Sleep loss and driver performance: Quantitative predictions with zero free parameters

Action editor: David Peebles

Glenn Gunzelmann [a,*], L. Richard Moore Jr. [b], Dario D. Salvucci [c], Kevin A. Gluck [a]

[a] Air Force Research Laboratory, United States
[b] Lockheed Martin at Air Force Research Laboratory, 6030 South Kent St., Mesa, AZ 85212, United States
[c] Department of Computer Science, Drexel University, 3141 Chestnut St., Philadelphia, PA 19104, United States

Available online 22 August 2010

Abstract

Fatigue has been implicated in an alarming number of motor vehicle accidents, costing billions of dollars and thousands of lives. Unfortunately, the ability to predict performance impairments in complex task domains like driving is limited by a gap in our understanding of the explanatory mechanisms. In this paper, we describe an attempt to generate a priori predictions of degradations in driver performance due to sleep deprivation. We accomplish this by integrating an existing account of the effects of sleep loss and circadian rhythms on sustained attention performance with a validated model of driver behavior. The predicted results account for published qualitative trends for driving across multiple days of restricted sleep and total sleep deprivation. The quantitative results show that the model's performance is worse at baseline and degrades less severely than human driving, and expose some critical areas for future research. Overall, the results illustrate the potential value of model reuse and integration for improving our understanding of important psychological phenomena and for making useful predictions of performance in applied, naturalistic task contexts.
Published by Elsevier B.V.

Keywords: Driver behavior; Fatigue; Computational model; Sustained attention; Sleep deprivation

1. Introduction

Accidents on roadways in the United States account for a distressingly high number of fatalities and substantial cost on an annual basis (Horne & Reyner, 1999; Klauer, Dingus, Neale, Sudweeks, & Ramsey, 2006; NTSB, 1995; Pack et al., 1995). According to a National Highway Transportation Safety Board report (NTSB, 1995), many of these accidents – 100,000 per year – may be wholly or partially attributable to the effects of drowsiness or fatigue on driver attention, judgment, and/or performance.

The alarmingly high cost of fatigued driving has been one motivation for studies to better understand changes in cognitive performance stemming from extended time awake (sleep deprivation), insufficient sleep (sleep restriction), and being awake at times of the day when the body is predisposed to sleep (circadian desynchrony: Dijk, Duffy, & Czeisler, 1992; Van Dongen & Dinges, 2005a, 2005b). This research has succeeded in identifying patterns of decline in cognitive performance related to time awake and circadian rhythms. However, there remain significant limitations in the capacity to make valid quantitative predictions about performance in novel task contexts based on historical information about sleep and wakefulness (Dinges, 2004; Van Dongen, 2004). Our computational modeling research has been targeted at addressing some of the existing limitations in predictive validity.

In the research presented here, we evaluate the capacity to make predictions about degradations in driver performance associated with an extended period of partial or total sleep deprivation. We discuss the implications of

* Corresponding author. Address: Air Force Research Laboratory – RHAC, 6030 South Kent Street, Mesa, AZ 85212-6061, United States. Tel.: +1 480 988 6561x674; fax: +1 480 988 2230.
  E-mail address: glenn.gunzelmann@us.af.mil (G. Gunzelmann).

our research in the context of potential applications of a performance prediction capability in the domain of driving. In the next section, we present some theoretical and methodological background, as well as a description of the model and mechanisms being utilized in this research. We then compare the model's predictions with data from the empirical literature, demonstrating that the predicted results in the model are aligned in crucial ways with those published results.

2. Requirements for a predictive theory

To make task-specific predictions about changes in human performance resulting from sleep loss, there are at least four necessary elements. These include: (1) a theory of the components of cognitive processing, (2) mechanisms to link fluctuations in alertness to processing changes in the various cognitive components, (3) an account of how those components are deployed in the performance of the task, and (4) an understanding of how overall alertness is influenced by factors like time awake and circadian rhythms. Here, we focus on the integration of these components. To set the stage, we describe how each of these requirements is addressed in this research, because they are all essential elements.

2.1. Components of cognitive processing

To predict how cognitive performance will change under conditions of fatigue, it is necessary to have a theory accounting for the cognitive processes are that are being affected. Abstract verbal theories have been presented in the literature, making reference to constructs like cognitive slowing, increased cognitive noise, or cognitive lapsing. Although such accounts provide useful descriptions of commonly observed effects, they do not formally specify the underlying mechanisms that give rise to such changes. Without a formal account of the underlying processes, it is impossible to generate quantitative performance predictions. To address this challenge, we utilize the Adaptive Control of Thought – Rational, or ACT-R, cognitive architecture (Anderson, 2007; Anderson et al., 2004). Cognitive architectures are intended as general theories of the foundational mechanisms of human cognition. As an instance of such a theory, ACT-R provides validated mechanisms representing many aspects of cognitive functioning. Additionally, ACT-R is implemented in software, which enables the development of models that simulate cognitive processing and behavior in particular task contexts and generate quantitative performance data that can be compared directly to human data. These features are critical to our research goals, and the results presented below illustrate some of the potential afforded by leveraging such existing theory.

ACT-R comprises a set of distinct processing modules (e.g., vision, declarative knowledge, motor), which are integrated through a production system that represents central cognition. The mechanisms of central cognition are of particular relevance in the research described here. Central cognition in ACT-R operates through a series of conflict resolution cycles to produce cognitive processing and behavior. During each cycle the subset of actions, or productions, whose conditions match the current system state is identified. In this case, the "system state" is represented by the contents of a set of buffers that provide limited-bandwidth communication between central cognition and the various information processing modules (e.g., a retrieval buffer for declarative memory and a visual object buffer for vision).

Once the applicable productions are identified, a "utility" value is computed for each. The utility computation is based upon reinforcement learning, and reflects the learned usefulness of the production for achieving the current goal. The production with the highest utility is selected and its actions are executed, unless its utility does not exceed a threshold for action (referred to as the "utility threshold" in ACT-R). The execution of a production typically leads to changes in the contents of one or more of ACT-R's buffers, creating a new system state. This state then serves as the starting point for the next cognitive cycle. The default duration for the completion of a single cognitive cycle in ACT-R is 50 ms. The utility value and the threshold for action play a critical role in our account of changes in cognitive processing stemming from fatigue, which is described next.

2.2. Impact of fatigue on cognitive components

The cognitive cycle in central cognition is foundational to all ACT-R models. Thus, changes to the functioning of this component of ACT-R stemming from variations in alertness would have general impacts on performance. Our explanatory mechanisms are based on the theoretical perspective that fluctuations in overall alertness or arousal can be associated with changes in utility values involved in selecting and executing actions in ACT-R's central production system. This is instantiated as proportional decreases in utility values as alertness declines. The primary impact of this mechanism is to decrease the likelihood that production utility values will exceed the threshold for action, leading to cognitive cycles where no actions are taken. We refer to these brief gaps in cognitive processing as *microlapses* (see Gunzelmann, Gross, Gluck, & Dinges, 2009).

While sleep loss results in degradations to performance, there is some evidence in the literature that people can stave off the negative effects of fatigue by increasing the effort they are putting into a task (e.g., Chua, Venkatraman, Dinges, & Chee, 2006; Portas et al., 1998). To account for the potential impacts of increased effort, a second parameter is manipulated – the threshold for action – which sets the minimum utility value required for a production to fire. Decreasing this threshold simulates greater *effort* by making it more likely that a production will successfully fire. However, this manipulation also increases

the probability that, under fatigue, a suboptimal action (a production with a low utility) will be executed instead (Gunzelmann et al., 2009). This is because proportional decreases in utility values lead to greater reductions for high-utility (i.e., better) alternatives. In contrast, Gaussian noise, which is added to utility values before selection, is not scaled. The resulting dynamics increase the likelihood that lower-utility options will be selected on a particular cycle under fatigued conditions.

The mechanisms just described instantiate a computational account of cognitive changes associated with sleep loss. We have validated the mechanisms using empirical data from a simple sustained attention task performed by participants every 2 h during 88 h of total sleep deprivation. The declines in performance of our model on this task, which result from the mechanisms described here, are closely aligned with the empirical results (Gunzelmann et al., 2009).

Note that to generate quantitative performance data, it was necessary to add task knowledge to ACT-R to enable it to perform the task. For the sustained attention task this was quite straightforward – the model needed only to wait for the stimulus, attend to it when it appeared, and then press a response button. Each of these actions can be represented by a production in ACT-R. Executing the appropriate action at the appropriate time is the major processing component of this model, aligning it nicely with central cognition in ACT-R. Understanding how cognitive resources are utilized in more complex tasks – like driving – can be much more challenging. Yet it is exactly these kinds of tasks that we are targeting as application contexts for this research. Fortunately, other research has addressed the challenge of modeling cognitive processes in driving, providing a well-validated model of human behavior in this context.

2.3. Cognitive processing in driving

There is an art to developing computational process models for specific tasks, which has led to criticism regarding the increased degrees of freedom that result when the implementation of the knowledge for performing the task is under the control of the modeler (e.g., Roberts & Pashler, 2000). To avoid this possible concern in the context of making predictions about the consequences of fatigue, we have opted to use a model that has already been developed, validated, and described in the literature to account for the cognitive processing required in driving (Salvucci, 2006). This model was developed using ACT-R, and was focused on modeling the cognitive requirements for lane keeping, lane changing, and passing other vehicles. The driver model is based on a control law of steering behavior (Salvucci & Gray, 2004) that visually encodes two salient points on the roadway: a near point in the lane center immediately in front of the vehicle; and a far point such as the vanishing point on a straight road, the tangent point on a curved road, or the lead vehicle when present. The control law describes how steering can be realized by keeping the far point stable while keeping the near point both stable and centered in the current lane.

The driver model that uses this control law relies on the production system that represents central cognition. This establishes an important link between the fatigue mechanisms described above and the driver model. The driver model uses successive iterations of four ACT-R productions to represent the control law of steering behavior. Specifically, these four rules compose a *control update cycle* during which the model: (1) encodes the near point, (2) encodes the far point, (3) updates steering and acceleration according to the control law, and (4) checks the vehicle's current stability as measured by the lateral velocity and position of the near and far points. If the vehicle is not yet stable, the model immediately initiates another control update; otherwise, the model waits approximately 500 ms to initiate the next control update.

The driver model has been shown to account well for driver behavior with respect to curve negotiation and lane changing (Salvucci, 2006). The most critical aspect of the model for our purposes is the execution time for a control update cycle: One of these cycles requires approximately 250 ms, including 50 ms for each of four production-rule firings plus some additional time for attention shifts and visual encoding. The update cycle time can increase, however, when attention is divided between driving and a secondary task, thus resulting in degradations in driver performance. For example, recent work has shown how dialing a phone (Salvucci, 2001; Salvucci & Taatgen, 2008) and rehearsing a memorized list of numbers (Salvucci & Beltowska, 2008) affects the driver model's performance; in both cases, concurrent execution of the secondary task interferes with processing of the driving task, thereby increasing the update cycle time and degrading performance (measured by, e.g., lateral deviation from lane center or brake response time to an external event). As we will describe, microlapses stemming from degraded alertness in ACT-R also can prolong or delay the update cycle, leading to similar degradations in driver performance.

2.4. Dynamics of alertness

With an account of how fatigue impacts particular components of cognition, and a model that accounts for how cognitive resources are utilized in a specific task context, it is possible to account for change in human performance as alertness declines. However, one critical aspect is still missing: How can we estimate the changes in parameter values over time to fit human empirical data as performance fluctuates with variations in alertness? For this aspect of our research, we draw, in part, upon extensive research on sleep, where understanding the dynamic fluctuations in alertness as a function of such factors as time awake, circadian rhythms, and even light exposure has been a central focus. Research in this area has led to the development of formal, mathematical accounts, which

quantify the influence of these factors on overall cognitive performance (e.g., Jewett & Kronauer, 1999; McCauley et al., 2009; Neri, 2004).

There are a variety of mathematical models that quantify the dynamics of alertness related to sleep loss and circadian rhythms, and a detailed description is beyond the scope of this paper. Neri (2004) contains an extensive review and evaluation of a number of particular examples. They are all based on a "two-process" model of alertness, emphasizing the roles of time awake and circadian rhythms on cognitive performance. In general, these models incorporate monotonic declines in alertness with time awake, combined with sinusoidal fluctuations in alertness associated with circadian rhythms. They differ in the details regarding the particular functions representing these processes (e.g., linear vs. sigmoidal declines in alertness with time awake), but they all tend to make qualitatively similar predictions about performance levels across extended periods of total sleep deprivation (Mallis, Mejdal, Nguyen, & Dinges, 2004; Van Dongen, 2004).

Mathematical models have been used to make performance predictions by scaling the generic output to existing empirical data, but are fundamentally limited in their capacity to make *a priori* performance predictions in novel task contexts. The reason is that they do not address the other required components of a predictive theory. Specifically, they do not represent the components of cognition nor how they are utilized in particular task contexts. Thus, they cannot anticipate the task-specific decrements that will be observed, although they tend to characterize effectively the qualitative trends associated with total sleep deprivation. Our research is focused on understanding the impact of fatigue on cognitive processing, while leveraging existing cognitive and mathematical theories and models to create an integrated account. Evaluating our integration approach is the focus of the remainder of this paper.

3. Integration

Our proposed mechanisms for fatigue instantiate a theory of changes in central cognitive processing resulting from fluctuations in alertness attributable to sleep loss and circadian rhythms (see Gunzelmann et al., 2009). Importantly, the foundation of the ACT-R driver model is procedural knowledge in central cognition that allows it to successfully keep the vehicle in its lane. As a result, an opportunity exists to bring together an existing model of driver behavior with an existing account of fatigue to explore the implications of fatigue on driving behavior. This opportunity represents an important step in the evolution of computational architectural accounts of cognitive phenomena, and illustrates the potential utility of unified theories that integrate theoretical insights from various domains of psychological research.

The integration of the driver model and fatigue mechanisms was a straightforward process because both were already implemented within ACT-R. The implementation

of the driver model was altered to run on a high performance computing (HPC) cluster but was not changed with respect to its behavioral performance. The driver model is similar to the sustained attention model in that neither makes extensive use of declarative memory, simplifying the account by minimizing the need to consider potential influences of fatigue on other components of cognitive functioning, like declarative knowledge access (e.g., Gunzelmann, Gluck, Kershner, Van Dongen, & Dinges, 2007). The fatigue mechanisms were taken directly from Gunzelmann et al. (2009) and applied to the driver model.[1] Thus, our procedural fatigue mechanisms alone provide the moderating effects in the driving model.

The actual effects of the fatigue mechanisms center on the production selection and execution phases of the production cycle in ACT-R. Proportional scaling of utility values for the production selection phase in the driver model creates situations in which the matching production with the highest utility fails to exceed the utility threshold, even with a reduced threshold reflecting *effort* or *compensation*. In those instances, no production is executed, which produces a microlapse as described above in the context of sustained attention. Parameter changes associated with fluctuations in alertness influence the frequency of microlapses, and microlapses lead to the performance changes exhibited by "tired" models.

We have demonstrated that this single computational consequence can explain phenomena in the sleep research community that have been associated with the combined influence of both *cognitive slowing* and *cognitive lapses* (e.g., Dinges & Kribbs, 1991). Cognitive slowing is produced when a relatively small number of microlapses interrupt the timely execution of task actions. As microlapses become more likely, they can also produce cognitive lapses lasting on the order of seconds, when many microlapses occur in a sequence.

The latter phenomenon – long sequences of microlapses producing more dramatic breakdowns in cognitive processing – is made more likely by a final mechanism in our account. Specifically, we further attenuate utility values as a consequence of a microlapse. From an implementation perspective, the microlapse condition is detected when there are no ACT-R events remaining to execute. When this happens, a follow-up conflict resolution is scheduled, while utility values are reduced progressively by 1.5% on each occasion. The noise component of the utility values allows the subsequent conflict resolution to potentially match a production and continue model execution. When

---

[1] A detail in this process is that Gunzelmann et al. (2009) discuss fatigue mechanisms in the context of procedural selection mechanism based upon expected utility, while the current model uses a reinforcement learning algorithm. The parameter values used in this model were established through a translation of the parameters from Gunzelmann et al. (2009) to the same model based in the reinforcement learning equation. Thus, while the parameter impacting utility is different (in name and value), it has been validated to produce the same results in the sustained-attention task as the published model.

the model successfully completes a control update cycle and updates its steering of the car, utility values are returned to the value prescribed by the current level of alertness (see below). However, this may not happen, and since each successive decline in alertness further reduces the possibility of utilities rising above the threshold, a model can quickly spiral into a state analogous to sleep. In the model described in Gunzelmann et al. (2009), this mechanism was critical in capturing the most substantial breakdowns in cognitive processing (i.e., sleep attacks).

In the next section, we evaluate the impact of our fatigue mechanisms on the driver model. Recall that the driver model realizes the continuous control law through four key productions. It is in this control update cycle that the fatigue mechanisms are most influential, since microlapses increase the overall update cycle time. As will be shown, even brief delays in cognitive activity can amount to significant and occasionally devastating impacts on driver performance.

4. Model evaluation

To evaluate the model, its behavior was assessed using the driving scenario described in Salvucci and Taatgen (2008). In the task, the driver steered down a single-lane highway, keeping the vehicle as centered as possible in the roadway. The vehicle moved at a constant speed that was not controlled by the driver, thus focusing the task particularly on lateral control. One key measure of performance in the task is *lateral deviation*: the root-mean-squared error between the lane center and the vehicle's lateral position within the lane. The baseline driver model navigating this environment exhibited an average lateral deviation of approximately 15 cm across a 10-min driving scenario (see Salvucci & Taatgen, 2008).

To produce predictions of driver behavior and performance, we used parameter values for the fatigue mechanisms that were estimated based upon our research on sustained attention (e.g., see Gunzelmann et al., 2009). Specifically, the model for that research was able to account for human sustained attention performance at 2 h intervals across 88 h of total sleep deprivation. This was accomplished by fitting a linear regression to map biomathematical model predictions of alertness to utility and threshold parameters in ACT-R. Using the same linear function, we generated predicted ACT-R parameter values for an experimental protocol described in Peters, Kloeppel, and Alicandri (1999). In that experiment, participants completed a 40-min driving scenario in a driving simulator once a day between 2:00 and 4:00 PM on four consecutive days. On the first day participants were well-rested. That night, however, they were restricted to 4 h of sleep, and were then completely deprived of sleep for the next two nights.

An additional detail that is important in our simulation results is that we modified the default cognitive cycle time in ACT-R to be 40 ms. This was done to completely parallel the parameter values we used in the sustained attention model, since there are interdependencies among our fatigue mechanisms and cognitive cycle time. This adjustment was made in transitioning the mechanism the reinforcement learning algorithm. While adjusting this parameter is generally discouraged, it is potentially relevant that Stewart, Choo, and Eliasmith (2010) found evidence that cycle times in this range may be appropriate for "simple" cognitive actions.

We generated estimated levels of alertness for each driving scenario using the Circadian Neurobehavioral Performance and Alertness (CNPA) model developed by Jewett and Kronauer (1999). This model is one of the biomathematical models of alertness mentioned above, and one of those that was used in Gunzelmann et al. (2009). The predicted alertness values were used to generate predicted parameter values for our ACT-R model, using the linear function estimated for the sustained attention task. Table 1 shows the results of this estimation process. CNPA produces numerical estimates of "cognitive throughput," which ranges from 0 to 1. These estimates were used to generate values for a production utility scaling parameter and the threshold for action. Those values, then, were used in the driver model to produce estimates of driving performance with 0 free parameters available for adjusting the model's behavior to fit the human data.

The model was run through 400, 40-min driving sessions using each of those parameter sets, leading to reliable measures of central tendency in the performance measures as well as evidence regarding the variability in fatigue effects. Importantly, on some of these runs, the model's performance degraded into an unrecoverable state where it swerved dramatically back and forth across the lane. This situation became more common for the later days of the study, illustrating the important role of the fatigue mechanisms in producing this situation. In cases where a lane violation duration was greater than 10 min, the model run was removed from the analyses presented below. The proportion of model runs removed were 0.075, 0.10, 0.19, and 0.25 for the Baseline day and 3 days of partial or total sleep loss respectively.

The number of runs removed clearly indicates that there are important limitations with regard to both the driver model and the fatigue mechanisms. We discuss this further in Section 5. They were removed for the analyses presented below because they represent degenerate model behavior

Table 1
CNPA predicted alertness values and the resulting ACT-R predicted parameter values for each day of total sleep deprivation (TSD). ACT-R parameters include utility scaling, which attenuates production utilities, and the utility threshold, which is a lower limit required for a production to fire.

Day	CNPA	Utility scaling	Utility threshold
Baseline (0 Days TSD)	0.95	1.00	2.01
Day 1 TSD	0.89	0.99	2.00
Day 2 TSD	0.51	0.90	1.91
Day 3 TSD	0.32	0.85	1.86

that significantly impacts the performance predictions. Interestingly, there were cases in Peters et al. (1999) where participants had accidents caused by straying from their lane. Unfortunately, the description of the results does not provide a specific definition of how crashes could occur, not does it indicate how the simulation recovered from such occurrences. So we are unable to recreate those details in our simulation. We do, however, perform a qualitative evaluation below.

To assess model performance, the lateral deviation of the model was recorded for each second during each model run. Fig. 1 shows the distribution of these deviation values for the model for each day of the study. Perhaps surprisingly, the distributions are not radically different. Note, however, that the proportion of time that the model spends near the middle of the lane (the left side of the distribution) decreases with more sleep loss. The overall trend is toward an increasingly skewed distribution: although performance remains normal most of the time, it diverges more often and to a greater extent as sleep deprivation increases. This pattern of results matches the data from the sustained attention task that we have used in developing the mechanisms applied to the driver model in this paper (see Gunzelmann et al., 2009).

While the distributions in the larger deviations (21–83 cm) are not very different, clear differences emerge in the categories representing the largest deviations. A lane violation ("LV" in the figure) represents a situation where some portion of the vehicle had crossed the lane line (i.e., the vehicle was partially outside its lane). Lane violations comprise deviations between 83 cm and 366 cm. The proportions of lane violations more than double for Days 2 and 3 of sleep deprivation as compared to the baseline day or a single night with inadequate sleep. Even more severe are lane shifts ("LS" in the figure), which represent cases where the vehicle has moved an entire lane's width

Fig. 2. Distribution of model update times as a function of number of days of total sleep deprivation (TSD).

laterally (a lane deviation greater than 366 cm) – a substantial and potentially catastrophic degree of driver performance error. Whereas the Baseline condition predicts about 2% lane shifts,[2] by Day 3 the model predicts a substantial increase to around 6%. This means that close to 6% of the time, the model is driving completely out of its intended lane (possibly off the road or possibly into oncoming traffic).

To better understand the nature of this performance in terms of the underlying processing, Fig. 2 shows the distribution of update times for the driver model in each condition — that is, the amount of time needed for the model to complete its four-production control update cycle. As was the case for lateral deviation, the distributions shift with increasing sleep deprivation such that update times reflecting cycles that are not interrupted (less than $\sim$250 ms) become less frequent and longer update times become more prevalent. The increase in update times arises because production rules are more likely to fall below threshold under the influence of fatigue mechanisms, thus missing an opportunity to fire during a particular cognitive cycle. Once again, the shift is relatively small, but the proportion of update cycles falling in the tail of the distribution increases noticeably. This shift illustrates greater inattention of the model to steering control, which is what leads to the performance changes shown in Fig. 1.

4.1. Comparison to human performance

There are interesting changes in the performance of the driver model with more severe sleep loss. To evaluate the model predictions in the context of actual human driver performance, we used data presented in Peters et al.
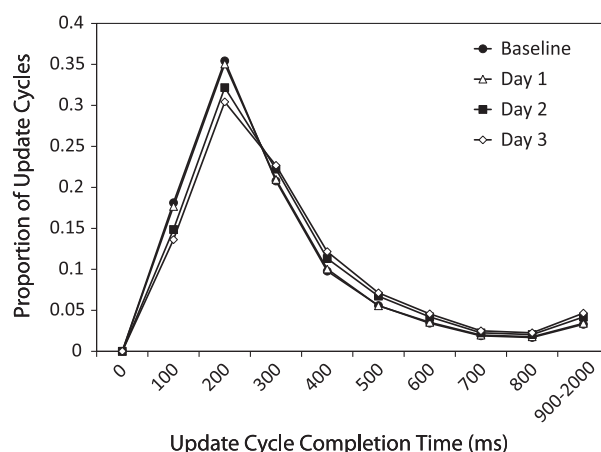

Fig. 1. Proportion of 1-s samples of lateral deviation falling into each of the specified bins. The last two categories represent instances where (1) the vehicle is partway out of the proper lane (a lane violation, "LV"), and (2) the vehicle's deviation is more than a full lane width off (a lane switch, "LS"). Separate lines represent 0, 24, 48, and 72 h of total sleep deprivation (TSD).

---

[2] Using default ACT-R parameters, the model does not produce any lane shifts under normal driving circumstances (Salvucci, 2006). This highlights the influence of the fatigue mechanisms even for baseline conditions, and exposes an aspect of the integration where additional research is required.
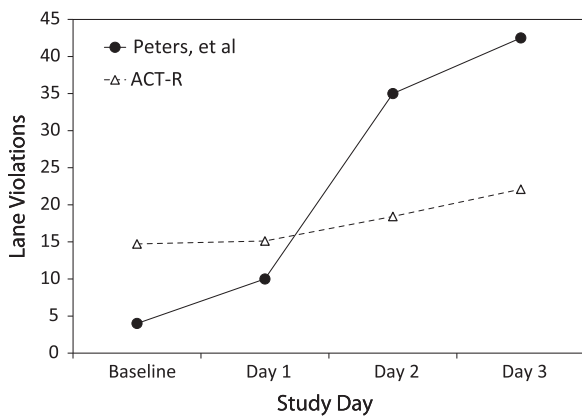
Fig. 3. Lane violations from Peters et al. (1999) compared to the lane violations in the model.

Fig. 4. Accidents from Peters et al. (1999) compared to the proportion of model runs that resulted in an unrecoverable state.

(1999) from the study briefly described above. The paper presented average frequency counts of lane violations observed in the participants in each of the four driving sessions in their study. Fig. 3 compares the reported lane violations from the Peters study to the lane violations produced by the model. The model captures the overall trend well ($r = .96$), although it does not degrade as severely as the human drivers. In addition, its performance is somewhat worse at baseline than human performance, owing to the influence of the fatigue mechanisms on a model that was fit to human data without those mechanisms in place.

Another measure reported by Peters et al. (1999) was accidents. In the paper, these are described simply as "off-road crashes" and presented as "crash rates" (p. 3). The measure appears to present a measure of the mean number of crashes per participant in a 40 min driving session, though it is not explicitly stated in the paper. This crash rate increases from 0 on the baseline day to nearly 8 on the last day of the study. While the lack of details on the crash measure make it impossible to compare the model directly, we can compare qualitatively the changes in crash likelihood to the proportion of model runs that were removed from our analysis because of it ending up in an unrecoverable state. This comparison is shown in Fig. 4. The correlation between these measures is .94. While this comparison is tenuous, it appears to reinforce the results presented in Fig. 3.

5. Discussion

Importantly, the model's behavior is produced with zero free parameters, providing an encouraging approximation of human performance in this complex task. Moreover, there are several known factors that may be contributing to the observed differences between model performance and those reported by Peters et al. First, there are a variety of potential differences in the experimental context such as lane width, road curvature, steering wheel gain, and lane violation threshold, all of which can play a significant role
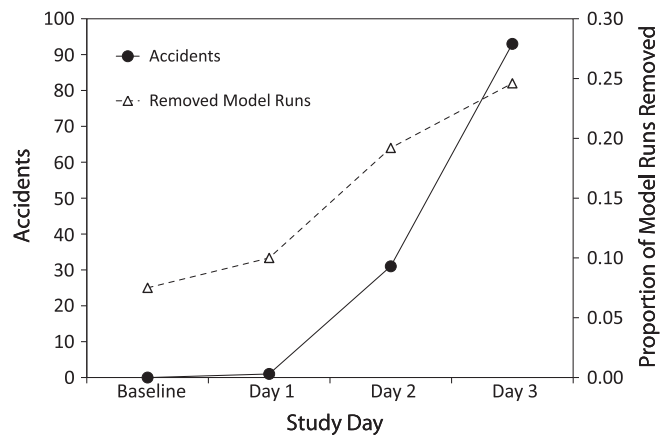
in the reported metrics. In addition, the road varied to a greater extent in Peters et al., including multiple speed limits (35 and 55 MPH), and different numbers of lanes (2 or 4) in different parts of the course. Participants in Peters et al. also controlled their own speed, and the duration of lane violations were not reported in that study. These task-related factors have unknown influences on driving performance, particularly under conditions of restricted or deprived sleep.

In addition to task-related issues, there is also a potential contribution of individual differences to the performance differences observed. Inter-individual differences in driving ability, as well as systematic differences in the impact of fatigue (e.g., Van Dongen, Baynard, Maislin, & Dinges, 2004), could play a substantial role in this case. Finally, the model's alertness levels do not change over the course of the 40-min driving sessions, either upon recognition of a lane violation, or as time on task increases. These dynamics are certainly at play in human driving, and expose important areas for future research, some of which we have begun to explore (Gunzelmann, Moore, Gluck, Van Dongen, & Dinges, 2010).

Future research is needed to validate our approach with participants that perform both the sustained attention task and the driving scenario, under conditions where all of the relevant task-related characteristics can be made equivalent for the humans and the model. Even with the discrepancies in the data shown in Figs. 3 and 4, however, we feel that the model performs quite reasonably under the constraints of a zero free parameter prediction.

6. Conclusions and future directions

The model described in this paper exhibits declines in performance when mechanisms are implemented to represent the deleterious effects of sleep loss on central cognitive functioning. The foundation is a validated model of skilled driver behavior (Salvucci, 2006), implemented in a validated theory of the human cognitive architecture

(Anderson, 2007). The model is augmented with a set of mechanisms that account for fatigue-related changes in central cognitive processing (Gunzelmann et al., 2009). The dynamics of those changes, in turn, are constrained by a biomathematical model capturing fluctuations in alertness associated with time awake and circadian rhythms (Jewett & Kronauer, 1999).

The primary contribution of this research is the demonstration that it is possible to make *a priori* predictions regarding the effects of extended wakefulness on performance in complex, dynamic tasks. The qualitative changes in the model's performance are similar to the performance changes observed in human participants attempting to drive after extended periods of partial or total sleep deprivation. The results go beyond intuitive notions regarding degradations in cognitive processing and performance as time awake increases. The scientific methodology described here produces quantitative estimates about the actual impact of those changes on performance in the driving task.

The discrepancies between the model and human performance are informative in that they expose some critical limitations of the model. Some of these have already been mentioned. For instance, our fatigue mechanisms do not represent finer-grained dynamics of alertness that are at play within a 40-min driving session. It is well established that extended time on task leads to degradations in performance (e.g., Davies & Parasuraman, 1982; Van der Hulst, Meijman, & Rothengatter, 2001). There is also evidence for systematic oscillations in alertness over relatively short periods of a couple of minutes (Arruda, Zhang, Amoss, Coburn, & Aue, 2009). In addition, when people note that they have drifted well off the road, it is likely that they experience a transient spike in alertness, which would allow them to reestablish control. Our model currently does not incorporate these dynamics, which surely impacts its performance relative to human participants.

It is likely that modifications to the steering control cycle would improve the model as well. The mechanism lacks the anticipation that is necessary to recover from very large lane deviations, which contributes to situations where the model swerves wildly back and forth across the lanes. Lastly, there is a real question regarding how to set parameter values for baseline performance. The values we used – taken directly from a model of sustained attention performance – lead to a level of performance that is misaligned on a quantitative level with the data from the human participants. The interaction of task characteristics, motivation, and alertness is a challenging issue that remains to be addressed in models of fatigue.

The current research effort represents a critical step in the process of using computational cognitive modeling to make predictions about human cognition and behavior in naturalistic task contexts. The modular design of ACT-R facilitates this convergence of research efforts by providing an infrastructure that allows new theoretical components

(like the account of fatigue) to be added seamlessly to the architecture. Once added, these new components, or modules, influence model behavior to the extent that the proper conditions arise to activate the mechanisms. In this case, the mechanisms for fatigue have a substantial impact on the model at more extreme levels of fatigue. Importantly, the impact appears to be appropriate based on human data from a similar task in the research literature.

An important question regarding this research effort relates to the generalizability of our findings to other tasks and domains. Our strategy of using a cognitive architecture is targeted explicitly at ensuring that the mechanisms we identify can be applied in other tasks and contexts to explain and predict human behavior. In the case of the mechanisms at play in the research described here, our expectation is that they will be influential in explaining the impact of fatigue across tasks and domains. This is because, as noted above, the cognitive cycle is a foundational component of all ACT-R models. At the same time, we have proposed mechanisms within other components of ACT-R, like declarative knowledge (Gunzelmann et al., 2007; Halverson, Gunzelmann, Moore, & Van Dongen, 2010) to explain effects of fatigue. These mechanisms should be influential in predicting performance changes stemming from fluctuations of alertness to the extent that those cognitive capacities are involved in performance on the particular task.

In other words, we believe that there are decrements with fatigue across components of cognitive functioning. To understand and predict how performance will change in a particular task, it will be necessary to understand both the interaction of fatigue with those components of cognition, as well as how, and to what extent, the information processing mechanisms are utilized in performing the task. Thus far, our research has focused primarily on tasks where we can effectively isolate the cognitive requirements to understand how they are impacted by fatigue. However, as our research progresses, it will become increasingly important to consider the interactions among various components of cognition to develop a more general and comprehensive account.

Finally, the integration achieved in this paper demonstrates both the scientific value of model reuse and the potential utility of cognitive architectures as vehicles for the integration and synthesis of scientific theories to produce cumulative scientific progress. Not only do these features reduce degrees of freedom in our model, but they will also be essential for making real predictions in naturalistic task contexts where data for model fitting is generally unavailable. This brings us back to a major goal of research on fatigue: To develop an understanding of the impact of sleep loss that is useful in making predictions regarding the consequences for performance in applied settings. At the outset, we cited the enormous cost of fatigue – both in dollars and lives – on highways in the United States. A better understanding of the relationship between fluctuations in alertness and changes in observable human

behavior has the potential to greatly reduce this cost, potentially saving thousands of lives. Moreover, driving is not the only area for which such potential benefits exist. In many applied settings – including military operations, commercial airline piloting, and nuclear power plant monitoring – lack of sleep and circadian desynchrony may lead to disastrous consequences (e.g., Caldwell, Caldwell, Brown, & Smith, 2004; Dinges, 1995). Accurate predictions of the consequences of fatigue could help to avert some of these potential tragedies.

Acknowledgments

References

Anderson, J. R. (2007). *How can the mind exist in a physical universe?* Oxford University Press.

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

Arruda, J. E., Zhang, H., Amoss, R. T., Coburn, K. L., & Aue, W. R. (2009). Rhythmic oscillations in quantitative EEG measured during a continuous performance task. *Applied Psychophysiology and Biofeedback, 34*(1), 7–16.

Caldwell, J. A., Caldwell, J. L., Brown, D. L., & Smith, J. K. (2004). The effects of 37 hours of continuous wakefulness on the physiological arousal, cognitive performance, self-reported mood, and simulator flight performance of F-117A pilots. *Military Psychology, 16*(3), 163–181.

Chua, Y. M. L., Venkatraman, V., Dinges, D. F., & Chee, M. W. L. (2006). The neural basis of interindividual variability in inhibitory efficiency after sleep deprivation. *The Journal of Neuroscience, 26*(27), 7156–7162.

Davies, D. R., & Parasuraman, R. (1982). *The psychology of vigilance.* New York, NY: Academic Press.

Dijk, D.-J., Duffy, J. F., & Czeisler, C. A. (1992). Circadian and sleep/wake dependent aspects of subjective alertness and cognitive performance. *Journal of Sleep Research, 1*, 112–117.

Dinges, D. F. (1995). An overview of sleepiness and accidents. *Journal of Sleep Research, 4*(2), 4–11.

Dinges, D. F. (2004). Critical research issues in development of biomathematical models of fatigue and performance. *Aviation, Space, and Environmental Medicine, 75*(3), A181–A191.

Dinges, D. F., & Kribbs, N. B. (1991). Performing while sleepy: effects of experimentally-induced sleepiness. In T. H. Monk (Ed.), *Sleep, sleepiness and performance* (pp. 97–128). Chichester: John Wiley & Sons.

Gunzelmann, G., Gluck, K. A., Kershner, J., Van Dongen, H. P. A., & Dinges, D. F. (2007). Understanding decrements in knowledge access resulting from increased fatigue. In D. S. McNamara & J. G. Trafton (Eds.), *Proceedings of the twenty-ninth annual meeting of the cognitive science society* (pp. 329–334). Austin, TX: Cognitive Science Society.

Gunzelmann, G., Gross, J. B., Gluck, K. A., & Dinges, D. F. (2009). Sleep deprivation and sustained attention performance: Integrating mathematical and cognitive modeling. *Cognitive Science.*

Gunzelmann, G., Moore, L. R., Gluck, K. A., Van Dongen, H. P. A., & Dinges, D. F. (2010). Fatigue in sustained attention: Generalizing mechanisms for time awake to time on task. In P. L. Ackerman (Ed.), *Cognitive Fatigue: Multidisciplinary perspectives on current research and future applications* (pp. 83–96). Washington, DC: American Psychological Association.

Halverson, T., Gunzelmann, G., Moore, L. R., & Van Dongen, H. P. A. (2010). Modeling the effects of work shift on learning in an orientation task. In D. D. Salvucci & G. Gunzelmann (Eds.), *Proceedings of the 10th international conference on cognitive modeling* (pp. 79–84). Philadelphia, PA: Drexel University.

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

Jewett, M. E., & Kronauer, R. E. (1999). Interactive mathematical models of subjective alertness and alertness in humans. *Journal of Biological Rhythms, 14*, 588–597.

Klauer, S. G., Dingus, T. A., Neale, V. L., Sudweeks, J. D., & Ramsey, D. J. (2006). The impact of driver inattention on near-crash/crash risk: An analysis using the 100-car naturalistic driving study data. *Technical report #DOT HS 810 594*. Washington, DC: National Highway Traffic Safety Administration.

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.

McCauley, P., Kalachev, L. V., Smith, A. D., Belenky, G., Dinges, D. F., & Van Dongen, H. P. A. (2009). A new mathematical model for the homeostatic effects of sleep loss on neurobehavioral performance. *Journal of Theoretical Biology, 256*, 227–239.

National Transportation Safety Board (1995). Factors that affect fatigue in heavy truck accidents. NTSB safety study: *Report no. NTSB//SS-95/01 (PB95-9170012)* (Vol. 1). Washington, DC: NTSB.

Neri, D. F. (Ed.). (2004). Fatigue and performance modeling workshop, June 13–14, 2002 [Special Issue]. *Aviation, Space, and Environmental Medicine, 75*(3, Suppl.).

Pack, A. I., Pack, A. M., Rodgman, E., Cucchiara, A., Dinges, D. F., & Schwab, C. W. (1995). Characteristics of crashes attributed to the driver having fallen asleep. *Accident Analysis & Prevention, 27*(6), 769–775.

Peters, R. D., Kloeppel, E., & Alicandri, E. (1999). Effects of partial and total sleep deprivation on driving performance. *Publication No. FHWA-RD-94-046*, Washington, DC: Federal Highway Administration, US Department of Transportation.

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.

Roberts, S., & Pashler, H. (2000). How persuasive is a good fit? A comment on theory testing. *Psychological Review, 107*(2), 358–367.

Salvucci, D. D. (2001). An integrated model of eye movements and visual encoding. *Cognitive Systems Research, 1*, 201–220.

Salvucci, D. D. (2006). Modeling driver behavior in a cognitive architecture. *Human Factors, 48*, 362–380.

Salvucci, D. D., & Beltowska, J. (2008). Effects of memory rehearsal on driver performance: Experiment and theoretical account. *Human Factors, 50*, 834–844.

Salvucci, D. D., & Gray, R. (2004). A two-point visual control model of steering. *Perception, 33*, 1233–1248.

Salvucci, D. D., & Taatgen, N. A. (2008). Threaded cognition: An integrated theory of concurrent multitasking. *Psychological Review, 115*, 101–130.

Stewart, T. C., Choo, X., & Eliasmith, C. (2010). Dynamic behavior of a spiking model of action selection in the basal ganglia. In D. D. Salvucci & G. Gunzelmann (Eds.), *Proceedings of the 10 international*

conference on cognitive modeling (pp. 235–240). Philadelphia, PA: Drexel University.

Van der Hulst, M., Meijman, T., & Rothengatter, T. (2001). Maintaining task set under fatigue: A study of time-on-task effects in simulated driving. *Transportation Research Part F, 4*, 103–118.

Van Dongen, H. P. A. (2004). Comparison of mathematical model predictions to experimental data of fatigue and performance. *Aviation, Space, and Environmental Medicine, 75*(3), A15–A36.

Van Dongen, H. P. A., Baynard, M. D., Maislin, G., & Dinges, D. F. (2004). Systematic interindividual differences in neurobehavioral impairment from sleep loss: Evidence of trait-like differential vulnerability. *Sleep, 27*(3), 423–433.

Van Dongen, H. P. A., & Dinges, D. F. (2005a). Circadian rhythms in sleepiness, alertness, and performance. In M. H. Kryger, T. Roth, & 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.