

Understanding Decrements in Knowledge Access Resulting from Increased Fatigue

Glenn Gunzelmann (glenn.gunzelmann@mesa.afmc.af.mil)

Kevin A. Gluck (kevin.gluck@mesa.afmc.af.mil)

Air Force Research Laboratory; 6030 South Kent Street; Mesa, AZ 85212 USA

Jeffrey Kershner (jeffrey.kershner@mesa.afmc.af.mil)

L3 Communications at Air Force Research Laboratory

6030 S. Kent St., Mesa, AZ 85212-6061 USA

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

Sleep and Performance Research Center, Washington State University Spokane

P.O. Box 1495; Spokane, WA 99210-1495, USA

David F. Dinges (dinges@mail.med.upenn.edu)

Division of Sleep and Chronobiology, University of Pennsylvania School of Medicine

1013 Blockley Hall; 423 Guardian Drive; Philadelphia, PA 19104-6021 USA

Abstract

Understanding the impact of fatigue on human cognition represents an important challenge in applying research in cognitive science to real-world situations. In this study, we explored the cognitive mechanisms responsible for performance decrements in people doing the Walter Reed Serial Addition/Subtraction Task (SAST) periodically during 88 hrs of total sleep deprivation. In our model, performance on the SAST relies heavily on declarative knowledge of mathematical facts, allowing us to extend fatigue mechanisms associated with procedural knowledge from previous research to include analogous parameters and mechanisms in declarative knowledge in the Adaptive Control of Thought-Rational (ACT-R) cognitive architecture. This research contributes to a comprehensive theory of how the human arousal system impacts cognition and performance.

Keywords: Fatigue; Memory; Arithmetic; Sleep Deprivation; Cognitive Model; ACT-R; Alertness.

Introduction

Across a variety of tasks and situations, human performance varies as a consequence of dynamic changes in general alertness. In addition to progressive declines associated with time awake, alertness varies throughout the day as a function of circadian rhythms (Borbély & Achermann, 1999; Van Dongen & Dinges, 2005a). This is important because modern society places increasing pressure on people to operate on inadequate sleep and at times of the day when the circadian system is at its low point (e.g., shift work & long-haul airline flights). Working under such conditions of reduced alertness has been identified as a contributing factor in a number of industrial and commercial disasters (e.g., Caldwell, 2003; Dinges, 1995).

Research on fatigue has been targeted at developing strategies for managing alertness levels to maximize performance and minimize the likelihood of fatigue-related errors. This research has taken a variety of forms, including the evaluation of a variety of psychopharmacological agents

to offset the negative effects of fatigue (e.g., Åkerstedt & Ficca, 1997, Bonnet et al., 2005), and the development of biomathematical models to predict general alertness as a function of circadian rhythms and sleep history (e.g., Hursh et al., 2004, Jewett & Kronauer, 1999).

Our research draws heavily on advances being made in these areas and provides a means of extending those efforts. The goal is to identify mechanisms within a cognitive architecture to represent the impact of general alertness on the cognitive system. In this research, we are using the Adaptive Control of Thought-Rational architecture (ACT-R; Anderson et al., 2004). Our approach is to use existing biomathematical models of fatigue to drive parameter changes in ACT-R, to produce decrements in the model's performance like those observed in human participants.

In previous papers (Gross, Gunzelmann, Gluck, Van Dongen, & Dinges, 2006; Gunzelmann, Gluck, Van Dongen, O'Connor, & Dinges, 2005), we presented a model that captured the deleterious effects of fatigue on a highly procedural task testing sustained attention – the psychomotor vigilance test (PVT; Dinges & Powell, 1985). Although that work allowed us to identify important mechanisms within ACT-R that appear to be impacted by fatigue, it was only the first step in a more ambitious research agenda. The goal is to develop a general account of fatigue that addresses performance changes across different components of the human cognitive system.

In this paper, we present efforts at extending our account of fatigue to tasks involving a more substantial declarative memory component. Specifically, we model changes in performance on the Walter Reed Serial Addition/Subtraction Task (SAST) (Thorne et al., 1985). This task and the experimental data we are using for validation are described next.

The Serial Addition/Subtraction Task

The SAST involves solving one-digit addition and subtraction problems. However, if the result of an addition

is greater than 10, the participant is asked to subtract 10 and respond with that result. Alternatively, if the result of a subtraction is less than 0, then the participant is instructed to respond with the result of adding 10 to that number. Thus, the SAST is a mod task, where the correct response is the result of the operation, *mod 10*.

The difficulty of the task is increased by presenting the components of the problem only briefly. For every trial, each element of the problem is presented for 200 ms, spaced by 200 ms with nothing on the screen. The first number is presented, followed by the second number, followed by the operator. The correct operation is $N1 <operator> N2$. Participants are instructed to respond as quickly and as accurately as possible. Participants were instructed to guess at the answer if for some reason the stimuli were not observed or the calculation could not be completed.

Empirical Data

Dinges and colleagues conducted a controlled laboratory study in which participants were kept awake for 88 continuous hrs (Van Dongen et al., 2001; Van Dongen & Dinges, 2005b). This sleep deprivation period followed 3 days of acclimation, where participants were given 8 hrs in bed (11:30 PM – 7:30 AM) per night for sleep. The acclimation period helped both to reduce any existing sleep debt and to stabilize the circadian rhythm. Subjects were awakened after the third night at 7:30 AM, and were kept awake for 88 hrs until 11:30 PM three days later (after missing 3 nights' sleep).

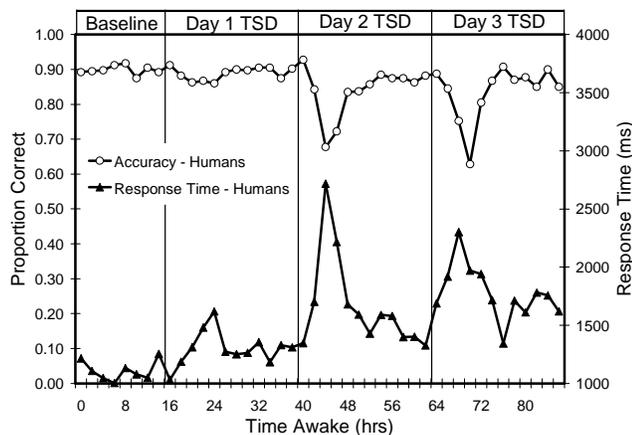


Figure 1: Human performance (accuracy and response times) on the SAST (N=8) across 88 hrs of total sleep deprivation (TSD).

While the participants were awake, they completed a set of neurobehavioral performance tasks, including the SAST, at 2 hr intervals. The full set of tasks required approximately 30 minutes to complete. In each session of this study, the SAST required participants to respond to 50 trials as quickly and accurately as possible. Participants received feedback on their responses to the first 10 trials in each of these sessions (either 'C' or 'E' was flashed for 200 ms following the response). Average response time and overall accuracy were recorded and are presented in Figure 1, averaged across 8 subjects. As is shown, accuracy decreased over the

course of the 88 hr period while response times increased. In addition, there was a daily rise and fall in these measures due to circadian rhythmicity interacting with homeostatic sleep drive.

A Theory and Model of Fatigue

Our approach to explaining the effects of fatigue on cognitive performance is to link predictions of general alertness, which can be obtained from existing biomathematical models of alertness, to specific mechanisms and parameters within the ACT-R cognitive architecture. Biomathematical models embody mechanisms that capture the influence of time awake and circadian rhythms on overall cognitive functioning (e.g., Hursh et al., 2004; Jewett & Kronauer, 1999). Although they have limitations, they perform well under conditions of total sleep deprivation (see Van Dongen, 2004). However, these models lack mechanisms for predicting *in situ* cognitive performance; predictions about performance must be scaled *post hoc* to dependent measures in particular tasks (e.g., Van Dongen, 2004).

We can enhance the predictive power of biomathematical models by linking them to a general theory of human cognition. ACT-R contains mechanisms for perception, cognition, and action, with a number of parameters influencing the speed and effectiveness of those processes (Anderson et al., 2004). We view the primary impact of fatigue as influencing the values of those parameters. In the next section, we describe the ACT-R architecture, focusing on the mechanisms in the architecture that are critical for our model of the SAST. Then, we illustrate how the negative impact of fatigue on human performance can be captured in a computational model by systematically varying the parameter values influencing those mechanisms.

ACT-R

ACT-R is a cognitive architecture that instantiates a general theory of human cognition. The theory posits a division between declarative and procedural knowledge, the existence of specialized information processing modules, and a serial bottleneck in central cognition, which in ACT-R is a serial production system (Anderson et al., 2004). Within the constraints of the architecture, accounts of human performance have been developed for tasks in a variety of domains of psychological research (see Anderson & Lebiere, 1998).

Performance on mathematics tasks is a central part of the history of the development of ACT-R, including accounts of how individuals solve simple addition and subtraction problems like those involved in the SAST (e.g., Anderson & Lebiere, 1998; Lebiere, 1999; Terao, Koedinger, Sohn, Anderson, & Carter, 2004). Thus, the SAST provides a useful context for exploring an understanding of how fatigue may impact mechanisms in the architecture.

Model Description

Models of mathematics performance in ACT-R focus on the acquisition and strengthening of declarative knowledge representing particular mathematical facts and problems.

The conceptualization of mathematics expertise in ACT-R is typically that, over a lifetime, these facts are rehearsed in a variety of situations and settings, resulting in fast and accurate access (Lebiere, 1999). The data presented in Figure 1 illustrate this. In the baseline period of the study, overall accuracy was 89.8%, and response times were 1.11 s on average. In contrast, by the third day of the 88 hrs of total sleep deprivation (TSD), accuracy was reduced to 83.7% overall, and response times had risen to 1.78 s. This is a substantial increase in both errors and response latency compared to performance during the baseline period.

As noted above, the SAST involves rapid presentation of two digits, followed by an operator. To do the task, the model encodes each of these elements of the problem as they are presented, generating a representation of the entire problem. The model then uses this information to probe memory for a chunk of declarative knowledge that encodes the solution. Because all of the participants in the study were adults with at least high school education, we did not include knowledge for more time-intensive alternative strategies, like counting, which may be used to compute, rather than retrieve, an answer (Siegler & Schrager, 1984). The model uses the answer to the math fact retrieved from memory to identify the solution to the problem. If it is greater than 10, the model retrieves the ones-digit of the solution and uses that to respond. If the answer is negative, the model probes memory a second time for a solution to the problem of the initial result *plus 10*. When the answer is determined, the model responds by making a virtual key press. The response and the response latency are recorded for each trial.

Sources of Errors The data in Figure 1 clearly illustrate that human performance is not perfect. Even during the baseline period, participants make errors on approximately 10% of the trials. This is probably due to the speeded presentation of the problem elements combined with an emphasis on responding quickly. There are multiple opportunities for errors in the model's performance, including encoding errors/omissions and retrieval errors. The first opportunities are when the elements of the problem are presented. The model may fall behind in encoding them, resulting in missing elements in the problem to be solved. In these cases, the model is forced to probe memory using an incomplete representation of the problem, which essentially means guessing at the correct answer.

In addition to failures to encode, the model can make errors as a result of retrieving inaccurate information from memory. This can occur when encoding problem elements from the screen (misidentifying what number is presented), when retrieving math facts from memory, and when retrieving the representation of the response digit from memory. This is possible because of a similarity-based partial matching mechanism in ACT-R (Anderson et al., 2004), which allows items in memory that are similar to the requested chunk to be retrieved. Due to Gaussian noise added to the calculations of the activation of declarative chunks in the retrieval process, the activation of partially matching chunks may exceed the activation of the appropriate chunk of knowledge. Following Lebiere (1999),

similarity values between numbers in our model are proportional to the ratio between them. Thus, numbers closer in value are more easily confused, as seen in empirical research (e.g., Siegler & Schrager, 1984).

Fatigue Mechanisms In addition to errors, there is variability in the response latencies to problems across the experimental protocol. In Gunzelmann et al. (2005), we presented a model of the psychomotor vigilance test (PVT), which tests sustained attention by having participants monitor a known location and respond to the onset of stimuli appearing at random intervals, for the duration of a 10-minute session. In the model for the PVT, response latencies increased under conditions of fatigue as a consequence of 'micro-lapses' in the functioning of the production system resulting from changes in numerical procedural parameters. We have evaluated using the same mechanism in the model for the SAST, and determined that, by itself, the occurrence of micro-lapses is unable to produce decrements in both accuracy and response times that mirror the effects observed in the human data.¹

Because the mechanisms we identified in earlier research are not adequate to account for performance changes in the SAST, we extended the set of mechanisms producing fatigue-related declines in the model's performance. Specifically, we incorporated parameters in ACT-R's declarative memory module, which are analogous to the parameters we have already identified as being influential in procedural knowledge. To ensure that this research is cumulative in the sense of generating a comprehensive and general account of fatigue, we vary the previously identified procedural parameters in this model as well, under the assumption that fatigue has global, task-independent effects on human information processing mechanisms.

The parameters we identified in Gross et al. (2006) included the parameter G , which is involved in selecting and executing a single production on each cognitive cycle, and the utility threshold, T_u , which impacts the likelihood of any cognitive action being performed at all on a cognitive cycle. In our account, we associated G with alertness within the procedural system, and we manipulated T_u to represent attempts by individuals to compensate for the deleterious effects of fatigue on attentional vigilance. Manipulations to both parameters were tied to predicted levels of general alertness stemming from existing biomathematical models (Gross et al., 2006).

The additional parameters we now manipulate in declarative memory to produce the model predictions presented below are the base-level activation, A_i , of declarative knowledge (particularly of numbers and math facts), combined with the retrieval threshold, T_r . The activation of knowledge plays a role for declarative knowledge similar to the role G plays in procedural knowledge. In addition, there is a direct correspondence in

¹ In fact, by the time the frequency of micro-lapses increases sufficiently to produce latencies like those seen in the human data, the model is reduced to near-chance levels of performance in terms of accuracy, because it almost always fails to accurately encode the elements of the problem.

function between T_r and T_u in ACT-R, in that both parameters control how ‘active’ or ‘useful’ information in memory must be in order to be accessed.

To represent decreased arousal within the declarative memory system, A_i is decremented. This has an immediate impact on the speed with which information is retrieved from memory because retrieval time in ACT-R is explicitly dependent on activation:

$$T_i = F * e^{-A_i}$$

In this equation, the time to retrieve chunk i from memory decreases as its activation, A_i , increases. F is a scaling parameter, which we leave at its default value of 1.

When retrievals are slowed there are two observable consequences in the model. First, failure to fully encode the problem elements becomes more likely, because retrieving the propositional representation of the visual information takes longer, leading to delays that can interfere with encoding subsequent items in the problem. Second, latencies increase directly due to longer retrieval times after the problem has been encoded.

In addition to decrementing A_i in the fits presented below, we also raise the retrieval threshold, T_r , as a consequence of fatigue. This suggests that additional effort is required to access declarative memory under conditions of fatigue, in addition to the decreased availability of that knowledge reflected in lower activation values. This role for T_r contrasts with the theoretical role of T_u in our model of the PVT described in Gross et al. (2006), where T_u was *decreased* under conditions of fatigue to represent explicit attempts at compensating for the negative consequences of reduced alertness. This is an issue to be addressed in future research.

In the next section, we evaluate the ability of these parameters to influence the model’s performance in a manner that is consistent with the declines associated with sleep deprivation in humans. To constrain this effort, we relate changes in ACT-R parameter values to predictions of general alertness in a biomathematical model of fatigue.

Evaluating the Model

Not surprisingly, manipulations to the overall activation of declarative knowledge and retrieval threshold, combined with the previously investigated parameters influencing procedural knowledge in central cognition, can capture the full range of human performance shown in Figure 1, in terms of both accuracy and response times. More important than fitting the data, however, is to understand the impact of fatigue on performance by assessing what is changing in the model’s information processing that is responsible for the changes in overall performance. This is a key issue for understanding the impact of fatigue that cannot be addressed with biomathematical models of alertness, illustrating an important contribution of using a cognitive architecture in this area.

In the case of the SAST, the activation level (A_i) appears to be the main driver of performance changes in the model. Decreasing activation reduces the availability of knowledge in declarative memory. This has complex effects on the

performance of the model. First, slower retrievals lead to increased response times by increasing the latency between when the problem is encoded and when the response is retrieved from memory. In addition, slower retrievals have the potential to impact the model’s encoding of the problem, as noted above. If encoding one element of the problem takes too long, then the model will fail to encode the next element, forcing the model to guess when responding.

Besides increasing retrieval latencies, decreased activation levels can lead to failures to retrieve knowledge from declarative memory, when chunks fail to exceed the retrieval threshold, T_r . A retrieval failure in this model leads to a need to guess about the correct response. When a retrieval failure occurs, another retrieval is attempted, but the context of the original request is lost, meaning that the chunk retrieved from memory is essentially random. However, because the information in this task is well-learned, retrieval failures rarely occur, except at the most extreme parameter values, which typically are not appropriate even under conditions of lowest general alertness as predicted by the biomathematical model. In contrast, for less well-learned information, the impact of retrieval failures could be critical, and exploring this issue is one direction of our current research.

Relating Parameters to General Alertness

One of the goals in this research is to develop a set of mechanisms to systematically map levels of general alertness to parameter values in the architecture, thereby reducing or eliminating the need to do unstructured searches through parameter spaces to fit the data. At this stage of our research, we are exploring the potential for such a systematic mapping by using biomathematical models of alertness to control how parameter values in ACT-R change as a consequence of fatigue. This limits the degrees of freedom and significantly constrains the model’s behavior.

Since we have no previous research to use as a basis for estimating the relationship between general alertness and model parameter values for declarative tasks, we used best-fitting parameter values for each session of the human data to infer a function mapping between alertness predictions from the biomathematical model and ACT-R parameter values. In this initial model evaluation, we limited our search to linear functions. Thus, using best-fitting parameters for each session, we estimated a function of the form:

$$V_x = ma_x + b$$

Where V_x is the value for the parameter at session x . We estimated the values for m and b for each of the four parameters in the model using the least squares method. To produce predicted parameter values for each session, then, the equation was solved using the predicted value of alertness for that session (a_x) to drive the prediction. Whereas we manipulated all four parameters in the model, it should be noted that the model is relatively insensitive to the values for G and T_u , because the effects are driven more by changes in the availability of knowledge, rather than the use of that knowledge in central cognition, as described above.

In this paper, we present the predictions based on this method for a single biomathematical model, namely the Circadian Neurobehavioral Performance and Alertness (CNPA) model developed by Jewett and Kronauer (1999). The CNPA model is typical of a class of biomathematical models of fatigue in that it incorporates processes to account for (a) circadian rhythms and (b) progressive declines in alertness as time awake increases (see Van Dongen, 2004, for a comparison of 7 such models). It is distinct in that it also includes components to account for sleep inertia (low alertness immediately upon awakening) and for the impact of light on the circadian component. See Jewett and Kronauer (1999) for a more detailed discussion of this model and how it is applied in understanding fatigue. Figure 2 presents the average performance of the model over 250 iterations using parameter values estimated using the alertness predictions from the CNPA model as described above. The resulting model performance roughly corresponds with the overall average pattern of the human data ($r=.62$, $RMSD=.071$ for accuracy; $r=.69$, $RMSD=278$ ms for response times).

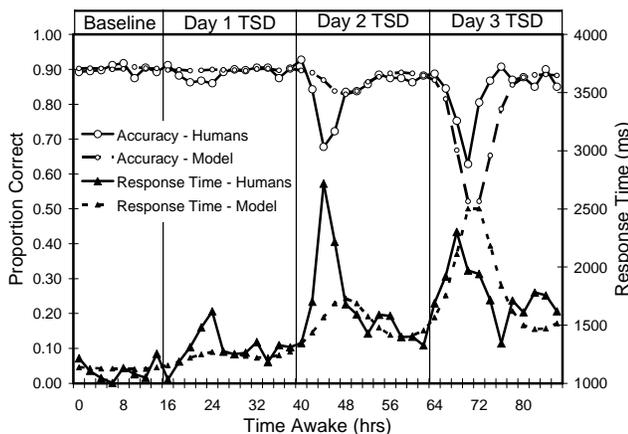


Figure 2: Model performance and human data, with model parameters estimated using a linear transformation of the biomathematical model (CNPA) predictions of general alertness.

The discrepancies between the model and the human data in Figure 2 illustrate several aspects of this modeling work. Since we are using predictions of alertness from an existing biomathematical model, we are tied to those predictions about the temporal dynamics of general alertness, including the high- and low-points in human performance over the course of the day. It appears in this case that those predictions from the biomathematical model are slightly misaligned relative to the performance dips found in the human data, particularly in the morning following the second day without sleep.

In addition to the misalignment between the predicted and observed performance low-points, the biomathematical model also predicts that the most severe decrements in performance will occur in the morning of the last day of the sleep deprivation period. Consequently, the ACT-R model does not fully capture the drop in human performance in the morning following the second night without sleep, which is

as severe as the decrements observed on the last morning of the study for the human participants. It is likely that using a non-linear function to map general alertness to model parameters could improve this prediction somewhat, and we intend to evaluate other potential functions as this research progresses. Although these discrepancies between the model and the data exist, we find the overall correspondence between the model and the data to be encouraging evidence for the utility of our approach. We sum up the important contributions of this research in the conclusion, and also identify some additional avenues for future research.

Conclusions

In this paper, we have demonstrated how parameters within a cognitive architecture can be manipulated to account for the impact of fatigue on cognitive performance. The research is an extension of previous efforts (e.g., Gross et al., 2006; Gunzelmann et al., 2005), increasing the breadth of our account to include the impact of fatigue on mechanisms in declarative knowledge. There is evidence that fatigue has global effects on cognitive functioning, and thus it is not surprising that we have identified mechanisms and parameters in multiple processing systems that are needed to account for the negative consequences of decrements in alertness across multiple tasks. As this research progresses, we anticipate that additional information processing mechanisms in ACT-R may be implicated in accounting for how fatigue changes cognitive performance.

The mechanisms implicated in performance impairment on the SAST relate to the availability of declarative knowledge, in conjunction with mechanisms associated with the selection and execution of procedural knowledge that we identified previously in our research on the PVT. The declarative parameters we manipulated are analogous to those manipulated in ACT-R's procedural system, providing a consistent account of the impact of fatigue across these two information processing systems, though we need to evaluate the contrasting roles of the threshold parameters.

As our research on how fatigue impacts human information processing progresses, we will continue to extend our account to include additional components of the cognitive system and seek to identify theoretically motivated parameters that reflect changes in cognitive functioning. A challenge of this research is to create a parsimonious, globally coherent account of the impact of fatigue that produces accurate predictions of performance across tasks. Specifically, manipulating new mechanisms in developing an account for performance in a new task can only be considered progress if the entire set of mechanisms and parameters can provide an account of performance across the full set of tasks. Such coherence is a focus of validation efforts. For instance, we plan to re-run our PVT model while including manipulations to the declarative parameters identified here to demonstrate that the account still holds.²

² It can be shown that manipulations of the declarative parameters will not impact the PVT model, since no information is retrieved from declarative knowledge while the model performs the task.

Our further research and validation efforts will ensure that the mechanisms we develop provide a robust and general account of how increased levels of fatigue degrade cognitive functioning. This will establish a foundation for making performance predictions in complex, dynamic tasks where empirical data regarding human performance may be unavailable. As noted above, the ultimate value of understanding fatigue using a cognitive architecture is in making predictions in applied settings, to avoid situations where errors stemming from fatigue have consequences that may be dramatic and catastrophic.

Acknowledgments

The views expressed in this paper are those of the authors and do not reflect the official policy or position of the Department of Defense or the U.S. Government. This research was sponsored by the Air Force Research Laboratory's (AFRL) Warfighter Readiness Research Division and by grant 04HE02COR from the Air Force Office of Scientific Research (AFOSR). Human data collection was sponsored partly by AFOSR grants F49620-95-1-0388 and F49620-00-1-0266, and by NIH grant RR00040. HVD was supported by AFOSR grant FA9550-05-1-0086, and DFD was supported by AFOSR grant FA9550-05-1-0293. We appreciate the thoughtful comments of three anonymous reviewers.

References

- Åkerstedt, T., & Ficca, G. (1997). Alertness-enhancing drugs as a countermeasure to fatigue in irregular work hours. *Chronobiology International*, *14*(2), 145-158.
- Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review*, *111*, 1036-1060.
- Anderson, J. R., & Lebiere, C. L. (1998). *The atomic components of thought*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- 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.
- Borbély, A. A., & Achermann, P. (1999). Sleep homeostasis and models of sleep regulation. *Journal of Biological Rhythms*, *14*, 557-568.
- Caldwell, J. (2003, Fall). Wake up to the importance of sleep for air safety! *Flightline*, 30-33.
- Dinges, D. F. (1995). An overview of sleepiness and accidents. *Journal of Sleep Research*, *4*(2), 4-11.
- 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.
- Gross, J. B., Gunzelmann, G., Gluck, K. A., Van Dongen, H. P. A., & Dinges, D. F. (2006). Computational modeling of the combined effects of circadian rhythm and sleep deprivation. In R. Sun & N. Miyake (Eds.), *Proceedings of the 28th Annual Meeting of the Cognitive Science Society* (p. 297-302). Mahwah, NJ: Lawrence Erlbaum Associates.
- Gunzelmann, G., Gluck, K. A., Price, S. C., Van Dongen, H. P. A., & Dinges, D. F. (2007). Incorporating the effects of sleep deprivation into a cognitive architecture. In W. D. Gray (Ed.), *Integrated Models of Cognitive Systems*. New York: Oxford University Press.
- 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 27th Annual Meeting of the Cognitive Science Society* (pp. 857-862). Mahwah, NJ: Lawrence Erlbaum Associates.
- Hursh, S. R., Redmond, D. P., Johnson, M. L., Thorne, D. R., Belenky, G., Balkin, T. J., et al. (2004). Fatigue models for applied research in warfighting. *Aviation, Space, and Environmental Medicine*, *75*(3), A44-60.
- Jewett, M. E., & Kronauer, R. E. (1999). Interactive mathematical models of subjective alertness and alertness in humans. *Journal of Biological Rhythms*, *14*, 588-597.
- Lebiere, C. (1999). The dynamics of cognition: An ACT-R model of cognitive arithmetic. *Kognitionswissenschaft*, *8*(1), 5-19.
- Siegler, R. S., & Shrager, J. (1984). Strategy choices in addition and subtraction: How do children know what to do? In C. Sophian (Ed.), *Origins of cognitive skills* (pp. 229-293). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Terao, A., Koedinger K., Sohn, M-H., Anderson, J. R., & Carter, C. S. (2004) An fMRI study of the interplay of visual-spatial systems in mathematical reasoning. In K. Forbus, D. Gentner, and T. Regier (Eds.), *Proceedings of the 26th Annual Conference of the Cognitive Science Society* (pp. 1327-1332). Mahwah, NJ: Lawrence Erlbaum Associates.
- Thorne, D. R., Genser, S. G., Sing, G. C., & Hegge, F. W. (1985). The Walter Reed performance assessment battery. *Neurobehaviora. 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., & Dinges, D. F. (2005a). Circadian rhythms in sleepiness, alertness, and performance. In M. H. Kryger, T. Roth, and W. C. Dement (Eds.), *Principles and Practice of Sleep Medicine*, 4th ed. (pp. 435-443). Philadelphia, PA: Elsevier Saunders.
- Van Dongen, H. P. A., & Dinges, D. F. (2005b). Sleep, circadian rhythms, and psychomotor vigilance. *Clinical Sports Medicine*, *24*, 237-249.
- 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.