

Decreased Arousal as a Result of Sleep Deprivation

The Unraveling of Cognitive Control

*Glenn Gunzelmann, Kevin A. Gluck, Scott Price,
Hans P. A. Van Dongen, & David F. Dinges*

This chapter discusses recent efforts at developing mechanisms for capturing the effects of fatigue on human performance. We describe a computational cognitive model, developed in ACT-R, that performs a sustained attentional task called the psychomotor vigilance task (PVT). We use neurobehavioral evidence from research on sleep deprivation, in addition to previous research from within the ACT-R community, to select and to evaluate a mechanism for producing fatigue effects in the model. Fatigue is represented by decrementing a parameter associated with arousal in ACT-R, while also reducing a threshold value in the architecture to capture attempts at compensating for the negative effects of decreased arousal. These parameters are associated with the production utility computation in ACT-R, which controls the selection/execution cycle to determine which production (if any) to execute on each cognitive cycle. In ACT-R, this mechanism is linked to the basal ganglia and the thalamus. In turn, portions of the thalamus show heightened activation in attentional tasks under conditions of sleep deprivation. The model we describe closely captures the performance of human participants on the PVT, as observed in a laboratory experiment involving 88 hours of total sleep deprivation.

Until recently, computational cognitive models of human performance were developed with little consideration of how factors such as emotions and alertness influence cognition. However, with increased sophistication in models of cognitive systems, advances in computer technology, and pressure for ever more realistic representations of human performance, cognitive moderators are emerging as an important area of research within the field of computational modeling (e.g., Gratch & Marsella, 2004; Hudlicka, 2003; Ritter, Reifers, Klein, Quigley, & Schoelles, 2004). There is a sense in which this development is both premature and long overdue. Evidence for its prematurity can be found in many of the other chapters in this volume. Cognitive science has yet to unravel many of the intricacies of “normal” human cognition. Therefore, adding additional complexity by including cognitive moderators that influence those thought processes constitutes a substantial challenge. However, cognitive moderators are pervasive in human cognition. It seems essential, therefore, that they be considered in attempts to

understand human cognitive functions. If cognitive architectures are to be viewed as “unified theories of cognition” (Newell, 1990), then they must include mechanisms to represent those factors that have substantial modulatory effects on cognitive performance.

This chapter describes an effort to introduce a theory of degraded cognitive functioning into the adaptive control of thought—rational, or ACT-R, cognitive architecture. In this case, the degradation arises from the combined effect of sleep deprivation and endogenous circadian variation. We describe a computational cognitive model that incorporates mechanisms to represent decreased alertness and describe the impact of those mechanisms on the model’s performance on the psychomotor vigilance task (PVT), a sustained attention task that has been extensively validated to be sensitive to variation in sleep homeostatic and circadian dynamics, while being relatively immune to the effects of aptitude and learning (Dorrian, Rogers, & Dinges, 2005). Our modeling effort draws on recent research on partial and total sleep deprivation (e.g., Van Dongen

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE

2007

2. REPORT TYPE

Book Chapter

3. DATES COVERED

00-10-2004 to 00-10-2006

4. TITLE AND SUBTITLE

Decreased Arousal as a Result of Sleep Deprivation The Unraveling of Cognitive Control

5a. CONTRACT NUMBER

F41624-97-D-5000 TO 23

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

62202F

6. AUTHOR(S)

Glenn Gunzelmann; Kevin Gluck; Scott Price; Hans Van Dongen; David Dinges

5d. PROJECT NUMBER

1123

5e. TASK NUMBER

AS

5f. WORK UNIT NUMBER

1123AS06;2313HA09

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Air Force Research Laboratory/RHA, Warfighter Readiness Research Division, 6030 South Kent Street, Mesa, AZ, 85212-6061

8. PERFORMING ORGANIZATION REPORT NUMBER

AFRL; AFRL/RHA

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Air Force Research Laboratory/RHA, Warfighter Readiness Research Division, 6030 South Kent Street, Mesa, AZ, 85212-6061

10. SPONSOR/MONITOR'S ACRONYM(S)

AFRL; AFRL/RHA

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

AFRL-RH-AZ-BC-2007-0001

12. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES

Published as Chapter 17 in W. Gray (Ed.), Integrated models of cognitive systems (pp. 243-253). New York NY: Oxford University Press

14. ABSTRACT

This chapter discusses recent efforts at developing mechanisms for capturing the effects of fatigue on human performance. We describe a computational cognitive model developed in ACT-R, that performs a sustained attentional task called the psychomotor vigilance (PVT). We use neurobehavioral evidence from research on sleep deprivation, in addition to previous research from within the ACT-R community, to select and to evaluate a mechanism for producing fatigue effects in the model. Fatigue is represented by decrementing a parameter associated with arousal in ACT-R, while also reducing a threshold value in the architecture to capture attempts at compensating for the negative effects of decreased arousal. These parameters are associated with the production utility computation in ACT-R, which controls the selection/execution cycle to determine which production (if any) to execute on each cognitive cycle. In ACT-R, this mechanism is linked to the basal ganglia and the thalamus. In turn, portions of the thalamus show heightened activation in attentional tasks under conditions of sleep deprivation. The model we describe closely captures the performance of human participants on the PVT, as observed in a laboratory experiment involving 88 hours of total sleep deprivation.

15. SUBJECT TERMS

Arousal; Sleep deprivation; Cognitive control; Fatigue; Human performance; Computational cognitive process modeling; Psychomotor vigilance task;

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified	Public Release	10	

et al., 2003), and leverages recent advances in understanding how sleep deprivation impacts neurobehavioral and brain functioning (e.g., Drummond et al., 1999, 2000; Drummond, Gillin, & Brown 2001; Habeck et al., 2004; Portas et al., 1998).

In the sections that follow, we describe relevant research related to sleep loss. This is followed by a description of the PVT and then the ACT-R model we have developed to perform it. We use the model to demonstrate the effectiveness of our approach for capturing performance decrements as a function of sleep deprivation. In describing the model, we suggest some alternative mechanisms to illustrate how the effects of sleep deprivation can be seen as resulting from impacts to either central control (Type 1 control) or the internal control of functional processes (Type 2 control), which includes processes like memory retrieval or programming motor movements. This distinction constitutes a major theme of this book. Although the mechanistic explanation for the effects of sleep deprivation we have developed is not explicitly defined in terms of Type 1 or Type 2 control, the discussion illustrates how the modeling effort is improved through consideration of this distinction.

Neuropsychological Research on Sleep Deprivation

Unquestionably, sleep deprivation has a negative effect on human performance across a wide array of tasks and situations. Determining the particular impacts of sleep deprivation, both behaviorally and physiologically, has been a significant topic of study in psychological and medical research for quite some time (e.g., Patrick & Gilbert, 1896; von Economo, 1930). Research originally focused on identifying the nature of neurobehavioral incapacitation but shifted to changes in cognitive performance when early studies did not provide conclusive evidence that sleep loss eliminated the ability to perform specific tasks (e.g., Kleitman, 1923; Lee & Kleitman, 1923). Current research directions have been motivated by the desire to uncover the neurophysiologic mechanisms that produce diminished alertness and decrements in cognitive performance, as well as any compensatory mechanisms. Research evaluating behavioral, pharmacological, and technological countermeasures to offset deficits of sleep deprivation has also been a long-standing focus of research (e.g., Bonnet et al., 2005; Caldwell, Caldwell, & Darlington 2003; Caldwell, Caldwell, Smith, & Brown, 2004; Dinges & Broughton, 1989).

At the cortical level, studies have shown inconsistent patterns of regional activation responses to sleep deprivation, depending on the type of cognitive task, its difficulty, and the method used to measure activation (e.g., Chee & Choo, 2004; Drummond et al., 1999, 2001; Habeck et al., 2004). At the subcortical level, a main area that consistently shows sensitivity to sleep deprivation is the thalamus (Chee & Choo, 2004; Habeck et al., 2004; Lin, 2000; Portas et al., 1998). The thalamus typically shows an increase in activation when individuals are asked to perform a task while sleep deprived, relative to performing the task when well rested. For instance, Portas et al. (1998) asked participants to perform a short-duration attention task while activity was measured using fMRI. They found that the thalamus showed increased activation while performing the attention task under conditions of sleep loss, while overall performance (response time) was not significantly different from baseline. From these results, they 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" (p. 8987). The possibility of such a compensatory mechanism involving the thalamus is discussed further in the section on the computational model later in this chapter.

Biomathematical Models of Sleep Deprivation

In addition to the significant progress that has been made in understanding the neurobehavioral mechanisms of sleep deprivation, researchers studying fatigue have also developed biomathematical models that reflect the influence of sleep history and circadian rhythms on overall cognitive performance, or alertness (Mallis, Mejdal, Nguyen, & Dinges, 2004). Such models provide a means for describing the dynamic interaction of these factors. For instance, Figure 17.1 shows the predictions for one of these models, the circadian neurobehavioral performance and alertness (CNPA) model (Jewett & Kronauer, 1999), for a protocol involving 88 hr of total sleep deprivation. The circadian rhythm component of the model is responsible for the cyclic nature of the predictions and increased sleep loss is responsible for the overall decline across days.

Although there is room for improvement in all current biomathematical models of performance (Van Dongen, 2004), the models have potential value for predicting global changes in alertness over time in a




FIGURE 17.1 Predictions of alertness from the circadian neurobehavioral performance and alertness (CNPA) model for a study involving 88 continuous hours awake, beginning at 7:30 a.m. on the baseline day.

variety of circumstances. However, a key limitation is that these models do not make predictions of how changes in alertness will affect performance on particular tasks (e.g., changes in response times or changes in types or frequencies of errors). The fits described in Van Dongen (2004) were produced by scaling the alertness predictions to minimize the deviation from the data. These values had to be computed post hoc. So, while the predictions from the models approximate relative changes in performance, they do not actually provide a priori estimates of how much response times will change in absolute magnitude or how errors will increase over time.

The computational cognitive modeling research described in this chapter will eventually allow us to bridge the gap between biomathematical models and complex cognitive task performance. Computational cognitive models make detailed predictions about human performance, including response times and errors. The goal of the project is to use the predictions from the biomathematical models to drive changes in mechanisms in the ACT-R cognitive architecture. In this way, the predictions of the biomathematical model can be used to produce parameter changes in the cognitive model, which can be used to make specific predictions about how human performance declines as a function of fatigue. Although this latter goal has not yet been reached, this chapter describes the progress we have made toward it, especially the determination of a set of mechanisms in ACT-R to account for changes in alertness. These mechanisms are demonstrated in the context of the PVT, which is described next.

Psychomotor Vigilance Task

The psychomotor vigilance task (PVT; Dinges & Powell, 1985) assesses vigilant/sustained attention and has been used frequently in sleep deprivation research. Its main advantages are that performance is both sensitive to the levels of sleep deprivation and relatively insensitive to either aptitude or learning (Dorrian et al., 2005). During a typical PVT trial, a stimulus appears in a prespecified location on a monitor at random intervals between 2 s and 10 s. The subject's task is to press a response button as fast as possible each time a stimulus appears but not to press the button too soon. When the response button is pressed, the visual stimulus displays reaction time in milliseconds to inform the subject of how well they performed. The duration of a test session is typically 10 min.

The data from a PVT session consist of approximately 90 responses, which can be classified to facilitate understanding how PVT performance changes as fatigue increases (Dorrian et al., 2005). The range for the first category, which we will refer to as "alert" responses, is between 150 ms and 500 ms after stimulus onset (median is typically around 250 ms), indicative of a participant that is responding about as rapidly as neurologically possible to each stimulus. Responses greater than 500 ms but less than 30,000 ms (i.e., 30 s) are considered to be "lapses" of attention (errors of omission; Dinges & Kribbs, 1991; Dorrian et al., 2005). These responses indicate that attention is wavering from the display, but that participants are recovering at some point to detect the stimulus. In some instances,

participants fail to respond even after 30s, which is a dramatic breakdown in performance that is classified as a “sleep attack” (Dorrian et al., 2005). In these cases, the experimenter intervenes to wake the participant. At the opposite end of the response-time continuum are “false starts” (errors of commission), which are responses that occur before the stimulus appears, or within 150ms of the stimulus onset (i.e., neurologically too fast to be a normal, alert response). These responses represent anticipation of the stimulus’s appearance.

As sleep deprivation increases, the proportion of alert responses decreases, and the distribution of reaction times shifts to the right, resulting in increased proportions of lapses and sleep attacks. As participants attempt to compensate based on feedback that they are lapsing (errors of omission) more frequently, the proportion of false starts (errors of commission) increases as well (Doran, Van Dongen, & Dinges, 2001). A sample set of data from the PVT is shown in Figure 17.2 (these data are from Van Dongen, 2004; Van Dongen et al., 2001). In the experiment that provided the data, participants first spent three nights in the laboratory to acclimate to a common sleep cycle of 8 hr for sleep per day. After this, participants were kept awake continuously for 88 hr, until near midnight on the fourth day. This is the same protocol that was used to generate the CNPA predictions in Figure 17.1, which shows alertness predictions for the last day of acclimation and for

the 88-hr sleep-deprivation period. The first day of this period, during which no actual sleep loss was yet incurred, was used as a baseline day. Beginning at 7:30 a.m. on the baseline day, participants completed a series of tasks, including the PVT, repeatedly in 2-hr cycles (the set of tasks took approximately 30 min to complete). Note that the PVT data shown in Figure 17.2 are averaged over sessions performed within each day of the protocol, whereas the CNPA data in Figure 17.1 illustrate the dynamic changes in alertness that occur within each of the days (circadian rhythms).

The next section describes the computational cognitive model. The model represents the first step in developing the capability to make detailed a priori predictions about changes in human performance on particular tasks as a function of increased levels of fatigue. The model performs the PVT, and parameter changes in the model impact performance in a manner similar to human performance under conditions of sleep deprivation.

Computational Cognitive Model

The computational model described in this chapter was developed in the ACT-R 5 cognitive architecture (Anderson et al., 2004). Here we will describe only the ACT-R mechanisms that are associated with the

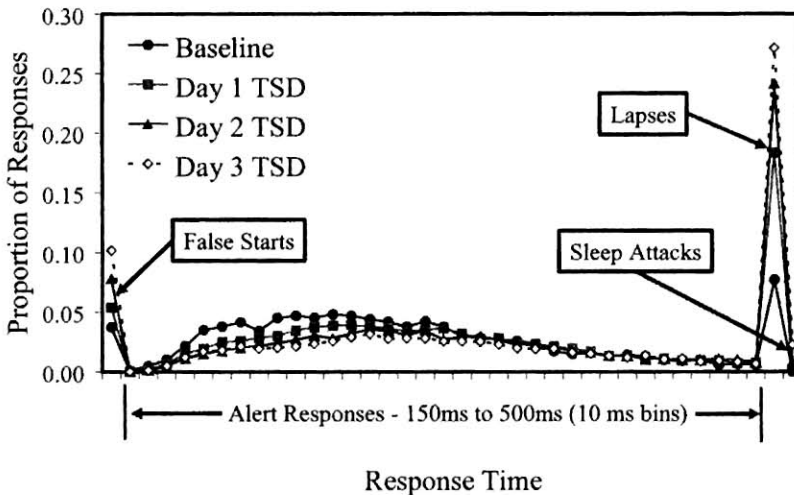


FIGURE 17.2 Human performance on the psychomotor vigilance task. Data are from a study where participants were kept awake for 88 continuous hours, while performing a battery of tests every 2 hr (data from Van Dongen et al., 2001; Van Dongen 2004). Averages across test sessions within each day are shown.

parameters that were manipulated to alter the architecture's level of alertness, which produces the performance decrements exhibited by the model.

We have constrained the selection of appropriate parameters and mechanisms for this effort in several ways. For instance, we have taken into account previous research in the ACT-R community (Belavkin, 2001; Jongman, 1998), and we have used the conclusions from neuropsychological research on the effect of sleep deprivation on the functioning of various brain areas, particularly the thalamus (Chee & Choo; Habeck et al., 2004; Portas et al., 1998). To use the conclusions from this work, we leveraged recent advances in the development of the ACT-R architecture, which have included mapping its components to brain areas (Anderson, chapter 4, this volume). This mapping establishes a "common space," where links between neuropsychological research on fatigue can be putatively linked with aspects of the architecture. The constraints imposed by this research implicate a mechanism in ACT-R that is related to the production selection/execution cycle as a candidate for being impacted by fatigue. This process is associated with the basal ganglia and the thalamus in the current conceptualization of ACT-R (Figure 17.3).

The production/execution cycle involves evaluating alternative productions and then selecting the "best" among them. During the selection process, productions are compared using a value called expected utility (U_i), which is calculated for each candidate production using the equation:

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

In this equation, P_i is the probability of success if production_{*i*} is used and C_i is the anticipated cost. In general, G has been termed *the value of the goal*. However,

Production Execution Cycle in ACT-R (including hypothesized mapping to brain areas)




FIGURE 17.3 Production execution cycle in the adaptive control of thought-rational (ACT-R) cognitive architecture, including hypothesized mapping to brain areas. The expected utility equation is associated with the selection component of this process, while execution is controlled by the utility threshold (T_u). Adapted from <http://actr.psy.cmu.edu/>.

the research cited above uses the G parameter to capture the influence of arousal on performance (Belavkin, 2001; Jongman, 1998). We use this conceptualization of G in our model as well. Noise (ϵ) is added to the calculation to add a stochastic component to the value. The noise is sampled from a Gaussian distribution with a mean of 0 and a variance of about 0.21.¹ A value for U_i is calculated for each production_{*i*} that matches the current state on each production cycle. The production_{*i*} with the highest value for U_i is selected.

Once a production is selected, the next step is execution. This process is associated with the thalamus in ACT-R (Figure 17.3). Production execution is controlled by a parameter called the *utility threshold*, T_u . The selected production is executed, provided that U_i exceeds T_u . If it does not, no production is executed and the model is "idle" for the duration of that production cycle (approximately 50 ms).² The neuropsychological data suggest that fatigue may indirectly affect this process, with individuals trying to offset the adverse effects through an attempt at compensation (Portas et al., 1998). As the behavior of the model illustrates, some compensation may be possible, but it does not completely offset the negative effects associated with sleep loss.

We find it encouraging that research on the neurobehavioral effects of fatigue and research within the ACT-R community both point to a common mechanism for capturing fatigue effects in ACT-R. The convergence of this research on the production selection/execution cycle in ACT-R indicates that one of the impacts of fatigue may be a decreased likelihood of successfully executing an appropriate sequence of productions. This entails both an increased likelihood of having cognitive cycles where the system is idle as well as the execution of inappropriate productions. Next we describe the model we constructed in ACT-R, which is based on this conceptualization of the impact of fatigue.

Model Design

Because the PVT is simple in design, the ACT-R model is relatively straightforward. Before the stimulus appears the model can (1) deliberately wait for the stimulus or (2) errantly make a response (a false start). Once the stimulus has appeared, the model can (1) attend to the stimulus and then respond (this is two productions) or (2) respond without attending the stimulus (a false start that happens to come after the stimulus appears and is therefore counted as an appropriate response). At any

point in the task, it is possible for the model to be idle for one or more cognitive cycles.

For nearly all the productions in the model, P_i was set to 1, meaning that the goal would be achieved successfully if that production was fired. The lone exception to this was the production that errantly responds. P_i for this production was 0, on the assumption that it is highly unlikely to result in achieving the goal of successfully responding to the stimulus. The consequence of this is a reduced likelihood of that production firing relative to the other, appropriate productions, since U_i becomes a negative value ($-C_i$). However, with noise added to the utility computation, this production is occasionally the one with the highest U_i and also rises above T_u .

To produce decrements in performance like those associated with sleep deprivation, we conceptualized increased sleep deprivation as resulting in decreased arousal. Therefore, we implement fatigue in ACT-R by decreasing the G parameter, in line with previous research within the ACT-R community (Belavkin, 2001; Jongman, 1998). Reducing the value of G decreases U_i for all productions, where P_i is greater than 0. This makes it more likely that the expected utilities for those productions will fall below T_u , thereby making them less likely to fire. In turn, this increases the probability of idle cognitive cycles. A key feature of this mechanism is that idle cycles actually become more likely than appropriate actions once U_i (before noise is added) falls below T_u , which happens after the first full day of total sleep deprivation.

In addition to the decreased arousal that is associated with sleep loss, we include a secondary process that can be viewed as an attempt to compensate for the negative effect of fatigue, as suggested by Portas et al. (1998). In ACT-R terms, the increased activity they observed in the thalamus can be seen as an attempt to make it easier for the selected production to fire successfully. The most natural way of representing this in ACT-R is by reducing T_u . With lower values for T_u , productions with increasingly lower U_i values are

able to fire. In cases where there is an idle cognitive cycle, new U_i values are computed at the beginning of the next cycle and the process repeats. The noise added to the calculation of U_i creates the possibility of all matching productions being below threshold on one cycle, followed by a cycle where at least one production rises above threshold.

There is one additional component to the model. A mechanism was implemented to represent the process of falling asleep. The idle cognitive cycles that result when no productions rise above threshold represent situations where arousal is so low that none of the available actions are executed. As an individual falls asleep, the probability of being in such a state should increase. In the model, this is represented by having the value of G decrease when idle cognitive cycles occur during the time when the stimulus is on the screen. On each occasion when this occurs, G is decremented by 0.035. This is limited by the architectural requirement that G remain positive. So, G is decremented by 0.035 on each idle cognitive cycle, unless that decrement results in a negative number, in which case G is set to a minimum value of 0.0001, where it remains stable until the start of the next trial. This progressively reduces the probability of a response with each passing cycle in the model. The value of G is reset to the starting value at the beginning of each trial, reflecting either a successful response or being awakened by the experimenter after a sleep attack. The values used for G and T_u for the four days of the experiment are presented in Table 17.1. As this table illustrates, G falls more rapidly than T_u , meaning that the model's performance deteriorates across the four days of the study. The model's performance is described in more detail in the next section.

Model Performance

Baseline performance in the model reflects the interplay of the knowledge in the system with the various mechanisms described above using the parameter values

TABLE 17.1 Parameter Values for G and T_u in the ACT-R Model for Each Day in the Sleep Deprivation Protocol

Day	G (Arousal)	T_u (Utility Threshold)
Baseline	1.98	1.84
Day 1 of total sleep deprivation	1.80	1.78
Day 2 of total sleep deprivation	1.66	1.70
Day 3 of total sleep deprivation	1.58	1.64

identified above and shown in Table 17.1. All other parameters were kept at their default ACT-R values. Baseline performance represents an alert, well-rested participant. The model's performance in this condition, along with the human data, is illustrated in panel A of Figure 17.4. The model closely captures all of the phenomena of interest. The largest discrepancy between the two data sets is that the human participants' alert responses tend to be somewhat faster than the model's. This may reflect an attentional strategy in use by the participants, but not implemented in the model.³ The accuracy of the model's predictions at baseline is matched by the predictions it makes for performance after one, two, and three days without sleep, as shown in panels B, C, and D of Figure 17.4, respectively. The model captures the increase in false starts, the shift in alert response times, and the increases in lapses and sleep attacks. For all four days, the performance of the model closely matches the human data. Overall, the correlation between the human data and the model predictions is 0.99 (root mean square deviation = 0.0047).

The impact of decreasing G in the model is to reduce the likelihood that the appropriate productions will have U_i values above T_u . We interpret this situation as reduced arousal (i.e., ACT-R gets sleepy), with the impact being that the appropriate productions are not executed. The reduction in arousal relative to the level of compensation immediately decreases the likelihood that a fast response will be made. Thus, there is a shift in the distribution of reaction times to the right, with progressively fewer fast responses as G gets lower. The decrease in frequency of fast responses has secondary effects. Each time an opportunity to respond is missed (idle cycles), G is decremented. This makes the model less likely to respond at the next opportunity. So the impact of lower arousal accumulates to produce many more lapses, along with more sleep attacks.

Finally, the false starts in the model increase similarly to the human data. This is primarily a side effect of the compensation in the model, reflected by the decreased values for T_u . As T_u decreases, it is more likely that the value for U_i for the production that produces a false start (i.e., *just click*) will have a U that

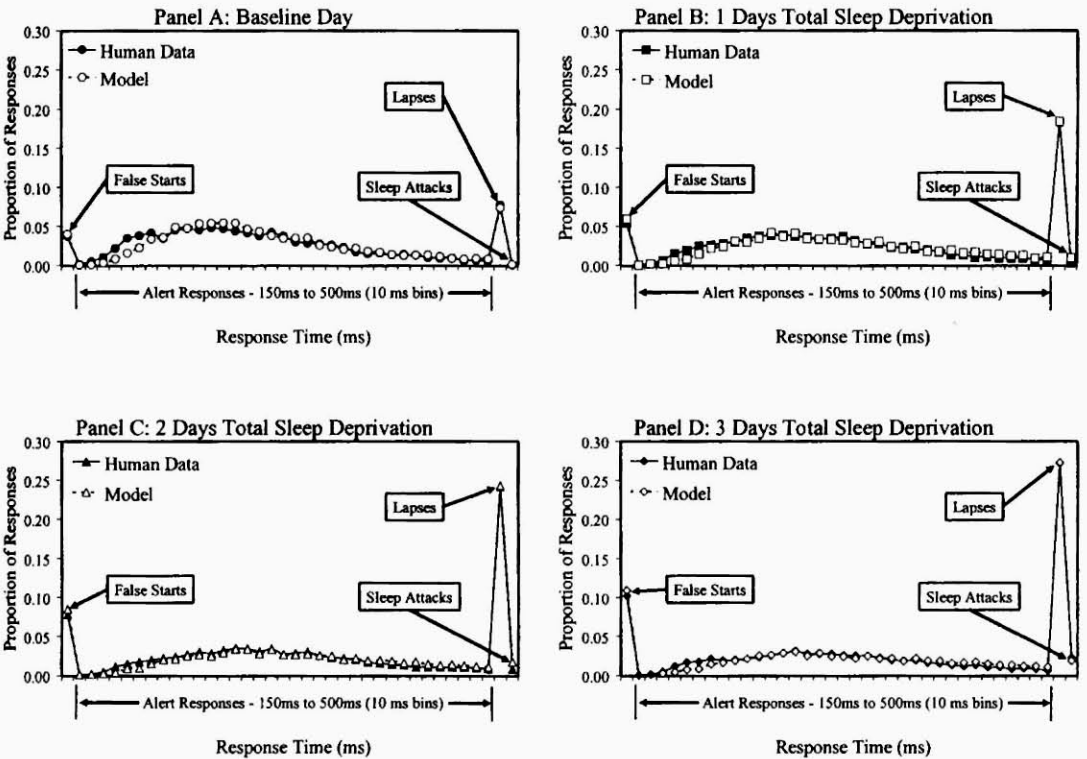


FIGURE 17.4 Human data and model performance for each day of the 88-hr total sleep deprivation protocol (data from Van Dongen et al., 2001; Van Dongen, 2004).

rises above threshold. Thus, the mechanism for compensation produces an undesirable behavioral consequence in the model that matches the human data; namely, an increased likelihood of responding before the stimulus appears. It is also the case that the decrease in G contributes to this effect (though to a lesser extent), since changes in G do not impact U for the *just-click* production (because $P = 0$ for this production). The combination of these two aspects of the model produces an increase in the probability of producing a false start that lines up with the data from the human participants.

General Discussion

The model described accurately captures the impact of sleep deprivation on human performance on the PVT. The mechanisms that were implemented are based on research within the ACT-R community, combined with neurobehavioral evidence concerning the impact of fatigue on brain activity. This illustrates the potential for using multiple constraints from diverse research communities to develop an understanding of cognitive mechanisms that impact performance. The ultimate goal of this work is to develop a general set of mechanisms for understanding how fatigue impacts the human cognitive system. The mechanisms must specify the relationship between fatigue and cognitive control, even if it is not done explicitly. ACT-R has components that correspond to the distinction between Type 1, or centralized, control and Type 2 control, or the internal control of functional processes.

The central production system represents the Type 1 control in ACT-R, while the other modules represent Type 2 control structures. The production cycle drives behavior in ACT-R models and directs the operation of the other modules. Other modules respond to requests from the production system by producing actions (e.g., motor movements or shifts of visual attention) or by performing some operation and making the results available to the production system (e.g., a retrieval from declarative memory).

At one extreme, alertness could be implemented as a separate module in ACT-R that operates as a Type 2 control structure to influence central cognition. At the other extreme, alertness could be viewed as a component of the Type 1 control system.⁴ We are not committed to a particular view of how sleep deprivation fits with cognitive control. Our model represents fatigue as a set of mechanisms that operate on the central

production system in ACT-R. Decreased arousal (G) and a lower utility threshold (T_u) both involve the central production system and influence the actions that are taken by the model. The performance decrements in the other modules stem from productions failing to fire to send the appropriate requests. Still, the model has little to say about whether those mechanisms are part of the Type 1 control structure or are a separate, Type 2, control structure. This is because the driving force behind the fatigue mechanisms is not yet integrated into the ACT-R architecture. Specifically, the biomathematical models of alertness, which we intend to use to drive the parameter changes, are currently outside the architecture. How they are represented in the architecture could have important implications for our understanding of how cognitive moderators like fatigue operate.

In the current model, the deficits in performance that develop with diminished alertness represent a breakdown in Type 1 control, where appropriate actions are no longer taken. Increases in the frequency of idle cognitive cycles produces a reduction and shift in alert response times. In addition, when the model is idle arousal continues to fall, making the model progressively less likely to respond. Essentially, the Type 1 control in the model disappears, as the central production system ceases to drive behavior forward. The model has, in essence, gone to sleep. Even though false starts are a consequence of a production becoming more likely to fire, this effect still represents a breakdown in control, since the production has a low probability of success. More generally, with the current mechanism for decreasing arousal and compensation, ACT-R becomes less able to differentiate between successful and unsuccessful productions. The effect is more erratic (i.e., unstable) performance, just as is found in human participants who are deprived of sleep (Doran, et al. 2001).

It is not the case that all of the effects must come through interaction with Type 1 control. We have considered an alternative mechanism for producing false starts that involves processes within the motor module. In ACT-R, motor actions can be preplanned so that they can be executed more rapidly. Type 1 control guides the execution of the planned action. It is possible that one of the effects of fatigue is to increase the likelihood that planned actions are inadvertently executed without direction from the central production system. There is evidence that increased levels of fatigue may reduce cognitive inhibition (e.g., Harrison & Home, 1998).

Thus, if planned actions are held back using a process like inhibition, then sleep deprivation could result in more frequent false starts through a breakdown in Type 2 control. We have not implemented this alternative, partly because the current ACT-R architecture does not include mechanisms for inhibiting motor movements that could be manipulated to produce this effect. However, it provides an interesting possibility for future research.

This alternative mechanism for false alarms raises an important issue for understanding fatigue. If it is the case that fatigue can directly influence one Type 2 control structure, like a motor module, then presumably other modules may be affected as well. Such a conclusion would suggest that sleep deprivation has effects that transcend the distinction between Type 1 and Type 2 control. That is, it is possible that sleep deprivation influences mechanisms associated with Type 1 control *and* Type 2 control.

Conclusion

We have presented a model that performs the PVT, combined with a set of mechanisms that produces decrements in performance on that task, which mirror the effect of sleep loss in human participants. The mechanisms implemented in the model were based on previous research, both in the ACT-R community and on the neurobehavioral impacts of sleep loss. The data were fit at the level of average performance for each day of the study. We have demonstrated that these mechanisms produce performance changes in the model that closely match the performance changes of human participants who have been denied sleep over a period of nearly four consecutive days. The next step is to fit the model to the experimental data from each 10-minute session, which occurred at 2-hour intervals throughout the study. This effort is already underway and has provided further support for our approach. We are using the results from this process to develop a mapping between alertness predictions generated by CNPA and parameter values in the ACT-R model. Once the appropriate mapping is identified, it will be used to validate our ability to make a priori predictions about human performance in other experimental protocols and for other tasks.

A major challenge for validating our model lies in its ability to account for findings from a variety of tasks. The model's performance on the PVT is almost entirely

dependent on procedural components of the architecture; the only declarative knowledge needed is a representation of the goal. It will be interesting to see whether additional mechanisms are required to account for changes in performance associated with tasks that require more declarative knowledge. For instance, we are currently developing a model for a serial addition/subtraction task (SAST), which relies on knowledge of simple (single-digit) addition and subtraction problems. Our initial examinations of the model for this task suggest that the procedural mechanisms that produced the effects in PVT may not be sufficient to account for the changes in response times and accuracy on the SAST. This may indicate that decreased alertness effects many aspects of the cognitive system simultaneously. Certainly, other areas of the brain show changes in activity as a function of sleep loss. As this research progresses and matures, we hope to develop a comprehensive, mechanistic explanation of the changes that occur at the level of the human cognitive architecture as a function of changes in alertness. To do this, we will continue to incorporate as many theoretical and empirical constraints as we can, to guide the identification of mechanisms and to inform our understanding of the processes involved.

Acknowledgments

Cognitive model development was sponsored partly by AFRL's Warfighter Readiness Research Division and partly by grant number 04HE02COR from the Air Force Office of Scientific Research (AFOSR). Human data collection was sponsored by AFOSR grants F49620-95-1-0388 and F49620-00-1-0266, and by NIH grant RR00040. We would like to thank Robert O'Connor for his role in the original experiment described here, and for his more recent assistance in providing the raw data for comparison with the model. We would also like to thank Wai-Tat Fu, Frank Ritter, and Mike Schoelles for their helpful comments on earlier drafts.

Notes

1. In ACT-R, noise is controlled by a parameter, s , whose relation to the variance (σ^2) is defined by the equation: $\sigma^2 = (\pi^2 * s^2) / 3$. In this model, s was set to the default ACT-R value of 0.25.

2. ACT-R does not produce idle cycles by default. This required an enhancement of this component of the architecture, and we are grateful to Dan Bothell for his help implementing this. In addition, cycle times in this model were noisy, which is an available option in ACT-R. The duration of a particular cycle in the model varied between 25 ms and 75 ms, according to a uniform distribution, which is the default function for this noise value in ACT-R.

3. The implemented model does not explicitly maintain attention once it has responded to the stimulus. As a result, when the next stimulus appears, an extra production is required to shift attention to it. With a probabilistic mixture of this process and one that maintains attention on the stimulus location between trials, the distribution of response times can be captured more accurately. However, the better fit comes at the cost of model complexity, which seemed unwarranted for the current modeling effort.

4. It is possible to consider a third option, which is that arousal is outside cognitive control altogether (i.e., it is not under control). However, arousal clearly seem to affect aspects of control in cognition, suggesting that it operates through these structures, if not as one, or part, of them.

References

- Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review*, *111*, 1036–1060.
- Belavkin, R. V. (2001). The role of emotion in problem solving. In C. Johnson (Ed.), *Proceedings of the AISB '01 Symposium on emotion, cognition, and affective computing* (pp. 49–57). York, UK: Heslington.
- Bonnet, M. H., Balkin, T. J., Dinges, D. F., Roehrs, T., Rogers, N. L., & Wesensten, N. J. (2005). The use of stimulants to modify performance during sleep loss: A review by the Sleep Deprivation and Stimulant Task Force of the American Academy of Sleep Medicine. *Sleep*, *28*(9), 1144–1168.
- Caldwell, J. A., Caldwell, J. L., & Darlington, K. K. (2003). Utility of dextroamphetamine for attenuating the impact of sleep deprivation in pilots. *Aviation, Space and Environmental Medicine*, *74*(11), 1125–1134.
- , Caldwell, J. L., Smith, J. K., & Brown, D. L. (2004). Modafinil's effects on simulator performance and mood in pilots during 37 h without sleep. *Aviation, Space and Environmental Medicine*, *75*(9), 777–784.
- Chee, M. W. L., & Choo, W. C. (2004). Functional imaging of working memory after 24 hr of total sleep deprivation. *Journal of Neuroscience*, *24*(19), 4560–4567.
- Dinges, D. F., & Broughton, R. J. (Eds.). (1989). *Sleep and alertness: Chronobiological, behavioral and medical aspects of napping*. New York: Raven Press.
- , & 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: Wiley.
- , & 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.
- 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.
- , Brown, G. G., Stricker, J. L., Buxton, R. B., Wong, E. C., & Gillin, J. C. (1999). Sleep deprivation-induced reduction in cortical functional response to serial subtraction. *Neuroreport*, *10*, 3745–3748.
- , 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.
- Gratch, J., & Marsella, S. (2004). A domain-independent framework for modeling emotion. *Cognitive Systems Research*, *5*, 269–306.
- Habeck, C., Rakitin, B., Moeller, C., Scarmeas, N., Zarahn, E., Brown, T., et al. (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.
- Harrison, Y., & Horne, J. A. (1998). Sleep loss impairs short and novel language tasks having a prefrontal focus. *Journal of Sleep Research*, *7*(2), 95–100.
- Hudlicka, E. (2003). Modeling effects of behavior moderators on performance: Evaluation of the MAMID methodology and architecture. In *Proceedings of the 12th Conference on Behavior Representation in Modeling and Simulation* (pp. 207–215). Orlando, FL: Institute for Simulation and Training.
- Jewett, M. E., & Kronauer, R. E. (1999). Interactive mathematical models of subjective alertness and

- alertness in humans. *Journal of Biological Rhythms*, 14, 588–597.
- Jongman, L. (1998). How to fatigue ACT-R? In *Proceedings of the Second European Conference on Cognitive Modelling* (pp. 52–57). Nottingham: Nottingham University Press.
- Kleitman, N. (1923). The effects of prolonged sleeplessness on man. *American Journal of Physiology*, 66, 67–92.
- Lee, M. A. M., & Kleitman, N. (1923). Studies on the physiology of sleep: II. Attempts to demonstrate functional changes in the nervous system during experimental insomnia. *American Journal of Physiology*, 67, 141–152.
- Lin, J. S. (2000). Brain structures and mechanisms involved in the control of cortical activation and wakefulness, with emphasis on the posterior hypothalamus and histaminergic neurons. *Sleep Medicine Reviews*, 4(5), 471–503.
- 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, A4–A14.
- Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.
- Patrick, G. T. W., & Gilbert, J. A. (1896). On the effects of loss of sleep. *Psychology Review*, 3, 469–483.
- 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., Reifers, A., Klein, L. C., Quigley, K., & Schoelles, M. (2004). Using cognitive modeling to study behavior moderators: Pre-task appraisal and anxiety. In *Proceedings of the Human Factors and Ergonomics Society* (pp. 2121–2125). Santa Monica, CA: Human Factors and Ergonomics Society.
- 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), 15–36.
- , 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.
- von Economo, C. (1930). Sleep as a problem of localization. *Journal of Nervous and Mental Disease*, 71, 249–259.
- Wu, J., 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, 155–162.