Title

A Unified Model of Fatigue in a Cognitive Architecture:Time-of-Day and Time-on-Task Effects on Task Performance

Permalink

https://escholarship.org/uc/item/8jn4g1z9

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 41(0)

Authors

Khosroshahi, Ehsan B.
Salvucci, Dario D.
Gunzelmann, Glenn
et al.

Publication Date

2019

A Unified Model of Fatigue in a Cognitive Architecture: Time-of-Day and Time-on-Task Effects on Task Performance

Ehsan B. Khosroshahi[a] **(ehsanebk@drexel.edu)**
Dario D. Salvucci[a] **(salvucci@drexel.edu)**
Glenn Gunzelmann[b] **(glenn.gunzelmann@us.af.mil)**
Bella Z. Veksler[c] **(bellav717@gmail.com)**

[a] Department of Computer Science, Drexel University, 3675 Market St.
Philadelphia, PA 19104, United States

[b] Warfighter Readiness Research Division, Air Force Research Laboratory, 2620 Q St.
Wright Patterson Air Force Base, OH 45433, United States

[c] Tier1 Performance Solutions, 100 E. Rivercenter Blvd., Suite 100
Covington, KY 41011, United States

Abstract

Capturing the effects of fatigue and, more generally, the effects of physical and mental states on human performance has been a topic of research for many years. Recent models, especially those developed in a cognitive architecture, have shown great promise in capturing these effects by providing insight into the specific cognitive and other components involved in task performance (like perception and motor movement). In particular, separate models have been developed to account for both time-of-day and time-on-task effects related to fatigue. In this paper, we present a novel unified model, developed in the ACT-R cognitive architecture, that captures both time-of-day and time-on-task effects with a single set of mechanisms and parameters. We demonstrate how this unified model accounts for quantitative and qualitative aspects of fatigued performance from two experiments, one focused on time-on-task effects under conditions of moderate fatigue, the other focusing on time-of-day effects under conditions of severe fatigue in a study of long-term (88-hour) sleep deprivation.

Keywords: Fatigue; sleep deprivation; cognitive architectures

Introduction

One of the most significant physiological states that affects human cognition is fatigue. Decades of research have investigated the effects of fatigue, sleep deprivation, and time-on-task in a number of important areas, including industrial disasters (e.g. Mitler et al.,1988), transportation accidents (e.g. Lauber & Kayten, 1988; Dinges, 1995), and motor vehicle crashes (e.g. Horne & Reyner, 1999; Pack et al., 1995). These studies have explored in depth the question of how fatigue modulates cognition and performance, and how we might quantify the effects of fatigue using mathematical or computational models and formalisms.

Of the many aspects of cognitive fatigue, there are two main factors that affect sustained attention and task performance: (1) sleep-related factors which are a function of sleep history and the time of the day when the task is being performed (circadian rhythm); and (2) task-related factors which are a function of the type of the task and how long the person has been doing the task, or *time-on-task* (Figure 1). Fatigue can also vary widely in its level of intensity: mild to moderate time-of-day or time-on-task effects may affect performance significantly (e.g., Pattyn et al, 2008; Bakan, 1955; Mackworth, 1948; Parasuraman, 1979), but severe fatigue that occur with long-term sleep deprivation can have even more drastic impacts on performance (e.g., Doran, Van Dongen, Dinges, 2001; Dorrian, Rogers, & Dinges, 2005).
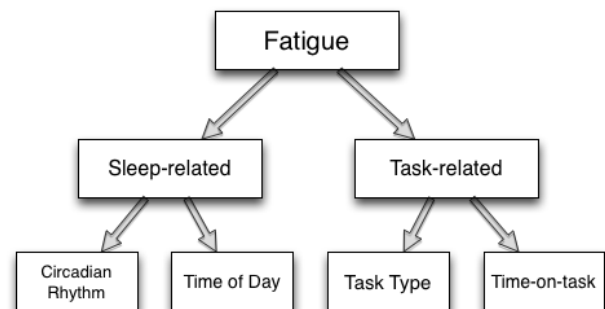

Figure 1: Main factors contributing to fatigue
in sustained-attention tasks.

Mathematical models of fatigue come in various forms, and can provide very good insight into the fluctuations in overall performance, accounting for moderate time-of-day and time-on-task effects (e.g. Fisk & Schneider, 1981; Giambra & Quilter, 1987; Mackworth, 1964) as well as effects of long-term sleep deprivation (e.g., Achermann, 2004; Borb & Achermann, 1999; Hursh et al., 2004; Jewett & Kronauer, 1999; McCauley et al., 2013). Such mathematical models aim to model the overall level of fatigue

at different points in time, but do not provide a detailed account of cognitive (and other) processes involved. Building on this work, recent models have focused on modeling fatigue within a computational cognitive architecture (e.g. French & Morris, 2003; Jones, Laird, & Neville, 1998; Gunzelmann, Gross, Gluck, & Dinges, 2009; Gunzelmann, Moore, Salvucci, & Gluck, 2011; Walsh, Gunzelmann, & Van Dongen, 2017; Veksler & Gunzelmann, 2018) to offer deeper insight into relationship between fatigue and the basic information processing mechanisms inherent to a task.

In this paper, we present a new unified model of fatigue that accounts for both sleep-related and task-related factors, and accounts for performance under both moderate and severe fatigue. In particular, we extend the work of Veksler and Gunzelmann (2018) and Walsh et al. (2017) by first examining the underlying theoretical foundations of both types of fatigue based on recent empirical work. We then utilize these ideas to propose an updated formulation of fatigue within the ACT-R cognitive architecture, testing its predictions against two data sets that demonstrate the benefits of a unified model.

Theoretical Foundations

The central idea of our modeling work is the concept of *microlapses*, introduced by Gunzelmann et al. (2009) to account for changes in behavioral performance related to fatigue. Microlapses can be viewed as an implementation of the "state instability" hypothesis (Doran et al., 2001): that a person's fatigue may be characterized as the switching between sleep and awake states, which may fluctuate second by second and can eventually progress to a physiological sleep state. Microlapses, however, incorporate the idea that switches between sleep and awake states may be more rapid (i.e., tens of ms), with transitions into, and remaining within, a "sleep" state becoming more likely as fatigue increases.

The concept of microlapses relies on the computational mechanisms of a procedural system in a cognitive architecture. A procedural system implemented as a production system is the central core of most well-established cognitive architectures like ACT-R (Anderson, 2007). ACT-R's production system implements a serial bottleneck in cognitive processing, representing cognition as a sequence of recognize-decide-act cycles that require about 50 ms each to execute. Under fatigue, microlapses cause the execution phase of the cycle to fail, leading to delays in completing, or even failure to complete, a task. As we will see, this theoretical foundation allows for an elegant model of fatigue that can account for both sleep- and task-related factors, and for performance across a range of degrees of fatigue.

Modeling Time-of-Day Effects

The first building block for our unified model is the model of sleep-related fatigue described in Walsh et al. (2017). Their model relied on ACT-R's concept of *utility*, namely that each

production rule (effectively a 50-ms unit of action) has an associated utility that determines its usefulness in being activated, and this utility can be compared to those of other rules to determine the next action. By manipulating the utility of the productions and the utility threshold, the system is able to produce microlapses: if the utility U_i of the selected production is less than a set utility threshold UT, a microlapse occurs. Because U_i values are noisy, changes in U_i and UT thus influence the probability of microlapses occurring.

To account for sleep-related factors, Walsh et al. (2017) used a biomathematical model to quantify the overall impact of time awake and circadian rhythms. First, let us assume that we have a biomathematical model value $B^S(t)$ that, given a sleep schedule S (i.e., the prior hours for which the person was asleep and awake), provides the level of fatigue at a given time of day t. As mentioned earlier, several such models have been developed in the past; Gunzelmann et al. use the formulation provided by McCauley et al. (2013), which we include here as well. Using this value, we can specify a fatigue scale factor $F_{bio}(t)$ that will scale a production's overall utility proportionally based on the biomathematical model's predictions:[1]

$$F_{bio}(t) = 1 - c_{bio} * B^S(t)$$

We also include a fatigue constant c_{bio} to scale the biomathematical value, and we will consider this constant as one parameter to estimate in our model fitting later.

The next component of the model represents the accumulated effect of microlapses, and incorporates the fact that when a microlapse occurs, another microlapse is more likely to occur immediately after. This component is formulated as follows:

$$F_{dec}(n) = (c_{dec})^n$$

Here, n is the number of consecutive microlapses that have occurred—thus, $n = 0$ after a normal production has fired, but would increase by 1 for each consecutive microlapse thereafter until another normal production firing. c_{dec} is assumed to be a constant between 0 and 1, and thus the value $F_{dec}(n)$ is also a value between 0 and 1 that decreases with larger values of n. As described by Walsh et al. (2017), $F_{dec}(n)$ can quickly decay to the point that will be too low to fire any production; however, there is a counterbalancing effect that resets $F_{dec}(n)$ by setting $n = 0$ (akin to awakening the model) when a stimulus is presented.

Integrating these factors together, following Gunzelmann et al. (2009), Walsh et al. (2017) defined a fatigued utility $FU_i(t, n)$ as a modified value of production i's base utility $U_i(t)$ scaled by both $F_{bio}(t)$ and $F_{dec}(n)$:

$$FU_i(t, n) = F_{bio}(t) * F_{dec}(n) * U_i(t) + \epsilon$$

The final term ϵ adds noise to the final fatigued utility, where the noise is sampled from a logistic distribution. This component is carried over from the standard utility function

---

[1] The names of some variables and constants have been changed from the original formulation for increased clarity.

in ACT-R, which includes this parameter to generate stochasticity in model behavior. Once this fatigued utility is computed, its value is compared to a utility threshold $UT(t)$, computed using the biomathematical model and a specified initial utility threshold UT_0:

$$UT_{bio}(t) = 1 - d_{bio} * B^S(t)$$

$$UT(t) = UT_{bio}(t) * UT_0$$

These equations introduce another constant, d_{bio}, that scales the biomathematical model value.

Modeling Time-on-Task Effects

As an extension to the above model of time-of-day effects, Veksler and Gunzelmann (2018) developed a model to capture the effects of time-on-task. Using the same core mechanisms as Walsh et al. (2017) described earlier, they replaced the biomathematical factor $F_{bio}(t)$ with a time-on-task factor $F_{tot}(T)$ defined as follows:

$$F_{tot}(T) = (1 + T)^{c_{tot}}$$

$$FU_i(t, T, n) = F_{tot}(T) * F_{dec}(n) * U_i(t) + \epsilon$$

Here, T represents the total time-on-task, or time spent performing the same task. Veksler et al. used a similar formulation to revise the computation of the utility threshold:

$$UT_{tot}(T) = (1 + T)^{d_{tot}}$$

$$UT(T) = UT_{tot}(T) * UT_0$$

The constants c_{tot} and d_{tot} are assumed to be between –1 and 0, and thus their respective functions decrease as the time-on-task T increases.

A Unified Model of Fatigue

The foundational components above provide the basis for our own unified model, and at first glance, one might expect that we could simply combine the equations and have a unified account directly. Unfortunately, a simple combination does not work well either theoretically or experimentally. We thus explore how we might combine these accounts and then proceed with a specification of the final unified model.

Developing a Unified Model

Examining the formulations for the time-of-day and time-on-task models above, the most straightforward approach to a unified model would be to simply multiple the respective factors together—that is, computing fatigued utility as:

$$FU_i(t, T, n) = F_{bio}(t) * F_{tot}(T) * F_{dec}(n) * U_i(t) + \epsilon$$

This approach multiples the biomathematical component $F_{bio}(t)$ with the time-on-task component $F_{tot}(T)$ to derive the total fatigued utility. In fact, this formulation has been tried with limited success in earlier work: Khosroshahi et al. (2016) used it to account for time-of-day effects on performance in psychomotor vigilance and driving.

Unfortunately, however, we have attempted to use this formulation to account for a broader set of time-of-day and time-on-task effects (discussed more later), and found this approach lacking for several reasons. Using this formulation, it was impossible to find a set of parameter values that produces acceptable results simultaneously for both time-on-task and time-of-day effects—especially when the latter is drawn out to long periods of sleep deprivation. For example, consider how the model might account for lapses in the psychomotor vigilance task (PVT), where participants simply see a visual stimulus and press a button in response, and where a *lapse* is defined as a response time greater than 500 ms. Using the formulation above, the model can nicely fit the number of lapses in the early stages of fatigue, namely during the first day or two without sleep; however, this produces a model that rarely suffers the sleep attacks (response times greater than 30 s) suffered by humans after 48-88 hours of sleep deprivation. On the flip side, if the model parameters were fitted to produce a human-like frequency of sleep attacks, the lapses under moderate fatigue would be much too large.

In summary, this was not an issue of parameter fitting—the model formulation itself was fundamentally flawed. Closer analysis of the model revealed its theoretical flaw: increasing values of the biomathematical model $B^S(t)$ over time would actually scale down the time-on-task effect—effectively making the time-on-task effects *smaller* as the model became more fatigued. This effect is counterintuitive, and indeed, we did not find any evidence to support it in our available data or in the literature. In addition, in their study of time-on-task effects, Veksler et al. (2018) found no correlation between either prior night's sleep or wake-up time and the difference in response times between the first and last blocks of a 35-minute task—indicating an additive, not multiplicative, relationship between time-of-day and time-on-task (see Kribbs & Dinges, 1994; Gunzelmann et al., 2010).

Yet another observation about the naïve combined model, and about the earlier time-on-task model, relates to the model's $F_{dec}(n)$ equation. Recall that this factor incorporates the idea of cascading microsleeps, such that when a microsleep occurs, another is more likely to happen in the subsequent cycle. In the original formulation, because $F_{dec}(n) = (c_{dec})^n$ and $0 < c_{dec} < 1$, there is a rapid initial drop for small n followed by a leveling off to an asymptote near zero. Instead, based on our observations of sleep attacks, a better formulation would allow for only a slight drop for small n, but as n gets larger, the microsleeps would rapidly deteriorate into a sleep attack.

The Unified Model

Given the reasoning above, we created our unified model based on the earlier models of time-of-day and time-on-task while reflecting the evidence above. In particular, we modified the formulations of several equations as follows. First, we changed the decrement factor to a negated exponential function to introduce a steep drop in fatigue as microsleeps accumulate:

$$F_{dec}(t, n) = -(e)^{c_{dec}*n} + 2$$

Next, we modified the biomathematical factor to eliminate the initial 1 in a way that forces it to reduce the overall utility:

$$F_{bio}(t) = -c_{bio} * B^S(t)$$

We then introduced the additive effect between time-of-day and time-on-task into the computation of fatigued utility:

$$F_{tot}(T) = (1 + T)^{c_{tot}}$$

$$FU_i(t, T, n) = F_{dec}(t, n) * [F_{bio}(t) + F_{tot}(T) + U_i(t) + \epsilon]$$

Analogous changes were applied to the utility threshold:

$$UT_{bio}(t) = -d_{bio} * B^S(t)$$

$$UT_{tot}(T) = (1 + T)^{d_{tot}}$$

$$UT(t, T) = UT_{bio}(t) + UT_{tot}(T) + UT_0$$

These changes all together represent our unified model that accounts for both time-of-day and time-on-task effects. The next section aims to validate this model across two experimental data sets.

Model Evaluation

To validate our model, we rely on two studies that employ arguably the most common task in fatigue-related studies, namely the psychomotor vigilance task (PVT: Dinges and Powell, 1985). As mentioned, the PVT involves an extremely simple stimulus-response. PVT has been used extensively in sleep-related studies because of its sensitivity to sleep and circadian-based fatigue and its procedural simplicity and the consistency of individual performance (e.g., Gunzelmann, Moore, Gluck, Van Dongen, Dinges, 2008; Dorrian et al., 2005). PVT is thus a highly sensitive sustained attention task which can be an independent measure of fatigue (Van Dongen et al., 2011).

A typical PVT trial lasts 10 minutes and requires a button response every 2-10 seconds. The visual stimulus is a millisecond counter displayed on the screen, which starts at 0 at stimulus onset and counts forward as time passes; when the person presses the response key, the counter stops, thus providing feedback for performance. The main dependent measure in the PVT is the number of *lapses*, where a lapse is defined as a reaction time of more than 500 ms. Researchers have also measured the *median response time* (RT) of alert responses (reaction times between 150 and 500 ms), *false starts* (incorrect keypresses or reaction times less than 150 ms), and *sleep attacks* where the participant does not respond for 30 seconds or more.

It is worth noting that we used a single set of parameter values for the models in both studies. Our unified model contains 7 free parameters in total (see Table 1). Another parameter that was treated as a free parameter in previous models is *cycle time*, which controls the amount of time to evaluate and select a production during each cognitive cycle.
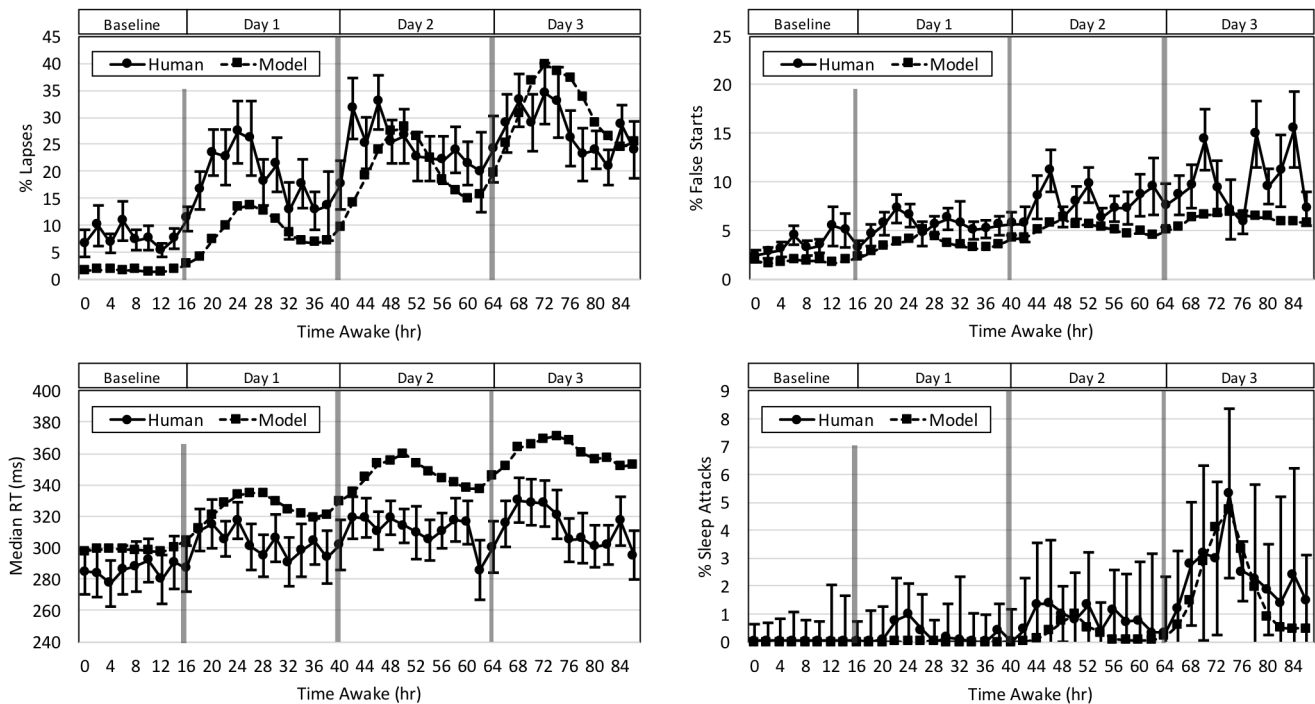

Figure 2: Human and model results for the PVT across 88 hours of sleep deprivation (Study 1).

We used the default value of 50 ms (Anderson, 2007) to keep the consistency with ACT-R theory. To reduce the chance of overfitting, we searched the parameter values to first fit parameters related to the time-on-task effect and then used the same values to fit parameters related to the time-of-day effect. Table 1 shows the list of free parameters and our best estimates for each parameter.

Table 1: ACT-R unified fatigue model free parameters and their estimations.

Parameter	Definition	Estimates
c_{dec}	Utility decrement factor	.006
c_{bio}	Utility biomathematical factor	.028
c_{tot}	Time-on-task decrement factor	.12
U_i	Base utility[2]	1.56
d_{bio}	Threshold biomathematical factor	.01
d_{tot}	Threshold time-on-task decrement factor	.04
UT_0	The initial threshold	1.15

PVT Model

Because the fatigue mechanisms described here are general to any production system or task, we require a model specifically of the PVT to test the fatigue mechanisms. For this purpose, we developed an ACT-R model that performs the PVT in as straightforward a manner as possible, with three main production rules, following the original model by Walsh et al. (2017):

1. *Attend*: shift visual attention to the stimulus
2. *Encode-and-Respond*: completes the visual encoding of the stimulus and initiates the response keypress
3. *Wait*: wait for the next stimulus

To capture the false starts in the PVT model, Walsh et al. (2017) used procedural partial matching: when enabled, productions whose conditions do not perfectly match the current state get a chance to be selected with a similarity difference (a negative value) added to their utility:

$$U_i' = U_i + SD_i + \epsilon$$

SD_i is the similarity difference which is added to the utility value when the conditions for the production are not met. At each cycle, the production with the greatest value U_i is selected when its utility exceeds the utility threshold. By enabling the procedural partial matching, Walsh et al. (2017)

---

[2] Base utility is defined as the standard utility value for all productions.

eliminated the need of a separate production (false-response); encode-and-respond can be selected at any time and when it is selected before the stimulus appears, false starts occur (which happens rarely because of the similarity difference added to it).

In the design of PVT in Walsh et al. (2017), U_i' was treated as a single free parameter meaning that one value was estimated and used for all the productions. The ACT-R's procedural learning (Anderson, 2007) was also disabled due to the nature of PVT and similar sustained attention tasks (Van Dongen et al., 2003) and the similarity difference was set to negative value of the production utility to simplify matters. Here we follow a similar design to stay consistent with earlier studies.

Study 1: Time-of-Day Experiment and Results

The first study for our model evaluation is a study of long-term sleep deprivation conducted by Doran et al. (2001). The study included 13 healthy participants who experienced 88 hours of total sleep deprivation. During periods of wakefulness for the duration of the study, participants
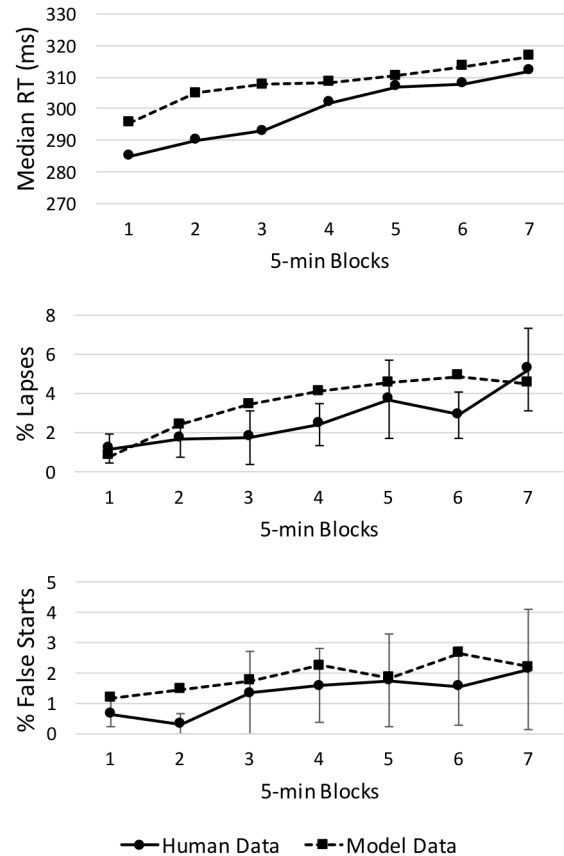
Figure 3: Human and model PVT results across the 5-min blocks of the 35-min experiment (Study 2).

completed a battery of performance evaluation tasks every 2 hours, including a 10-minute PVT. Although Gunzelmann et al. (2009) modeled the same experiment, our effort here is different in two ways: (1) the older model used different parameter values for each day, whereas we are constraining our model to a single set of parameters; and (2) the older model did not include the time-on-task factor; although time-on-task was not a focus of the study, it is important for us to know that the model can produce a good fit while incorporating this factor (see also Gunzelmann et al. 2011).

To model this study, we ran iterations of the PVT model for 10-minute periods and matched the sleep schedule of the model to the 88-hour sleep deprivation experimental protocol. Parameters were estimated to produce the best fit across the four PVT measures; a single set of parameter values was used across the entire experiment. Figure 2 shows the human data and model's performance for all four measures: lapses ($R^2 = 0.68$, $RMSE = 7.92$), median reaction times ($R^2 = 0.55$, $RMSE = 34.41$), false starts ($R^2 = 0.56$, $RMSE = 3.42$), and sleep attacks ($R^2 = 0.77$, $RMSE = 0.64$). Overall, the model accounted for all the major aspects of the data; it slightly underpredicted lapses in days 1-2, and slightly overpredicted RT in days 4-5, but in general, the model captured most of the fluctuations in performance across all four measures.

Study 2: Time-on-Task Experiment and Results

The second study for our model evaluation is a study examining time-on-task effects conducted by Veksler and Gunzelmann (2018). In the study, 20 participants performed a 35-minute PVT instead of the usual 10 minutes; by extending the typical PVT duration, they were able to draw out how the effects of time-on-task on PVT are similar to those of sleep loss. As mentioned earlier, Veksler and Gunzelmann modeled the time-on-task effects in this experiment, but at the time did not incorporate the biomathematical model, and used a different set of parameters than earlier models. To include biomathematical modeling in our simulations, we assumed 8 hours the night before the experiment, waking at 7:30am and performing the experiment at 10:00am.

The results of the model compared to the human data are shown in Figure 3. For this evaluation, we compared the performance of the model with the experimental results across seven 5-minute blocks of PVT. The model was able to capture the changes across the blocks for median reaction times ($R^2 = 0.85$, $RMSE = 9.67$), lapses ($R^2 = 0.53$, $RMSE = 1.27$), and false starts ($R^2 = 0.55$, $RMSE = 0.69$). The model shows a slight overprediction of lapses in the middle blocks, but in general, the model performs well for these three measures, especially considering that this is the same model with the same parameters as the previous study.

General Discussion

In this paper, we introduce a unified computational model that accounts for two of the most important aspect of fatigue, namely time-of-day and time-on-task effects on behavior and performance. Our result once again accounts for the microlapse hypothesis (Gunzelmann et al. 2009) and the fact that microlapses could account for both sleep loss and time-on-task effects in sustained attention (following Veksler et al., 2018). We were also able to capture both the time-of-day and time-on-task effects with the same parameters; going forward, we are interested in understanding how these parameters might generalize to other tasks, and how they might vary across individuals. It is also notable that the mechanisms here are complex, with a number of free parameters that are sensitive to changes in setting. Nevertheless, we believe that as we continue to fit additional experiments with this unified model, we can reduce the space of free parameters and can find parameter values that cut across a variety of task domains, providing an even more general model with easier estimation of parameters.

In conclusion, by validating that the unified model can account for the negative consequences in behavioral performance of both time-of-day and time-on-task effects, we have demonstrated that both phenomena have similar natures and as a result could be modeled with a single set of mechanisms. Although PVT as a testbed for our modeling seems to be a simple task, this research will give us a strong foundation to expand the model to more complex domains. We are also interested in extending this model beyond the sleep-loss and time-on-task to moderate levels of fatigue (e.g., sequential sleep limitation), which would further bolster the model's generalizability to complex real-world task domains.

References

Achermann, P. (2004). The two-process model of sleep regulation revisited. *Aviation, Space, and Environmental Medicine*, *75*, A37-A43.

Anderson, J. R. (2007). *How Can the Human Mind Occur in the Physical Universe?* New York: OUP.

Bakan, P. (1955). Discrimination decrement as a function of time in a prolonged vigil. *Journal of Experimental Psychology*, *50*, 387.

Borb, A. A., & Achermann, P. (1999). Sleep homeostasis and models of sleep regulation. *Journal of Biological Rhythms*, *14*, 559-570.

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*, 652-655.

Dinges, D. F. (1995). An overview of sleepiness and accidents. *Journal of Sleep Research*, *4*, 4-14.

Doran, S. M., Van Dongen, H. P. A., & Dinges, D. F. (2001). Sustained attention performance during sleep deprivation:

evidence of state instability. *Archives Italiennes de Biologie, 139*, 253-267.

Dorrian, J., Rogers, N. L., & Dinges, D. F. (2005). Psychomotor vigilance performance: Neurocognitive assay sensitive to sleep loss. In C. A. Kushida (Ed.), *Sleep Deprivation: Clinical Issues, Pharmacology and Sleep Loss Effects* (pp. 39–70). New York: Marcel Dekker.

Fisk, A. D., & Schneider, W. (1981). Control and automatic processing during tasks requiring sustained attention: A new approach to vigilance. *Human Factors*, *23*, 737-750.

French, J., & Morris, C. S. (2003). Modeling fatigue degraded performance in artificial agents. In *Proc. of the Human Factors and Ergonomics Society Annual Meeting* (Vol. 47, No. 3, pp. 307-310). Los Angeles, CA: SAGE.

Giambra, L. M., & Quilter, R. E. (1987). A two-term exponential functional description of the time course of sustained attention. *Human Factors*, *29*, 635-643.

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*, *33*, 880-910.

Gunzelmann, G., Moore, L. R., Gluck, K. A., Van Dongen, H. P., & Dinges, D. F. (2008). Individual differences in sustained vigilant attention: Insights from computational cognitive modeling. In *Proc. of the 30th Annual Meeting of the Cognitive Science Society* (pp. 2017-2022).

Gunzelmann, G., Moore, L. R., Gluck, K. A., Van Dongen, H. P., & Dinges, D. F. (2010). Fatigue in sustained attention: Generalizing mechanisms for time awake to time on task. *Cognitive fatigue: Multidisciplinary perspectives on current research and future applications* (pp. 83-101). Washington, DC: American Psychological Association.

Gunzelmann, G., Moore Jr., L. R., Salvucci, D. D., & Gluck, K. A. (2011). Sleep loss and driver performance: Quantitative predictions with zero free parameters. *Cognitive Systems Research*, *12*, 154-163.

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

Hursh, S. R., Redmond, D. P., Johnson, M. L., Thorne, D. R., Belenky, G., Balkin, T. J., ... & Eddy, D. R. (2004). Fatigue models for applied research in warfighting. *Aviation, Space, and Environmental Medicine*, *75*, A44-A53.

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

Jones, R. M., Laird, J. E., & Neville, K. (1998). Modeling pilot fatigue with a synthetic behavior model. In *Proc. of the 7th Conference on Computer Generated Forces and Behavioral Representation* (pp. 349-356).

Khosroshahi, E. B., Salvucci, D. D., Veksler, B. Z., & Gunzelmann, G. (2016). Capturing the effects of moderate

fatigue on driver performance. In *Proc. of the 14th Intl. Conference on Cognitive Modeling* (pp. 163-168).

Kribbs, N. B., & Dinges, D. (1994). Vigilance decrement and sleepiness. In R. D. Ogilvie & J. R. Harsh (Eds.), *Sleep onset: Normal and abnormal processes* (pp. 113-125). Washington, DC: APA.

Lauber, J. K., & Kayten, P. J. (1988). Sleepiness, circadian dysrhythmia, and fatigue in transportation system accidents. *Sleep: Journal of Sleep Research & Sleep Medicine*, *11*, 503-512.

McCauley, P., Kalachev, L. V., Mollicone, D. J., Banks, S., Dinges, D. F., & Van Dongen, H. P. (2013). Dynamic circadian modulation in a biomathematical model for the effects of sleep and sleep loss on waking neurobehavioral performance. *Sleep*, *36,* 1987–1997.

Mackworth, N. H. (1948). The breakdown of vigilance during prolonged visual search. *Quarterly Journal of Experimental Psychology*, *1*, 6-21.

Mackworth, J. F. (1964). Performance decrement in vigilance, threshold, and high-speed perceptual motor tasks. *Canadian Journal of Psychology*, *18*, 209.

Mitler, M. M., Carskadon, M. A., Czeisier, C. A., Dement, W. C., Dinges, D. F., & Graeber, R. C. (1988). Catastrophes, sleep, and public policy: Consensus report. *Sleep*, *11*, 100-109.

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*, 769-775.

Parasuraman, R. (1979). Memory load and event rate control sensitivity decrements in sustained attention. *Science*, *205*, 924-927.

Pattyn, N., Neyt, X., Henderickx, D., & Soetens, E. (2008). Psychophysiological investigation of vigilance decrement: boredom or cognitive fatigue? *Physiology & Behavior*, *93*, 369-378.

Van Dongen, H. P. A., Belenky, G., & Krueger, J. M. (2011). A local, bottom-up perspective on sleep deprivation and neurobehavioral performance. *Current Topics in Medicinal Chemistry*, *11*, 2414-2422.

Van Dongen, H. P. A., Maislin, G., Mullington, J. M., & Dinges, D. F. (2003). The cumulative cost of additional wakefulness: Dose-response effects on neurobehavioral functions and sleep physiology from chronic sleep restriction and total sleep deprivation. *Sleep, 26*, 117–126.

Veksler, B. Z., & Gunzelmann, G. (2018). Functional equivalence of sleep loss and time on task effects in sustained attention. *Cognitive Science, 42*, 600-632.

Walsh, M. M., Gunzelmann, G., & Van Dongen, H. P. A. (2017). Computational cognitive models of the temporal dynamics of fatigue from sleep loss. *Psychonomic Bulletin & Review, 24*, 1785-1807.