributions generated by the ACT-R and diffusion model approaches.

Both the ACT-R model and the diffusion model described here are based on putative principles of cognitive functioning and would be expected to generalize to tasks other than the PVT (Gunzelmann et al., 2011; Ratcliff & Van Dongen, 2011). Although encompassing different levels of abstraction, the two models both describe how the temporal dynamics of fatigue from sleep loss affect components of cognition. These components are evoked in the PVT, and they are evoked in various other tasks. As such, these computational modeling approaches have considerable predictive potential – because the same components of cognition are used in different tasks and contexts, theories of how fatigue affects those components can be leveraged to predict performance in novel circumstances. This potential has already been demonstrated for the ACT-R model (Gunzelmann et al., 2009a; Gunzelmann et al., 2011). Even so, the extent of generalizability of the two distinct accounts discussed here needs to be documented, and we may find that this ultimately differentiates them.

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. This research was supported by the AFRL's Warfighter Readiness Research Division and by NIH grant R01HL105768, and in part by AFOSR grant FA9550-09-1-0136 and ONR grant N00014-13-1-0302. Matthew Walsh held a National Research Council Research Associateship Award with the AFRL while conducting this research.

References

- Anderson, J. R. (2007). *How can the human mind occur in the physical universe*. New York, NY: Oxford University Press.
- Borbély, A. A., & Achermann, P. (1999). Sleep homeostasis and models of sleep regulation. *Journal of Biological Rhythms*, 6, 559-570.
- Chee, M. W. L., & Asplund, C. L. (2013). Neuroimaging of attention and alteration of processing capacity in sleep-deprived persons. In E. Nofzinger, P. Maquet, & M. J. Thorpy (Eds.), *Neuroimaging of sleep and sleep disorders* (pp. 137-144). Cambridge: Cambridge University Press.
- Dinges, D. F., (1995). An overview of sleepiness and accidents. *Journal of Sleep Research*, 4, 4-14.
- Gunzelmann, G., Byrne, M.D., Gluck, K.A., & Moore, L.R. (2009a). Using computational cognitive modeling to predict dual-task performance with sleep deprivation. *Human Factors*, 51, 251-260.
- Gunzelmann, G., Gross, J. B., Gluck, K. A., & Dinges, D. F. (2009b). Sleep deprivation and sustained attention performance. Integrating mathematical and cognitive modeling. *Cognitive Science*, 33, 880-910.
- Gunzelmann, G., Moore, L.R., 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 applications* (pp. 83-96). Washington, DC: American Psychological Association.
- Gunzelmann, G., Moore, L. R., Salvucci, D. D., & Gluck, K. A. (2011). Sleep loss and driver performance: Quantitative predictions with zero free parameters. *Cognitive Systems Research*, 12, 154-163.
- Harris, J. (2008). MindModeling@Home: a large-scale computational cognitive modeling infrastructure. In *Proceedings of the 6th annual conference on systems engineering research*.
- Hublin, C., Kaprio, J., Partinen, M., & Koskenvuo, M. (2001). Insufficient sleep: a population-based study in adults. *Sleep*, 24, 392-400.
- Jackson, M. L., Gunzelmann, G., Whitney, P., Hinson, J. M., Belenky, G., Rabat, A., & Van Dongen, H. P. A. (2013). Deconstructing and reconstructing cognitive performance in sleep deprivation. *Sleep Medicine Reviews*, 17, 215-225.
- Lim, J., & Dinges, D. F. (2008). Sleep deprivation and vigilant attention. *Annals of the New York Academy of Sciences*, 1129, 305-322.
- McCauley, P., Kalachev, L. V., Mollicone, D. J., Banks, S., Dinges, D. F., & Van Dongen, H. P. A. (2013). Dynamic circadian modulation in a biomathematical model for the effects of sleep and sleep loss on waking neurobehavioral performance. *Sleep*, 36, 1987-1997.
- Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation*, 20, 873-922.
- Ratcliff, R., & Van Dongen, H. P. A. (2011). Diffusion model for one-choice reaction-time tasks and the cognitive effects of sleep deprivation. *Proceedings of the National Academy of Sciences*, 108, 11285-11290.
- Tuerlinckx, F., Maris, E., Ratcliff, R., & De Boeck, P. (2001). A comparison of four methods for simulating the diffusion process. *Behavioral Research Methods, Instruments, & Computers*, 33, 443-456.
- Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, 108, 550-592.
- Van Dongen, H. P. A., Belenky, G., & Krueger, J. M. (2011). A local, bottom-up perspective on sleep deprivation and neurobehavioral performance. *Current Topics in Medicinal Chemistry*, 11, 2414-2422.