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.

Approaches to Modeling the Effects of Fatigue on Cognitive Performance

Glenn Gunzelmann
Kevin A. Gluck
Air Force Research Laboratory
Warfighter Readiness Research Division
6030 S. Kent St.
Mesa, AZ 85212
480-988-6561 x-677
glenn.gunzelmann@mesa.afmc.af.mil;
kevin.gluck@mesa.afmc.af.mil

Keywords:
Fatigue, Computational Model, Biomathematical Model, Cognitive Architecture

ABSTRACT: *Behavior representations in simulation environments have become increasingly sophisticated as computational power and understanding of human cognitive processes have increased. Despite improvements in the breadth and depth of models of human behavior, comparatively little research has explored how human cognition is moderated by external and internal factors. In this review paper, we focus on fatigue, a pervasive cognitive moderator that has dramatic and sometimes tragic effects on human cognitive functioning. We discuss prior and current research investigating models of how fatigue impacts human cognition and behavior, sometimes leading to dramatic degradations. The paper ends with a description of our on-going efforts to identify mechanisms for fatigue within the Adaptive Control of Thought – Rational (ACT-R) cognitive architecture. Our approach involves integrating diverse research areas, especially biomathematical modeling of the effects of fatigue on neurobehavioral functioning, and computational cognitive modeling of human cognition and behavior.*

1. Motivation and Background

Increasingly sophisticated human behavior representations are rapidly becoming a standard component of modern modeling and simulation capabilities. These entities can serve in a variety of roles, including adversaries, teammates, and neutral forces. They engage in complex tasks and have become critical to serving the variety of goals associated with immersive environments, simulations, synthetic tasks, and games, including training, education, analysis, and entertainment.

As the tools, techniques, and implementations of human behavior representations have begun to mature, increasing emphasis has been placed on ensuring that they accurately reproduce the cognitive capacities, limitations, and processes of the human participants they replace. One particular area of increased focus is the impact of cognitive moderators (e.g., Gluck, Gunzelmann, Gratch, Hudlicka, & Ritter, 2006). Cognitive moderators are factors that influence human cognition, but are not necessarily part of the goal-directed thought that is the focus of information processing. In this paper we focus on mathematical and computational implementations of fatigue as a cognitive moderator.

Psychological research has only begun to uncover the details of how cognitive moderators affect thought. In the case of fatigue, however, there is an extensive empirical literature documenting the negative consequences of sleep restriction, sleep deprivation, and the intrinsic circadian clock (e.g., Durmer & Dinges, 2005; Dinges, Baynard, & Rogers, 2005). In contrast, models of these processes, and the concomitant changes in cognitive functioning associated with these processes, are relatively few. Our goal here is to review significant prior and on-going efforts and accomplishments in the development of mathematical and computational theories of how fatigue impacts cognitive performance.

The approaches are varied, but they are linked by a common overarching goal: The desire to make principled, quantitative predictions about how increased levels of fatigue will impact performance in applied settings. The applications of this research are extensive. A surprising number of industrial and commercial disasters have been attributed at least partially to fatigue (e.g., Caldwell, 2003; Mitler et al., 1988), and so have a sizable proportion of other tragedies like automobile accidents (e.g., Horne & Reyner, 1999). One application of this research is to make predictions regarding an individual's readiness to engage in a particular, cognitively demanding activity. In fact, some research on fatigue has led to transitions of tools to industry to support decision making concerning work-rest schedules (e.g., Dean, Fletcher, Hursh, & Klerman, 2007).

Another potential application of this research is in the development of synthetic entities to populate virtual environments. In training contexts it is valuable for the trainee experience to match as closely as possible the cognitive demands that will be present in the operational context. In many operational environments, restricted sleep and periods of acute sleep deprivation are the norm, meaning that most or all individuals will be operating at reduced levels of alertness. This is certain to have consequences for their performance. Thus, trainees should benefit from interacting with synthetic teammates that can display the appropriate consequences of fatigue, resulting in better preparedness for the real world context.

In the next section, we review a number of research efforts aimed at developing quantitative models of fatigue. We begin with efforts directed specifically at developing behavior representations that exhibit fatigue-related changes in performance, and then consider several projects that have focused on generating algorithms to quantify the dynamics of fatigue, but which are not connected directly to cognitive mechanisms to simulate human performance. Following this overview, we describe efforts at integrating these approaches, including our research, which represents a synthesis of these two broad approaches to understanding fatigue.

2. Approaches to Modeling Fatigue

Fatigue has been an important topic of research for decades within psychology. Behavioral studies of the effects of fatigue have a long history, and have documented the negative effects associated with lack of sleep across a broad range of tasks and domains (e.g., Durmer & Dinges, 2005; Dinges et al., 2005). As technological advances enabled more sophisticated neuroimaging research, investigations of fatigue focused more extensively on neurophysiological and neuropsychological data, uncovering a variety of changes in brain functioning associated with fatigue and circadian rhythms (e.g., Drummond, Brown, Salamat, & Gillin, 2004; Drummond, Gillan, & Brown, 2001; Portas et al., 1998). Most recently, some effort has been invested in using the extensive empirical literature to develop computational accounts of the dynamics of fatigue and circadian rhythms and the impact of those factors on cognitive functioning.

2.1 Architected Models of Fatigue

There have been two prior efforts over the last 10 years to incorporate mechanisms for fatigue into architecture-based behavior representations. In both cases these appear to have been single-model investigations of this topic. The goal in each case was to increase the cognitive validity of the model, to better reflect the limitations of human arousal and motivation.

2.1.1 Modeling Within-Task Fatigue in ACT-R

Jongman (1998; Jongman & Taatgen, 1999) conducted research to understand how to represent the impact of extended time on task within the ACT-R cognitive architecture (Anderson et al., 2004). Although this is related to the topic of sleep- and circadian-influenced fatigue, it seems worth emphasizing that the cognitive fatigue associated with extended time on task may be distinct from fatigue caused by long periods of wakefulness. Jongman and Taatgen developed a task requiring participants to determine the weights of packets of coffee (under a set of constraints), and identified two general strategies for completing the task. One of the strategies required more effort to execute effectively, but led to better performance compared to the alternative strategy, which was easier to execute but less successful. They found that participants tended to shift from the more effective to the less effective strategy when mentally fatigued.

The model developed by Jongman and Taatgen instantiates the two strategies for performing the task. The transition that occurs as mental fatigue increases is achieved by manipulating a parameter that controls the likelihood of selecting one strategy versus another. Using this mechanism, Jongman and Taatgen were able to account for strategy shifts for individual participants, including relatively good fits to performance on individual trials.

2.1.2 Modeling Pilot Fatigue with Soar

Using the Soar cognitive architecture, Jones, Laird, & Neville (1998) added mechanisms for fatigue to enhance the realism of behavior representations within the TacAir-Soar model of human pilot performance (Laird et al., 1998), by adding mechanisms to model the effects of cognitive moderators on performance. To accomplish this, Jones et al. (1998) made reference to a mathematical model developed by Neville (1997) and colleagues (Neville et al., 2000) to represent the dynamics of fatigue (the SILCS model described below). The model includes mechanisms to produce both a slowdown in cognitive functioning and an increased prevalence of “lapses” in cognitive functioning, representing two popular hypotheses regarding the mechanisms through which fatigue impacts performance (e.g., Dinges & Kribbs, 1991).

To introduce fatigue effects into the TacAir-Soar model, Jones et al. (1998) modified Soar’s mechanisms to allow for the introduction of artificial delays in processing. Once these mechanisms were implemented, Jones and colleagues compared the performance of "fatigued" pilot models to "non-fatigued" pilot models, however their comparisons were based only upon preliminary data and anecdotal evidence regarding the validity of the model's behavior. In addition, it is not clear that the model from Neville (1997) was integrated into TacAir-Soar to generate differential predictions across different levels of fatigue or at different points in the circadian cycle.

Despite the limitations, this research does represent perhaps the first serious attempt to develop robust behavior simulations that show performance decrements as a consequence of sleep-related fatigue. Indeed, the methodology pursued by Jones et al. (1998) shares similarities with our own approach to developing mechanisms for fatigue, which we describe below. One important commonality is the interest in integrating a mathematical model representing the dynamics of fatigue. Such models have been developed independently from direct attempts to produce behavior models. Some of these efforts are described in the next section.

2.2 Mathematical Algorithms for Fatigue

2.2.1 The Sleepiness-Induced Lapsing and Cognitive Slowing (SILCS) Model

SILCS is an empirically-derived model for describing the behavioral dynamics of human performance (Neville et al., 2000) under conditions of total sleep deprivation. Interesting strengths include a solid theoretical basis for the components of the model (hemi-circadian, circadian, and sleep homeostat) and fits to multiple dependent measures. Neville and colleagues explicitly state that the intended application for SILCS is in synthetic warfighters for high fidelity simulation systems. SILCS itself is comprised of three functions representing the three empirical findings of interest to its developers: lapse frequency, lapse duration, and general response slowing. Each of these three functions includes parameters representing hemi-circadian, circadian, and sleep homeostat components, as well as a baseline performance parameter.

The SILCS modeling methodology involves finding the best-fitting mathematical parameters for the empirical result relevant to that function. For instance, Neville et al. (2000) computed the number of lapses occurring per minute in their empirical results, and then fit the "Lapse frequency per min" function to that result. They did the same, separately, for mean lapse durations and also for mean response times, with different best-fitting parameter values for the different dependent measures. Neville et al. did not provide quantitative fit metrics, such as correlations or deviations, but they did plot their model predictions and the empirical human data from two studies, which allows for a qualitative visual assessment of the goodness of their account. For many measures, the fit appears to capture qualitative trends, although there are several cases where the model predictions stray considerably from the observed data.

2.2.2 The Fatigue Degradation (FADE) Algorithm

French and Morris (2003) describe a development and validation approach for their FAtigue DEgradation (FADE) algorithm that is similar in spirit to that used for SILCS, but extends the methodology employed by Neville et al. (2000) in a couple of interesting ways. The fundamental similarity is that the FADE algorithm parameters are derived from a sample of human subjects in a sleep deprivation experiment. The algorithm itself is simpler, however, comprised of a linear "time awake" component and a cosine "circadian oscillation" component. French and Morris fit the FADE algorithm to data from a 52-hour sleep deprivation experiment and found that it correlated .92 with the human subjects data. They then did a zero-free-parameter fit of the same model to data from a different task in the same 52-hour sleep deprivation experiment and found that it correlated .87 with those data. This zero-free-parameter prediction is one way in which French and Morris went beyond the methodology of Neville et al.

French and Morris also extended their model by adding a parabolic recovery component that allows them to make predictions about fatigue levels over longer periods of time and varying levels of sleep restriction. With this extension in place, French and Morris then compared the fatigue levels predicted by two maritime work-rest schedules: the centuries old 4-hour watch standard that is in common use today, which FADE predicts results in unacceptably high fatigue levels after just 8 days of deployment, and an alternative schedule that appears to maintain acceptable fatigue levels indefinitely. There is no validation of these predictions against human data, but the effort speaks to an intended application of fatigue models of this sort, which is to guide policy and practice.

2.2.3 Biomathematical Models of Fatigue

The approaches to modeling fatigue described so far have been developed by researchers whose focus is on developing behavior representations. In contrast, biomathematical models represent the main approach to modeling the effects of fatigue that have come out of the sleep research community. There are a number of these models (see Klerman & St. Hilaire, 2007; Mallis, Mejdal, Nguyen, & Dinges, 2004; Van Dongen, 2004), which all share a common theoretical origin in a two-process theory of human arousal and alertness (Acherman, 2004; Borbely & Acherman, 1999). This theory posits that there are two main influences on overall alertness, a sleep homeostatic pressure that increases with time awake and a circadian rhythm that varies over the course of approximately 24 hrs.

Biomathematical models have been developed on the basis of extensive empirical investigations of human neurobehavioral performance under conditions of sleep deprivation. Participants in these studies are monitored continuously (often in a hospital environment) and perform a wide array of experimental tasks over the course of sometimes extensive laboratory stays. In addition to performance data from experimental tasks, physiological measures are also obtained, including core body temperature, hormone levels, and often neurophysiological data (e.g., EEG, fMRI, etc.). From these data, models of the dynamics of the sleep homeostatic pressure and circadian rhythm have been developed that accurately capture the relative changes in human neurobehavioral functioning observed under conditions of sleep deprivation.

Biomathematical models have been influential in the sleep research community, and have been transitioned to applied settings where they are used in planning and scheduling to maximize overall effectiveness. Perhaps the most impressive case of transition for this class of models to date is the example set by the Sleep, Activity, Fatigue, and Task Effectiveness (SAFTE™) model (Hursh et al., 2003). SAFTE™ is based on the same general two-process model as the others cited above, and also includes a sleep inertia component that accounts for the well-documented period of sleepiness that persists briefly after awakening. It makes predictions about cognitive effectiveness, relative to a well-rested baseline. SAFTE™ has been incorporated into a software package called the Fatigue Avoidance Scheduling Tool (FAST™). To their credit, Hursh et al. performed a quantitative evaluation of the validity of the cognitive effectiveness predictions of SAFTE™ using empirical data from two studies, resulting in very high correlations.

All models have their weaknesses and the biomathematical models are no exception. They suffer from at least two important limitations. First, on the whole they have been developed with little attention to results from sleep restriction studies of fatigue. Thus, whereas they make fairly accurate predictions regarding the relative changes in overall functioning resulting from sleep deprivation, they are much less accurate in predicting the dynamics associated with long-term restricted sleep.

The other main limitation of biomathematical models of alertness is that they produce only predictions of overall level of functioning. That is, these models do not generate predictions of performance in particular tasks. In many cases, the output from these models is scaled to particular dependent measures to produce quantitative estimates of performance in specific contexts (e.g., Van Dongen, 2004). While this may be a useful approach in some circumstances, it is not ideal in complex, real-world settings, where the relevant data regarding human performance may be difficult or impossible to collect ahead of time. In these cases, models are needed that can produce accurate, *a priori*, quantitative estimates of human performance. In the next section, we describe integrated approaches to modeling fatigue, which are aimed at achieving this goal.

2.3 Hybrid Approaches

2.3.1 PMFServ

PMFServ is a human behavior emulation system that synthesizes dozens of performance moderator functions (PMFs) into a unifying framework (Silverman, 2007). PMFs are mathematical characterizations of the influences of individual factors (sleep, affect, role in group) on performance. They are organized into seven modules: perception, biology/stress, personality/culture/emotion, memory, social, decision making, and expression. The implementation of the effects of fatigue is located in the biology/stress module, as a component of the Gillis and Hursh (1999) model of stressed performance, which is a linear additive model that includes event stress, time pressure, and effective fatigue (Silverman, Johns, Cornwell, & O'Brien, 2006). Effective fatigue is described in Silverman et al. as "a normalized metric based on current level of many of the physiological reservoirs" (p. 147) Regarding the validity of the fatigued performance predictions generated by PMFServ, we are unable to find an explicit comparison with quantitative human performance data.

2.3.2 IMPRINT Pro

IMPRINT Pro is the latest in the line of IMProved Research INTegration (IMPRINT) task network modeling tools, which allow analysts to develop human performance models for estimating manpower, personnel, and training (MPT) requirements and constraints for new weapon systems very early in the acquisition process. IMPRINT Pro includes a variety of new features, one of which is the incorporation of fatigue as a moderator for its performance predictions (Alion Science and Technology, personal communication). The approach that was adopted for achieving this was to build the cognitive effectiveness algorithms from SAFTE™ into a plugin (.dll) that can be loaded into the IMPRINT Pro software. With this feature loaded, as cognitive effectiveness predictions deviate from maximal, those changes affect task completion times. The incorporation of this feature into IMPRINT Pro is a very new development and there are not yet any validation studies completed.

2.3.3 ACT-R/F

Our research on fatigue is in very much the same spirit as the efforts described above. We aim to develop a psychologically valid theory of fatigue that can be used to understand and predict changes in cognitive functioning resulting from extended periods of wakefulness, restricted sleep, and circadian rhythms. Part of what is distinct about our research, however, is that we bring together mathematical modeling research and cognitive architectures in the development of behavior representations that attempt to capitalize on the strengths of both approaches in a flexible and generalizable way.

We have adopted an approach that begins with careful validation of mechanisms in detailed laboratory tasks, followed by applying those mechanisms to account for human behavior in more complex, naturalistic tasks. The result, we believe, is a more detailed and powerful account of fatigue that will facilitate making predictions about human performance in applied contexts. Our goal in this paper is not to review in detail our modeling results, but rather to describe our approach and its contribution to research in this area. For details on the modeling results, the interested reader is directed to the publications that are cited throughout this section.

At the core of our research is a focus on understanding the details of human performance, and how that performance changes as fatigue levels increase. We use the ACT-R cognitive architecture to represent "normal" human cognitive performance (Anderson et al., 2004), and identify mechanisms within ACT-R that can be manipulated to produce performance decrements associated with fatigue (e.g., Gunzelmann, Gluck, Van Dongen, O'Connor, & Dinges, 2005). This work has allowed us to identify mechanisms in multiple components of ACT-R that appear to be sensitive to the negative effects of sleep deprivation (e.g., Gross, Gunzelmann, Gluck, Van Dongen, & Dinges, 2006; Gunzelmann et al., 2005; Gunzelmann, Gluck, Kershner, Van Dongen, & Dinges, 2007).

Identifying mechanisms within ACT-R that are impacted by fatigue is a critical step in our research. However, to bring us closer to the goal of making a priori predictions about the impact of fatigue, it is important to have some constraint on selecting parameter values within those mechanisms and for identifying the dynamics of those parameters as time awake increases and time of day changes. For this purpose, we leverage existing biomathematical models of alertness (e.g., Gross et al., 2006; Gunzelmann et al., 2007; Gunzelmann, Gluck, Kershner, Van Dongen, & Dinges, under review). Specifically, we use a function to map biomathematical model predictions of alertness to parameters in ACT-R. Thus far, we have used linear functions of the form:

$$P_x = s \cdot A_t + I$$

In this equation, P_x represents a particular parameter x in ACT-R, s_x represents the slope of the linear equation, and I_x is the intercept. A_t is the predicted level of alertness at time t from the biomathematical model. Using this straightforward approach, we have been able to successfully account for changes in human performance as a consequence of fatigue in multiple tasks emphasizing different components of ACT-R (see Gross et al., 2006; Gunzelmann et al., 2007; Gunzelmann, Gross, et al., under review). These results are described in the following subsection.

In our research efforts, we have focused on developing accounts for changes in human performance on relatively simple laboratory tasks, which rely differentially on various components of cognitive functioning. This strategy allows us to get "close to the architecture" (Newell, 1990), which is helpful when trying to find architectural mechanisms that explain empirical results, and it allows us to rely on academic collaborators for human subjects experiment data, which is helpful for evaluating the validity of our theory. We began by investigating a sustained attention task, emphasizing central cognition (e.g., Gunzelmann et al., 2005; Gunzelmann, Gross, et al., under review). Since then, we have expanded this research to a serial addition/subtraction task requiring the retrieval of simple math facts (Gunzelmann et al., 2007) and to a dual-task reaction time paradigm (Gunzelmann, Byrne, Gluck, & Moore, under review).

2.3.3.1 Sustained Attention Performance

Sustained attention performance has been shown to be highly sensitive to the negative consequences of fatigue (e.g., Dinges & Powell, 1985; Dorrian, Rogers, & Dinges, 2005). To explore computational mechanisms associated with declines on this task, we developed a model to perform the Psychomotor Vigilance Test, or PVT (Dinges & Powell, 1985). This task requires participants to monitor a location on a computer monitor for the presentation of a stimulus. When the stimulus appears, participants respond by pressing a response button. Stimuli are presented with moderate frequency, at delays distributed uniformly from 2-10 s. However, the duration of a session, 10 minutes, makes the task challenging to perform, particularly under conditions of fatigue (e.g., Dinges & Powell, 1985; Doran, Van Dongen, & Dinges, 2001).

To model these effects in ACT-R, we considered task constraints, previous theoretical work in ACT-R (e.g., Belavkin, 2001; Jongman, 1998), and neuropsychological findings on attention and fatigue (e.g., Drummond, et al., 2001; 2004; Portas et al., 1998) to target our search for appropriate mechanisms in the architecture. All of these constraints suggested that ACT-R's production system, representing a central cognitive bottleneck in processing, was the appropriate mechanism to target to understand degraded performance on the PVT.

Indeed, manipulations to parameters that influence this mechanism result in degraded model performance on the PVT that corresponds to the kinds of decrements observed in human participants (e.g., Gunzelmann et al., 2005). Further, the dynamics of these parameter changes can be constrained using biomathematical models of alertness, resulting in an account that is both strongly theoretically motivated and also able to produce quantitative performance predictions that correspond closely to observed empirical data (see Gross et al., 2006; Gunzelmann, Gross, et al., under review).

2.3.3.2 Serial Addition/Subtraction

To understand degradations in cognitive performance on tasks requiring significant use of declarative knowledge, we investigated changes in performance on the Walter Reed Serial Addition/Subtraction Task, or SAST (Thorne Genser, Sing, & Hegge, 1985). With increased fatigue, performance on this task degrades both in terms of average response time and accuracy (Van Dongen et al., 2001). An initial investigation showed that the procedural mechanisms used to model fatigue effects in the PVT were not sufficient to capture changes in performance on the SAST. Thus, to model these effects we added mechanisms in ACT-R's declarative memory module that are analogous to the mechanisms manipulated in our account of the PVT within ACT-R's procedural knowledge. In both cases, the mechanisms are associated with the availability of knowledge and the likelihood that it will be used in a particular context. Once again, the decrements in the model corresponded to the effects observed in human participants. We were also able to constrain the dynamics of the parameters in this account using biomathematical model predictions of alertness (see Gunzelmann et al., 2007).

It was not surprising that our procedural mechanisms were not sufficient to account for fatigue effects in a more complex task with a more substantial reliance on declarative knowledge. Physiological changes in brain functioning stemming from fatigue are widespread. This leads to an expectation that mechanisms in multiple components of ACT-R will be required to capture the broad effects that are observed in human participants.

Importantly, in our model of human performance on the SAST, we included the mechanisms implemented in ACT-R's procedural memory in our account of fatigue. Our motivation for this was based upon the recognition that fatigue has consistent effects on cognitive performance regardless of the task being performed. Differences in how severely performance is degraded is a function of how dependent task performance is on particular components of cognitive functioning, rather than variable effects of fatigue. Thus, for a truly general account, all of the mechanisms we identify must be applied across tasks, rather than picking and choosing which components to use on a task by task basis.

2.3.3.3 Dual Task Performance

Most recently, we have attempted to generalize the mechanisms developed to account for fatigue effects in the PVT to a dual-task paradigm and the psychological refractory period (Gunzelmann, Byrne, et al., under review). The task involves an auditory two-choice reaction time task combined with a visual two-choice reaction time task. The delay, or stimulus onset asynchrony (SOA), between the onset of the auditory stimulus (Task 1) and the presentation of the visual stimulus (Task 2) is varied. Reaction time on Task 2 when the SOA is long (i.e., 1000 ms) is faster than when the SOA is short (i.e., 50 ms), because individuals are done responding to the Task 1 stimulus before the Task 2 stimulus appears at the longer SOA. The difference in Task 2 reaction time from short SOA's to long SOA's is referred to as the psychological refractory period (PRP).

In this effort, we have demonstrated that our account of fatigue in the PVT task also provides a sufficient account of changes in performance on each task in the dual task paradigm and changes in the magnitude of the PRP effect resulting from increased fatigue (Gunzelmann, Byrne, et al., under review). We continue to use the biomathematical model predictions of alertness to constrain the dynamics of the parameters in this model. This model has also provided evidence regarding theoretical arguments about the causes of performance decrements, suggesting that a general slow-down in cognitive processing is not necessary to account for changes in human performance resulting from fatigue.

3. Challenges & Future Directions

Thus far, we have been encouraged by the results of our efforts to develop computational mechanisms for fatigue within the ACT-R cognitive architecture. In particular, the integration of a cognitive architecture with biomathematical model predictions of alertness has added real explanatory power while reducing degrees of freedom in fitting our models to empirical data. We have been able to build upon the strong foundation ACT-R provides as a validated theory of human cognition, and make connections to carefully developed and validated models of the dynamics of the human arousal system.

Despite the encouraging progress made thus far, there remain many important and challenging questions, which will only be answered through the integration of research efforts using multiple methodologies in multiple domains. For instance, we have only begun to explore mechanisms within the ACT-R cognitive architecture that may be impacted by fatigue. Thus far, we have focused on the production system in central cognition and ACT-R's declarative memory, the two components of ACT-R with the longest history and most well-validated mechanisms. However, ACT-R comprises other functional capabilities as well, including multiple learning mechanisms, as well as perceptual and motor modules. There is evidence that these aspects of cognitive functioning may degrade with fatigue as well (e.g., Atkinson, Coldwells, & Reilly, 1993, Heuer, Spijkers, Kiesswetter, & Schmidtke, 1998; Tassi, Pellerin, Moessinger, Eschenlauer, & Muzet, 2000). Thus, important extensions to our account are necessary before we can claim to have a general theory that can be applied across tasks and domains.

A related issue to exploring mechanisms in other cognitive functions is extending accounts of fatigue to larger and more complex tasks. Above we described research that implemented mechanisms for fatigue within a complex aircraft pilot model (Jones et al., 1998). This is an example of the kind of application where quantitative models of fatigue are needed. On the one hand, Jones et al. (1998) did not include careful empirical validation in their modeling efforts. On the other hand, our efforts are a long way from scaling up to such complex, knowledge-rich environments, although the architectural mechanisms themselves should scale unproblematically.

There are also limitations of the biomathematical models that have been developed to represent the dynamics of alertness in relation to time awake and circadian rhythms (Klerman & St. Hilaire, 2007). These models make relatively accurate predictions regarding the dynamics of alertness under conditions of total sleep deprivation. However, their accuracy diminishes substantially for making predictions when sleep is restricted, such as obtaining only 4 or 6 hours of sleep per night for several weeks continuously (Van Dongen, 2004). To be useful in applied contexts, these models must be able to make predictions for such scenarios, since restricted sleep is prevalent in applied contexts and operational settings.

Within our own research efforts, there is need to fully integrate the biomathematical models into the ACT-R cognitive architecture. Such an integration would allow modelers to vary the time of day and the sleep history of the model in a simulation and produce differential predictions about performance. This functionality would provide a potentially powerful tool for understanding the trade-offs intrinsic to planning operations and work-rest schedules.

The challenges, however, run much deeper than issues of integration and scope. In addition to understanding the relationship between alertness and human information processing mechanisms, there is also the tricky issue of how the dynamics of alertness are impacted by interaction with the environment. There is evidence that the deleterious effects of fatigue are modulated by the task environment, which can have an arousing effect on the cognitive system when the task is engaging (e.g., Caldwell & Ramspott, 1998; Drummond et al., 2004; Pilcher, Band, Odle-Dusseau, & Muth., 2007). In addition to these interactions, there are the complicated relationships of fatigue with other cognitive moderators, like stress, emotions, and drugs. Drugs may be an especially important topic, given the prevalence of substances like caffeine, Dexedrine, and modafinil for offsetting the negative effects of sleep loss (e.g., Bonnet et al., 2005).

Finally, the ultimate achievement of this research is not to predict changes in performance under conditions of fatigue at the aggregate level, but rather to develop a tool that would allow the relevant parameters to be tailored to the capabilities and sensitivity to fatigue of particular individuals. Such a tool would allow decision makers in applied settings to judge the suitability of schedules and assignments on the basis of quantitative predictions of performance for a particular person based upon particular assumptions regarding that individual's history of sleep and wakefulness. We are still far from such a vision. To achieve it will require advances in understanding how fatigue impacts cognitive performance, as well as improved models of the mechanisms of human cognition more generally. There are myriad challenges associated with these areas of research, but the research that has been conducted in this area has begun to lay the foundation for success.

4. Acknowledgements

The views expressed in this paper are those of the authors and do not reflect the official policy or position of the Department of Defense or the U.S. Government. The research was sponsored by the Air Force Research Laboratory's (AFRL) Warfighter Readiness Research Division and by grant 07HE01COR from the Air Force Office of Scientific Research (AFOSR).

5. References

Achermann, P. (2004). The two-process model of sleep regulation revisited. *Aviation, Space, and Environmental Medicine, 74(3)*, A37-A43.

Alion Science and Technology (2008). The SAFTE plugin for IMPRINT Pro: Incorporation of a model of fatigue. *Personal Communication*.

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y . (2004). An integrated theory of the mind. *Psychological Review 111,* (4). 1036-1060.

Atkinson, A., Coldwells, A., & Reilly, T. (1993). A comparison of circadian rhythms in work performance between physically active and inactive subjects. *Ergonomics, 36,* 273–281.

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, England: 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.

Borbely, A. A., & Acherman, P. (1999). Sleep homeostasis and models of sleep regulation. *Journal of Biological Rhythms, 14*, 557-568.

Caldwell, J. A. (2003, Fall). Wake up to the importance of sleep for air safety! *Flightline*, 30-33.

Caldwell, J. A., & Ramspott, S. (1998). Effects of task duration on sensitivity to sleep deprivation using the multi-attribute task battery. *Behavior Research Methods, Instruments, & Computers, 30(4)*, 651-660.

Dean, D. A. II, Fletcher, A., Hursh, S. R., & Klerman, E. B. (2007). Developing Mathematical Models of Neurobehavioral Performance for the "Real World." *Journal of Biological Rhythms, 22(3),* 246-258.

Dinges, D. F., Baynard, M., & Rogers, N. L. (2005). Chronic sleep deprivation. In M. H. Kryger, T. Roth, and W. C. Dement (Eds.), *Principles and Practice of Sleep Medicine (4th edition)* (pp. 67-76). Philadelphia, PA: W. B. Saunders.

Dinges, D. F., & Kribbs, N. B. (1991). Performing while sleepy: effects of experimentally-induced sleepiness. In Monk, T.H. (Ed.) *Sleep, sleepiness and performance* (pp. 97-128). Chichester: 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 Kushida, C.A. (Ed.), *Sleep Deprivation: Clinical Issues, Pharmacology and Sleep Loss Effects* (pp. 39-70). New York, NY: Marcel Dekker, Inc.

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. (2003). Modeling fatigue degraded performance in artificial agents. *Proceedings of the Human Factors and Ergonomics Society 47th Annual Meeting* (pp. 307-310).

Gillis, P. D., & Hursh, S. R. (1999). Using behavior moderators to influence CGF command entity effectiveness and performance. *Proceedings of the Eighth Conference on Computer Generated Forces and Behavioral Representation,* Orlando, Florida, 237–251.

Gluck, K. A., Gunzelmann, G., Gratch, J., Hudlicka, E., & Ritter, F. E. (2006, July). The Impact of Cognitive Moderators on Human Cognition & Performance (Co-organizer, with K. A. Gluck). In R. Sun & N. Miyake (Eds.), *Proceedings of the Twenty-Eighth Annual Meeting of the Cognitive Science Society* (p. 2658). Mahwah, NJ: Lawrence Erlbaum Associates.

Gross, J. B., Gunzelmann, G., Gluck, K. A., Van Dongen, H. P. A., & Dinges, D. F. (2006). Prediction of circadian performance during sleep deprivation. In R. Sun & N. Miyake (Eds.), *Proceedings of the Twenty-Eighth Annual Meeting of the Cognitive Science Society* (p. 297-302). Mahwah, NJ: Lawrence Erlbaum Associates.

Gunzelmann, G., Byrne, M. D., Gluck, K. A., & Moore, L. R. (prep). Using Computational Cognitive Modeling to Evaluate Alternative Theories: Implications for Understanding Fatigue. *Human Factors*.

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 and J. G. Trafton (Eds.), *Proceedings of the Twenty-Ninth Annual Meeting of the Cognitive Science Society* (pp. 329-334). Austin, TX: Cognitive Science Society.

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, and M. Bucciarelli (Eds.), *Proceedings of the Twenty-*

Seventh Annual Meeting of the Cognitive Science Society (pp. 857-862). Mahwah, NJ: Lawrence Erlbaum Associates.

Gunzelmann, G., Gross, J. B., Gluck, K. A., & Dinges, D. F. (under review). A computational account of changes in psychomotor vigilance as a consequence of fatigue. *Cognitive Science*.

Heuer, H., Spijkers, W., Kiesswetter, E., & Schmidtke, V. (1998). Effects of sleep loss, time of day, and extended mental work on implicit and explicit learning of sequences. *Journal of Experimental Psychology: Applied, 4(2)*, 139-62.

Horne, J., & Reyner, L. (1999). Vehicle accidents related to sleep: A review. *Occupational and Environmental Medicine, 56*, 289-294.

Hursh, S., et al. (2003). The DOD Sleep, Activity, Fatigue, and Task Effectiveness Model. *Proceedings of the Twelfth Conference on Behavior Representation in Modeling and Simulation* (pp. 189-196), Orlando, FL: Institute for Simulation and Training.

Jones, R., Laird, J. and Neville, K., (1998). Modeling pilot fatigue with a synthetic behavior model. *Proceedings of the Seventh Conference on Computer Generated Forces and Behavioral Representation*, Orlando, FL.

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.

Jongman, L. & Taatgen, N. A. (1999). An ACT-R model of individual differences in changes in adaptivity due to mental fatigue. In *Proceedings of the Twenty-First Annual Meeting of the Cognitive Science Society* (pp. 246-251). Mahwah, NJ: Lawrence Erlbaum Associates.

Klerman, E. B., & St. Hilaire, M. (2007). On Mathematical Modeling of Circadian Rhythms, Performance, and Alertness. *Journal of Biological Rhythms, 22(2)*, 91-102.

Laird, J. E., Coulter, K. J., Jones, R. M., Kenny, P. G., Koss, F. V., & Nielsen, P. E. (1998). Integrating intelligent computer generated forces in distributed simulation: TacAir-Soar in STOW97. *Proceedings of the 1998 Simulation Interoperability Workshop*. Orlando, FL.

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.

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.

Neville, K. (1997). Progress report: Fatigue model development and insertion project. *Unpublished manuscript*.

Neville, K., Takamoto, N., French, J., Hursh, S. R., Schiflett, S. G., & Storm, W. S. (2000). The sleepiness-induced lapsing and cognitive slowing model: Predicting fatigue effects on warfighter performance.

Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.

Pilcher, J. J., Band, D., Odle-Dusseau, H. N., & Muth, E. R. (2007). Human performance under sustained operations and acute sleep deprivation conditions: Toward a model of controlled attention. *Aviation, Space, and Environmental Medicine, 78(5-II Suppl.)*, B15-B24.

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.

Silverman, B. G. (2007). Human terrain data – What should we do with it? In S. G. Henderson, B. Biller, M.-H. Hsieh, J. Shortle, J. D. Tew, & R. R. Barton (Eds), Proceedings of the 2007 Winter Simulation Conference.

Silverman, B.G., Johns, M., Cornwell, J., & O'Brien, K. (2006). Human behavior models for agents in simulators and games: Part I: Enabling science with PMFServ. *Presence, 15*(2), 139-162.

Tassi, P., Pellerin, N., Moessinger, M., Eschenlauer, R., & Muzet, A. (2000). Variation of visual detection over the 24-hour period in humans. *Chronobiology International, 17*, 795–805.

Thorne, D. R., Genser, S. G., Sing, G. C., & Hegge, F. W. (1985). The Walter Reed performance assessment battery. *Neurobehavioral Toxicology and Teratology, 7*, 415-418.

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., 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.

Author Biographies

KEVIN GLUCK is a Senior Research Psychologist at the Air Force Research Laboratory's Warfighter Readiness Research Division, where he is the Science and Technology Advisor for the Cognitive Models and Agents Branch and the Team Lead for the Performance and Learning Models (PALM) research team.

GLENN GUNZELMANN is a Research Psychologist with the Air Force Research Laboratory. He is the technical lead for the Performance and Learning Models (PALM) Team's basic research, which is focused on developing psychologically valid computational accounts of human cognition and performance.