

Computational cognitive modeling of the temporal dynamics of fatigue from sleep loss

Matthew M. Walsh¹ · Glenn Gunzelmann² · Hans P. A. Van Dongen^{3,4}

Published online: 16 February 2017
© Psychonomic Society, Inc. 2017

Abstract Computational models have become common tools in psychology. They provide quantitative instantiations of theories that seek to explain the functioning of the human mind. In this paper, we focus on identifying deep theoretical similarities between two very different models. Both models are concerned with how fatigue from sleep loss impacts cognitive processing. The first is based on the diffusion model and posits that fatigue decreases the drift rate of the diffusion process. The second is based on the Adaptive Control of Thought – Rational (ACT-R) cognitive architecture and posits that fatigue decreases the utility of candidate actions leading to microlapses in cognitive processing. A biomathematical model of fatigue is used to control drift rate in the first account and utility in the second. We investigated the predicted response time distributions of these two integrated computational cognitive models for performance on a psychomotor vigilance test under conditions of total sleep deprivation, simulated shift work, and sustained sleep restriction. The models generated equivalent predictions

of response time distributions with excellent goodness-of-fit to the human data. More importantly, although the accounts involve different modeling approaches and levels of abstraction, they represent the effects of fatigue in a functionally equivalent way: in both, fatigue decreases the signal-to-noise ratio in decision processes and decreases response inhibition. This convergence suggests that sleep loss impairs psychomotor vigilance performance through degradation of the quality of cognitive processing, which provides a foundation for systematic investigation of the effects of sleep loss on other aspects of cognition. Our findings illustrate the value of treating different modeling formalisms as vehicles for discovery.

Keywords ACT-R · Biomathematical fatigue models · Cognitive architecture · Diffusion model · Leaky accumulator · Local sleep · Psychomotor vigilance test (PVT)

Introduction

Fatigue due to sleep loss is associated with a wide range of deficits in cognitive performance (Jackson & Van Dongen, 2011; Killgore, 2010; Lim & Dinges, 2010) and an increased risk of errors and accidents (e.g., Dinges, 1995). The effects of fatigue are particularly profound for tasks involving sustained attention (Lim & Dinges, 2008). However, the effects of fatigue on cognitive performance vary considerably across tasks (Jackson et al., 2013). This variability reflects in part the differential effects of fatigue on different cognitive processes (Tucker, Whitney, Belenky, Hinson, & Van Dongen, 2010). Therefore, to better understand the effects of fatigue from sleep loss, it is necessary to consider the range of cognitive processes evoked during task performance (Whitney & Hinson, 2010). One way to do so is through

Electronic supplementary material The online version of this article (doi:10.3758/s13423-017-1243-6) contains supplementary material, which is available to authorized users.

✉ Matthew M. Walsh
mmw188@gmail.com

¹ TiER1 Performance Solutions, Covington, KY, USA

² Air Force Research Laboratory, Dayton, OH, USA

³ Sleep and Performance Research Center, Washington State University, Spokane, WA, USA

⁴ Elson S. Floyd College of Medicine, Washington State University, Spokane, WA, USA

the use of computational cognitive models (Gunzelmann, Gross, Gluck, & Dinges, 2009; Gunzelmann et al., 2015; Jackson et al., 2013).

Computational cognitive models provide a way to examine task performance in terms of underlying processing mechanisms (Newell, 1990). Models vary in terms of their focus (e.g., linguistics, learning, perception, motor planning; see Gray, 2007) and in terms of the level of analysis or representation they employ (e.g., connectionist models, cognitive architectures, Bayesian models; see Sun, 2008). Despite the prevalence of computational models in psychological science, quantitative models that account for the effects of fatigue from sleep loss on cognitive processes have only recently appeared (e.g., Gunzelmann et al., 2009a; Gunzelmann, Gluck, Moore, & Dinges, 2012; Ratcliff & Van Dongen, 2009, 2011; Veksler & Gunzelmann, in press). This reflects a longstanding tendency for models developed in psychology to focus on empirical phenomena and data associated with favorable physiological and psychological states that are conducive to cognitive processing.

Representing the detrimental effects of fatigue in cognitive models remains a critical challenge for developing unified theories of cognition (Newell, 1990). Representing these effects in cognitive models also creates an opportunity to evaluate theories from the sleep and sleep deprivation literature. Computational models enhance the transparency of theories – including theories of fatigue – and enable quantitative predictions. Recent modeling work suggests that fatigue mainly affects central cognition rather than perceptual and motor processes (Gunzelmann et al., 2009a; Ratcliff & Van Dongen, 2011). Here we build on that work and examine how fatigue affects cognitive processes in two distinct computational models. We show that the quantitative predictions of both models correspond closely to human performance data. More importantly, we demonstrate that the models implement ostensibly different but actually similar theoretical accounts of fatigue, despite being realized in distinct and substantively different conceptual and computational frameworks. In this way, two disparate models of fatigue, implemented computationally, map onto one theory. The juxtaposition of models provides an opportunity to pursue broader understanding of how fatigue impacts cognition, and illustrates how convergence can be achieved across alternative modeling approaches.

Psychomotor vigilance deficits due to fatigue

In this paper, we focus on deficits in sustained attention, one of the most profound and extensively studied aspects of fatigue from sleep loss. These deficits are frequently measured with the psychomotor vigilance test (PVT), a performance task in which stimuli are presented with random inter-trial intervals (Dinges & Powell, 1985; Lim & Dinges, 2008). The most

commonly used version of the task is 10 min in duration and has inter-trial intervals ranging from 2 to 10 s.¹ Participants watch a blank computer monitor for the sudden onset of a visual millisecond counter after which they respond with a key press. The counter starts from zero and continuously increments until the participant responds, or until 30 s have passed. Elapsed time remains on the screen for 1 s after the participant responds, serving as performance feedback. Participants are instructed to respond as quickly as possible, but to avoid responding before the stimulus appears.

PVT responses are frequently divided into three categories. *False starts* occur before the stimulus appears or within 150 ms of stimulus presentation; *alert responses* occur from 150 ms to 500 ms after stimulus presentation; and *lapses* occur more than 500 ms after stimulus presentation.² A fourth category, *non-responses*, is sometimes included for cases when no response occurs before the trial automatically terminates after 30 s.

Increased performance instability is a hallmark effect of fatigue on PVT performance (Doran, Van Dongen, & Dinges, 2001). The distribution of response times on the PVT, which has a long right tail even when participants are well rested, becomes increasingly skewed as participants become more fatigued. Fatigued participants make more extremely slow responses (i.e., more lapses), and they respond before the counter appears more often (i.e., more false starts). The progressive increase in performance variability typified by extremely slow responses and premature responses, and the general slowing of responses with time awake, are standard effects of sleep deprivation on the PVT and other vigilance tasks (Dorrian, Rogers, & Dinges, 2005; Kleitman, 1963). This performance instability contributes to an elevated risk of errors and accidents in fatigued individuals (Van Dongen & Hursh, 2010).

The PVT is among the most widely used cognitive assays of fatigue from sleep loss. PVT performance is highly sensitive to acute total sleep deprivation, sustained sleep restriction, circadian rhythm, time on task, and a range of fatigue countermeasures (Dorrian, Rogers, & Dinges, 2005; Lim & Dinges, 2008). Aptitude and practice effects on the PVT are negligible, making the task suitable for repeated administration and comparisons both within and between individuals (Dorrian, Rogers, & Dinges, 2005; Horne & Wilkinson, 1985). PVT performance deficits due to fatigue from sleep loss vary systematically among participants, reflecting stable

¹ Though we focus on the 10-min PVT, our description of PVT results and the integrated cognitive models we present apply to the more recently developed 3-min PVT as well (Basner & Dinges, 2011; Basner, Mollicone, & Dinges, 2011).

² Various sources define false starts as responses within 100 ms, 130 ms, or 150 ms of stimulus onset. These slight variations did not impact the cognitive modeling results presented later.

inter-individual differences in vulnerability to sleep loss (Van Dongen, Baynard, Maislin, & Dinges, 2004).

Biomathematical models of the temporal dynamics of fatigue

Biomathematical models of fatigue describe and predict changes in the level of fatigue over the course of hours and days. Different models account for different sets of fatigue factors, but all instantiate two primary processes of sleep regulation (Hursh & Van Dongen, 2010; Mallis, Mejdal, Nguyen, & Dinges, 2004). A homeostatic process controls the increasing drive for sleep with continuous time awake, and a circadian process promotes wakefulness during the day and sleep at night (Borbély & Achermann, 1999). Together these two primary processes produce dynamic changes in the level of fatigue. The homeostatic process is modulated by chronic sleep restriction across days (Hursh et al., 2004; McCauley et al., 2009, 2013).

Biomathematical models of fatigue are not specifically concerned with cognitive mechanisms of task performance. Rather, model outputs are typically scaled to a generic summary measure of performance such as an effectiveness score, or a selected outcome measure such as the number of lapses on the PVT. The basic trends produced by the models are observed in a variety of tasks. However, because the outputs of biomathematical models are fitted to specific tasks and measures using existing data, their ability to predict performance in novel tasks is limited. In addition, biomathematical models of fatigue are generally silent about underlying changes in cognitive processing caused by sleep loss and circadian rhythmicity.

Computational cognitive models of fatigue

Two classes of computational models have recently been used to account for the cognitive effects of fatigue. The first is based on a diffusion process that simulates the accumulation of evidence during simple and two-alternative forced-choice tasks (Laming, 1968; Ratcliff & McKoon, 2008; Stone, 1960). In this model, evidence accumulates stochastically until reaching a decision threshold, at which time a response is initiated. By varying parameter values in the diffusion model, researchers have reproduced the effects of various experimental manipulations on choice accuracy and response time distributions (e.g., Ratcliff & McKoon, 2008).

Ratcliff and Van Dongen (2011) showed that fatigue may be simulated by adjusting a composite diffusion model parameter, drift rate divided by drift rate variability. Varying this composite parameter allowed the diffusion model to capture changes in the response time distribution for individuals performing the PVT with increasing time awake (Ratcliff & Van Dongen, 2011). The same approach accounted for the

effects of fatigue in a two-choice numerosity discrimination task (Ratcliff & Van Dongen, 2009).

The second class of models uses a cognitive architecture to account for how fatigue affects specific components of cognition (e.g., Gunzelmann et al., 2009a). Cognitive architectures are general theories of cognition that specify foundational information processing mechanisms and how they interact with one another (Gluck, 2010). In the cognitive architecture used here, Adaptive Control of Thought – Rational (ACT-R), information enters the system through perceptual modules, affects processing throughout a collection of internal modules, and ultimately causes the manual module to issue a motor response (Anderson, 2007).

Gunzelmann et al. (2009a) developed an account of how fatigue selectively impairs information processing mechanisms in ACT-R. Briefly, fatigue reduced the utility of candidate behaviors, causing them probabilistically to fall below the threshold for action (cf. Gartenberg et al., 2014). This integrated account captured changes in the complete PVT response time distribution, including false starts, alert responses, lapses, and sleep attacks (Gunzelmann et al., 2009a). The same account made realistic predictions of the effects of fatigue on dual-task performance (e.g., Gunzelmann, Byrne, Gluck, & Moore, 2009) and lane deviation in driving (Gunzelmann, Moore, Salvucci, & Gluck, 2011).

Overview

Despite the many experiments that have documented the effects of fatigue on cognitive performance, and the increasingly widespread use of biomathematical models of fatigue in real-world applications, computational accounts of how fatigue affects cognitive processing mechanisms remain limited (cf., Gunzelmann et al., 2015). Further, there has been no attempt at a comparative analysis of the few available computational cognitive models of fatigue. Several basic questions remain unresolved: How does fatigue impact cognitive processes? How do different cognitive models account for the effects of fatigue? And how do these models relate to one another and to theories of fatigue from the sleep research literature?

We developed and compared two computational cognitive accounts of the effects of fatigue on cognition – a diffusion model and an ACT-R model – and we used the PVT as a test bed for evaluating these accounts. The diffusion model is based on that of Ratcliff and Van Dongen (2011), which we augmented with a leaky accumulator (Usher & McClelland, 2001) to capture false starts. Additionally, we merged the diffusion model with a biomathematical model of fatigue (McCauley et al., 2013) to systematically vary parameters in the diffusion model that are affected by sleep loss. The ACT-R model is based on that of Gunzelmann et al. (2009a). We updated their model for the latest release of the architecture, ACT-R 7.0 (Anderson, 2007), and fitted it to new data sets.

The same biomathematical model (McCauley et al., 2013) was used to systematically vary parameters in the ACT-R model that are affected by fatigue.

To date, the two accounts are the only computational cognitive models that have been shown to predict the complete distribution of response times on the PVT. Yet, aside from this shared capability, the diffusion model and ACT-R have few obvious similarities. As such, the question arose whether the two would support a consistent picture of how fatigue affects psychomotor vigilance performance. The focus of the research described here is not to evaluate the viability of either ACT-R or diffusion models as theories of cognitive processes per se. Nor is the objective to declare one of the models superior in accounting for the effects of fatigue. Rather, we evaluate what each model tells us about how fatigue impacts cognitive processing. In this way, we take seriously McClelland's (2009) statement that models are "vehicles for scientific discovery" (p. 16) and employ the two models' different levels of abstraction to look for converging evidence about the nature of fatigue, not tied to a specific modeling formalism, in the context of sustained attention.

Integrated models

Diffusion and ACT-R models have previously been proposed to account for the cognitive effects of fatigue on PVT performance, but the models did not focus on the same observed deficits. We modified and expanded both in order to be able to compare them directly.

Diffusion model

The diffusion model is based on a random walk sequential sampling process (Ratcliff & McKoon, 2008). In the diffusion model, evidence is accumulated over time (Fig. 1). Incoming information drives the process toward one of two decision boundaries during two-alternative forced choice tasks, or one decision criterion (A) during one-choice reaction time tasks. The rate of evidence accumulation, called drift rate, varies across trials according to a normal distribution with mean V and standard deviation η . The process terminates when accumulated evidence reaches a decision criterion, at which point the decision is made and the response initiated. Decision time is the elapsed time from when the diffusion process begins until it reaches the decision criterion. Other, non-central cognitive processes involved in task performance, such as perceptual and motor processes, make up non-decision time, which is represented as a uniformly distributed variable with mean T_{ND} and spread S_T . Response time equals the sum of decision time and non-decision time.

The diffusion model is typically fitted to accuracy rates and response time distributions for correct and error

responses. The model parameters capture a wide range of empirical effects (Ratcliff, 2006; Ratcliff & McKoon, 2008). For example, mean drift rate (V) nominally corresponds to the signal-to-noise ratio in the evidence accumulation process. Low drift rates produce slower, less accurate responses, and high drift rates produce faster, more accurate responses. Variations in mean drift rate can account for the effects of stimulus quality on response time and accuracy (Ratcliff & McKoon, 2008; Voss, Rothermund, & Voss, 2004). The decision criterion (A) controls whether responses are conservative or liberal. When the decision criterion is high, responses are slower but more accurate, and when the decision criterion is low, responses are faster but less accurate. Variations in the decision criterion can capture, for example, the effects of incentives or instructions to prioritize speed versus accuracy on performance (Ratcliff & McKoon, 2008).

The diffusion model has been applied to performance on many different tasks, such as perceptual discrimination (Ratcliff & Rouder, 2000), signal detection (Ratcliff, Van Zandt, & McKoon, 1999), and lexical decision making (Ratcliff, Gomez, & McKoon, 2004). In the context of fatigue research, the model has been applied to numerosity discrimination and the PVT (Ratcliff & Van Dongen, 2009, 2011; Patanaik, Zagorodnov, Kwoh, & Chee, 2014). For our purposes, the diffusion model is particularly suitable for studying the effects of fatigue on the PVT because it predicts complete response time distributions, and because its parameters can be manipulated to produce continuous performance decrements due to fatigue.

Task model

Ratcliff and Van Dongen (2011) developed a diffusion model of the PVT. Because the PVT is a one-choice reaction time task, the model includes a single, positive decision criterion, A (Fig. 1). Although accumulated evidence could become negative during the accumulation process, the process terminates only upon reaching the positive decision criterion.

Ratcliff and Van Dongen (2011) fitted this diffusion model to data from participants who completed the PVT every 2 h while staying awake for 36 h. They found that sleep deprivation mainly affected evidence accumulation, whereas the decision criterion and non-decision time were only marginally influenced by fatigue. As fatigue increased, drift rate (V) decreased and between-trial variability in drift rate (η) increased. The temporal dynamics of the effect of fatigue on PVT performance could be described by changing a single (composite) diffusion model scaling parameter, the drift ratio V/η (Ratcliff & Van Dongen, 2011).




Fig. 1 Three sample trajectories from a one-criterion diffusion process with mean drift rate V . Evidence accumulation starts at 0 and terminates upon reaching criterion A . Histograms show example densities of response times

Task model extension

The model of Ratcliff and Van Dongen (2011) makes predictions about the complete response time distribution. However, because the model pertains to processing after the stimulus appears, it does not produce false starts. Increased numbers of false starts under conditions of sleep loss are an important phenomenon, as false starts indicate that fatigue-related performance impairments cannot be explained solely by general slowing or reduced motivation (Doran, Van Dongen, & Dinges, 2001). Here we expanded the diffusion model to also predict the proportion of false starts.

In a typical diffusion model, the decision process is initiated when a stimulus is presented. This is problematic because, in theory, it would necessitate a separate decision process to detect trial onset and to initiate the primary decision process associated with the trial. Further complicating matters, in tasks where different types of stimuli necessitate different decision processes (e.g., numerosity judgments vs. lexical decisions), the pre-trial decision process would need to include multiple boundaries associated with deciding which of the tasks to perform. To avoid this basic dilemma, we expanded the model so that the diffusion process began immediately after feedback from the preceding trial disappeared (i.e., in anticipation of the next stimulus). The diffusion process persisted throughout the pre-stimulus interval and after the stimulus appeared. Although this may seem counterintuitive, we believe that it represents a plausible characterization of decision processes. The change allows multiple processes to operate in parallel, with irrelevant processes remaining essentially dormant until the appropriate stimuli move their drift rates from zero.

Conceptually, the signal-to-noise ratio in the PVT should be zero during the pre-stimulus interval and greater than zero once the stimulus appears. To produce these dynamics, we allowed V to take on positive values once the

stimulus appeared, and we fixed it to zero during the pre-stimulus interval. Due to within-trial stochastic variability, the diffusion process is equally likely to move upward or downward at each point in time when V is zero. The noisy accumulation of evidence during the pre-stimulus interval allows the diffusion process to occasionally reach the decision criterion before the stimulus appears and thereby cause a false start.

Preliminary simulations with this model extension confirmed that stochasticity alone caused the diffusion process to nearly always reach the decision criterion prematurely, and the predicted number of false starts far exceeded what is typically seen. Therefore, we implemented an additional process based on the intrinsic decay present in neural integrator models (Abbott, 1991; Cain & Shea-Brown, 2012; Goldman, Compte, & Wang, 2009; Robinson, 1989; Smith & Ratcliff, 2009). We used the leaky competing accumulator model (LCA; Ossmy et al., 2013; Usher & McClelland, 2001), in which the diffusion process depends on two opposing forces: accumulation of evidence from a stimulus drives the process toward the decision criterion, while decay (or leakage) pulls the process back to zero with a certain decay rate, λ .

Decay can be seen as controlling response inhibition. When the value of λ is high, the diffusion process tends to return to zero and is thus more robust against noise (greater inhibition), whereas when the value of decay is low, the diffusion process is more sensitive to noise (less inhibition). Because decay rate depends on the simulation step size, we sometimes refer to the diffusion process's integration time constant instead, which does not depend on the simulation step size. The integration time constant is the time it takes for the diffusion process to reach $1 - 1/e$ of its final (asymptotic) value given a signal of constant intensity.

The frequency of false starts is jointly determined by the value of λ and by the decision criterion, A (Fig. 2). With the LCA process implemented, the diffusion model

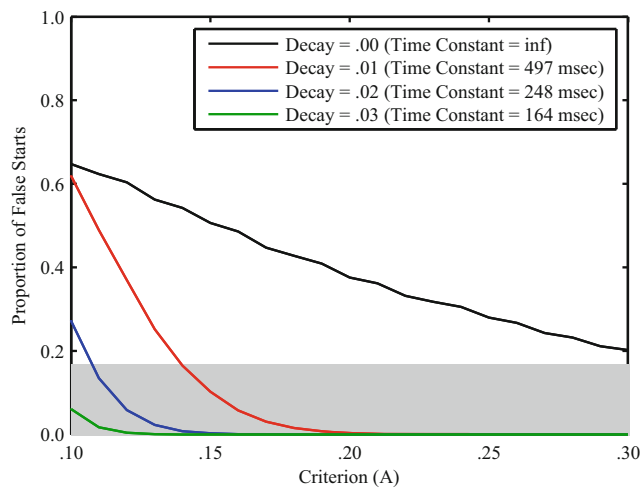


Fig. 2 Proportion of false starts in simulations as a function of leaky competing accumulator model (LCA) decay rate, λ , and diffusion model decision criterion, A . The gray band shows the range of the proportion of false starts committed by participants during the third day of total sleep deprivation (see Experiment 1 below)

of the PVT is able to produce false starts within the range of what is typically seen (Doran, Van Dongen, & Dinges, 2001).

The task model was implemented using a random walk approximation of the diffusion process with a step size $\tau = 5$ ms (Tuerlinckx, Maris, Ratcliff, & De Boeck, 2001). At each step j , a displacement Δ_j occurs with probability p , and a displacement $-\Delta_j$ occurs with probability $1-p$. The size of the displacement is determined by τ and the within-trial stochastic component of the accumulation process $s = 0.1$ (Tuerlinckx et al., 2001)³:

$$\Delta_j = s\sqrt{\tau}. \quad (1)$$

The probability of the positive displacement is given by:

$$p = \left(1 + \frac{V\sqrt{\tau}}{s}\right) / 2. \quad (2)$$

where V corresponds to drift rate. A within-trial stochastic component of the process is instantiated in Eq. 2, and a between-trial stochastic component is produced by drawing a value v from a normal distribution with a mean of V and a standard deviation of η for each trial.

To implement decay, we modified the displacement term in (1), as follows⁴:

$$\Delta_j = s\sqrt{\tau} - \lambda E_j, \quad (3)$$

³ Taking the square root of τ ensures that v meaningfully corresponds to drift rate in the continuous diffusion model and in a way not dependent on time-step size in its discrete approximation (Tuerlinckx et al., 2001).

⁴ For related expressions, see Goldman et al. (2009) and Usher and McClelland (2001).

where E_j is the evidence accumulated to that time in the diffusion process:

$$E_j = \sum_{k=1}^{j-1} \Delta_k, \quad (4)$$

with $E_0 = 0$. The first step of the diffusion process, $j = 1$, comes immediately after the offset of feedback from the preceding trial.

Biomathematical model of fatigue

We combined the extended diffusion model of the PVT with the biomathematical model of fatigue described by McCauley et al. (2013). This biomathematical model is based on the two-process model of sleep regulation (Borbély & Achermann, 1999), in which a homeostatic process increases the drive for sleep with time awake, and a circadian process promotes wakefulness during the day and sleep at night. The homeostatic and circadian processes interact in a non-linear manner such that the amplitude of the circadian process increases in sleep-deprived individuals (McCauley et al., 2013). A third process adjusts the set point of the homeostatic process in response to long-term sleep/wake history. The biomathematical model is sensitive to fatigue due to total sleep deprivation, circadian misalignment during shift work, and sustained sleep restriction (Fig. 3).⁵

Integration

We integrated the biomathematical model of fatigue with the diffusion model of the PVT. Ratcliff and Van Dongen (2011) showed that the dynamic changes in the response time distribution on the PVT across 36 h of total sleep deprivation were captured by varying the drift ratio V/η . Furthermore, they found a high correlation between the fitted values of V/η and the output of an earlier version of the biomathematical model of fatigue (McCauley et al., 2009). We captured this dynamic by fixing η across sessions and treating V as a linear function of the predictions of the biomathematical model of fatigue, F :

$$V = a_V F + b_V, \quad (5)$$

where a_V and b_V represent slope and intercept. When the slope term a_V is negative, drift rate decreases with fatigue, producing longer and more variable response times.

Dynamic changes in false starts could result from temporal changes in the decision criterion, A , or in the decay rate of the LCA process, λ . Ratcliff and Van Dongen

⁵ For a complete description of the model and its equations, see Table 1 of McCauley et al. (2013). A DOS executable of the model is available at <https://pounce.spokane.wsu.edu/data/public/9d87c1>

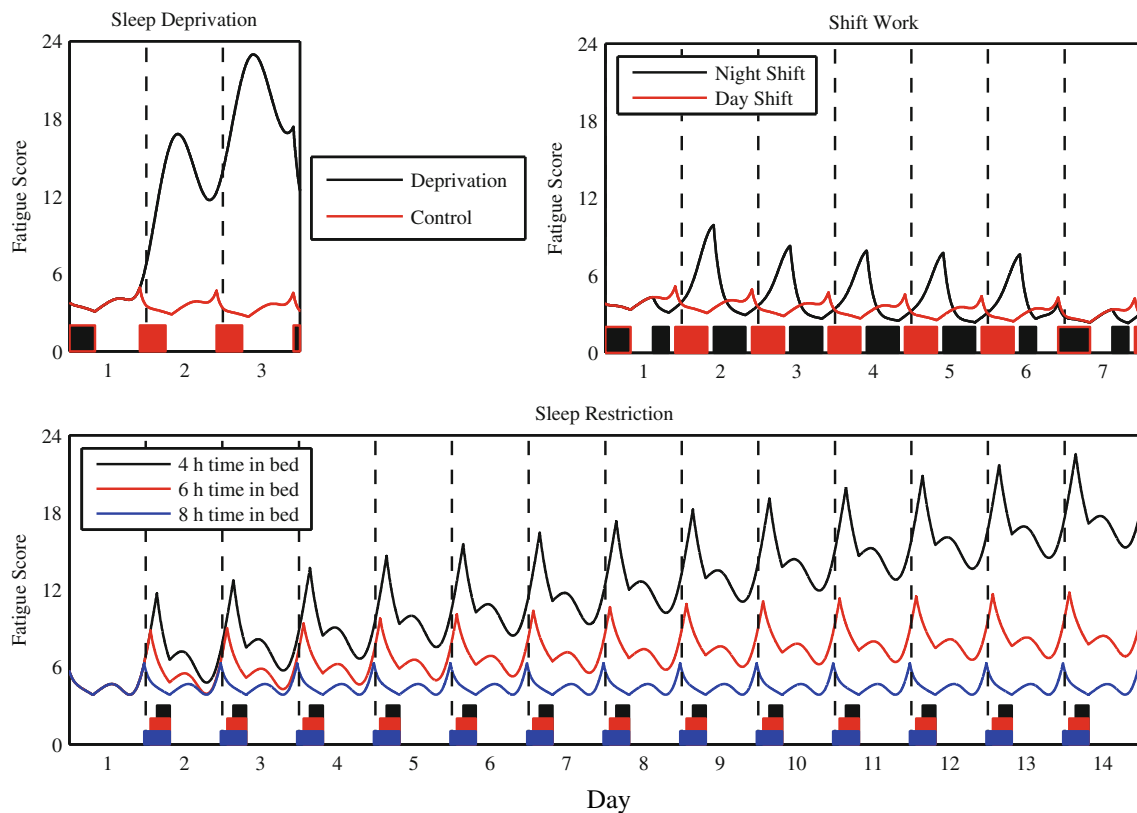


Fig. 3 Biomathematical model predictions of fatigue. **Top-left panel:** Predicted fatigue across 3 days of total sleep deprivation versus a control condition with daily time in bed (TIB) from 22:00 until 08:00. **Top-right panel:** Predicted fatigue across 5 days of a shift schedule with daily TIB from 22:00 until 08:00 (day shift) versus a shift schedule with daily TIB from 10:00 until 20:00 beginning on day 2 (night shift).

Bottom panel: Predicted fatigue across 14 days for schedules with 4 h, 6 h, or 8 h TIB daily, with TIB ending at 07:30 for all three conditions. Colored bars at the base of each panel show TIB for the corresponding conditions. In this model (McCauley et al., 2013), Fatigue Score is the predicted number of lapses in a 10-min Psychomotor Vigilance Test (PVT)

(2011) reported that A did not change significantly in response to sleep deprivation. We therefore fixed A , and treated λ as a linear function of the predictions of the biomathematical model of fatigue, F :

$$\lambda = a_\lambda F + b_\lambda, \tag{6}$$

where a_λ and b_λ represent slope and intercept.⁶ When the slope term a_λ is negative, the decay rate decreases with fatigue. This partially offsets the effect of fatigue on drift rate. However, this also offsets the dampening effect of decay during the pre-stimulus interval, which potentiates false starts.

In total, the extended, integrated diffusion model contained eight free, participant-specific parameters (Table 1). Ratcliff and Van Dongen (2011) reported that three diffusion model parameters (V , η , and A) were not uniquely determinable for the PVT. Rather, the two ratios A/V and V/η were approximately constant over different

model fits. In our extended model, the interplay between decay (λ) and drift rate make the values of all three parameters consequential. In our simulations, A and η were estimated for individuals and held constant across sessions, while V was estimated for individuals and allowed to vary across sessions by Eq. 5.

Table 1 Diffusion model parameters

Parameter	Definition
a_V	Drift rate scaling slope
b_V	Drift rate scaling intercept
a_λ	Decay rate scaling slope
b_λ	Decay rate scaling intercept
η	Drift rate variability
A	Decision criterion
T_{ND}	Mean of non-decision time
SD	Spread of non-decision time

Note. The first and second parameters (a_V , b_V) scale the effects of fatigue on drift rate, and the third and fourth parameters (a_λ , b_λ) scale the effects of fatigue on decay. The remaining parameters and corresponding processes are not involved in capturing fatigue in the model

⁶ In a later set of reported simulations, we examined the results of instead allowing A to vary with fatigue while holding λ constant. Although the fits were nearly as good, they were slightly better when λ was allowed to vary with fatigue.

Adaptive Control of Thought – Rational (ACT-R)

The cognitive architecture ACT-R (Anderson, 2007) contains a set of specialized information processing modules (Fig. 4). These include a visual module for locating and identifying objects in the visual field, a manual module for producing motor responses, a declarative module for storing and retrieving information in memory, an imaginal module for holding current problem representations, a goal module for maintaining information about context and intent, and a procedural module for coordinating other modules' behavior. Buffers allow information and commands to pass from the specialized information processing modules to the central procedural module and back.

In ACT-R, procedural knowledge is represented in the form of production rules. Each rule has a set of conditions that must be met for it to be selected, and a set of actions that modify the external world and the internal state of the architecture. Each rule also has real-valued utility, corresponding to the reward the rule is expected to lead to with respect to task performance and completion. Cognitive performance unfolds across a sequence of production cycles lasting on the order of tens of milliseconds. During each cycle, conditions for different productions are compared against the contents of the buffers, and the production with the highest utility is selected and enacted. The resulting state of the world and architecture, represented by the contents of the buffers, serves as the starting point for the next production cycle.

ACT-R has been used to model cognitive performance on dozens of laboratory tasks (see Anderson, 2007) and to

simulate complex skills such as air traffic control, algebra problem solving, and driving (Anderson et al., 2004; Salvucci, 2006). In the context of fatigue research, ACT-R has been applied to dual-tasking (Gunzelmann et al., 2009b), arithmetic retrieval (Gunzelmann et al., 2012), driving (Gunzelmann et al., 2011), flying (Gunzelmann & Gluck, 2009), and the PVT (Gunzelmann et al., 2009a). Like the diffusion model, ACT-R is particularly suitable for studying the effects of fatigue on the PVT because it predicts complete response time distributions, and because its parameters can be manipulated to produce continuous performance decrements due to fatigue.

Task model

Gunzelmann et al. (2009a) developed an ACT-R model of the PVT. The model contains four productions: (1) *wait* for the stimulus to appear, (2) *attend* to the stimulus, (3) *respond*, and (4) *press key*. Each production is eligible for selection in certain conditions. *Wait* is eligible when nothing is present on the screen, *attend* is eligible when the stimulus is present on the screen but has not yet been attended, *respond* is eligible when the stimulus is present on the screen and has been attended, and *press key* is always eligible. Thus, during each production cycle a choice is made between one of the first three productions, and the *press key* production. Logistically distributed noise is added to production utilities (U_i),

$$U'_i = U_i + \varepsilon_i, \quad (7)$$

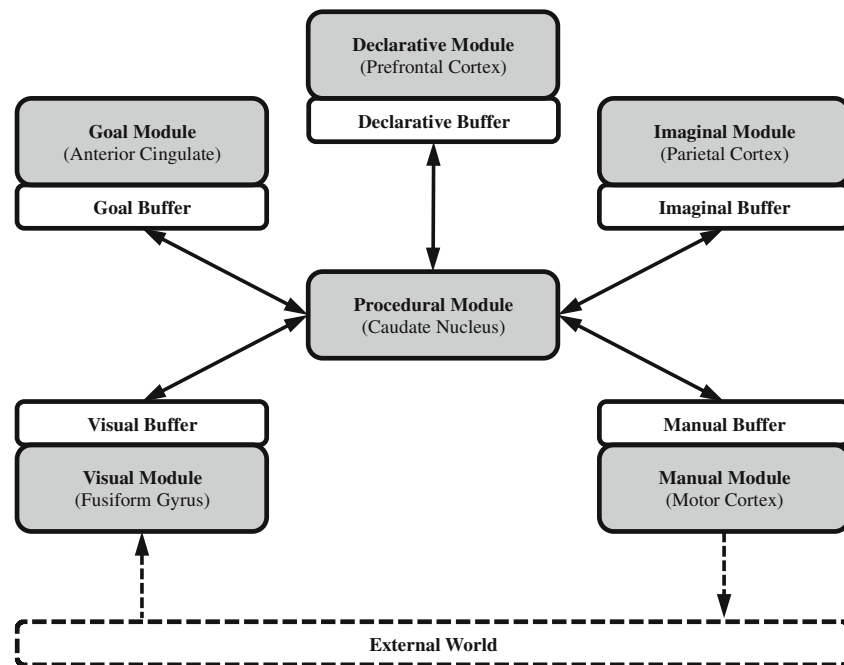


Fig. 4 Schematic representation of the modules and buffers of the Adaptive Control of Thought – Rational (ACT-R) cognitive architecture and their associations with brain regions

and the production with the greatest resulting utility value U'_i is selected. The occasional selection of *press key* before the stimulus appears permits false starts.

The production with the greatest value is enacted if its utility exceeds the threshold,

$$\text{Production} = \max(U'_i); \text{enacted if } \max(U'_i) > T. \quad (8)$$

When no production's utility exceeds the threshold, the model becomes briefly inactive before initiating the next production cycle – a microlapse. The period of inactivity equals the duration of one production cycle (i.e., tens of milliseconds).

Gunzelmann et al. (2009a) fitted this ACT-R model to data from participants who completed the PVT every 2 h while staying awake for 88 h. They found that sleep deprivation affected production utilities and the utility threshold; as time awake increased, production utilities and the utility threshold decreased. Changes in the duration of perceptual and motor processes, production execution time, and production noise were not needed to account for the effects of fatigue.

Task model update

Subsequent to the development of the ACT-R model of the PVT, a mechanism called *production partial matching* was introduced in ACT-R. With production partial matching, productions whose conditions are not perfectly met are eligible for selection, but their utility values are penalized:

$$U'_i = (U_i - MMP_i) + \varepsilon. \quad (9)$$

MMP_i is the mismatch penalty charged if the conditions for the production are not perfectly met. The production with the greatest value U'_i is selected and enacted provided its utility exceeds the threshold (Eq. 8). This is true even when the production with highest utility does not perfectly match the conditions, but exceeds the threshold.

The addition of production partial matching to ACT-R obviated the need for the *press key* production in the PVT model of Gunzelmann et al. (2009a); with this addition, *respond* can be selected at any time. Consequently, the updated model we advance contains a total of three productions: (1) *wait*, (2) *attend*, and (3) *respond* to the stimulus. When *respond* is selected before the stimulus appears, a false start occurs. However, because the utility of *respond* is penalized before the stimulus appears, this happens infrequently.

Productions' baseline utilities (U_i) were treated as a single free parameter – one value of U was estimated and used for all productions. U could be acquired from experience using ACT-R's procedural learning equation (Anderson, 2007; Fu & Anderson, 2006), but we disabled utility learning because learning effects on the PVT are negligible (Van Dongen et al., 2003). To simplify matters, the mismatch penalty (MMP , Eq. 9)

was set to the value of production utility. The resulting payoff matrix was symmetric with zero assigned to mismatches and U assigned to matches (see Eq. 9).⁷ The mean and standard deviation of the logistically distributed noise added to these values during each production cycle were set to the default ACT-R values of 0.0 and 0.453 (Anderson et al., 2004).

Integration

We integrated the biomathematical model of fatigue (McCauley et al., 2013) with the updated ACT-R model of the PVT. In the integrated model, the effects of fatigue play out through three component interactions. First, fatigue reduces productions' utility values (Fig. 5):

$$U''_i = FP \cdot (U_i - MMP_i) + \varepsilon_i, \quad (10)$$

where FP (i.e., *Fatigue in Procedural Knowledge*) is a linear function of the predictions of the biomathematical model of fatigue, F :

$$FP = a_p F + 1. \quad (11)$$

The parameter a_p is a slope parameter. In the absence of fatigue, FP equals one (i.e., utilities are unchanged). When the slope parameter a_p is negative, production utilities decrease with fatigue.⁸ Consequently, selections are increasingly driven by noise. Also, productions increasingly fall below the utility threshold, causing more microlapses.

Second, fatigue lowers the utility threshold, T (Fig. 5):

$$T' = FT \cdot T. \quad (12)$$

FT (i.e., *Fatigued Threshold*) is a linear function of the predictions of the biomathematical model of fatigue:

$$FT = a_T F + 1. \quad (13)$$

The parameter a_T is a slope parameter. In the absence of fatigue ($F = 0$), FT equals one. When the slope term a_T is negative, the utility threshold decreases with fatigue. This partially compensates for the effect of fatigue on utility values. However, this also reduces the inhibitive effect of the mismatch penalty on the *respond* production, allowing more false starts.

⁷ The mismatch penalty between the conditions for the *wait* and *respond* productions could conceivably be set to a different value than the mismatch penalty between the conditions for the *attend* and *respond* productions. However, such adjustments were not necessary to capture the results of the experiments modeled here.

⁸ For large negative values of a_p , FP could conceivably become negative, and U'' could become negative as well. Based on estimates of a_p from the reported simulations, however, FP did not approach negative values even after 62 h of total sleep deprivation. Future work could explore bounding FP at zero, or using a non-linear function to map the output of the biomathematical model to cognitive model parameter estimates to eliminate this possibility.




Fig. 5 In the integrated Adaptive Control of Thought – Rational (ACT-R) model of the Psychomotor Vigilance Test (PVT), fatigue reduces utility values (distributions), and to a lesser extent the utility threshold (vertical lines)

Third, when no production's utility exceeds the threshold and a microlapse occurs, FP is decreased by a small amount:

$$FP \leftarrow FP \cdot FP_{dec}, \quad (14)$$

where $0 < FP_{dec} < 1$. This small change makes it more likely that microlapses will occur in subsequent production cycles. Across such a series of cycles, the probability of responding progressively decreases, leading to behavioral lapses. The value of FP is restored when a stimulus next appears.

In total, the updated, integrated ACT-R model contained six free, participant-specific parameters (Table 2). The one parameter not yet discussed, *cycle time*, controls the amount of time to evaluate and select a production during each production cycle, and has a default value of 50 ms (Anderson, 2007). We allowed cycle time to vary across individuals, consistent with the notion that this parameter reflects stable differences in processing speed (Deary, Der, & Ford, 2001; Larson & Alderton, 1990). The interplay between these parameters and the biomathematical model of fatigue account for the complete RT distribution in the PVT, including false starts and lapses (Supplementary Fig. 1). Although the durations of events in the ACT-R model are on the order of tens of milliseconds, the summation of time across events and the

Table 2 ACT-R parameters

Parameter	Definition
U	Utility
a_P	Utility scaling slope
T	Threshold
a_T	Threshold scaling slope
FP_{dec}	Utility decrement after lapse
$Cycle$	Cycle time

Note. The second parameter (a_P) scales the effects of fatigue on production utilities, and the third parameter (a_T) scales the effects of fatigue on the utility threshold. The remaining parameters and corresponding processes are not involved in capturing fatigue in the model

millisecond-level variability in event durations produces continuous reaction-time distributions.

Experiments and results

We investigated whether the integrated diffusion and ACT-R models could each account for the effects of fatigue stemming from total sleep deprivation, simulated shift work, and sustained sleep restriction on PVT performance. We compared simulations to observations from three experiments (Van Dongen, Belenky, & Vila, 2011; Van Dongen, Maislin, Mullington, & Dinges, 2003; Whitney, Hinson, Jackson, & Van Dongen, 2015).

Fitting procedure

In each of the experiments reported below, we collapsed data across multiple 10-min PVT sessions to form probability density functions. We binned response times corresponding to the 0.1-interval quantiles of responses that occurred after 150 ms for each participant. The 0.1 quantile contained the fastest 10% of responses after 150 ms, and the 1.0 quantile contained the slowest 10% of responses after 150 ms (including the few trials with no response after 30 s). We then calculated the overall proportion of responses that occurred before 150 ms (i.e., false starts), and the overall proportions of responses within the 10 quantiles.⁹ This provided 11 proportion values against which to compare the models for each participant.

We used the models to create corresponding expected probability density functions. To do so, we simulated participants' performance during each 10-min PVT session, and collapsed predictions across sessions as was done with the observations. We then calculated the proportion of predicted responses that occurred before 150 ms, and the proportions of responses within the 10 quantiles defined by boundaries derived from observed response times. Predictions were based on 1,000 simulations for each PVT session.

The observed and expected probability density functions were used to compute the likelihood ratio chi-square (G^2), which is asymptotically equivalent to the chi-square,

$$G^2 = 2 \sum_{i=1:day} \sum_{j=1:bin} N_{ij} \cdot \log \left(\frac{p_{ij}}{\pi_{ij}} \right). \quad (15)$$

N_{ij} is the observed number of responses in the j^{th} bin on the i^{th} day, p_{ij} is the predicted proportion of responses in that bin for that day, and π_{ij} is the observed proportion of responses (Smith & Ratcliff, 2009). A simplex search algorithm with

⁹ Each quantile contained 10% of responses that occurred after 150 ms of stimulus onset. This is slightly less than 10% of all responses, which also include false starts (responses that occurred before or within 150 ms of stimulus onset).

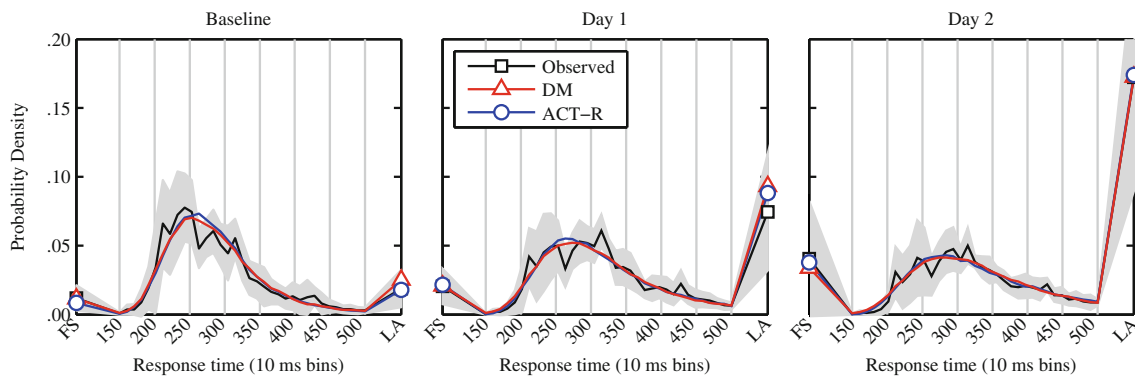


Fig. 6 Psychomotor Vigilance Test (PVT) response time distributions across 62 h of total sleep deprivation, aggregated over each day of Experiment 1. The first bin shows the proportion of false starts (FS), the final bin shows the proportion of lapses (LA), and the middle bins show the proportion of responses occurring in 10 ms intervals from 150–

500 ms. The gray area shows ± 1 SD around the mean for the observations. The red and blue curves show the predictions of the diffusion model (DM) and the Adaptive Control of Thought – Rational (ACT-R) model, respectively

multiple start points was used to find parameter values that minimized G^2 for each participant and model. Simulations were conducted using large-scale computational resources (Harris, 2008).¹⁰ The supplementary material contains additional information about model fitting procedures.

We used two criteria to assess model fit: the G^2 statistic and the Bayesian Information Criterion (BIC). The BIC is calculated from the G^2 statistic according to

$$BIC = G^2 + m \cdot \log(n), \quad (16)$$

where m is the number of free parameters and n is the total number of observations per participant aggregated across all PVT sessions.

Experiment 1: Acute total sleep deprivation

The first experiment involved a 62-h total sleep deprivation condition and a well-rested control condition in a laboratory (Whitney et al., 2015). Participants in the sleep deprivation condition ($n = 13$) remained awake for 62 h, starting at 08:00 after two baseline days, whereas participants in the control condition ($n = 13$) received 10 h time in bed (TIB; 22:00–08:00) each night (Fig. 3, top left). Participants performed the PVT approximately once every 2 h during scheduled wakefulness. Because each session contained relatively few response time observations (mean \pm SD: 94 ± 8), we collapsed data across sessions that occurred within each 24-h period of the experiment (22:00–22:00) for each participant. This yielded aggregate data sets for day 0 (0–15 h awake, baseline), day 1 (16–39 h awake in the sleep deprivation condition), and day 2 (40–62 h awake in the sleep deprivation condition).

Figure 6 shows the average proportions of responses that occurred before or within 150 ms of stimulus presentation (i.e.,

false starts), from 150–500 ms after stimulus presentation (i.e., *alert responses*), and more than 500 ms after stimulus presentation (i.e., *lapses*) during each day of the experiment. As time awake increased, participants in the sleep deprivation condition responded more slowly, committed more false starts, and experienced more lapses (see Supplementary Fig. 2 for individuals' data). None of these effects appeared for participants in the control condition (see Supplementary Fig. 3 for individuals' data).

Figure 6 also shows the fits of the diffusion and ACT-R models to the observations. Fits for individual participants are shown in Supplementary Figs. 2 and 3, and cumulative distributions based on quantile response times are shown in Supplementary Figs. 44 and 5. Table 3 contains measures of model fit to the quantile response times of the individuals in the sleep deprivation and control conditions, and Tables 4 and 5 contain the parameter estimates. The G^2 statistic was lower for the diffusion model in both conditions; however, this measure does not take model complexity into account. The BIC, which does take model complexity into account, favored the ACT-R model in the control condition and the diffusion model in the sleep deprivation condition, but the absolute differences between the model fits were small. Comparison of the models' outputs to one another reinforces this conclusion (Supplementary Table 3).

The diffusion model closely matched the observed response time distributions (Fig. 6). The sum of squared errors for each participant ranged from .007 to .017 with a mean (\pm SE) of .010 (\pm .001). Correlations between the predicted and observed proportions of responses in 10-ms bins for each participant ranged from $r = .90$ to .97 (Supplementary Fig. 2). As time awake increased, the diffusion model responded more slowly and produced more lapses. This occurred because of changes in the rate of evidence accumulation. The drift rate scaling slope (a_V) was significantly less than zero, $t(12) = 16.86$, $p < .001$ (Table 4, Experiment 1). Consequently, the response time distribution shifted to the right and the distribution became more skewed with increasing time awake.

¹⁰ See <http://www.mindmodeling.org>

Table 3 Average model fits to individuals in Experiments 1 and 2 (± 1 standard error across individuals)

	Model	Model df	Experiment df	G^2	Observations	BIC
Experiment 1:	Diffusion	8	29 \pm .1	120 \pm 11	2623 \pm 15	182 \pm 11
Sleep deprivation	ACT-R	6		137 \pm 9		184 \pm 9
Experiment 1:	Diffusion	8	28 \pm .3	89 \pm 8	1698 \pm 3	148 \pm 8
Control	ACT-R	6		94 \pm 11		138 \pm 11
Experiment 2:	Diffusion	8	78 \pm 1	431 \pm 33	7700 \pm 21	502 \pm 33
Night shift	ACT-R	6		451 \pm 52		505 \pm 52
Experiment 2:	Diffusion	8	79 \pm .6	366 \pm 25	7692 \pm 31	437 \pm 25
Control	ACT-R	6		394 \pm 27		447 \pm 26

The diffusion model also committed more false starts across days. This occurred because of changes in stability during the pre-stimulus interval. The decay rate scaling slope (a_λ) was significantly less than zero, $t(12) = 3.80$, $p < .01$. Reduced decay rate partially compensated for the effect of fatigue on drift rate once the stimulus appeared. However, reduced decay rate also inadvertently allowed noise to drive the diffusion process beyond the decision criterion during the pre-stimulus interval, causing more false starts. The average value of the decay rate scaling slope ($-.0010$; Table 4) is small because the term is multiplied by relatively large values from the biomathematical model (up to 24; Fig. 3), and because small changes in decay have a large impact on performance. With an average value of $-.0010$ for reduced decay, the integration time constant changes from 264 ms at baseline to 192 ms after 3 days of total sleep deprivation.

The ACT-R model closely matched the observed response time distributions as well (Fig. 6). The sum of squared errors for each participant ranged from .008 to .014 with a mean (\pm SE) of .011 (\pm .001). Correlations between the observed and predicted response time distributions in 10-ms bins for each participant ranged from $r = .91$ to .96 (Supplementary Fig. 2). The ACT-R model responded more slowly and generated more lapses with increasing time awake. This was because of the rising frequency of microlapses. The utility intercept

Table 4 Diffusion model parameters estimates for individual participants (± 1 standard error across individuals) for Experiments 1 and 2

Parameter	Experiment 1: Sleep deprivation	Experiment 2: Shift work
a_V	$-.024 \pm .001$	$-.027 \pm .003$
b_V	$1.35 \pm .03$	$1.34 \pm .02$
a_λ	$-.0010 \pm .0001$	$-.0016 \pm .0001$
b_λ	$.040 \pm .002$	$.043 \pm .002$
η	$.23 \pm .01$	$.21 \pm .01$
A	$.101 \pm .002$	$.107 \pm .003$
T_{ND}	$.146 \pm .006$	$.136 \pm .005$
SD	$.101 \pm .007$	$.077 \pm .006$

was greater than the threshold intercept, $t(12) = 6.98$, $p < .001$ (Table 5; Experiment 1). Consequently, production utilities predominantly exceeded the utility threshold during early sessions, minimizing microlapses. However, the utility scaling slope (a_P) was more negative than the threshold scaling slope (a_T), $t(12) = 11.93$, $p < .001$. Consequently, production utilities fell below the threshold during later sessions with increasing probability, resulting in more microlapses. Microlapses slowed alert responses, and sometimes delayed responses beyond 500 ms, causing lapses.

The ACT-R model also committed more false starts across days. This occurred because of changes in the threshold. The threshold scaling slope (a_T) was significantly less than zero, $t(12) = 7.02$, $p < .001$. Reducing the utility threshold partially compensated for the effect of fatigue on production utilities. However, lowering the threshold also reduced the influence of the mismatch penalty on the *respond* production, leading to more false starts.

Analyzing the response time distributions by day emphasizes the homeostatic process of sleep regulation; performance declines across days of sleep deprivation. Performance also varies within days in accordance with the circadian process. To examine the dynamics across time of day, we calculated the proportions of false starts and lapses and the median times of alert responses for each session during the 3 days (Fig. 7). Due to the interaction between the homeostatic and circadian processes, participants committed more false starts and lapses and

Table 5 ACT-R parameter estimates for individual participants (± 1 standard error across individuals) for Experiments 1 and 2

Parameter	Experiment 1: Sleep deprivation	Experiment 2: Shift work
U	$5.07 \pm .05$	$4.98 \pm .11$
a_P	$-.014 \pm .001$	$-.022 \pm .001$
T	$4.59 \pm .06$	$4.51 \pm .07$
a_T	$-.010 \pm .001$	$-.015 \pm .001$
FP_{dec}	$.984 \pm .003$	$.983 \pm .006$
$Cycle$	$.041 \pm .003$	$.034 \pm .002$

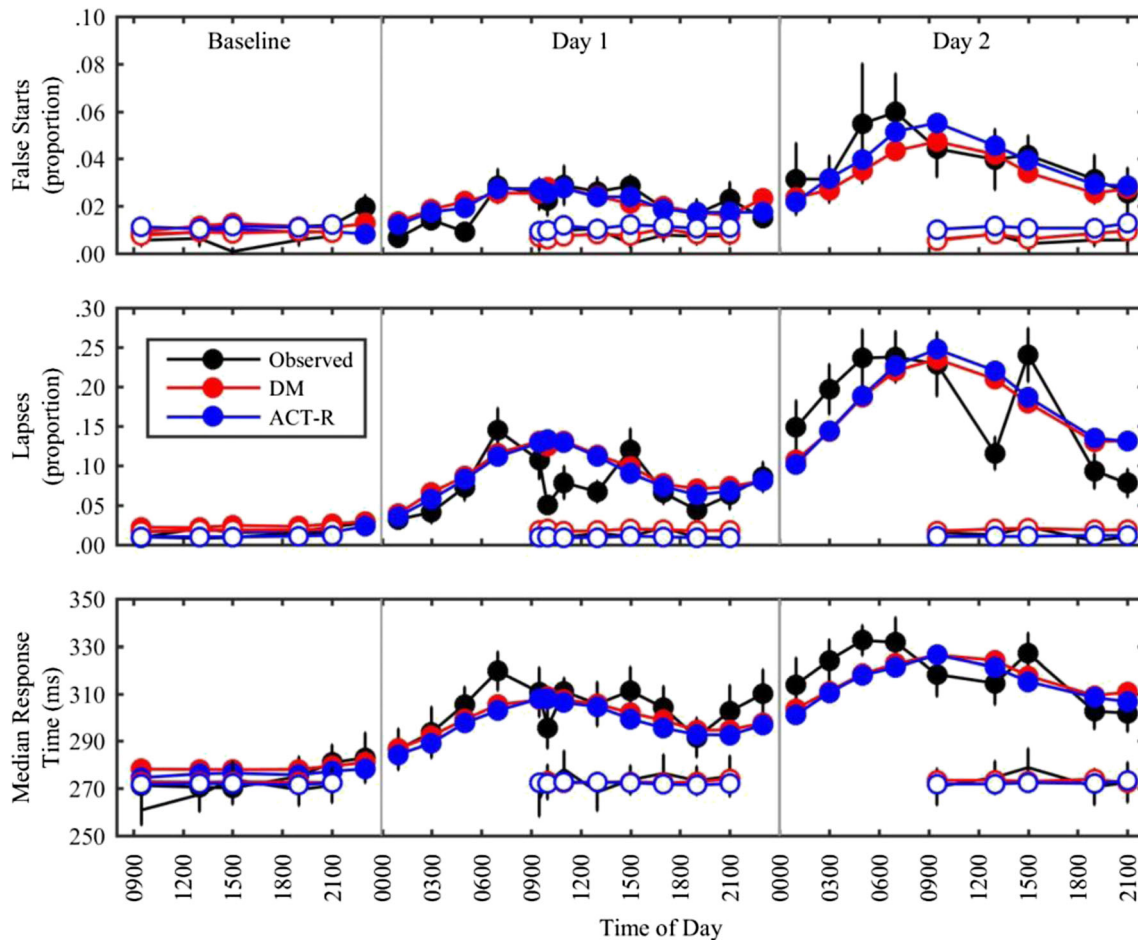


Fig. 7 Proportion of false starts (**top panel**), proportion of lapses (**middle panel**), and median times of alert responses (**bottom panel**) for each session in Experiment 1. Filled shapes correspond to the 62-h

total sleep deprivation condition; open shapes correspond to the control condition (for which there was no testing during the nights). Error bars indicate ± 1 standard error

they responded more slowly during the early morning hours. As instantiated in the biomathematical model of McCauley et al. (2013), the circadian process interacted nonlinearly with the homeostatic process, such that time-of-day effects were greater in the sleep deprivation condition than in the control condition.

The diffusion and ACT-R models reproduced the effects of the circadian process and its interaction with the homeostatic process. Fits to the proportion of false starts (diffusion model: $r = .92$ ACT-R: $r = .92$), lapses (diffusion model: $r = .91$; ACT-R: $r = .91$) and median response times (diffusion model: $r = .95$; ACT-R: $r = .95$) were close to the observations for both models.

Experiment 2: Simulated night shift work

Night shift work is associated with increased fatigue and deficits in cognitive performance due to circadian misalignment (Åkerstedt, 1988; Van Dongen, 2006). The second experiment we modeled involved a simulated night shift condition and a control condition in a laboratory (Van Dongen et al., 2011c).

Participants in the night shift condition ($n = 13$) completed two 5-day night-time duty cycles with duty time spanning from 20:00 until 10:00 (Fig. 3, top right). The two duty cycles were separated by a 34-h break from the primary task of driving a high-fidelity driving simulator.¹¹ The 34-h break included a 5-h nap opportunity from 10:00 until 15:00, a night of sleep from 22:00 until 08:00, and another 5-h nap opportunity from 15:00 until 20:00. Participants in the control condition ($n = 14$) completed two 5-day daytime duty cycles with duty time spanning from 08:00 until 22:00. The two duty cycles were separated by a 34-h break from the primary task of driving as well, but the sleep schedule was unaltered.

Participants performed the PVT eight times per duty day. Because each session contained relatively few observations (mean \pm SD: 96 ± 1), we combined data from sessions that occurred at the same time of day across duty days and duty cycles. In the control condition, sessions occurred at 09:05,

¹¹ Figure 3 only shows the output of the biomathematical model through the first duty cycle and break. The output for the second duty cycle is nearly identical.

09:55, 12:05, 12:55, 15:05, 15:55, 18:05, and 18:55. In the night shift condition, session times were offset from these times by 12 h. The fitting procedure and evaluation metrics were identical to those used in Experiment 1.

In Experiment 2, performance varied primarily by time of day (Van Dongen et al., 2011c). Performance remained relatively constant across time of day in the control condition, but degraded across duty hours in the night shift condition (Fig. 8). The number of false starts increased significantly, albeit slightly, from the initial to the final testing session in the night shift condition, but not in the control condition. The number of lapses and the median response times also increased across time of day in the night shift condition, but not in the control condition.

Table 3 contains measures of model fit to the quantile response times from the daytime and night-time conditions. The G^2 statistic was slightly lower for the diffusion model in both conditions, as was the BIC. However, model fits were skewed by the poor correspondence between their outputs and the data of one participant (Participant 13, Supplementary Fig. 6). Neither model was fully able to account for the peakedness of the participant's RT distributions. Excluding this participant reduced the average G^2 values (DM = 439; ACT-R = 422), and the resulting BIC scores (DM = 491; ACT-R = 451).

Model parameter estimates for Experiment 2 were similar to those for Experiment 1 (Tables 4 and 5), indicating generality of the models across different experimental manipulations inducing fatigue. We used the best fitting parameter estimates to generate response time distributions for the diffusion and ACT-R models for each participant (Supplementary Figs. 6 and 7). We calculated the expected proportions of false starts and lapses, and the median times of alert responses during each of the scheduled testing sessions. Both models predicted an effect of duty hour in the nighttime condition only (Fig. 8), characterized by progressively worse performance throughout the night shift. Fits to the proportions of false starts (diffusion model: $r = .85$; ACT-R: $r = .82$), lapses (diffusion model: $r = .89$; ACT-R: $r = .66$) and median response times (diffusion model: $r = .92$; ACT-R: $r = .93$) were comparable between the models.

Experiment 3: Sustained sleep restriction

Sleep restriction, when sustained across multiple nights, results in cumulative deficits in cognitive performance (Belenky et al., 2003; Van Dongen et al., 2003). The homeostatic process in the McCauley et al. (2013) bi mathematical model tracks cumulative fatigue due to sustained sleep restriction,

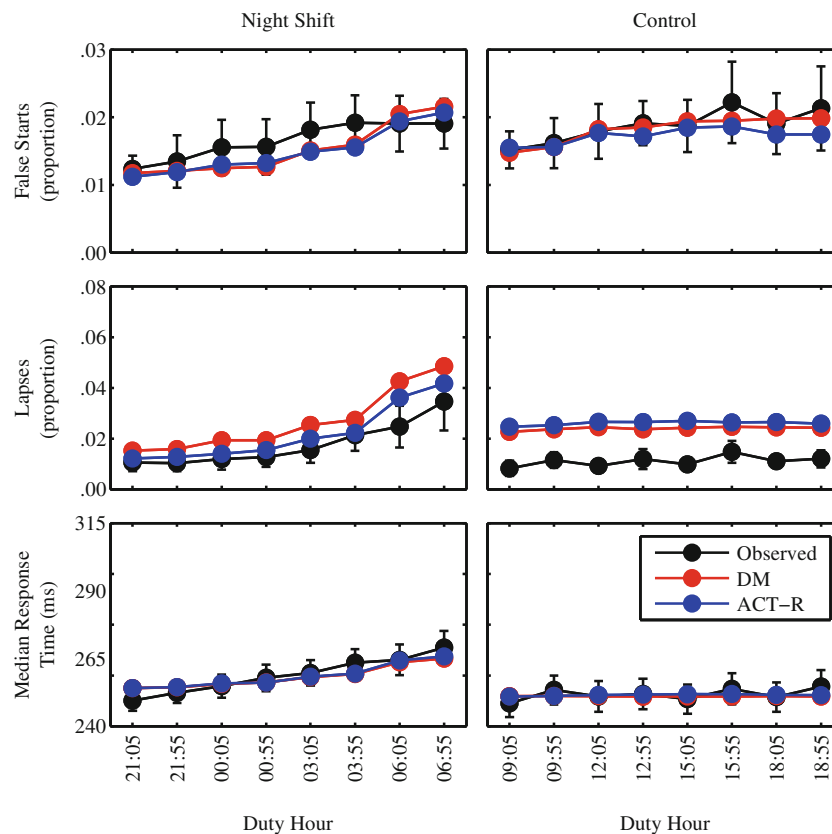


Fig. 8 Proportion of false starts (**top panel**), proportion of lapses (**middle panel**), and median times of alert responses (**bottom panel**) for each session in Experiment 2. Black circles show observations, red

circles show diffusion model predictions, and blue circles show Adaptive Control of Thought – Rational (ACT-R) predictions. Error bars indicate ± 1 standard error

and can thus capture the resulting deficits. The third experiment we modeled involved restricting sleep to 4, 6, or 8 h TIB each night over 14 days (Van Dongen et al., 2003). The experiment began after 3 days of baseline adaptation, and continued for 14 consecutive days (Fig. 3, bottom).

To examine whether the integrated diffusion and ACT-R models extend to sustained sleep restriction, we generated predictions using the parameter sets recovered from the total sleep deprivation condition of Experiment 1. We followed this approach because the published data on sustained sleep restriction only included the number of lapses, which was not adequate to fit the model. We collapsed predictions across PVT sessions that occurred within each 24-h period of Experiment 3, and computed the expected proportions of false starts and lapses, and the median response times, for the three experiment conditions (Fig. 9). In agreement with the results from Van Dongen et al. (2003), both models predicted dose-dependent, mounting impairment across days for participants in the 4- and 6-h conditions, but not for participants in the 8-h condition (see also Belenky et al., 2003).

Alternate parameterizations

In the preceding sections, we considered theoretically-constrained parameterizations of each model. The model fits provided evidence that allowing two parameters (ν and λ in

the diffusion model and *Utility* and *Threshold* in the ACT-R model) to vary with fatigue was sufficient to capture the effects of fatigue on PVT performance. We next examined whether these parameterizations were adequate by fitting variants of the diffusion and ACT-R models where each parameter was allowed to vary alone with fatigue, or all combinations of two parameters were allowed to vary with fatigue.

For this purpose we used the data from the sleep deprivation condition of Experiment 1, where the effects of fatigue were most substantial. We fitted 21 versions of the diffusion model (six where one parameter varied and 15 where all combinations of two parameters varied) and 11 versions of the ACT-R model (four where one parameter varied and seven where all combinations of two parameters varied). We compared fits among diffusion model variants and ACT-R model variants using BIC scores calculated from the G^2 values (Supplementary Tables 4 and 5).

Of the diffusion models, the best fitting variant was the one used in the previous sections, in which drift rate and decay varied with fatigue (BIC = 182). Other variants that performed well included ones where fatigue affected drift rate alone (BIC = 190), drift rate and drift rate variability (BIC = 191), and drift rate and decision criterion (BIC = 193). Of the ACT-R models, the best fitting variant was the one used in the previous sections as well, in which utility and threshold varied with fatigue (BIC = 184). All

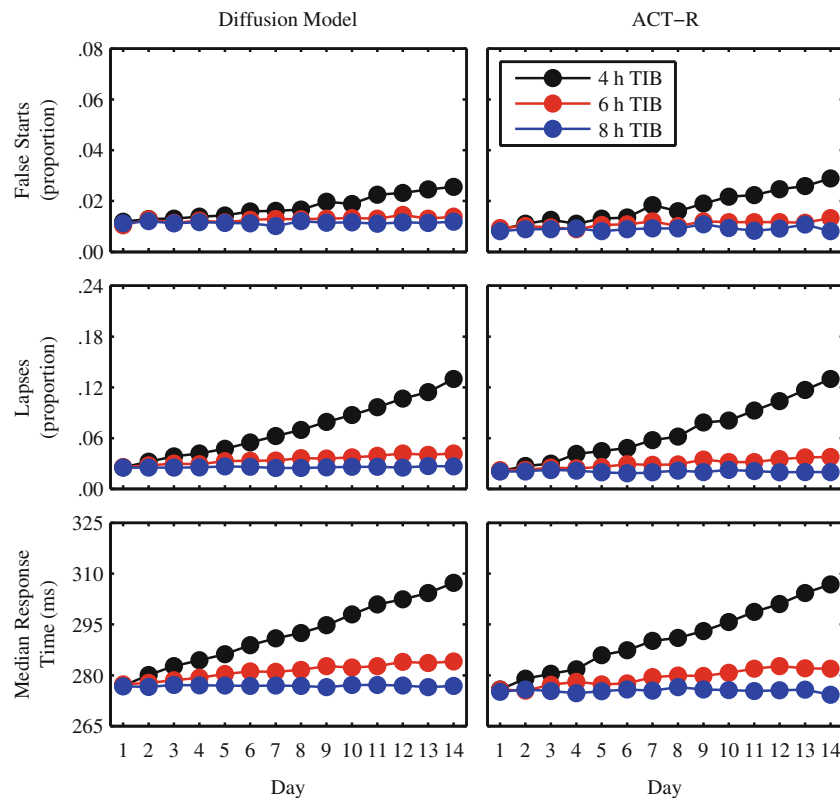


Fig. 9 Proportion of false starts (**top panel**), proportion of lapses (**middle panel**), and median time of alert responses (**bottom panel**) for each day in Experiment 3

other variants fitted substantially more poorly. In the next best-fitting variants, fatigue affected utility and FP_{DEC} (BIC = 226), or utility and cycle time (BIC = 226). In sum, the *a priori* model variants explored in the previous sections, which were based on existing computational theories of fatigue (Gunzelmann et al., 2009a; Ratcliff & Van Dongen, 2011), also provided the best fits to the experimental data.

One might also wonder whether a different implementation of the diffusion model in which RTs arise from the convolution of a distribution from the diffusion process and a separate contaminant distribution based on random guesses better fits the data (Ratcliff & Van Dongen, 2009; Ratcliff & Tuerlinckx, 2002). According to this account, *false starts* could be a consequence of fatigue increasing the probability of random guesses rather than decreasing inhibition. This possibility implies two ancillary assumptions: (1) Guesses are distributed over some interval both preceding and following stimulus onset, and (2) the probability of random guesses increases with fatigue, producing more false starts. We explored this possibility by replacing the LCA process with a random guesses process (see Ratcliff & Van Dongen, 2009). To capture false starts, which occur before the stimulus is presented, we treated the interval of assumption 1 as elapsed time from the offset of feedback from the previous trial to the length of the slowest response after the stimulus – that is, the ITI plus the duration of the slowest response. We fitted the model to data from participants in the total sleep deprivation condition of Experiment 1, and found that the estimated proportion of random guesses increased with fatigue, contributing to the rise in false starts. Still, the goodness-of-fit was lower in the diffusion model with random guesses than in the model with the LCA process on average (BIC = 204 vs. BIC = 182), and for nine of 12 participants. As such, we favored the diffusion model with the LCA process.

Sources of individual differences

Inter-individual differences in PVT performance may be due to baseline performance differences and/or differences in the dynamic changes across time awake and time of day (Van Dongen et al., 2004; Van Dongen, Bender, & Dinges, 2012). We investigated which parameters in the integrated diffusion and ACT-R models accounted for these sources of inter-individual variability.

We first identified parameters that produced inter-individual differences in baseline performance, measured as the proportion of lapses during Day 0 of Experiment 1. Across the total sleep deprivation and control conditions, non-decision time (T_{ND}) from the diffusion model significantly correlated with the proportion of lapses ($r = .61$, $p < .01$), as did cycle time ($Cycle$) from the ACT-R model ($r = .51$, $p < .01$). The estimates of individuals' non-decision and cycle times were correlated ($r = .79$, $p < .001$) (Fig. 10), indicating that these parameters produced similar effects.

We then identified parameters that produced inter-individual differences in vulnerability to fatigue, measured as the increase in lapses and false starts from Day 0 to Day 3 in Experiment 1. Drift rate scaling slope (a_V) from the diffusion model significantly correlated with the increase in lapses ($r = .80$, $p < .001$), as did the utility scaling slope (a_P) from the ACT-R model ($r = .70$, $p < .01$). Across individuals, drift rate and utility slopes were correlated ($r = .63$, $p < .05$) (Fig. 10), indicating that these parameters produced similar effects. Further, the decay rate scaling slope (a_λ) in the diffusion model and the threshold scaling slope (a_T) in the ACT-R model were correlated with the increase in false starts from Day 0 to Day 3 (decay slope: $r = .66$, $p < .05$; threshold slope: $r = .70$, $p < .01$), but not with the increase in lapses. Across individuals, decay and threshold slopes were correlated ($r = .59$, $p < .05$), again indicating that these parameters produced similar effects.

The inter-individual differences in non-decision time (T_{ND}), drift rate scaling slope (a_V) and decay rate scaling slope




Fig. 10 Model parameters in the diffusion model (DM) and the Adaptive Control of Thought – Rational (ACT-R) model capturing inter-individual differences in baseline performance (left) and vulnerability to fatigue (middle and right). Each gray point in the left scatter plots represents

an individual participant from the control condition of Experiment 1, and each number in plots represents an individual from the total sleep deprivation condition, whose data are plotted in Supplementary Fig. 2

(a_λ) of the diffusion model were all moderately interrelated (T_{ND} vs. a_V : $r = .45$; T_{ND} vs. a_λ : $r = .05$; a_V vs. a_λ : $r = .17$). The individual differences in cycle time (*Cycle*), utility scaling slope (a_P) and threshold scaling slope (a_T) of the ACT-R model were also moderately interrelated, with the exception of a_P and a_T , which were strongly interrelated (*Cycle* vs. a_P : $r = .52$; *Cycle* vs. a_T : $r = .42$; a_P vs. a_T : $r = .92$).

Taken together, these observations suggest that individual differences in baseline performance are fundamentally distinct from individual differences in vulnerability to performance impairment from sleep loss (Van Dongen et al., 2004).¹² The relatively weak interrelationships among the individual subjects' diffusion model parameters also suggest that individual differences in degradation of the decision process (after stimulus presentation) and degradation of inhibition (before stimulus presentation) during sleep deprivation are distinct, which may be indicative of different mechanistic pathways. This would not seem to be confirmed by the ACT-R model, in which individual differences in degradation of the decision process and degradation of inhibition are highly correlated. However, strong interrelationships between individual differences in parameter estimates may also be caused by intrinsic correlation among the parameter estimates in the model fitting. An experimental manipulation of the PVT that deliberately dissociates the decision process from the inhibition process will be needed to resolve this issue.

Discussion

Fatigue from sleep loss degrades cognitive performance. The effects of fatigue are especially pronounced for tasks that involve sustained attention, such as the PVT (for others, see e.g., Killgore, 2010). The PVT is commonly used in sleep research because of its high sensitivity to fatigue from sleep loss and circadian rhythms. This sensitivity suggests that understanding the mechanisms associated with performance degradation on the PVT may provide insight regarding the impact of sleep loss and circadian rhythms on cognition more generally.

We used two computational cognitive models to study how fatigue affects cognitive performance: the first is based on the diffusion model (Ratcliff & Van Dongen, 2011), and the second on ACT-R (Gunzelmann et al., 2009a). We integrated each cognitive model with a biomathematical model of fatigue (McCauley et al., 2013). We then investigated the performance of the integrated models across three PVT experiments that measured the effects of fatigue arising from total sleep deprivation, simulated shift work, and sustained sleep restriction. The integrated diffusion and ACT-R models reproduced

three key phenomena in the PVT under conditions of fatigue: increased lapses, increased false starts, and slower alert responses. The two models were also able to account for the complete response time distribution in detail.

We did not reject either model; that was neither the aim nor were we prompted by our findings to do so. Both models provided excellent fits to a wide range of empirical data. They involve fundamentally different levels of abstraction, but we found that they account for the effects of fatigue in surprisingly equivalent ways. This result may have broader implications than would rejecting either model. Juxtaposing the two models and examining why they produce such similar outcomes despite their substantive differences is informative both theoretically and from an applied computational point of view.

Mechanistic effects of fatigue

That the same biomathematical model was used to induce fatigue in the diffusion and ACT-R models did not mean, *prima facie*, that the two integrated accounts would yield converging results. Each account's predictions emerged from interactions among (1) the biomathematical model, (2) the cognitive processing mechanisms instantiated in the accounts, and (3) the manner in which fatigue impacted these mechanisms. Thus, although using different biomathematical models would produce varying results, using the same model does not guarantee convergence. This further depends on how fatigue impacted processing mechanisms in each account, as we discuss next.

With regard to their theoretical basis, the diffusion and ACT-R models are quite distinct. The diffusion model is devoid of specifics about the performance task at hand. The onset of a stimulus drives the continuous accumulation of evidence toward a decision criterion. The ACT-R model, in contrast, contains a set of cognitive processes thought to be involved in performing a given task. The onset of a stimulus alters which processes occur and in what order.

The diffusion and ACT-R models are also quite distinct in terms of their computational implementations. The models differ in whether they treat the decision to respond as a unitary or repeated event. Decision time in the diffusion model is determined by the duration of the diffusion process, which is implemented as one ongoing process. Slow responses arise from low values of the drift rate, which prolong decision time. In contrast, decision time in the ACT-R model corresponds to aggregate time across multiple, short-duration production cycles preceding selection of the *attend* and *respond* productions. The duration of one production cycle is relatively brief (between 30 and 50 ms, Table 5). But because a production can be enacted only when its value exceeds the utility threshold, many production cycles may occur before the model responds. Slow responses arise from reduced utility values,

¹² Even so, it has been reported that baseline diffusion model drift rate can be used to predict PVT performance during total sleep deprivation (Patanaiik et al., 2015).

which tend to cause microlapses and thereby increase the number of production cycles.

Related to this, the diffusion and ACT-R models differ in whether they treat evidence accumulation as continuous or discrete. In the diffusion model, a response is initiated when accumulated evidence exceeds the decision criterion. If a response has not yet occurred and drift rate is positive, the probability of responding increases over time. In the ACT-R model, a response is initiated when the utility value of the *respond* production exceeds that of all other productions and the utility threshold. If a response has not yet occurred and the state has not changed, the probability of responding during the next production cycle essentially remains the same (Eq. 10).

Against the backdrop of these differences, the accounts capture the detrimental effects of fatigue through two essentially identical component interactions (Table 6). First, in both models, fatigue reduces the signal-to-noise ratio in the decision process, albeit in computationally different ways. In the diffusion model, the signal-to-noise ratio in the decision process is reduced because of the decreasing mean drift ratio, which nominally corresponds to the signal-to-noise ratio in the evidence accumulation process (Ratcliff & Van Dongen, 2011). In ACT-R, the signal-to-noise ratio in the decision process is reduced because of the decreasing production utilities relative to the utility threshold. In both models, dynamic decreases in the signal-to-noise ratio produce increasingly skewed response time distributions with longer right tails.

Second, in both models, fatigue reduces response inhibition – the ability to suppress actions that are inappropriate in the current context and that interfere with goal-driven behavior (Mostofsky & Simmonds, 2008) – but again in computationally different ways. In the diffusion model, the reduction in response inhibition arises from the decreasing decay rate in the LCA component. As decay decreases, so too does the suppression of responses prior to stimulus onset. In ACT-R, the reduction in response inhibition arises from the decreasing utility threshold. This allows actions that were previously suppressed on the basis of their low utility to be enacted. In both models, dynamic changes to response inhibition partially compensate for the primary effect of fatigue and increase the probability that a response will eventually be made after the stimulus appears. These changes also cause more responses to occur before the stimulus appears.

The sufficiency of these two mechanisms in accounting for the effects of fatigue on cognitive performance is shown by our simulations of experiments involving total sleep deprivation, circadian misalignment, and sleep restriction (see above, Experiments 1, 2, and 3). The necessity of the two mechanisms is corroborated by our exploration of alternate model variants (see above, *Alternate parameterizations*). Allowing only a single parameter or any other combination of two parameters in each model to vary with fatigue reduced the goodness of fit. It is also noteworthy that the relationship between

how the models implemented fatigue, evident at the group level, held at the level of individual participants. The positive correlations between parameter values in the two models across individuals (drift rate slope and utility slope; decay slope and threshold slope) provide further evidence for equivalent effects in the two modeling frameworks.

The ACT-R model can be thought of as approximating a discrete diffusion process with 38-ms time steps τ (i.e., the mean duration of a production cycle), a decision criterion of $0.1\sqrt{\tau}$, and a decay of 1.0.¹³ For all combinations of *utility* and *threshold* in the ACT-R model, there is a corresponding drift rate in the diffusion model that yields an identical probability of the process terminating after one step (Fig. 11). As seen in the figure, decreasing *utility* with fatigue in the ACT-R model has the same effect as decreasing drift rate in the diffusion model, a point confirmed in an earlier model mimicry simulation study (Fisher, Walsh, Blaha, & Gunzelmann, 2015). This is not to say that ACT-R is merely a special case of the diffusion model (or vice versa). The two approaches are motivated by entirely different considerations, and they are implemented in completely different ways. Yet despite these differences, the conceptual relationship between *utility* and *threshold* in the ACT-R model, and *drift rate* in the diffusion model underlies their similar behavior with respect to the PVT.

There is also a relationship between *threshold* in the ACT-R model and *decay* in the diffusion model. Decay dampens accumulated evidence, which arises from both signal and noise. When decay is low, the probability of noise driving the decision process beyond the decision criterion increases. Likewise, utilities in ACT-R reflect a production's underlying value in addition to noise (Eq. 7). When threshold is low, the probability of noise causing a non-matching production to exceed the threshold increases.

Although the diffusion and ACT-R models describe the mechanisms underlying PVT performance at different levels of abstraction and in distinct ways, they capture the detrimental effects of fatigue through essentially equivalent component interactions. The insight gained here from comparing the two modeling frameworks – that fatigue reduces the signal-to-noise ratio in the decision process (after stimulus presentation) *as well as* response inhibition (both before and after stimulus presentation) during performance of the PVT – is a new finding.

Relationship to neuronal theories of fatigue

The substantial convergence between models helps to constrain possibilities for the neuronal mechanisms underlying the effects of fatigue on PVT performance. The wide

¹³ Setting decay to 1.0 renders the process memoryless, and setting the decision criterion to $0.1\sqrt{\tau}$ causes the process to terminate immediately once positive evidence is accumulated.

Table 6 Primary and secondary effects of fatigue in diffusion model and ACT-R model

	Diffusion model	ACT-R	Behavioral outcome
Primary effect	Decreased drift rate	Decreased utility	Slower alert responses Increased lapses
Secondary effect	Decreased decay	Decreased threshold	Increased false starts

applicability of the diffusion model across different performance tasks as documented in the literature (see Ratcliff & McKoon, 2008) implies that a single generic underlying mechanism may be able to account for the impact of fatigue on performance. However, ACT-R simulations of fatigue effects across different task platforms underline the importance of differentiating the cognitive processes involved (Gunzelmann et al., 2005, 2009b, 2012; Halverson, Gunzelmann, Moore, & Van Dongen, 2010; see also Jackson et al., 2013). At the neuronal level, these considerations point to mechanisms in which fatigue degrades cognitive processing in a generic fashion (i.e., common to many neuronal pathways) that is nonetheless process-specific (i.e., in neuronal pathways involved in select cognitive processes). Current paradigms positing that subcortical brain mechanisms induce global cortical changes responsible for the effects of fatigue on cognitive performance (Aston-Jones, Chen, Zhu, & Oshinsky, 2001; Doran et al., 2001; Saper et al., 2005; Thomas et al., 2000) fail to explain this process-specificity of fatigue effects (Jackson et al., 2013).

An emerging theoretical view of how sleep deprivation may affect cognitive task performance posits that while the brain as a whole is awake, individual cortical columns involved in task performance may independently “fall asleep” (Jackson et al., 2013; Van Dongen, Belenky, & Krueger, 2011a). Based on the concept of local, use-dependent sleep (Krueger et al., 2008), this paradigm postulates that as a consequence of prior use, cortical columns may temporarily fail to

process information, effectively reducing functional connectivity and thereby degrading the quality of cognitive processing (Krueger, Huang, Rector, & Buysse, 2013; Van Dongen, Belenky, & Krueger, 2011a). Prior use is a function of time awake and is further modulated by task load (Van Dongen, Belenky, & Krueger, 2011b), which is determined by stimulus density and time on task (i.e., cumulative cognitive processing requirement) and is particularly high in repetitive, attention-demanding tasks such as the PVT. The effects of local sleep on performance depend on the number of functional neuronal circuits available to process information for a given task – that is, level of redundancy, or cognitive capacity – which may vary across tasks and among individuals (Chee & Van Dongen, 2013).

The concept of local, use-dependent sleep is consistent with the results of the computational cognitive models considered here, and fits well with the notion of reduced signal-to-noise ratio in the decision process. The transient loss of a subset of neural columns involved in task performance would be expected to reduce the quality of stimulus processing and evidence accumulation (as in the diffusion model), or to produce microlapses (as in the ACT-R model). The idea of local, use-dependent sleep also fits with the notion of reduced response inhibition, provided that inhibition is viewed as an *active* process that is also susceptible to local sleep. The implication that PVT performance relies not only on the ability to sustain attention but also on the ability to maintain inhibition, and that these are distinct aspects of cognition that may each separately instill vulnerability to PVT performance impairment due to fatigue, is a novel insight derived from our computational model comparison.

Predictive generalizability

By examining how fatigue impacts specific underlying mechanisms in each model, the accounts allow exploration of how fatigue may impact cognitive processing in other task contexts. The integrated diffusion model we developed explains the effects of fatigue from sleep loss on performance in terms of temporal changes in degradation of information processing in central cognition. Dynamic changes in drift ratio during sleep deprivation or night work are associated with reduced signal-to-noise ratio and, consequently, degraded quality of cognitive processing. This perspective is supported by neuroimaging data, which indicate that sleep deprivation is

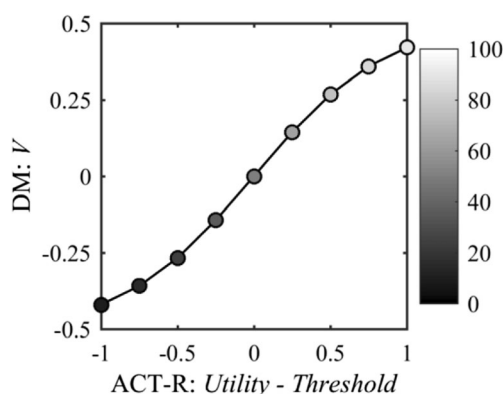


Fig. 11 Difference between utility and threshold in the Adaptive Control of Thought – Rational (ACT-R) model (x-axis) and value of drift rate that produces identical probability of reaching the decision criterion after one 38-ms time step (y-axis). Grayscale shows probability of reaching the criterion after one step

associated with a reduction in neuronal connectivity (Verweij et al., 2014) or available functional neuronal circuits, especially those that are most intensively used for the task at hand (Chee & Asplund, 2013).

However, the integrated diffusion model does not elucidate which circuits are most intensively used during performance of a given task. Thus, it is *a priori* unclear to what extent the model's predictions may generalize from one task to another. The integrated ACT-R model, on the other hand, is explicit regarding which aspects of cognition (i.e., ACT-R modules) are assumed to be involved in task performance, and how intensively. Furthermore, ACT-R modules have been linked to specific brain regions (Borst & Anderson, 2013), suggesting which neuronal circuits may be involved in performance of a given task. Most relevant for the PVT, production rules are thought to be instantiated by networks involving basal ganglia structures including the striatum, the pallidum, and the thalamus (Anderson, 2007). The utility threshold and the compensatory response to fatigue have been posited to be associated with the thalamus (Gunzelmann et al., 2009a), as supported by findings of decreased thalamic activation during sleep deprivation (Chee et al., 2008; Thomas et al., 2000).

As such, it is reasonable to assume that, while the *a priori* predictive generalizability of the integrated diffusion model is limited to generic changes in scaled performance outcomes over time, the integrated ACT-R model could generalize to novel tasks and contexts in terms of absolute performance predictions (see Gunzelmann et al., 2015). Additionally, unlike the diffusion model, the scope of ACT-R extends beyond one- and two-alternative forced-choice tasks. As a first step toward demonstrating these capabilities, we recently combined an ACT-R account of fatigue with validated ACT-R models of multi-tasking and driving behavior to make *a priori* predictions about the effects of extended wakefulness on task performance (Gunzelmann et al., 2009b; Gunzelmann et al., 2011; Khosroshahi, Salvucci, Veksler, & Gunzelmann, 2016). This is not to say the diffusion model cannot play an important role in simulations of complex tasks as well. For example, in the driving domain, brake light detection can be modeled as a signal detection process. In this way, the diffusion model can be used to simulate braking, one component of driving performance (Ratcliff & Strayer, 2014).

Integrated theories of cognition

Computational models have been applied to myriad topics in cognitive science. Integrative and comparative approaches such as those used here provide a pathway towards unification and the development of a coherent whole (Newell, 1990). In this paper, we integrated between cognitive capacities and a cognitive moderator, fatigue. To achieve this integration, we leveraged existing cognitive computational models (Gunzelmann et al., 2009a; Ratcliff & Van Dongen, 2011) and a

biomathematical model of fatigue (McCauley et al., 2013). Such model reuse has been recommended as a practice to accelerate cognitive architecture research (see Gluck, 2010).

The constituent models had previously been validated in isolation, and many of the constraints that shaped their development, though typically unrelated to the PVT, limited the number of assumptions we needed to make in order to create an account of the effects of fatigue on PVT performance. Integrating existing models thus allowed us to reduce the danger of the irrelevant-specification problem (Newell, 1990) – that is, needing to make a large number of under-constrained design decisions to allow the simulation to run.

Rather than adding new knowledge to the cognitive computational models (i.e., constructing new agents) to capture performance under conditions of fatigue, we used the same computational cognitive models and adjusted the settings of architectural parameters. Others have used this approach. For example, Ritter et al. (2007) adjusted the values of declarative memory and motor parameters to simulate the effects of caffeine and anxiety on serial subtraction performance. Likewise, published accounts of arousal capture the effects of fatigue by manipulating aspects of ACT-R's utility calculation (Belavkin, 2001; Jongman, 1998; Gonzalez, Best, Healy, Kole, & Bourne, 2011). Our approach goes a step further by linking parameter values with an underlying physiological account – the effects of fatigue on architectural parameters vary continuously over time and in the manner specified by a validated biomathematical model (McCauley et al., 2013). In doing so, the cumulative models get “further down the list” (Newell, 1990, p. 16) of areas to be covered by a unified theory of cognition.

The practice of implementing moderators by directly adjusting parameter values, though suitable for studying fatigue in isolation, may prove to be impractical for studying the combined effects of multiple moderators. One promising direction for future work is to integrate physiological models of the body – models that represent the combined effects of multiple moderators – with cognitive architectures (Dancy, Ritter, Berry, & Klein, 2015). A similar approach has been used to model the combined and often conflicting effects of emotions such as fear, anger, sadness, and happiness on architectural parameters (Hudlicka, 2007).

Integration is a potentially fruitful approach for leveraging multiple non-overlapping models. This was the case for the diffusion and ACT-R models, and the biomathematical model. However, the diffusion model and the ACT-R model account for the same decision process. When such a “zone of contention” exists, the typical approach is to try to falsify one of the models (McClelland, 2009). In this regard, our results do not point to a clear victor. At the same time, the models share an underlying theoretical interpretation, emphasizing the complementary rather than contradictory nature of their mechanisms (see also, Lebiere, Gonzalez, & Warwick,

2009). Because both theories have utility in advancing our understanding of the mechanisms of fatigue, we chose to focus on their shared perspective rather than their individual limitations. Only by determining how to best account for the effects of fatigue on PVT performance using both modeling formalisms did we recognized their theoretical correspondence.

Conclusion

The adverse effects of fatigue from sleep loss on cognitive performance are substantial, yet most computational models of performance do not include fatigue as a cognitive moderator. We leveraged existing models to explore how fatigue from sleep loss affects cognitive processes. Integrating a biomathematical model of fatigue with computational cognitive models produced a more comprehensive account than either approach alone: the integrated diffusion and ACT-R models captured in detail how fatigue impairs psychomotor vigilance performance. Juxtaposition of the integrated models, which provide accounts of cognitive performance at fundamentally different levels of abstraction, revealed a surprisingly consistent picture of how fatigue affects central cognition during PVT performance: (1) by reducing the signal-to-noise ratio in decision processes, and (2) by reducing response inhibition. Further, by considering response inhibition as an active process, both of these effects can be seen as arising from the loss of processing resources due to local sleep. These findings advance our theoretical understanding of fatigue and illustrate the synergy that can be achieved by comparing computational cognitive modeling at different levels of abstraction, focusing not only on how they differ, but also on how and why they converge.

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 US Government. This research was performed while Matthew Walsh held a National Research Council Research Associateship Award with the Air Force Research Laboratory's Cognitive Models and Agents Branch. Data collection was supported by NIH grant R01HL105768, FAA award DTFAAC-11-A-00003, and FMCSA contract DTMC75-07-D-00006. Hans Van Dongen was supported by CDMRP grant W81XWH-16-1-0319. The authors thank Christopher Fisher and Christopher Myers for their extremely constructive comments on earlier versions of this paper.

References

- Abbott, L. F. (1991). Firing rate models for neural populations. In O. Benhard, C. Bosio, P. Guidice, & E. Tabet (Eds.), *Neural networks: From biology to high energy physics* (pp. 179–196). Pisa, Italy: ETS Editrice.
- Åkerstedt, T. (1988). Sleepiness as a consequence of shift work. *Sleep, 11*, 17–34.
- Anderson, J. R. (2007). *How can the human mind occur in the physical universe*. New York, NY: 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.
- Aston-Jones, G., Chen, S., Zhu, Y., & Oshinsky, M. L. (2001). A neural circuit for circadian regulation of arousal. *Nature Neuroscience, 4*, 732–738.
- Basner, M., & Dinges, D. F. (2011). Maximizing sensitivity of the psychomotor vigilance test (PVT) to sleep loss. *Sleep, 34*, 581–591.
- Basner, M., Mollicone, D., & Dinges, D. F. (2011). Validity and sensitivity of a brief psychomotor vigilance test (PVT-B) to total and partial sleep deprivation. *Acta Astronautica, 69*, 949–959.
- Belavkin, R. V. (2001). Modelling the inverted-U effect in ACT-R. In E. M. Altmann, A. Cleeremans, C. D. Schunn, & W. D. Gray (Eds.), *Proceedings of the 4th International Conference on Cognitive Modeling* (pp. 275–276). Mahwah, NJ: Erlbaum.
- Belenky, G., Wesensten, N. J., Thorne, D. R., Thomas, M. L., Sing, H. C., Redmond, D. P., ... Balkin, T. J. (2003). Patterns of performance degradation and restoration during sleep restriction and subsequent recovery: A sleep dose-response study. *Journal of Sleep Research, 12*, 1–12.
- Borbély, A. A., & Achermann, P. (1999). Sleep homeostasis and models of sleep regulation. *Journal of Biological Rhythms, 14*, 557–568.
- Borst, J. P., & Anderson, J. R. (2013). Using model-based functional MRI to locate working memory updates and declarative memory retrievals in the fronto-parietal network. *Proceedings of the National Academy of Sciences of the USA, 110*, 1628–1633.
- Cain, N., & Shea-Brown, E. (2012). Computational models of decision making: Integration, stability, and noise. *Current Opinion in Neurobiology, 22*, 1047–1053.
- Chee, M. W. L., & Asplund, C. L. (2013). Neuroimaging of attention and alteration of processing capacity in sleep-deprived persons. In E. Nofzinger, P. Maquet, & M. J. Thorpy (Eds.), *Neuroimaging of Sleep and Sleep Disorders* (pp. 137–144). Cambridge, U.K.: Cambridge University Press.
- Chee, M. W. L., Tan, J. C., Zheng, H., Parimal, S., Weissman, D. H., Zagorodnov, V., & Dinges, D. F. (2008). Lapsing during sleep deprivation is associated with distributed changes in brain activation. *The Journal of Neuroscience, 28*, 5519–5528.
- Chee, M. W. L., & Van Dongen, H. P. A. (2013). Functional imaging of inter-individual differences in response to sleep deprivation. In E. Nofzinger, P. Maquet, & M. J. Thorpy (Eds.), *Neuroimaging of Sleep and Sleep Disorders* (pp. 154–162). Cambridge, U.K.: Cambridge University Press.
- Dancy, C. L., Ritter, F. E., Berry, K. A., & Klein, L. C. (2015). Using a cognitive architecture with a physiological substrate to represent effects of a psychological stressor on cognition. *Computational and Mathematical Organization Theory, 21*, 90–114.
- Deary, I. J., Der, G., & Ford, G. (2001). Reaction times and intelligence differences: A population-based cohort study. *Intelligence, 29*, 389–399.
- Dinges, D. F. (1995). An overview of sleepiness and accidents. *Journal of Sleep Research, 4*, 4–14.
- 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.
- Doran, S. M., Van Dongen, H. P. A., & Dinges, D. F. (2001). Sustained attention performance during sleep deprivation: Evidence of state instability. *Archives of Italian Biology: Neuroscience, 139*, 253–267.
- Dorrian, J., Rogers, N. L., & Dinges, D. F. (2005). Psychomotor vigilance performance: Neurocognitive assay sensitive to sleep loss. In C. Kushida (Ed.), *Sleep Deprivation*. New York: Marcel Dekker.
- Fisher, C. R., Walsh, M. M., Blaha, L. M., & Gunzelmann, G. (2015). ACT-R and LBA model mimicry reveals similarity across modeling formalisms. In D. Noelle et al. (Eds.), *Proceedings of the 37th Annual Cognitive Science Society Meeting* (pp. 710–715). Austin, TX: Cognitive Science Society.

- Fu, W. T., & Anderson, J. R. (2006). From recurrent choice to skill learning: A reinforcement-learning model. *Journal of Experimental Psychology: General*, *135*, 184–206.
- Gartenberg, D., Veksler, B., Gunzelmann, G., & Trafton, J. G. (2014). An ACT-R process model of the signal duration phenomenon of vigilance. In *Proceedings of 58th annual meeting of the Human Factors and Ergonomics Society*.
- Gluck, K. A. (2010). Cognitive architectures for human factors in aviation. In E. Salas & D. Maurino (Eds.), *Human factors in aviation* (2nd ed., pp. 375–400). New York: Elsevier.
- Goldman, M. S., Compte, A., & Wang, Z.-J. (2009). Neural integrator models. In L. R. Squire (Ed.), *Encyclopedia of Neuroscience* (pp. 165–178). Oxford: Academic Press.
- Gonzalez, C., Best, B., Healy, A. F., Kole, J. A., & Bourne, L. E. (2011). A cognitive modeling account of simultaneous learning and fatigue effects. *Cognitive Systems Research*, *12*, 19–32.
- Gray, W. D. (Ed.) (2007). *Integrated models of cognitive systems*. New York: Oxford University Press.
- Gunzelmann, G., Byrne, M. D., Gluck, K. A., & Moore, L. R. (2009). Using computational cognitive modeling to predict dual-task performance with sleep deprivation. *Human Factors*, *51*, 251–260.
- Gunzelmann, G., & Gluck, K. A. (2009). An Integrative approach to understanding and predicting the consequences of fatigue on cognitive performance. *Cognitive Technology*, *14*, 14–25.
- Gunzelmann, G., Gluck, K. A., Moore, L. R., & Dinges, D. F. (2012). Diminished access to declarative knowledge with sleep deprivation. *Cognitive Systems Research*, *13*, 1–11.
- 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, & M. Bucciarelli (Eds.), *Proceedings of the 27th Annual Meeting of the Cognitive Science Society* (pp. 857–862). Mahwah, NJ: Lawrence Erlbaum Associates.
- Gunzelmann, G., Gross, J. B., Gluck, K. A., & Dinges, D. F. (2009). Sleep deprivation and sustained attention performance. Integrating mathematical and cognitive modeling. *Cognitive Science*, *33*, 880–910.
- Gunzelmann, G., Moore, 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.
- Gunzelmann, G., Veksler, B. Z., Walsh, M. M., & Gluck, K. A. (2015). Understanding and predicting the cognitive effects of sleep loss through simulation. *Translational Issues in Psychological Science*, *1*, 106–115.
- Halverson, T., Gunzelmann, G., Moore, L. R., & Van Dongen, H. P. A. (2010). Modeling the effects of work shift on learning in a mental orientation and rotation task. In D. D. Salvucci & G. Gunzelmann (Eds.), *Proceedings of the 10th International Conference on Cognitive Modeling* (pp. 79–84). Philadelphia, PA: Drexel University.
- Harris, J. (2008). MindModeling@Home: A large-scale computational cognitive modeling infrastructure. In *Proceedings of the 6th Annual Conference on Systems Engineering Research 2008* (pp. 246–252). Los Angeles, CA.
- Horne, J. A., & Wilkinson, S. (1985). Chronic sleep reduction: Daytime vigilance performance and EEG measures of sleepiness, with particular reference to “practice” effects. *Psychophysiology*, *22*, 69–78.
- Hudlicka, E. (2007). Reasons for emotions: Modeling emotions with integrated cognitive systems. In W. D. Gray (Ed.), *Integrated Models of Cognitive Systems* (pp. 263–278). New York: Oxford University Press.
- 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.
- Hursh, S. R., & Van Dongen, H. P. A. (2010). Fatigue and performance modeling. In M. H. Kryger, T. Roth, & W. C. Dement (Eds.), *Principles and Practice of Sleep Medicine* (5th ed., pp. 745–752). Elsevier Saunders: St. Louis, Missouri.
- Jackson, M. L., Gunzelmann, G., Whitney, P., Hinson, J. M., Belenky, G., Rabat, A., & Van Dongen, H. P. A. (2013). Deconstructing and reconstructing cognitive performance in sleep deprivation. *Sleep Medicine Reviews*, *17*, 215–225.
- Jackson, M. L., & Van Dongen, H. P. A. (2011). Cognitive effects of sleepiness. In M. J. Thorpy & M. Billiard (Eds.), *Sleepiness* (pp. 72–81). Cambridge, UK: Cambridge University Press.
- Jongman, L. (1998). How to fatigue ACT-R? In *Proceedings of the Second European Conference on Cognitive Modelling* (pp. 52–57). Nottingham, England: Nottingham University Press.
- Khosroshahi, E. B., Salvucci, D. D., Veksler, B. Z., & Gunzelmann, G. (2016). Capturing the effects of moderate fatigue on driver performance. In *Proceedings of the 14th International Conference on Cognitive Modeling* (pp. 163–168). University Park, PA: The Pennsylvania State University.
- Killgore, W. D. (2010). Effects of sleep deprivation on cognition. *Progress in Brain Research*, *185*, 105–129.
- Kleitman, N. (1963). Deprivation of sleep. In N. Kleitman (Ed.), *Sleep and Wakefulness* (pp. 215–229). Chicago: University of Chicago Press.
- Krueger, J. M., Huang, Y. H., Rector, D. M., & Buysse, D. J. (2013). Sleep: A synchrony of cell activity-driven small network states. *European Journal of Neuroscience*, *38*, 2199–2209.
- Krueger, J. M., Rector, D. M., Roy, S., Van Dongen, H. P. A., Belenky, G., & Panksepp, J. (2008). Sleep as a fundamental property of neuronal assemblies. *Nature Reviews Neuroscience*, *9*, 910–919.
- Laming, D. R. J. (1968). *Information theory of choice-reaction times*. London: Academic Press.
- Larson, G. E., & Alderton, D. L. (1990). Reaction time variability and intelligence: A —worst performance analysis of individual differences. *Intelligence*, *14*, 309–325.
- Lebiere, C., Gonzalez, C., & Warwick, W. (2009). Convergence and constraints revealed in a qualitative model comparison. *Journal of Cognitive Engineering and Decision Making*, *3*, 131–155.
- Lim, J., & Dinges, D. F. (2008). Sleep deprivation and vigilant attention. *Annals of the New York Academy of Sciences*, *1129*, 305–322.
- Lim, J., & Dinges, D. F. (2010). A meta-analysis of the impact of short-term sleep deprivation on cognitive variables. *Psychological Bulletin*, *136*, 375–389.
- Mallis, M. M., Mejdal, S., Nguyen, T. T., & Dinges, D. F. (2004). Summary of the key features of seven biomathematical models of human fatigue and performance. *Aviation, Space, and Environmental Medicine*, *75*, A4–A14.
- McCauley, P., Kalachev, L. V., Mollicone, D. J., Banks, S., Dinges, D. F., & Van Dongen, H. P. A. (2013). Dynamic circadian modulation in a biomathematical model for the effects of sleep and sleep loss on waking neurobehavioral performance. *Sleep*, *36*, 1987–1997.
- 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.
- McClelland, J. L. (2009). The place of modeling in cognitive science. *Topics in Cognitive Science*, *1*, 11–38.
- Mostofsky, S. H., & Simmonds, D. J. (2008). Response inhibition and response selection: Two sides of the same coin. *Journal of Cognitive Neuroscience*, *5*, 751–761.
- Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Cambridge University Press.
- Ossmy, O., Moran, R., Pfeffer, T., Tsetsos, K., Usher, M., & Donner, T. H. (2013). The timescale of perceptual evidence integration can be adapted to the environment. *Current Biology*, *23*, 981–986.
- Patanaiik, A., Kwok, C. K., Chua, E. C., Gooley, J. J., & Chee, M. W. (2015). Classifying vulnerability to sleep deprivation using baseline measures of psychomotor vigilance. *Sleep*, *38*, 723–734.

- Patanaik, A., Zagorodnov, V., Kwok, C. K., & Chee, M. W. L. (2014). Predicting vulnerability to sleep deprivation using diffusion model parameters. *Journal of Sleep Research, 23*, 576–584.
- Ratcliff, R. (2006). Modeling response signal and response time data. *Cognitive Psychology, 53*, 195–237.
- Ratcliff, R., Gomez, P., & McKoon, G. (2004). Diffusion model account of lexical decision. *Psychological Review, 111*, 159–182.
- Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation, 20*, 873–922.
- Ratcliff, R., & Rouder, J. N. (2000). A diffusion model account of masking in letter identification. *Journal of Experimental Psychology: Human Perception and Performance, 26*, 127–140.
- Ratcliff, R., & Strayer, D. (2014). Modeling simple driving tasks with a one-boundary diffusion model. *Psychonomic Bulletin & Review, 21*, 577–589.
- Ratcliff, R., & Tuerlinckx, F. (2002). Estimating parameters of the diffusion model: Approaches to dealing with contaminant reaction times and parameter variability. *Psychonomic Bulletin & Review, 9*, 438–481.
- Ratcliff, R., & Van Dongen, H. P. A. (2009). Sleep deprivation affects multiple distinct cognitive processes. *Psychonomic Bulletin & Review, 16*, 742–751.
- Ratcliff, R., & Van Dongen, H. P. A. (2011). Diffusion model for one-choice reaction-time tasks and the cognitive effects of sleep deprivation. *Proceedings of the National Academy of Sciences, 108*, 11285–11290.
- Ratcliff, R., Van Zandt, T., & McKoon, G. (1999). Connectionist and diffusion models of reaction time. *Psychological Review, 106*, 261–300.
- Ritter, F. E., Reifers, A. L., Klein, L. C., & Schoelles, M. J. (2007). Lessons from defining theories of stress. In W. D. Gray (Ed.), *Integrated Models of Cognitive Systems* (pp. 254–262). New York: Oxford University Press.
- Robinson, D. A. (1989). Integrating with neurons. *Annual Review of Neuroscience, 12*, 33–45.
- Salvucci, D. D. (2006). Modeling driver behavior in a cognitive architecture. *Human Factors, 48*, 362–380.
- Saper, C. B., Scammell, T. E., & Lu, J. (2005). Hypothalamic regulation of sleep and circadian rhythms. *Nature, 437*, 1257–1263.
- Smith, P. L., & Ratcliff, R. (2009). An integrated theory of attention and decision making in visual signal detection. *Psychological Review, 116*, 283–317.
- Stone, M. (1960). Models for choice reaction time. *Psychometrika, 25*, 251–260.
- Sun, R. (Ed.). (2008). *Cambridge handbook of computational psychology*. New York: Cambridge University Press.
- Thomas, M., Sing, H., Belenky, G., Holcomb, H., Mayberg, H., Dannals, R., ... Redmond, D. (2000). Neural basis of alertness and cognitive performance impairments during sleepiness. I. Effects of 24 h of sleep deprivation on waking human regional brain activity. *Journal of Sleep Research, 9*, 335–352.
- Tucker, A. M., Whitney, P., Belenky, G., Hinson, J. M., & Van Dongen, H. P. A. (2010). Effects of sleep deprivation on dissociated components of executive functioning. *Sleep, 33*, 47–57.
- Tuerlinckx, F., Maris, E., Ratcliff, R., & De Boeck, P. (2001). A comparison of four methods for simulating the diffusion process. *Behavioral Research Methods, Instruments, & Computers, 33*, 443–456.
- Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review, 108*, 550–592.
- Van Dongen, H. P. A. (2006). Shift work and inter-individual differences in sleep and sleepiness. *Chronobiology International, 23*, 1139–1147.
- 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*, 423–433.
- Van Dongen, H. P. A., Belenky, G., & Krueger, J. M. (2011a). A local, bottom-up perspective on sleep deprivation and neurobehavioral performance. *Current Topics in Medicinal Chemistry, 11*, 2414–2422.
- Van Dongen, H. P. A., Belenky, G., & Krueger, J. M. (2011b). Investigating the temporal dynamics and underlying mechanisms of cognitive fatigue. In P. L. Ackerman (Ed.), *Cognitive Fatigue: Multidisciplinary Perspectives on Current Research and Future Applications* (pp. 127–147). Washington, D.C.: American Psychological Association.
- Van Dongen, H. P. A., Belenky, G., & Vila, B. J. (2011). The efficacy of a restart break for recycling with optimal performance depends critically on circadian timing. *Sleep, 34*, 917–929.
- Van Dongen, H. P. A., Bender, A. M., & Dinges, D. F. (2012). Systematic individual differences in sleep homeostatic and circadian rhythm contributions to neurobehavioral impairment during sleep deprivation. *Accident Analysis and Prevention, 45S*, 11–16.
- Van Dongen, H. P. A., & Hursh, S. R. (2010). Fatigue, performance, errors and accidents. In M. H. Kryger, T. Roth, & W. C. Dement (Eds.), *Principles and Practice of Sleep Medicine* (5th ed., pp. 753–759). Philadelphia: Elsevier Saunders.
- 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.
- Verweij, I. M., Romeijn, N., Smit, D. J., Piantoni, G., Van Someren, E. J. W., & van der Werf, Y. D. (2014). Sleep deprivation leads to a loss of functional connectivity in frontal brain regions. *BMC Neuroscience, 15*, 88.
- Voss, A., Rothermund, K., & Voss, J. (2004). Interpreting the parameters of the diffusion model: An empirical validation. *Memory & Cognition, 32*, 1206–1220.
- Whitney, P., & Hinson, J. M. (2010). Measurement of cognition in studies of sleep deprivation. *Progress in Brain Research, 185*, 37–48.
- Whitney, P., Hinson, J. M., Jackson, M. L., & Van Dongen, H. P. A. (2015). Feedback blunting: Total sleep deprivation impairs decision making that requires updating based on feedback. *Sleep, 38*, 745–754.