

Sleep Deprivation and Sustained Attention Performance: Integrating Mathematical and Cognitive Modeling

Glenn Gunzelmann,^a Joshua B. Gross,^b Kevin A. Gluck,^a David F. Dinges^c

^a*Air Force Research Laboratory, Mesa*

^b*The Pennsylvania State University, University Park*

^c*University of Pennsylvania School of Medicine, Philadelphia*

Received 30 October 2007; received in revised form 27 August 2008; accepted 1 September 2008

Abstract

A long history of research has revealed many neurophysiological changes and concomitant behavioral impacts of sleep deprivation, sleep restriction, and circadian rhythms. Little research, however, has been conducted in the area of computational cognitive modeling to understand the information processing mechanisms through which neurobehavioral factors operate to produce degradations in human performance. Our approach to understanding this relationship is to link predictions of overall cognitive functioning, or alertness, from existing biomathematical models to information processing parameters in a cognitive architecture, leveraging the strengths from each to develop a more comprehensive explanation. The integration of these methodologies is used to account for changes in human performance on a sustained attention task across 88 h of total sleep deprivation. The integrated model captures changes due to time awake and circadian rhythms, and it also provides an account for underlying changes in the cognitive processes that give rise to those effects. The results show the potential for developing mechanistic accounts of how fatigue impacts cognition, and they illustrate the increased explanatory power that is possible by combining theoretical insights from multiple methodologies.

Keywords: Fatigue; Attention; Computational model; ACT-R cognitive architecture; Alertness; Psychomotor vigilance test; Sleep deprivation

1. Introduction

Sleep is a necessary part of human functioning. There is an extensive literature documenting the negative cognitive consequences associated with less-than-adequate amounts of

Correspondence should be sent to Glenn Gunzelmann, Air Force Research Laboratory — RHAC, 6030 South Kent Street, Mesa, AZ 85212-6061. E-mail: glenn.gunzelmann@us.af.mil

sleep (e.g., Dinges, Baynard, & Rogers, 2005; Durmer & Dinges, 2005), which has greatly expanded our understanding of the effects of fatigue. Research has shown that cognitive performance in a variety of functional domains gets worse as sleep deprivation increases. With increasing degrees of sleep loss, experimental participants exhibit marked decreases in vigilant attention (e.g., Doran, Van Dongen, & Dinges, 2001; Dorrian, Rogers, & Dinges, 2005), slower response time and increased errors in performing mathematical operations (e.g., Hursh et al., 2004; Thomas et al., 2000; Van Dongen, Maislin, Mullington, & Dinges, 2003), impaired memory (e.g., Drummond & Brown, 2001; Drummond et al., 2000; Drummond, Gillin, & Brown, 2001; Habeck et al., 2004; Thomas et al., 2000; Van Dongen, Baynard, Maislin, & Dinges, 2004; Van Dongen et al., 2003), and decrements in performance on naturalistic tasks where errors can and do have serious consequences outside the laboratory (Caldwell, 2003; Caldwell, Caldwell, Brown, & Smith, 2004; Dinges, 1995; Landrigan et al., 2004).

The effects of sleep loss and circadian rhythmicity have an impact on individuals' overall level of cognitive functioning, which we refer to as *alertness* in this paper. Alertness varies according to a circadian rhythm and tends to decline as the length of time awake increases (e.g., Achermann, 2004; Borbely & Achermann, 1999). To quantify these processes in this research, we use estimates derived from biomathematical models of the dynamics of alertness (see below) as an index of fatigue, with lower levels of alertness reflecting higher levels of fatigue. Characterized in this way, fatigue has been implicated in a variety of industrial disasters (see Mitler et al., 1988), including serious accidents in all transportation modalities (Caldwell, 2003; Dinges, 1995) and a significant number of motor vehicle crashes (Horne & Reyner, 1999; Klauer, Dingus, Neale, Sudweeks, & Ramsey, 2006; National Transportation Safety Board, 1995; Pack et al., 1995; Royal 2002). Despite the extensive empirical literature, the ability to accurately predict performance decrements resulting from decreased levels of alertness is sharply limited. The goal of the current effort is to address this challenge using an integration of mathematical and computational process modeling, as described in the next section.

2. Research approach

Existing modeling methodologies for understanding the impact of fatigue are inadequate, in isolation, for generating precise, in situ, quantitative performance predictions. The means of addressing this challenge presented here involves linking predictions of alertness from biomathematical models to information processing mechanisms in a cognitive architecture. This integration draws upon a rich history of research exploring the dynamics of human performance as a function of time awake and circadian rhythms, and translates those dynamics into detailed performance predictions produced through the operation of a validated set of cognitively plausible information processing mechanisms. This not only helps in understanding changes in human performance as alertness declines but also provides an account of the underlying cognitive mechanisms that are responsible for producing those changes. To evaluate our theory, we use detailed empirical data from an intense investigation

examining the impact of total sleep deprivation (TSD) on human sustained attention performance (Doran et al., 2001). As we illustrate below, our model closely captures the degradation in performance resulting from sleep loss that has been observed in individuals deprived of sleep for an extended period. First, however, we comment on the value and limitations of both biomathematical models and cognitive architectures. This is followed by a description of the task under consideration and an integrated model that has been developed to provide a more detailed account of the impact of fatigue than would be possible using either modeling approach in isolation.

2.1. Biomathematical models of alertness

Biomathematical models of alertness have been developed over many years, based upon empirical studies investigating the changes associated with time awake and circadian rhythms that occur across a variety of dependent measures (see the special issue of *Aviation, Space, and Environmental Medicine*—Neri, 2004, for an extensive review). The models reflect current knowledge regarding neurophysiologic processes and provide a quantitative estimate of the overall state of the cognitive system, which we refer to as *alertness* here. Existing biomathematical models all instantiate a two-process theory of alertness, embodying the claim that the human arousal system consists of two primary components: a circadian system and a sleep homeostasis system (e.g., Achermann, 2004; Borbely & Achermann, 1999). The circadian system oscillates in an approximately 24-h cycle, resulting in variations in alertness within a single day. The sleep homeostat system, in contrast, produces continuous declines in alertness as time awake increases, which then recovers with sleep. One of the existing models includes a detailed account of how light exposure influences alertness by affecting the phase and amplitude of the circadian pacemaker (Kronauer, Forger, & Jewett, 1999), and many provide mechanisms to shift the circadian cycle as a function of changes to the light/dark schedule (e.g., traveling across time zones; Hursh et al., 2004; Jewett & Kronauer, 1999; Kronauer et al., 1999). Fig. 1 illustrates the predictions about alertness for two of the models across 88 h of TSD.

As a result of the common underlying perspective, current biomathematical models of alertness all make qualitatively similar predictions about how increased sleep deprivation should impact overall cognitive performance (c.f., Mallis, Mejdal, Nguyen, & Dinges, 2004; Van Dongen, 2004), even though they differ somewhat in the particular dynamics of the circadian system as well as the recovery mechanism associated with sleep in the homeostat system (see Klerman & St. Hilaire, 2007). For example, the predictions from the two models in Fig. 1 are highly correlated ($r = .94$). To show the potential of our approach in evaluating alternative biomathematical models of alertness and to illustrate the existing qualitative similarities, we show model performance using parameter values constrained by the alertness predictions from each of the models shown in Fig. 1 in the results presented below. The two models are (a) the Circadian Neurobehavioral Performance and Alertness (CNPA) model (Jewett & Kronauer, 1999) and (b) the Sleep, Activity, Fatigue, and Task Effectiveness (SAFTE) model (Hursh et al., 2004). Additional details regarding the mechanisms in the




Fig. 1. Predictions of alertness for two mathematical models of fatigue: the Circadian Neurobehavioral Performance and Alertness (CNPA) model, and the Sleep, Activity, Fatigue, and Task Effectiveness (SAFTE) model.

models can be found in Jewett and Kronauer (1999) and Hursh et al. (2004) for CNPA and SAFTE, respectively.¹

The parameters in the biomathematical models can be manipulated, and the output they generate can be scaled to empirical data, which helps to illustrate the ability of these models to capture the dynamics of human performance in particular tasks. Van Dongen (2004) performed a comparison of seven of these models, using least-squares optimal scaling to assess the fit of the biomathematical models to several dependent measures for a set of fatigue protocols. In fact, one of the datasets used by Van Dongen (2004) was the 88 h TSD protocol described above. Though the results of this evaluation showed promise, there are some inherent limitations of this approach, including that the scaling process is done posthoc, it relies on existing empirical data to drive the process, and it must be performed separately for each dependent measure and task of interest. Of course, in many naturalistic contexts the behavioral data are not available and may not be possible to collect. This exposes a major drawback to these models, since those real-world contexts are precisely the situations and tasks where quantitative predictions are the most valuable. More importantly, biomathematical models provide no insight regarding the cognitive mechanisms responsible for observed declines in performance, effectively eliminating the ability to generalize the quantitative predictions to other tasks, contexts, or dependent measures.

2.2. Cognitive architectures

In contrast to biomathematical models of human alertness, computational cognitive modeling represents an approach to understanding human cognition that focuses explicitly on producing performance predictions in particular tasks and contexts (e.g., Gluck & Pew, 2005; Ritter et al., 2003). Cognitive models have been developed in a variety of areas, providing insight into the basic information processing mechanisms that support human cognition and performance (e.g., Anderson, 2007; Gray, 2007; Newell, 1990; Ritter et al., 2003). However, there have been only limited efforts within the cognitive modeling or behavior

representation community to understand how those basic information processing mechanisms are impacted by dynamic changes in alertness associated with sleep history and circadian rhythms (see Gunzelmann & Gluck, 2008, for a review). In addition, in cases where the impact of fatigue has been incorporated into such models, the predictions typically have not been compared rigorously to empirical data (e.g., French, Morris, & Hancock, 2003; Jones, Laird, & Neville, 1998).

To link the biomathematical models of alertness to a computational theory of cognition involves determining which information processing mechanisms in the architecture are impacted by decreased alertness, and how. A critical assumption of our approach is that fatigue causes existing mechanisms in the architecture to operate less effectively or efficiently, which can be represented by changing their associated parameters. Using the biomathematical models to drive the dynamics of those parameters over time serves to embed a theory of alertness in the architecture, allowing a model's performance to vary as a function of time awake and circadian rhythms. Besides allowing for the generation of in situ performance predictions, using a cognitive architecture provides a window into the information processing mechanisms in the cognitive system and how they may be affected by decreased alertness. The next section describes the task we have used to develop our theory, as well as the particular experimental study that provided the data we used to test the mechanisms.

3. Experimental protocol and task of interest

To evaluate the theory embodied in our computational model, detailed performance data from human participants across a wide range of fatigue levels is needed. For this purpose, we use the data from a study described in Doran et al. (2001), which measured human performance on a set of tasks across 88 continuous hours of TSD. Participants were 13 healthy men between the ages of 22 and 37 years with normal sleep habits. To ensure that participants had established a common sleep schedule and that they were well rested, they were asked to adhere to a set schedule involving 8 h sleep per night for 1 week prior to entering the lab, and they were monitored for compliance using a combination of actigraph data, diary entries, and time-stamped phone-ins. They entered the lab 3 days before the extended wakefulness period. For each of the three nights preceding the 88 TSD, participants were given 8 h time in bed (2330 to 0730) in a dark room. When they awoke at 0730 after the third night, they remained awake for the next 88 h. During periods of wakefulness for the duration of the study, participants completed a battery of performance evaluation tasks, lasting approximately 30 min, every 2 h. The tasks included in this battery were intended to assess various aspects of cognitive functioning and performance. In the current research, we are concerned with performance on one of these tasks, referred to as the psychomotor vigilance test (PVT).

3.1. Psychomotor vigilance test

The PVT is a sustained attention task developed to assess changes in human functioning as a consequence of fatigue (Dinges & Powell, 1985). The task has been

extensively validated to be sensitive to fatigue from partial and total sleep loss, as well as circadian misalignment (Dorrian et al., 2005). The task measures vigilant attention by requiring participants to monitor a known location on a computer monitor for a stimulus. When the stimulus appears they react to it by pressing a button. Human performance on the PVT changes systematically in response to variations in alertness (Dinges & Powell, 1985), as characterized by the biomathematical models described above. In fact, changes in PVT performance follow a similar qualitative pattern as has been identified in other dependent measures, like response times for solving simple addition and subtraction problems (Van Dongen et al., 2001; Van Dongen & Dinges, 2005), and physiological measures like core body temperature and hormone levels (e.g., el-Hajj Fuleihan et al., 1997; Khalsa, Jewett, Duffy, & Czeisler, 2000). In addition to its sensitivity to sleep and circadian-based fatigue, the PVT is commonly used because of its procedural simplicity and the consistency of individual performance apart from systematic changes associated with fatigue (e.g., Gunzelmann, Moore, Gluck, Van Dongen, & Dinges, 2008).

In the experiment conducted by Doran et al. (2001), the interstimulus interval for the PVT varied randomly (according to a uniform linear distribution) between 2 and 10 s after the previous response. The stimulus consisted of a counter that displayed the number of milliseconds elapsed since trial onset. Participants were asked to respond as quickly as possible, but to avoid making any false starts (responding before the stimulus appeared). After each response, participants were given feedback about how long it took them to respond (i.e., the number of milliseconds on the counter). Each session lasted approximately 10 min and was performed once every 2 h as part of the battery of tasks in the study. Thus, each of the 13 participants completed 44 separate 10-min PVT sessions during the period of extended wakefulness. Each session consisted of approximately 60 to 90 trials, depending on the duration of reaction times.

3.2. Human performance on the psychomotor vigilance test

To characterize PVT performance, reaction times are typically categorized into one of four groups: false starts, alert responses, lapses, and sleep attacks. In the data we present, *false starts* are defined as responses before the stimulus appears, or within 150 ms of stimulus onset.² Based on the original study, responses between 150 ms and 500 ms are identified as *alert responses*, reflecting response times associated with appropriately noticing the onset of the stimulus and eliciting a response. Reactions to the stimulus that took longer than 500 ms (up to 30 s) were considered *lapses*, suggesting that some interruption or delay in the response process interfered with responding optimally. Finally, there are some occasions where participants failed to respond even after 30 s. These are classified as *sleep attacks*. In Doran et al. (2001), trials were halted at this point and participants were alerted by a beep from the computer to alert them for the beginning of the next trial. Throughout the experiment, participants were monitored behaviorally and “repeatedly admonished to perform to the best of their ability” (Doran et al., 2001, p. 256), to ensure they maintained motivated effort during all PVT sessions. One reason for the monitoring was to make certain that

participants' eyes were open and looking at the monitor throughout each session. Thus, lapses and, to a greater extent, sleep attacks appear to represent major breakdowns in cognitive processing.

Under increasing levels of fatigue, performance can be characterized as becoming progressively more variable from stimulus to stimulus (Doran et al., 2001). In general, the distribution of responses tends to shift to the right, producing three characteristic effects (Fig. 2). There is a decrease in the proportion of alert responses, with a small corresponding increase in median reaction time for responses in this category. Along with this shift, there is an increase in lapses, suggesting that participants' vigilant attention is unstable and frequently interrupted. As sleep loss becomes severe, there is often a small but notable increase in the proportion of trials classified as sleep attacks. Sleep attacks are important because their increased prevalence, though small on average, illustrates a dramatic breakdown in information processing activity. In conjunction with these three effects that illustrate a general trend to longer responses, the prevalence of false starts as a proportion of responses increases as well.

Performance decrements on the PVT reflect a common perspective on the causes of fatigue-related decrements in performance, specifically that they stem from a combination of cognitive slowing and lapses in cognitive performance (e.g., Dinges & Kribbs, 1991; Horne, Anderson, & Wilkinson, 1983; Kjellberg, 1977; Lisper & Kjellberg, 1972; Wilkinson, 1968). From this perspective, the shift to the right in the distribution of alert responses is hypothesized to result from a slowdown in cognitive and/or perceptual-motor processes, leading to longer execution times for the task. The increases in lapses and sleep attacks are attributable to more frequent and longer lapses in cognitive performance (Williams, Lubin, & Goodnow, 1959), lasting from seconds to tens of seconds. This latter phenomenon has also been cast as evidence for "state instability" (Doran et al., 2001). In this view, increased fatigue results in "a fundamentally unstable state that cannot be characterized as either fully awake or asleep; that fluctuates within seconds; and that can rapidly progress to




Fig. 2. Human performance on the psychomotor vigilance test across 88 h of total sleep deprivation (TSD), averaged over each day for the protocol ($n = 13$).

physiological sleep if not resisted'' (Doran et al., 2001; p. 264). Finally, the increase in False starts has been attributed to attempts at compensation leading to unintentional responses before the target appears (e.g., Doran et al., 2001; Horne & Pettitt, 1985).

4. Task model and Fatigue mechanisms

In this section, we present a model to perform the PVT and a set of mechanisms to account for changes in performance associated with fatigue. The model is available online as supplementary material for the article, which can be accessed here: <http://www.cogsci.rpi.edu/CSJarchive/supplemental/index.html>. The mechanisms reflect a variation of the lapse and state instability views of fatigue, combined with a mechanism to represent attempts to compensate for the negative impact of decreased alertness. Our variation of lapsing/state instability only differs in that the state of our model fluctuates within tens of milliseconds, not seconds as proposed by Doran et al. (2001). With this minor adaptation of the theory, we have found it unnecessary to posit a separate cognitive slowing mechanism to accurately account for changes in human behavior on the PVT, as has been proposed in the past (e.g., Dinges & Kribbs, 1991; Doran et al., 2001). The model is described in this section, followed by an illustration of how it captures the changes in human performance on the PVT over the course of 88 h without sleep. In the general discussion, we use the model to assess the prospects for developing a comprehensive account of fatigue within a cognitive architecture, and discuss future research aimed at that general objective.

4.1. Adaptive control of thought-rational

We use the adaptive control of thought-rational (ACT-R) cognitive architecture in this research to represent a theory of the mechanisms of human cognition (Anderson et al., 2004). It is a theory of human cognition that has been implemented as a running simulation (Anderson et al., 2004). Models developed using ACT-R have been used to provide quantitative accounts of human performance across a broad range of tasks and research areas (see Anderson & Lebiere, 1998, for a review of some of these). At its core, ACT-R is a production system, with a number of specialized modules linked to it. The production system is serial, representing cognition as a sequence of recognize-decide-act cycles. On each cycle, all productions are compared against the state of the system, represented by the contents of a set of buffers containing memories, goals, and perceptions. From the set of productions matching the current state, a single production is selected (using an expected utility function described below) and then executed (fired). This process results in some modification to the state (a change in one or more buffers), and potentially an action (such as pressing a key or shifting visual attention). Once a cycle has completed, the process begins again with the new system state.

The buffers in ACT-R serve as the interface between specialized processing modules and the central production system. For example, there are modules for vision, motor action, and declarative memory. Activity in the modules is driven mostly by requests

passed by the production system through the buffers. The results of the processing in the modules are deposited back in the buffers, where they can be accessed by central cognition. Like the central production system, each module operates in a serial manner—each module can only process a single request at a time, and each buffer can only hold a single piece of information (*chunk*) at a time. Still, there is parallelism in ACT-R, since all of the modules are able to operate simultaneously. A schematic view of the current architecture is presented in Fig. 3.

There are a number of particular features of ACT-R that facilitate the development of a computational account of fatigue. First, ACT-R has perceptual and motor modules, which allow it to interact directly with computer-based tasks. In the case of the PVT, this is essential since human performance relies extensively on perceptual-motor processes. Second, ACT-R’s symbolic level of representation is supported by an underlying subsymbolic layer, which instantiates neural-like computations that give ACT-R the ability to show graded variations in behavior. Graded degradations are typically what are observed in humans under conditions of fatigue (e.g., Fig. 2; Lee & Kleitman, 1923). Third, Anderson et al. (2004; Anderson, 2007) provide a mapping of ACT-R mechanisms to brain regions (see Fig. 3), which is supported by extensive research. While this mapping is still an active area of research, the theoretical commitment to linking cognition and neuropsychology provides a natural link between ACT-R and existing research on fatigue, which has examined in detail the brain’s response to sleep deprivation (e.g., Drummond & Brown, 2001; Drummond et al., 2000, 2001; Habeck et al., 2004; Portas et al., 1998; Thomas et al., 2000). Next, we describe some of ACT-R’s mechanisms in more detail, and then discuss how we are manipulating those mechanisms to account for the deleterious effects of fatigue on human performance on the PVT.

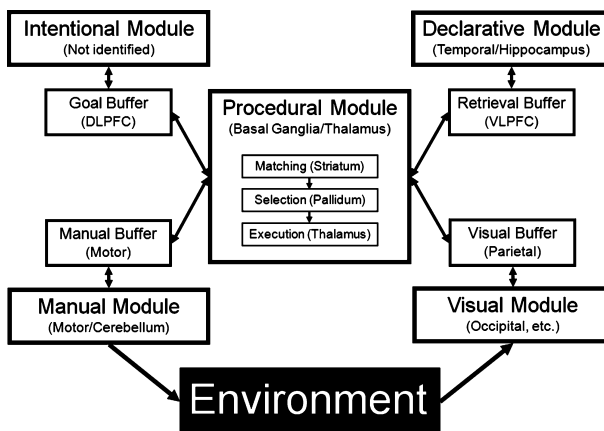


Fig. 3. Schematic illustration of the adaptive control of thought-rational cognitive architecture, including proposed mapping to brain regions. In the figure, VLPFC refers to “Ventrolateral prefrontal cortex” and DLPFC stands for “Dorsolateral prefrontal cortex.”

4.2. A model for the psychomotor vigilance test

For ACT-R to perform any given task, it is necessary to construct a model containing the knowledge needed to accomplish it. For the PVT, the knowledge is straightforward, because at any point in a trial there is a single correct action that can be taken. During the delay period, the correct action is to wait. Once the stimulus has appeared, the model must first recognize its presence, and then make a key press to respond. In the model, there are individual productions that execute each of these actions at appropriate points in the task:

1. The first production (*wait*) explicitly waits during the delay.
2. The second production (*attend*) shifts attention to the stimulus.
3. The third production (*respond*) executes a key press once the stimulus has been attended.

The conditions of these productions constrain them to fire only when they are appropriate, based on the model's representation of the current state of the trial. Thus, the *wait* production only fires when nothing is visible on the screen, the *attend* production fires only when something is on the screen that has not been attended, and the *respond* production is executed only when the stimulus has been attended. These productions allow the model to correctly perform the task. The architectural mechanisms in ACT-R that control the timing and execution of these processes result in performance that is consistent with alert human performance.

These productions are nearly sufficient to enable the model to exhibit the full range of human behavior that is observed in the PVT. However, because the *respond* production requires that a stimulus be attended on the screen, it is unable to produce false starts. To capture these anticipatory actions, an additional production was included in the model, which generates a key press regardless of whether the stimulus is present on the screen (*press-key*). Randomly making a response represents a possible alternative approach to the task, albeit highly ineffective. We recognize that there are other possible explanations for the occurrence of false start in the human data, such as diminished inhibitory control of planned motor actions (e.g., Chua, Venkatraman, Dinges, & Chee, 2006), though there are data suggesting the opposite effect (e.g., Manganotti, Palermo, Patuzzo, Zanette, & Fiaschi, 2001). We have chosen to implement the alternative described here to minimize the modifications to the ACT-R architecture itself.

The four productions represent the procedural knowledge implemented in ACT-R to perform the task. In addition to this, only a single chunk of declarative knowledge is needed. This chunk represents the task in the goal buffer that is part of ACT-R's intentional module. In essence, this chunk establishes the task context (i.e., "do the PVT") and enables ACT-R to do the minimal amount of tracking of the task state that is required (e.g., noting when a response has been made and preparing for the next trial). The knowledge in the model does not change as a function of fatigue. Instead, it is the availability of that knowledge, or the likelihood of it being used, that is altered as a consequence of sleep loss. This is consistent with the effects of fatigue on human cognition and performance. In the case of the PVT, this

applies specifically to the application of production rules to performing the task. ACT-R's subsymbolic procedural mechanisms influence the availability and use of procedural knowledge, providing a means of manipulating this process to produce graded changes in the model's performance. The set of mechanisms we propose in our account is the focus of the next subsection.

4.3. Computational mechanisms for fatigue

The research presented here describes the first step in the process of identifying information processing mechanisms that are impacted by fatigue, which is to address this challenge in the context of a specific task that avoids much of the complexity associated with more dynamic, naturalistic tasks (i.e., the PVT). The mechanisms described here represent the outcome of a process that included consideration of psychological theory (e.g., Dinges & Kribbs, 1991; Doran et al., 2001; Heuer, Kleinsorge, Klein, & Kohlisch, 2004; Williams et al., 1959), neuropsychological data from the literature (e.g., Portas et al., 1998; Wu et al., 1991), as well as theoretical commitments that have guided the development of ACT-R (Anderson et al., 2004). The first constraint, however, is the task itself. Cognitively, the PVT relies on perceptual and motor activities, and taxes the selection and execution of procedural knowledge in ACT-R, which is driven by the production execution cycle. Thus, manipulations to many existing parameters, for instance within the declarative memory module, would have no impact on the model's performance. However, the production cycle is critical since a rapid sequence of operations, driven by production firings, is needed to produce a fast response. Thus, our account focuses on this component of ACT-R, which is shown at the center of Fig. 3.

4.3.1. Decreased alertness

To constrain our search for parameters in the production execution cycle, we used existing neurophysiologic and psychological research on fatigue to identify particular components of this process likely to be impacted by changes in alertness. Prior research in ACT-R has identified a parameter used in the selection phase of the production cycle as being related to "arousal" or "motivation" (Belavkin, 2001; Jongman, 1998). Production selection occurs once productions that match the current system state have been identified. The selection process is controlled by calculating an expected utility (U) for each candidate production (i) and selecting the production that has the highest utility. The equation for expected utility in ACT-R is:

$$U_i = P_i G - C_i + \varepsilon$$

In this equation, P_i represents the probability that the goal will be achieved with production i (ranging from 0 to 1), and C_i represents the anticipated cost of achieving the goal, should that production be fired. We use the default value of 50 ms for C_i for all productions, which corresponds to the well-established default cognitive cycle time in ACT-R (e.g., Anderson et al., 2004; Anderson & Lebiere, 1998). Noise (ε) is added to the calculation of utility to produce stochasticity in the selection process. It is selected from an approximately normal

distribution with a mean of 0 and a standard deviation of approximately 0.453, based upon a commonly used noise parameter in ACT-R.³ The parameter G is the one used by Belavkin (2001) and Jongman (1998) to represent motivation or arousal, and which is conceptualized as alertness in the current work. The only constraint on G is that it must be positive, but it has been cast as “the value of the goal” in seconds, which provides some practical constraint on values. The lower bound on U_i , then, is $-C$ (when $P_i = 0$ or as G approaches 0). There is no explicit upper bound.

Though we manipulate the same parameter, the ways in which manipulations to G impact the model’s performance are different in some respects relative to previous ACT-R models. For instance, Jongman (1998) decremented G to represent decreasing motivation. The consequence of this was to increasingly bias the utility mechanism to lower cost options. In Jongman’s model, this produced a strategy shift to alternatives that were lower in cost, but less effective. This idea corresponds with findings from the sleep research community, where similar transitions have been observed (e.g., Horowitz, Cade, Wolfe, & Czeisler, 2003). This process is related to the likelihood of producing a false start as a consequence of the *press-key* production firing in the PVT model. In the model, $P_i = 0$ for the *press-key* production, reflecting the idea that it is very unlikely to lead to success on the task. For each of the other productions, which are appropriate actions for the task, P_i was set to the default value of one. The consequence of this is that decrements to G have no impact on U_i for the *press-key* production because $P_i = 0$. However, decrements to G do produce lower U_i -values for the other productions. Thus, one consequence of reduced alertness in this model (lower G -values) is that ACT-R is less able to discriminate between the appropriate *wait* production and the less-appropriate *press-key* production during the delay period, leading to a higher probability of producing a false start.

In contrast to Jongman’s work, Belavkin (2001) decremented G continuously as a model performed a task to account for “giving-up” behavior. In ACT-R, the utility (U_i) of the selected production must exceed the utility threshold, T_u , to be executed. In Belavkin’s work, when G was reduced to a level where no productions exceeded T_u , the model stopped running and effectively gave up on the task. However, whereas participants in Belavkin’s work eventually gave up, participants in sleep deprivation studies report being motivated and trying to perform the tasks (e.g., Doran et al., 2001), despite being unable to maintain performance as sleep deprivation increases. To allow ACT-R to persevere despite low levels of alertness, we modified the production cycle in ACT-R slightly to allow the system as a whole to continue operating despite the occurrence of cycles where no productions exceed T_u . In our modification, this situation produces a “micro-lapse,” which lasts for the duration of an ACT-R cognitive cycle (50 ms, plus noise⁴). Following a micro-lapse, a new production cycle is initiated. The noise (ϵ) added to the utility computation for each production on each cycle makes it possible for a micro-lapse to be followed by a cycle in which a production successfully fires. Isolated micro-lapses are responsible for the small increases reaction times for the task. Longer sequences of micro-lapses generate more substantial delays (lapses), or even complete failures to respond (sleep attacks).

Micro-lapses represent an implementation of the state instability view of fatigue, with the model transitioning between wakefulness and what amounts to sleep (micro-lapses) in rapid

succession. As noted above, however, our account involves transitions that are potentially much more frequent (each cognitive cycle, or approximately 50 ms) as compared to the prevailing view in the sleep research community (lapses lasting at least several seconds; Dinges & Kribbs, 1991; Doran et al., 2001; Williams et al., 1959), although the possibility of very brief lapses in cognitive processing has been considered by some (e.g., Heuer et al., 2004). This is why we refer to them as “micro-lapses” in our model. In this view, PVT lapses and sleep attacks are the result of long sequences of micro-lapses, rather than monolithic breakdowns in cognitive performance.

The value of G is reduced as sleep deprivation increases to represent decreased alertness. However, the range of human performance within each session suggests that a person’s level of alertness varies within trials and sessions as well. To capture this, we decrement G within a trial (by a small, prespecified unit) after each micro-lapse. The idea here is that a lapse in cognitive processing provides an opportunity for alertness levels to decrease even further, representing the process of falling asleep. This small drop in G means that it is more likely that a micro-lapse will occur on the subsequent cycle. Across a series of such cycles, it gradually becomes very unlikely that a response to the stimulus will be made, which contributes to the occurrence of lapses and, ultimately, sleep attacks. Another way of stating this is that experiencing a micro-lapse (a transition into a sleep state) increases the likelihood of remaining in that state or transitioning back into that state at a later point in the trial. This reflects the view that state instability “can rapidly progress to physiological sleep if not resisted” (Doran et al., 2001, p. 264). To complete the model description, a mechanism to represent resistance to sleep (i.e., compensation) is described next.

4.3.2. *Compensation for reduced alertness*

The final component of our account relates to the utility threshold (T_u). There is converging evidence to suggest a role for this parameter in compensatory behavior in sustained attention tasks like the PVT. First, T_u is critically related to production execution in ACT-R. Changes to T_u affect the probability that the selected production will be fired on a given cognitive cycle. Thus, reducing T_u can serve as a compensatory mechanism by allowing productions with lower utilities to fire, which can offset some of the negative effects of lower G -values. ACT-R’s production rules have been hypothesized as being instantiated in the cortico-striatal-thalamic loop (e.g., Anderson, 2007; Anderson, Fincham, Qin, & Stocco, 2008), with production execution relying on the thalamus (Anderson et al., 2004). This theoretical link can be combined with empirical evidence on fatigue, which shows evidence for compensatory behavior in response to decreased alertness, both behaviorally and neurophysiologically (e.g., Chee & Choo, 2004; Chua et al., 2006; Doran et al., 2001; Drummond, Brown, Salamat, & Gillin, 2004; Drummond et al., 2000; Portas et al., 1998). The research illustrates that efforts at compensation can help to insulate individuals from some of the performance decrements associated with decreased alertness for short periods (e.g., Drummond et al., 2004; Portas et al., 1998; Thomas et al., 2000) but may also lead to distinct changes in performance, like the increased occurrence of false starts in the PVT (e.g., Doran et al., 2001). Importantly, researchers have associated this compensation process with the thalamus (e.g., Portas et al., 1998). In discussing changes in activity found in the thalamus, Portas

et al. (1998) concluded, “This process may represent a sort of compensatory mechanism... We speculate that the thalamus has to ‘work harder’ in conditions of low arousal to achieve a performance that is equal to that obtained during normal arousal” (Portas et al., 1998, p. 8987).

Reflecting these findings, compensation is represented in our model by decreasing the value of T_u . With a lower threshold, it is easier for productions to fire, since lower values for U_i will still exceed T_u . Thus, it becomes more likely that the appropriate action will occur given a particular level of alertness. As we describe next, however, this compensatory mechanism has the secondary effect of increasing the likelihood of executing the inappropriate action as well, contributing to the increase in false starts. In addition to the correspondence with neurophysiologic data, this relationship is aligned with the explanation for increased false starts under sleep deprivation reported by Doran et al. (2001).

5. Evaluating the model and mechanisms

5.1. Impact of proposed mechanisms

As an initial evaluation of the model, its behavior is compared to the empirical data shown in Fig. 2. The corresponding model data is presented in Fig. 4. This figure illustrates the capacity for the model to capture the four primary phenomena of interest in the PVT. Specifically, the model produces more false starts, lapses, and sleep attacks, combined with increased median reaction times for alert responses, as sleep deprivation increases. The magnitudes of these effects in the model are also similar to those seen in the empirical data. The overall correspondence between the data in the two figures is quite good ($r = .99$, $\text{RMSD} = 0.0047$), which further supports the theoretical account instantiated in the model.




Fig. 4. Adaptive control of thought-rational model capturing the four phenomena of interest in the psychomotor vigilance test, averaged over each day for the protocol.

The data shown in Fig. 4 illustrate the sufficiency of the mechanisms in the model for accounting for the performance decrements in humans under conditions of fatigue. To produce these performance predictions, we explicitly manipulated G , T_u , and the magnitude of the decrement imposed on G as a consequence of a micro-lapse. The best-fit values for G and T_u for each day of the study are presented in Table 1 based upon the root mean squared deviation (RMSD) between the model’s performance and the human data shown in Fig. 2. The G -values in Table 1 represent the “starting values” for G on each trial to model human performance for the particular day of the study. From these starting values, G was decremented as a consequence of micro-lapses. The magnitude of this decrement was kept constant across all days at 0.035. In the most extreme cases, G was reduced to a minimum of 0.0001 after a series of micro-lapses, due to the ACT-R constraint that G remain positive. Once this value was reached, G was maintained at this level until the start of the next trial (at which point it was reset to the starting value based upon the day, as listed in Table 1). Next we describe the model’s explanation for the changes observed in each of the four dependent measures of interest.

5.1.1. False starts

False starts in the model increase as a consequence of the interaction of two processes. As noted above, lower G -values decrease the discriminability between the appropriate action (*wait*) and the inappropriate one (*press-key*) since changes to G do not impact U_i for *press-key*. The more influential mechanism, however, is the reduction of T_u to reflect compensation, because this manipulation makes it more likely that U_i will rise above T_u for the *press-key* production. Thus, our explanation for the rise in false starts is that it is partly an outcome of decreased alertness, but that it is primarily an unintended consequence of attempts by individuals to compensate for worsening performance. An earlier version of the model that does not incorporate compensatory reductions in T_u , presented in Gunzelmann, Gluck, Van Dongen, O’Connor, and Dinges (2005), illustrates this. While the performance of the model overall was in line with human performance, it failed to capture the magnitude of the increase in false starts. Fig. 4, in contrast, shows that the model with this compensatory mechanism captures the effect quite well.

Table 1
Best-fitting parameter values for G and T_u for capturing the impact of fatigue after 0, 1, 2, and 3 days of TSD

Day	G	T_u
Baseline (0 days TSD)	1.98	1.84
Day 1 TSD	1.80	1.78
Day 2 TSD	1.66	1.70
Day 3 TSD	1.58	1.64

Notes: TSD, total sleep deprivation.

5.1.2. Increased median reaction times

Longer alert reaction times in human performance on the PVT are a major piece of evidence used to argue for cognitive slowing in response to decreased alertness (e.g., Dinges & Kribbs, 1991). However, generalized slowing is unnecessary for explaining changes in median reaction time on the PVT in our model. The shift in median reaction time in the model actually arises because of the immediate impact of lower values of G on the model's performance, which results in a decreased probability of successfully responding to the stimulus at the first opportunity after it appears. That is, the probability of a micro-lapse increases for the cognitive cycle immediately following the presentation of the stimulus. A similar idea has been raised in the sleep research community as an alternative to general cognitive slowing (e.g., Heuer et al., 2004). In the model, the effect is compounded by the decrements that are made to G when micro-lapses occur. The consequence for alert responses is that the overall proportion decreases, and the distribution is increasingly skewed toward longer responses (see Fig. 4), resulting in an increase in median reaction time within the alert response range.

5.1.3. Increased likelihood of lapses and sleep attacks

Lower values for G have the initial impact of decreasing the probability of the *fastest* alert responses. However, as G is decremented during the course of a trial, it becomes increasingly unlikely that a response will be made at all, which extends the tail of the response distribution in the model. The consequence is that a higher proportion of responses become lapses. In extreme cases, G is decremented to its lowest possible value and the appropriate productions for completing the task fail to fire for 30 s, producing sleep attacks. This does not happen frequently in the model, but the increase is similar to the increase observed in humans (Fig. 4).

5.2. Integrating biomathematical models with ACT-R

Fig. 4 shows the general ability of the mechanisms we have proposed to capture performance characteristics related to fatigue at the granularity of whole days. In this section, we evaluate the viability of the model in more detail by examining the human data at the level of individual sessions, which reflects the impacts of both sleep deprivation and circadian rhythm on human performance.

5.2.1. Using biomathematical models to predict parameters

Using biomathematical models, quantitative conclusions regarding human behavior are only possible by scaling the output to existing performance data using regression techniques to align the predictions with the observed data. As an alternative to this approach, we used the biomathematical models to constrain the dynamics of changes to parameter values in ACT-R (G and T_{ii}) to capture performance changes across the sleep deprivation protocol. We applied the same technique in parallel using CNPA and SAFTE. We began by identifying the best and worst sessions in the human performance data, based upon the average rank-orderings of performance on each of the dependent measures of interest in the data.

Using this criterion, the best human performance was observed at 1200 on the baseline day, while the worst performance was at 0600 following the third without sleep. Because of the lack of relevant prior research to draw upon for estimating parameter values, the process began by explicitly fitting the data for those two sessions by manipulating the G and T_u parameters in ACT-R. The decrement applied to G as a consequence of a micro-lapse was maintained at 0.035, which was estimated in generating the model fits shown in Fig. 4 to the aggregate human data shown in Fig. 2. For all evaluations, fit was assessed using RMSD values comparing the model to human data based upon binning the data for individual sessions, as illustrated in Figs. 2 and 4 for whole days. The best-fitting parameter values, along with the resulting correlations and RMSD for those two sessions are shown in Table 2.

Table 3 presents the minimum and maximum alertness values predicted by the two biomathematical models for the 88 h TSD period. We used these values in conjunction with the best-fitting parameter values in Table 2 to generate a function to map between the two. To simplify the process of integrating the predictions of the biomathematical models with ACT-R parameters, we assumed a linear relationship between them. Thus, by using the best-fitting ACT-R parameter values from just two sessions, we were able to derive values for G and T_u for all 44 sessions by using a simple linear scaling.⁵ As described above, G and T_u both decrease with lower levels of alertness, reflecting lower levels of alertness and increasing

Table 2
Parameter values involved in estimating linear functions to map biomathematical model predictions of alertness to G and T_u values in ACT-R (TSD)

	Best-Fitting Parameters		Fit to Human Data	
	G	T_u	r	RMSD
Best performance (day 0 @1200)	2.12	1.94	.93	0.035
Worst performance (day 3 @0600)	1.52	1.60	.99	0.021

Notes: ACT-R, adaptive control of thought-rational; TSD, total sleep deprivation; RMSD, root mean squared deviation.

Table 3
Predicted range of alertness across 88 h of TSD for the CNPA and SAFTE models

	CNPA (Ranges From 0–1)	SAFTE (Ranges From 0–100)
Highest alertness	0.950 (1800, baseline day)	99.87 (2000, baseline day)
Lowest alertness	0.375 (0800, day 3 TSD)	15.17 (0400, day 3 TSD)
	CNPA	SAFTE
Equation for G	$G_s = 1.52 + \left[\left(\frac{A_s - 0.375}{0.575} \right) \times 0.60 \right]$	$G_s = 1.52 + \left[\left(\frac{A_s - 15.17}{84.70} \right) \times 0.60 \right]$
Equation for T_u	$(T_u)_s = 1.60 + \left[\left(\frac{A_s - 0.375}{0.575} \right) \times 0.34 \right]$	$(T_u)_s = 1.60 + \left[\left(\frac{A_s - 15.17}{84.70} \right) \times 0.34 \right]$

Notes: TSD, total sleep deprivation; CNPA, Circadian Neurobehavioral Performance and Alertness; SAFTE, Sleep, Activity, Fatigue, and Task Effectiveness.

efforts at compensation, respectively. Consequently, we can apply the same approach to deriving values for both G and T_u . To calculate G for session s , we used the equation:

$$G_s = G_w + \left[\left(\frac{A_s - A_m}{A_r} \right) \times G_r \right]$$

where G_w is the best-fit value for G for the worst observed human performance. A_s is the predicted level of alertness for session s , while A_m is the minimum predicted alertness value in the protocol and A_r is the range in predicted alertness values in the protocol. Finally, G_r is the range of G (the best-fit G -value for the session with the worst human performance subtracted from the best-fit value at the observed high point). This accomplishes a simple linear mapping between biomathematical model predictions of alertness and values of G within ACT-R. The same approach was used to predict T_u for session s :

$$(T_u)_s = (T_u)_w + \left[\left(\frac{A_s - A_m}{A_r} \right) \times (T_u)_r \right]$$

where $(T_u)_r$ is the range of T_u (once again, the best-fitting value at the observed performance peak minus the best-fitting value at the observed low point in human performance) and $(T_u)_w$ is the best-fitting value for T_u for the session where participants performed worst. The equations for each parameter using each biomathematical model are included in Table 3.

Using the equations in Table 3 and the predicted alertness values from the biomathematical models, we generated a predicted value for G and for T_u for each session in the experimental protocol. Fig. 5 shows the predicted values of G and T_u for both CNPA and SAFTE across the protocol. Because of stochasticity in the model, it was run using these parameter values through 250 simulated 10-min PVT sessions to obtain stable performance predictions. This produced two sets of model results for each of the 44 experimental sessions: one for the model using parameter values derived from CNPA, and one using parameter values estimated using SAFTE.

5.2.2. Integrated model performance

For each session in the experiment, we considered the four crucial performance phenomena associated with the PVT (described above): false starts, lapses, and sleep attacks measured in proportion of responses, and alert responses measured as median alert reaction time for responses between 150 ms and 500 ms. We compared each of these measures for each of the 44 human data points, giving 176 points of comparison (44 test sessions \times 4 measures) for the model predictions. In addition, we performed these comparisons for the model based on parameter values prescribed by each of the biomathematical models. Table 4 shows qualitative (r) and quantitative (RMSD) fit of the model to the human data. Fig. 6 illustrates the model's performance and human data for each of the dependent measures. The correspondence between the model and human performance is quite good given the variability of the human data, with correlations ranging from 0.70 to 0.88. RMSDs are also relatively low. Before concluding, we briefly consider several alternative conceptualizations of the impact of fatigue on cognitive performance in the context of ACT-R and the PVT.

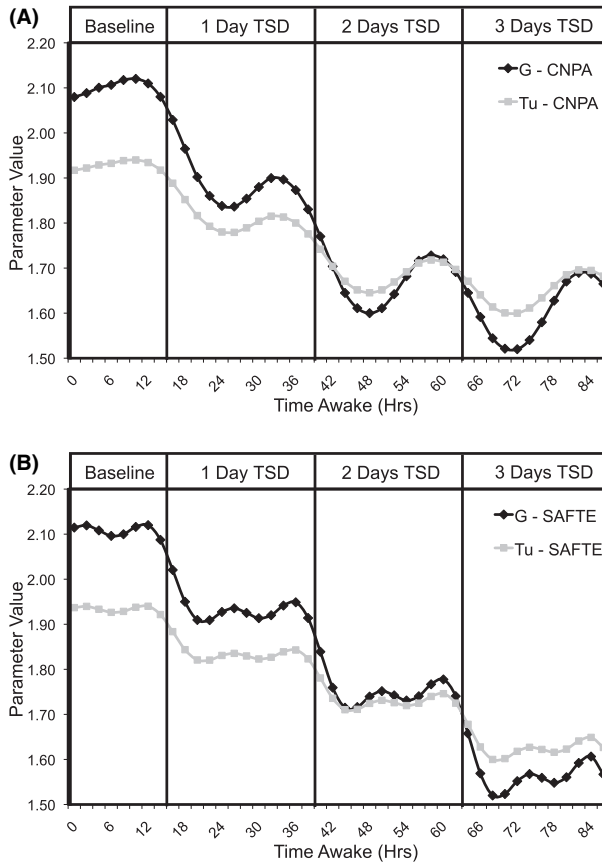


Fig. 5. Predicted values for G and T_u , based upon Circadian Neurobehavioral Performance and Alertness (CNPA) (A) and Sleep, Activity, Fatigue, and Task Effectiveness (SAFTE) (B) across 88 h of total sleep deprivation (TSD).

5.3. Potential alternative mechanisms

Our modeling results help to demonstrate that the theoretical account is consistent with the available evidence concerning how fatigue exerts its influence on human performance in the PVT. However, we have also considered alternative accounts based upon existing theories in the sleep research literature (e.g., Bratzke, Rolke, Ulrich, & Peters, 2007; Dinges & Kribbs, 1991; Horowitz et al., 2003). Whereas the performance of the model illustrated below speaks to the sufficiency of the mechanisms it embodies, the limitations of alternative accounts briefly considered here add further support to the model we have developed.

5.3.1. Increased cognitive noise

Increased levels of fatigue result in increased variability in human performance on the PVT (e.g., Fig. 2; Doran et al., 2001). A direct way of instantiating this in ACT-R would be

Table 4

Comparison of correlation data (r and RMSD) for each of the four principal measures of performance. For false starts, lapses, and sleep attacks, the unit is proportion of responses; for median alert reaction time, the unit is milliseconds

	ACT-R Model Fit to Human Data			
	CNPA-Based Parameters		SAFTE-Based Parameters	
	r	RMSD	r	RMSD
False starts	.72	0.023	.78	0.023
Median alert time reaction	.76	11.69 ms	.70	11.99 ms
Lapses	.88	0.052	.81	0.067
Sleep attacks	.77	0.008	.82	0.007

RMSD, root mean squared deviation; ACT-R, adaptive control of thought-rational; CNPA, Circadian Neuro-behavioral Performance and Alertness; SAFTE, Sleep, Activity, Fatigue, and Task Effectiveness.

to increase noise within ACT-R's production system (ε in the utility equation). Indeed, increasing noise can result in more instances of false starts and longer reaction times to the stimulus. However, in the context of ACT-R's production cycle, this explanation is sharply limited. Specifically, increased levels of noise can reduce the likelihood of performing the appropriate action to chance levels, but not lower. In the production system, this means that increased noise will make production selection and execution more variable for each cognitive cycle, which occur at approximately 50 ms intervals throughout the task. Thus, the primary impact of higher noise values on the model is to increase the likelihood that a false start will occur, since there will be an average of 40–200 cognitive cycles during the 2–10 s delay period. Each cycle provides an opportunity for the model to execute the inappropriate action, and noise increases the likelihood that the inappropriate action will be executed on each occasion. In addition, the likelihood of a lapse increases only modestly because noise will not adequately suppress the appropriate actions. Indeed, the likelihood of a sleep attack remains essentially unchanged since increasing randomness in the production cycle makes it quite likely that a response will be made at some point in the 30 s following the stimulus onset (approximately 600 cognitive cycles). Thus, the impact of noise on the model's performance is uncharacteristic of observed changes in human behavior on the PVT.

5.3.2. Cognitive slowing

The production system in ACT-R has a default time for cognitive actions to be performed (i.e., for productions to be executed), which is 50 ms. An alternative that was mentioned above is that one impact of fatigue is to increase how long cognitive actions take to perform (e.g., Bratzke et al., 2007; Dinges & Kribbs, 1991), which can be instantiated in ACT-R by increasing the cognitive cycle time parameter. This hypothesis is typically used to explain effects like the small shift in median reaction time that is observed in the PVT. Unfortunately, by itself cognitive slowing has no way to simultaneously increase the occurrence of false starts, lapses, or sleep attacks. For cognitive processing to be slowed to the point where




Fig. 6. Model predictions and human performance for the four phenomena of interest in the psychomotor vigilance test across 88 h of sleep deprivation: (A) False starts, (B) Median alert response time, (C) Lapses, and (D) Sleep attacks.

sleep attacks were occurring would require cycle times of many *seconds*. Such issues are likely reasons why cognitive slowing is typically proposed in conjunction with the lapse hypothesis to capture the range of changes that are observed in human performance as fatigue levels increase. A key contribution of our model is to demonstrate that cognitive slowing does not appear to be necessary to capture the subtle changes in reaction time that are observed on this task.

5.3.3. Increased noise and cycle time

Besides considering increased noise and cognitive slowing in isolation, one could hypothesize that the dual-mechanism account could be instantiated by using cognitive cycle time to produce cognitive slowing and noise to produce lapses in cognitive performance. Unfortunately, this account does not overcome the limitations of the two perspectives in isolation. With the probability of executing the appropriate action at chance levels, cognitive cycle time would still have to increase to multiple seconds to produce an appreciable increase in sleep attacks. And once again, false starts would increase dramatically before significant changes in lapses were observed. Thus, the combination of these mechanisms in ACT-R also fails to capture the observed trends in the human data.

5.3.4. *Compensation as the inverse of decreased alertness*

Aside from alternative mechanisms for decreased alertness, alternatives for characterizing compensatory efforts are possible as well. For instance, perhaps participants are successful in their attempts at compensating for lower alertness, which would suggest that the compensatory mechanism described by researchers serves to offset directly the negative consequences of fatigue. Regardless of the set of mechanisms, this would serve only to counteract the negative effects of decreased alertness. This idea, however, is clearly unsupported. None of the potential mechanisms described here for reducing alertness are able to produce increases in reaction times, lapses, and sleep attacks while capturing the other key finding in the PVT, which is the increase in false starts over longer periods of sleep deprivation. This has been described as being a side effect of efforts at compensation (Doran et al., 2001), suggesting that compensation for fatigue does not simply “undo” the negative effects but may have unintended consequences for cognitive performance, such as increasing the occurrence of false starts in the PVT.

5.3.5. *Optimal combination of plausible mechanisms*

One final option is to consider how all of the plausible mechanisms identified here, including the mechanisms in the model as well as cognitive noise and cognitive slowing, could be combined to fit optimally the empirical findings. Although this may result in quantitatively better fits, we demonstrate that the full set of mechanisms is not necessary to capture changes in human performance on the PVT with relatively high accuracy. Thus, we conclude that mechanisms and parameters in addition to what we propose currently represent unnecessary complications to the explanation of how cognitive processing changes under conditions of fatigue. As we expand the breadth of this account, we will strive to strike a balance between parsimony and explanatory power to produce a comprehensive account of performance changes observed in humans.

In line with this aspiration, we note that it is mathematically possible to use a combination of the utility threshold (T_u) and cognitive noise (ε), without manipulating G , to obtain model performance on the PVT that is essentially identical to the one described here (J. R. Anderson, personal communication, June 10, 2008). Despite the formal equivalence, however, we believe that the set of mechanisms proposed above offers at least two advantages. First, there is a straightforward theoretical mapping between the parameters that are manipulated and alertness and compensatory effort, including support from neuropsychological findings. Second, to capture the within-trial decrements that currently involve a constant-magnitude reduction in G when a micro-lapse occurs, both T_u and ε would need to be increased by different, exponentially increasing amounts. This would add both complexity and parameters to the model.

6. General discussion

We have presented a computational account of the effects of sleep loss and circadian rhythms on sustained attention performance. The mechanisms that underlie this account

have been linked to both cognitive psychological and neuropsychological findings of the effects of these factors on neurobehavioral performance. At first glance, this account may appear complicated. In ACT-R, we have modified the default production execution cycle to allow for micro-lapses, implemented a dynamic process for decrementing G over time, and manipulated two global parameters to account for performance changes on the PVT. In addition, we have incorporated biomathematical modeling predictions of alertness to drive parameter changes. However, these mechanisms are rooted in straightforward theoretical assumptions about the relationship between fatigue and cognitive functioning. In addition, we have described several alternative theories, none of which are able to capture the complex dynamics of human performance in this task in the context of our ACT-R model.

Fundamentally, we have made two proposals. One is that the global parameter G in ACT-R can be associated with an overall level of alertness. The second is that compensation for reduced alertness in ACT-R can be represented by decrementing T_u , which makes it easier for a given production to fire on any production cycle. The remaining components of the explanation serve to tie those parameter manipulations to theoretical evidence gathered in research on the effects of fatigue. This includes linking parameter changes to biomathematical models that attempt to capture the systematic influence that time awake and circadian rhythms have on overall cognitive functioning. This helps to connect our theory to the substantial literature that has documented the dynamics of those processes. The result is an account that is consistent with existing research on fatigue and that provides a close fit to the empirical data.

The theory underlying the mechanisms we implemented provides a synthesis of two related hypotheses in the sleep research community—the “state instability” hypothesis (Doran et al., 2001) and the “lapse” hypothesis (Dinges & Kribbs, 1991; Williams et al., 1959). The state instability hypothesis suggests that fatigue results in increased pressure to sleep, eventually causing uncontrolled switches into a sleep state. Relatedly, the lapse hypothesis states that fatigued performance is characterized by lapses in cognitive processing that delay performance and lead to errors in paced tasks (Dinges & Kribbs, 1991). The micro-lapses in our account can be interpreted as representing gaps in cognitive processing resulting from sleep-wake instability in the cognitive system. The key feature of our account relates to the duration of micro-lapses. Whereas previous research has described the processes as operating on a timescale of seconds, the micro-lapses described here last for only tens of milliseconds each. As noted above, one advantage of this proposal is that there is no need to posit an additional “cognitive slowing” mechanism to capture subtle changes in median alert reaction time in the PVT, as suggested in Doran et al. (2001; see also Bratzke et al., 2007). Our model demonstrates that dynamic, rapid (i.e., single cognitive cycles—50 ms) vacillations between wakefulness and sleep provide a parsimonious account for the range of changes in reaction times observed in human performance on the PVT.

By integrating predictions of alertness from biomathematical models, we have been able to constrain the proposed mechanisms in ACT-R to model performance across 88 h of total sleep deprivation. Notably, this does not improve the “fit” of the model to the empirical data. Specifically, it would be straightforward to fit each session by explicitly varying G and T_u , with results that would be a much closer fit to human performance. This is the approach that was taken to generate the fits shown in Fig. 3 to the data aggregated over whole days.

However, fitting in such a manner lacks theoretical constraint, and across 44 experimental sessions it results in 88 free parameters. By incorporating the biomathematical models, we have been able to reduce the number of free parameters to 4—a slope and an intercept mapping predicted alertness (from the biomathematical model) to both G and T_u . This reduces complexity in the model and leverages a well-validated approach to understanding the dynamics of alertness in human functioning.

An additional consideration in this context involves assessing the value that is added by using ACT-R above and beyond the biomathematical models alone. Once again, the value does not come in the form of better fits to the empirical data. As an illustration, Table 5 presents a statistical comparison of the biomathematical models by themselves to the empirical data from Doran et al., (2001) using a linear regression, similar to the evaluation by Van Dongen (2004). One important point is that the degree of correspondence between the biomathematical models and the human data is about the same as, and generally even slightly better than, the fits obtained when the biomathematical models are used in conjunction with our ACT-R cognitive model of the task.

What, then, is added through the use of ACT-R? We see two key advantages to the ACT-R model. First, by virtue of the fact that the ACT-R model *performs* the tasks and generates *behavior*, it produces predictions for all of the dependent measures simultaneously. This contrasts with the fits from the biomathematical models in Table 5, where the biomathematical models were scaled to each dependent measure independently. That is, for each measure a linear regression was performed in each case to maximize the correspondence between the model and the data. Not only does this add parameters to the fitting process, but it also exposes the second important contribution of the cognitive model presented here. Specifically, the central limitation of the biomathematical modeling approach in the context of predicting behavior is that they are unable to account for in situ cognitive performance. Our ACT-R model draws upon the strengths of biomathematical models in characterizing the dynamics of alertness and translates them into behavioral predictions based upon the impact they have on information processing mechanisms in ACT-R. This capability is essential to

Table 5
Comparison of correlation data (r and RMSD) for each of the four principal measures of performance

	Biomathematical Model Fit to Human Data			
	CNPA		SAFTE	
	r	RMSD	r	RMSD
False starts	.74	0.021	.79	0.019
Median alert reaction time	.77	8.77 ms	.70	9.78 ms
Lapses	.90	0.030	.85	0.036
Sleep attacks	.76	0.008	.80	0.007

Note: RMSD, root mean squared deviation; CNPA, Circadian Neurobehavioral Performance and Alertness; SAFTE, Sleep, Activity, Fatigue, and Task Effectiveness.

achieving a more ambitious goal, which involves extending this account so that the mechanisms and parameter values can be applied across *tasks* as well.

The mechanisms that we describe here relate to ACT-R's production system—a component of the architecture that is centrally involved in every ACT-R model. Thus, our account of fatigue in the PVT should generalize to other task contexts, even if additional mechanisms are added as other cognitive capabilities are addressed in our research. In fact, we have shown the effectiveness of the mechanisms described here in the context of a dual-task paradigm investigating the Psychological Refractory Period (Gunzelmann, Byrne, Gluck, & Moore, in press). As in the PVT, we have been able to accurately capture performance changes in participants deprived of sleep without the need for additional mechanisms for cognitive slowing, which has been presented as an explanation for the observed effects (Bratzke et al., 2007).

Our use of ACT-R is to provide an integrated account of the mechanisms of human cognition, which we augment by incorporating biomathematical models representing the dynamic impact of alertness on the information processing mechanisms that drive ACT-R's behavior. Biomathematical models will never overcome the need for scaling the output to particular dependent measures and tasks; they lack a theory of the information processing architecture that is responsible for human cognitive performance. Similarly, ACT-R can represent human cognition and behavior, but until now it has lacked a theory of human alertness to moderate those mechanisms. To make useful predictions in this area requires that both kinds of theories be integrated. We have presented a method for achieving this integration and have explored the promise of the approach for understanding changes in human performance in the PVT under conditions of fatigue.

6.1. Directions for future work

We have focused on the PVT because the task is commonly used in sleep research and is useful for diagnosing architectural changes in procedural functioning. As we continue this research, we will focus on other important aspects of human cognition and performance. For instance, looking at human data for the PVT in more detail reveals that there are sequential relationships in the likelihood of producing a false start on a particular trial, as well as more general time on task effects, which are not captured by the current model (e.g., Doran et al., 2001). These effects may reflect more dynamic changes in alertness that depend on both internal cognitive activity and external environmental stimulation. These are important areas that we plan to explore as we extend our account.

In addition to pursuing a more detailed understanding of the dynamics of, and influences on, alertness, we are also pursuing efforts to extend our account to other task domains by exploring the impact of changes in alertness on other mechanisms within ACT-R. For example, we have been working to extend our account to other areas of human functioning, including a task with significant declarative memory requirements (e.g., Gunzelmann, Gluck, Kershner, Van Dongen, & Dinges, 2007), and a visual search task (Moore, Gunzelmann, & Gluck, 2008). Initial work with these tasks has shown that the mechanisms developed to account for performance changes on the PVT are not sufficient to capture

performance changes in tasks with more complicated cognitive requirements. We do not find this result surprising. Research examining the neuropsychological effects of fatigue has shown that there are changes in activity across the cortex and in a number of subcortical areas resulting from sleep deprivation. Thus, it is to be expected that fatigue would have effects across different components of human functioning, which translates into decrements associated with multiple modules within the ACT-R architecture. This was a motivation for focusing first on a relatively simple task that specifically targets a component of ACT-R. This approach facilitates the process of isolating the impact of fatigue on different cognitive mechanisms. The challenge will be to integrate theoretical commitments in ACT-R with existing behavioral and neuropsychological research to identify precisely how the various components of ACT-R are impacted by fatigue. Though challenging, this approach has proven fruitful thus far.

In addition to providing an overall account for performance decrements arising from fatigue and circadian rhythmicity, the mechanisms we have implemented may provide a basis for understanding individual differences in human susceptibility to the negative effects of these processes. Research on fatigue has uncovered significant individual differences, in terms of both how much sleep is required and in the severity of the decrements associated with lack of sleep. Under the assumption that sleep deprivation impacts all individuals through the same mechanisms, individual differences can be represented through different-magnitude effects on those mechanisms. This implies that all individuals experience decrements in alertness as a consequence of sleep deprivation. However, it is entirely possible that the rate of decline in alertness across a particular period without sleep will differ among individuals. The mechanisms we have described provide a means of quantifying those differences, and one avenue of current research is to test our ability to account for them using consistent parameter manipulations for individual participants across multiple tasks. Of course, any individual differences in the impact of fatigue are layered on top of individual differences in baseline performance and ability. Thus, success with this effort will require addressing both sources of inter-individual variability. We have begun to address some of these issues, with initially promising results (Gunzelmann et al., 2008).

We are optimistic about the prospects for developing a comprehensive account of fatigue within ACT-R. By incorporating the capacity for representing individual differences as well, it will be possible to make a priori predictions about how human performance will change for a specific person, on a particular task, given a specific history of sleep. While we are still a long way from making these kinds of predictions for dynamic naturalistic tasks, the promise of this approach is clear. The potential benefits of such a capability for predicting readiness and averting dangerous situations are substantial.

Acknowledgments

The views expressed in this article are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government. Cognitive model development was sponsored by the Air Force Research Laboratory's (AFRL) Warfighter

Readiness Research Division and by grants 04HE02COR and 07HE01COR from the Air Force Office of Scientific Research (AFOSR). Human data collection was sponsored partly by AFOSR grants F49620-95-1-0388 and F49620-00-1-0266, and by NIH grant RR00040. Joshua Gross's participation was sponsored by an AFRL internship administered by the Oak Ridge Institute for Science and Education. Portions of this research were presented at the 27th and 28th Annual Meetings of the Cognitive Science Society. We would like to thank Steve Hursh and Elizabeth Klerman for their assistance in understanding and using the biomathematical modeling tools mentioned in this paper. We would also like to thank Hans Van Dongen for his comments on this research, and Don Lyon and Frank Ritter for their comments on earlier drafts of this paper.

Notes

1. As of writing, information about CNPA can be found here: <http://sleep.med.harvard.edu/research/faculty-research/tools/CPSS>. Information regarding SAFTE is available at http://www.novasci.com/index_files/page0003.htm.
2. Reaction times faster than 150 ms are unlikely to represent intentional responses to stimulus, due to limitations of the human nervous system.
3. We set ACT-R's utility noise parameter (s) to 0.25, which translates into a standard deviation through the following function:

$$\sigma = \sqrt{\frac{\pi^2 \times s^2}{3}}$$

4. We have enabled an option in ACT-R to make this parameter noisy. As a result, the duration of a cognitive cycle varies randomly between 25 ms and 75 ms, according to a uniform linear distribution.
5. Selecting the appropriate method for deriving parameter values across sessions is not an obvious decision. It would be reasonable to find best-fitting values for each session and use linear regression to optimally map ACT-R parameter values to alertness predictions from CNPA and SAFTE. While this may have resulted in a modest improvement of the model's fit to the data across sessions, our goal here is to illustrate some of the potential of this approach for making predictions about performance. One may argue that this parameter estimation method still requires data about the high- and low-points of human performance across the protocol; however, it still allows for the identification of a systematic relation between biomathematical predictions of alertness and particular parameters in the ACT-R architecture. This, we believe, is a critical step in advancing this approach and we are willing to accept some reduction in fit statistics to make this claim a little more strongly.

References

- Achermann, P. (2004). The two-process model of sleep regulation revisited. *Aviation, Space, and Environmental Medicine*, 74(3), A37–A43.
- Anderson, J. (2007). *How can the mind exist in a physical universe?* New York: Oxford University Press.
- Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review*, 111, 1036–1060.
- Anderson, J. R., Fincham, J. M., Qin, Y., & Stocco, A. (2008). A central circuit of the mind. *Trends in Cognitive Science*, 12(4), 136–143.
- Anderson, J. R., & Lebiere, C. L. (1998). *The atomic components of thought*. Hillsdale, NJ: Erlbaum.
- Belavkin, R. V. (2001). Modelling the inverted-U effect in ACT-R. In E. M. Altmann, A. Cleeremans, C. D. Schunn, & W. D. Gray (Eds.), *Proceedings of the 2001 Fourth International Conference on Cognitive Modeling* (pp. 275–276). Mahwah, NJ: Erlbaum.
- Borbely, A. A., & Achermann, P. (1999). Sleep homeostasis and models of sleep regulation. *Journal of Biological Rhythms*, 14, 557–568.
- Bratzke, D., Rolke, B., Ulrich, R., & Peters, M. (2007). Central slowing during the night. *Psychological Science*, 18(5), 456–461.
- Caldwell, J. A. (2003, Fall). Wake up to the importance of sleep for air safety! *Flightline*, 30–33.
- Caldwell, J. A., Caldwell, J. L., Brown, D. L., & Smith, J. K. (2004). The effects of 37 hours of continuous wakefulness on the physiological arousal, cognitive performance, self-reported mood, and simulator flight performance of F-117A pilots. *Military Psychology*, 16(3), 163–181.
- Chee, M. W., & Choo, W. C. (2004). Functional imaging of working memory after 24 hr of total sleep deprivation. *The Journal of Neuroscience*, 24, 4560–4567.
- Chua, Y. M. L., Venkatraman, V., Dinges, D. F., & Chee, M. W. L. (2006). The neural basis of interindividual variability in inhibitory efficiency after sleep deprivation. *The Journal of Neuroscience*, 26(27), 7156–7162.
- Dinges, D. F. (1995). An overview of sleepiness and accidents. *Journal of Sleep Research*, 4(2), 4–11.
- Dinges, D. F., Baynard, M., & Rogers, N. L. (2005). Chronic sleep deprivation. In M. H. Kryger, T. Roth, & W. C. Dement (Eds.), *Principles and practice of sleep medicine* (4th ed.) (pp. 67–76). Philadelphia: W. B. Saunders.
- Dinges, D. F., & Kribbs, N. B. (1991). Performing while sleepy: Effects of experimentally-induced sleepiness. In T. H. Monk (Ed.), *Sleep, sleepiness and performance* (pp. 97–128). Chichester, England: John Wiley & Sons.
- Dinges, D. F., & Powell, J. W. (1985). Microcomputer analyses of performance on a portable, simple visual RT task during sustained operations. *Behavior Research Methods, Instruments, & Computers*, 17(6), 652–655.
- Doran, S. M., Van Dongen, H. P., & Dinges, D. F. (2001). Sustained attention performance during sleep deprivation: Evidence of state instability. *Archives of Italian Biology: Neuroscience*, 139(3), 253–267.
- Dorrian, J., Rogers, N. L., & Dinges, D. F. (2005). Psychomotor vigilance performance: Neurocognitive assay sensitive to sleep loss. In C. A. Kushida (Ed.), *Sleep deprivation: Clinical issues, pharmacology and sleep loss effects* (pp. 39–70). New York: Marcel Dekker, Inc.
- Drummond, S. P. A., & Brown, G. G. (2001). The effects of total sleep deprivation on cerebral responses to cognitive performance. *Neuropsychopharmacology*, 25(S5), S68–S73.
- Drummond, S. P. A., Brown, G. G., Gillin, J. C., Stricker, J. L., Wong, E. C., & Buxton, R. B. (2000). Altered brain response to verbal learning following sleep deprivation. *Nature*, 403, 655–657.
- Drummond, S. P., Brown, G. G., Salamat, J. S., & Gillin, J. C. (2004). Increasing task difficulty facilitates the cerebral compensatory response to total sleep deprivation. *Sleep*, 27, 445–451.
- Drummond, S. P. A., Gillin, J. C., & Brown, G. G. (2001). Increased cerebral response during a divided attention task following sleep deprivation. *Journal of Sleep Research*, 10, 85–92.
- Durmer, J. S., & Dinges, D. F. (2005). Neurocognitive consequences of sleep deprivation. *Seminars in Neurology*, 25(1), 117–129.

- French, J., Morris, C. S., & Hancock, P. A. (2003). Modeling fatigue degraded performance in artificial agents. In *Proceedings of the Human Factors and Ergonomics Society* (pp. 307–310). Santa Monica, CA: Human Factors and Ergonomics Society.
- Gluck, K. A., & Pew, R. W. (Eds.) (2005). *Modeling human behavior with integrated cognitive architectures: Comparison, evaluation, and validation*. Mahwah, NJ: Erlbaum.
- Gray, W. D. (2007). *Integrated models of cognitive systems*. New York: Oxford University Press.
- Gunzelmann, G., Byrne, M. D., Gluck, K. A., & Moore, L. R. (in press). Using computational cognitive modeling to evaluate alternative theories: Implications for understanding fatigue. *Human Factors*, 51(2).
- Gunzelmann, G., & Gluck, K. A. (2008). Approaches to modeling the effects of fatigue on cognitive performance. In J. Hansberger (Ed.), *Proceedings of the seventeenth conference on behavior representation in modeling and simulation* (pp. 136–145). Orlando, FL: Simulation Interoperability Standards Organization.
- 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. In D. S. McNamara & J. G. Trafton (Eds.), *Proceedings of the twenty-ninth annual meeting of the Cognitive Science Society* (pp. 329–334). Mahwah, NJ: Erlbaum.
- Gunzelmann, G., Gluck, K. A., Van Dongen, H. P. A., O'Connor, R. M., & Dinges, D. F. (2005). A neurobehaviorally inspired ACT-R model of sleep deprivation: Decreased performance in psychomotor vigilance. In B. G. Bara, L. Barsalou, & M. Bucciarelli (Eds.), *Proceedings of the 27th annual meeting of the Cognitive Science Society* (pp. 857–862). Mahwah, NJ: Erlbaum.
- Gunzelmann, G., Moore, L. R., Gluck, K. A., Van Dongen, H. P. A., & Dinges, D. F. (2008). Individual differences in sustained vigilant attention: Insights from computational cognitive modeling. In B. C. Love, K. McRae, & V. M. Sloutsky (Eds.), *Proceedings of the thirtieth annual meeting of the Cognitive Science Society* (pp. 2017–2022). Austin, TX: Cognitive Science Society.
- Habeck, C., Rakitin, B. C., Moeller, J., Scarmeas, N., Zarahn, E., Brown, T., & Stern, Y. (2004). An event-related fMRI study of the neurobehavioral impact of sleep deprivation on performance of a delayed-match-to-sample task. *Cognitive Brain Research*, 18, 306–321.
- el-Hajj Fuleihan, G., Klerman, E. B., Brown, E. N., Choe, Y., Brown, E. M., & Czeisler, C. A. (1997). The parathyroid hormone circadian rhythm is truly endogenous — a general clinical research center study. *Journal of Clinical Endocrinology and Metabolism*, 82(1), 281–286.
- Heuer, H., Kleinsorge, T., Klein, W., & Kohlisch, O. (2004). Total sleep deprivation increases the costs of shifting between simple cognitive tasks. *Acta Psychologica*, 117, 29–64.
- Horne, J. A., Anderson, N. R., & Wilkinson, R. T. (1983). Effects of sleep deprivation on signal detection measures of vigilance: Implications for sleep function. *Sleep*, 6, 347–358.
- Horne, J. A., & Pettitt, A. N. (1985). High incentive effects on vigilance performance during 72 hours of total sleep deprivation. *Acta Psychologica*, 58, 123–139.
- Horne, J., & Reyner, L. (1999). Vehicle accidents related to sleep: A review. *Occupational and Environmental Medicine*, 56, 289–294.
- Horowitz, T. S., Cade, B. E., Wolfe, J. M., & Czeisler, C. A. (2003). Searching night and day: A dissociation of effects of circadian phase and time awake on visual selective attention and vigilance. *Psychological Science*, 14(6), 549–557.
- Hursh, S. R., Redmond, D. P., Johnson, M. L., Thorne, D. R., Belenky, G., Balkin, T. J., Storm, W. F., Miller, J. C., & Eddy, D. R. (2004). Fatigue models for applied research in warfighting. *Aviation, Space, and Environmental Medicine*, 75(3), A44–A60.
- Jewett, M. E., & Kronauer, R. E. (1999). Interactive mathematical models of subjective alertness and alertness in humans. *Journal of Biological Rhythms*, 14, 588–597.
- Jones, R. M., Laird, J. E., & Neville, K. (1998). Modeling pilot fatigue with a synthetic behavior model. In *Proceedings of the seventh conference on computer generated forces and behavioral representation*. (pp. 349–356). Orlando, FL: Simulation Interoperability Standards Organization.
- Jongman, L. (1998). How to fatigue ACT-R? In *Proceedings of the second European conference on cognitive modelling* (pp. 52–57). Nottingham, England: Nottingham University Press.

- Khalsa, S. B. S., Jewett, M. E., Duffy, J. F., & Czeisler, C. A. (2000). The timing of the human circadian clock is accurately represented by the core body temperature rhythm following phase shifts to a three-cycle light stimulus near the critical zone. *Journal of Biological Rhythms*, *15*(6), 524–530.
- Kjellberg, A. (1977). Sleep deprivation and some aspects of performance: II, Lapses and other attentional effects. *Waking and Sleeping*, *1*, 145–148.
- Klauer, S. G., Dingus, T. A., Neale, V. L., Sudweeks, J. D., & Ramsey, D. J. (2006). *The impact of driver inattention on near-crash/crash risk: An analysis using the 100-car naturalistic driving study data*. Technical Report #DOT HS 810 594. Washington, DC: National Highway Traffic Safety Administration.
- Klerman, E. B., & St. Hilaire, M. (2007). On mathematical modeling of circadian rhythms, performance, and alertness. *Journal of Biological Rhythms*, *22*(2), 91–102.
- Kronauer, R. E., Forger, D. B., & Jewett, M. E. (1999). Quantifying human circadian pacemaker response to brief, extended, and repeated light stimuli over the photopic range. *Journal of Biological Rhythms*, *14*, 500–515.
- Landrigan, C. P., Rothschild, J. M., Cronin, J. W., Kaushal, R., Burdick, E., Katz, J. T., Lilly, C. M., Stone, P. H., Lockley, S. W., Bates, D. W., & Czeisler, C. A. (2004). Effect of reducing interns' work hours on serious medical errors in intensive care units. *New England Journal of Medicine*, *351*(18), 1838–1848.
- Lee, M. A. M., & Kleitman, N. (1923). Attempts to demonstrate functional changes in the nervous system during experimental insomnia. *American Journal of Physiology*, *67*, 141–152.
- Lisper, H., & Kjellberg, A. (1972). Effects of 24-hour sleep deprivation on rate of decrement in a 10-minute auditory reaction time task. *Journal of Experimental Psychology*, *96*, 287–290.
- Mallis, M., Mejdal, S., Nguyen, T., & Dinges, D. (2004). Summary of the key features of seven biomathematical models of human fatigue and performance. *Aviation, Space, and Environmental Medicine*, *75*(3), A4–A14.
- Manganotti, P., Palermo, A., Patuzzo, S., Zanette, G., & Fiaschi, A. (2001). Decrease in motor cortical excitability in human subjects after sleep deprivation. *Neuroscience Letters*, *304*(3), 153–156.
- Mitler, M. M., Carskadon, M. A., Czeisler, C. A., Dement, W. C., Dinges, D. F., & Graeber, R. C. (1988). Catastrophes, sleep, and public policy: Consensus report. *Sleep*, *11*(1), 100–109.
- Moore, L. R., Gunzelmann, G., & Gluck, K. A. (2008). Evaluating mechanisms of fatigue using a digit symbol substitution task [Abstract]. In B. C. Love, K. McRae & V. M. Sloutsky (Eds.), *Proceedings of the thirtieth annual meeting of the Cognitive Science Society* (p. 1522). Austin, TX: Cognitive Science Society.
- National Transportation Safety Board (1995). *Factors that affect fatigue in heavy truck accidents (volume 1)*. NTSB Safety Study: Report no. NTSB//SS-95/01 (PB95-9170012). Washington, DC: NTSB.
- Neri, D. F. (ed.) (2004). Fatigue and performance modeling workshop, June 13–14, 2002 [Special Issue]. *Aviation, Space, and Environmental Medicine*, *75*(Suppl. 3), A1–A199.
- Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.
- Pack, A. I., Pack, A. M., Rodgman, E., Cucchiara, A., Dinges, D. F., & Schwab, C. W. (1995). Characteristics of crashes attributed to the driver having fallen asleep. *Accident Analysis & Prevention*, *27*(6), 769–775.
- Portas, C. M., Rees, G., Howseman, A. M., Josephs, O., Turner, R., & Frith, C. D. (1998). A specific role for the thalamus in mediating the interaction of attention and arousal in humans. *The Journal of Neuroscience*, *18*, 8979–8989.
- Ritter, F. E., Shadbolt, N. R., Elliman, D., Young, R., Gobet, F., & Baxter, G. D. (2003). *Techniques for modeling human and organizational behaviour in synthetic environments: A supplementary review*. Wright-Patterson Air Force Base, OH: Human Systems Information Analysis Center.
- Royal, D. (2002). *National Survey of Distracted and Drowsy Driving Attitudes and Behaviors: 2002*. Volume 1-Findings Report. Washington, DC: National Highway Traffic Safety Administration. Available at http://www.nhtsa.dot.gov/people/injury/drowsy_driving1/survey-distractive03/index.htm
- Thomas, M., Sing, H., Belenky, G., Holcomb, H., Mayberg, H., Dannals, R., Wagner, H. Jr, Thorne, D., Popp, K., Rowland, L., Welsh, A., Balwinski, S., & Redmond, D. (2000). Neural basis of alertness and cognitive performance impairments during sleepiness. I. Effects of 24 h of sleep deprivation on waking human regional brain activity. *Journal of Sleep Research*, *9*, 335–352.

- Van Dongen, H. P. A. (2004). Comparison of mathematical model predictions to experimental data of fatigue and performance. *Aviation, Space, and Environmental Medicine*, 75(3), A15–A36.
- Van Dongen, H. P. A., Baynard, M. D., Maislin, G., & Dinges, D. F. (2004). Systematic interindividual differences in neurobehavioral impairment from sleep loss: Evidence of trait-like differential vulnerability. *Sleep*, 27, 423–433.
- Van Dongen, H. P. A., & Dinges, D. F. (2005). Sleep, circadian rhythms, and psychomotor vigilance. *Clinical Sports Medicine*, 24, 237–249.
- Van Dongen, H. P. A., Maislin, G., Mullington, J. M., & Dinges, D. F. (2003). The cumulative cost of additional wakefulness: Dose-response effects on neurobehavioral functions and sleep physiology from chronic sleep restriction and total sleep deprivation. *Sleep*, 26, 117–126.
- Van Dongen, H. P. A., Price, N. J., Mullington, J. M., Szuba, M. P., Kapoor, S. C., & Dinges, D. F. (2001). Caffeine eliminates sleep inertia: Evidence for the role of adenosine. *Sleep*, 24, 813–819.
- Wilkinson, R. T. (1968). Sleep deprivation: Performance tests for partial and selective sleep deprivation. In L. E. Abt & B. F. Riess (Eds.), *Progress in clinical psychology: Dreams and dreaming* (pp. 28–43). New York: Grune and Stratton.
- Williams, H. L., Lubin, A., & Goodnow, J. J. (1959). Impaired performance with acute sleep loss. *Psychological Monographs*, 73(14), 1–26.
- Wu, J. C., Gillin, J. C., Buchsbaum, M. S., Hershey, T., Hazlett, E., Sicotte, N., & Bunney, W. E. (1991). The effect of sleep deprivation on cerebral glucose metabolic rate in normal humans assessed with positron emission tomography. *Sleep*, 14(2), 155–162.