

ACT-R and LBA Model Mimicry Reveals Similarity Across Modeling Formalisms

Christopher R. Fisher (christopher.fisher.27.ctr@us.af.mil)

Matthew M. Walsh (matthew.walsh.15.ctr@us.af.mil)

Leslie M. Blaha (leslie.blaha@us.af.mil)

Glenn Gunzelmann (glenn.gunzelmann@us.af.mil)

Air Force Research Laboratory, 711th Human Performance Wing
Wright-Patterson Air Force Base, OH 45433

Abstract

Adaptive Control of Thought-Rational (ACT-R) and the Linear Ballistic Accumulator (LBA) were compared in a model mimicry simulation of the Psychomotor Vigilance Task (PVT), a simple, reaction time (RT) task requiring sustained attention. The models use different formalisms to capture the full response profile of the PVT. The parameters were varied systematically to illustrate the ranges of the models' predictions, to assess the models' estimation properties, and to determine which parameters in the models correspond with each other. Both models produced skewed RT distributions typical of empirical data, including false starts and lapses. The simulation study demonstrated that both models and their parameters are recoverable. Lastly, isolated parameters in the LBA model captured the effects of varying parameters in the ACT-R model, but the reverse was not always true. These interesting correspondences across different modeling formalisms suggest the possibility of integrating ACT-R and the LBA in future work.

Keywords: ACT-R, LBA, PVT, reaction time, fatigue, model comparison

Introduction

The ability to detect a single stimulus is fundamental to cognition. Although this skill is basic, the study and modeling of stimulus detection is worthwhile for several reasons. Stimulus detection has been extensively examined in laboratory tasks involving vigilance and simple reaction time (RT; Luce, 1986). Additionally, this ability underlies successful performance in applied contexts that require sustained attention, such as driving. Finally, intuition suggests that the cognitive processes involved in stimulus detection should be involved in more-complex multi-alternative choices as well.

Despite the simplicity of detection tasks, the RT distributions they produce are complex and empirically rich. This is well-illustrated by the psychomotor vigilance task (PVT; Dinges & Powell, 1985), a 10-minute detection task in which stimuli are presented at random inter-trial intervals ranging from 2 to 10 seconds. Participants are instructed to respond as quickly as possible once the stimulus appears while avoiding premature responses. The PVT response profile consists of three categories: *false starts* occur before

or within 150 ms of stimulus presentation, *alert responses* occur between 150 and 500 ms of the stimulus onset, and *lapses* occur 500 ms after of the stimulus onset. The RT distribution on the PVT, which has a long right tail even when participants are well rested, becomes increasingly skewed to the right with greater fatigue from sleep loss, as reflected in increased lapses (Lim & Dinges, 2008). Additionally, participants commit more false starts. These features of the response profile reflect stable individual differences, both at baseline and following sleep loss (Van Dongen, Baynard, Maislin, & Dinges, 2004).

A complete model of the PVT should explain the full response profile, yet most biomathematical accounts from the sleep research literature only predict aggregate measures of performance such as the proportion of lapses (for a review, see Van Dongen, 2004). More recent work has attempted to use statistical functions to characterize the full RT distribution (Lim & Dinges, 2008), but those efforts still fail to explain why the particular distributions arise. A promising alternative is to use computational cognitive models, which specify the cognitive processes underlying task performance, to simulate behavior in the PVT (e.g., Gunzelmann, Veksler, Walsh, & Gluck, 2015).

In this paper, we compared two PVT models derived from very different formalisms. The first model is based on the integrated-cognitive architecture Adaptive Control of Thought-Rational (ACT-R), in which RTs are determined by the durations of a sequence of discrete cognitive events. The second model is based on the Linear Ballistic Accumulator (LBA; Brown & Heathcote, 2008), an analytically tractable member of the class of sequential sampling models. In the LBA, RTs are determined by the combined durations of a decision process in which evidence accumulates continuously, and an overall non-decision time attributed to perceptual and motor processes.

The PVT is an ideal test bed for comparing ACT-R and the LBA because (1) the PVT is simple, yet (2) it provides empirically rich data for inferring cognitive processes, and (3) both ACT-R and the LBA can be applied to the PVT. Rather than attempting to falsify one account, we sought to compare and contrast these differing formalisms.

We addressed three primary questions in this research. First, can both ACT-R and LBA generate the complete RT profiles, including false starts and lapses, observed in PVT studies? ACT-R models have predominantly been used to predict mean RTs, and attempts to account for full RT distributions have been rare (but see Walsh et al., 2014). The LBA has only been used to model the correct and error responses in multi-alternative choice tasks (Brown & Heathcote, 2008), and it was unclear whether it could also account for the full response profile observed in the PVT, especially the occurrence of false starts and lapses. Second, how well can ACT-R and LBA recover their own parameters from simulated PVT data? Both models are complex, and the estimation properties of their parameters have not been assessed in the PVT. As such, it was unknown whether model parameters could be reliably estimated from PVT data, or whether the models could even be distinguished from one another based on data from the PVT. Third, what are the relationships between core parameters in the two models? Although the models are distinct, it was unclear which of their parameters are conceptually and/or functionally linked.

Models

LBA

The LBA is a sequential sampling model that is similar to the drift diffusion model (DDM) in terms of parameter interpretation (Brown & Heathcote, 2008; Donkin et al., 2011). In both models, information is sampled from a stimulus and accumulates over time. When accumulated evidence in favor of an alternative reaches a threshold, a decision occurs. Sources of variation in the DDM, such as intra-trial variability in evidence accumulation and inter-trial variability in non-decision time, are absent from the LBA. These simplifications come with no loss of generality, making LBA a more parsimonious, complete account of basic empirical RT phenomena (Brown & Heathcote, 2008).

In the standard LBA, the stimulus onset triggers an evidence accumulation process. Accumulated evidence begins from a variable starting point between 0 and the response threshold, and proceeds towards the response threshold in a linear and deterministic fashion. The speed of the accumulation process is controlled by the drift rate. Between-trial variability in the drift rate and starting point of the evidence accumulation process contribute to the shape and spread of the RT distribution. The drift rate is normally distributed across trials with a mean of V , and a standard deviation of 1. The starting point is uniformly distributed with an adjustable maximum starting point, A . Other processes such as encoding and motor execution are combined into a composite measure of non-decision time, t_0 .

Several modifications were necessary to apply the LBA to the PVT (Fig. 1). Our modified LBA model involves two accumulation processes that occur in succession rather than one accumulation process. First, an inter-stimulus interval (ISI) accumulation process starts at the beginning of the

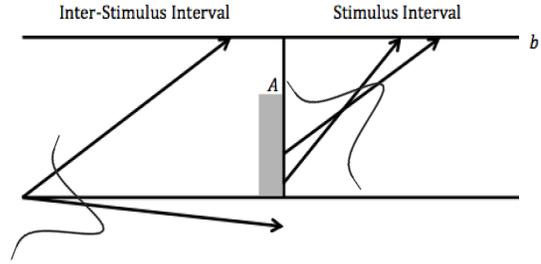


Figure 1. The modified LBA has separate accumulators for the inter-stimulus and stimulus intervals. A denotes spread of start points for stimulus interval, and b denotes threshold for both intervals. The vertical bar marks stimulus onset.

trial. Although this process has a negative drift rate on average, stochasticity occasionally results in a positive drift rate and, consequently, a false start. Once the stimulus appears, the ISI accumulation process halts and a separate stimulus interval (SI) accumulation process starts. The trial ends once a response is given.

The ISI and SI accumulation processes are identical, except for mean drift rate, V , and the maximum starting point, A . The ISI mean drift rate, V_{ISI} , is constrained to be negative, indicating that false starts are rare and produced randomly. Additionally, the ISI maximum starting point, A_{ISI} , is set to zero to reflect bias toward not responding. The threshold, b , is the same for the ISI and SI accumulation processes, as is non-decision time, t_0 . In total, the modified LBA model contains five free parameters: b , A_{SI} , V_{ISI} , V_{SI} , and t_0 .

ACT-R

ACT-R contains a set of specialized information-processing modules (e.g., a vision module, a declarative memory module, a motor module). These modules are connected to, and controlled by, a central procedural module (Anderson, 2007). Procedural knowledge is represented in the form of production rules, which consist of selection criteria and actions that modify the internal state of the architecture and the external state of the world when the selection criteria are met. The temporal dynamics of cognition unfold across a sequence of production cycles. During each cycle, the conditions for each production are compared against the conditions of the current state, and a production is selected and enacted if its conditions are met. The resulting state serves as the starting point for the next production cycle.

We adopted an ACT-R model of the PVT that consists of three productions: (1) wait for the stimulus to appear, which represents task engagement, (2) attend to the stimulus, and (3) respond to the stimulus (Walsh et al., 2014). Partial production matching allows productions whose conditions are not perfectly met to be selected in a stochastic fashion, producing occasional false starts. The probability that a production is selected is modulated by two adjustable parameters—a utility scalar (U_S) and a utility threshold (U_T). Formally, production utility can be expressed as:

$$(1) \quad U_{ij} = U_S(U_i - MMP_{ij}) + \epsilon_i$$

where U_{ij} is the utility of production i in state j , U_S is the utility scalar, U_i is the stored utility for production i , MMP_{ij} is the mismatch penalty for production i in state j , and ϵ_i is logistically distributed noise. The resulting payoff matrix is symmetric with 0 assigned to mismatches and 1 assigned to matches. The mismatch penalty ensures that productions whose conditions are not perfectly met will be selected with low probability.

The production with highest utility is selected and enacted if its utility exceeds the utility threshold, U_T ,

$$(2) \quad \text{Production} = \max(U_{ij}) \text{ if } \max(U_{ij}) > U_T$$

If no production's utility exceeds the utility threshold, a microlapse occurs and no production is enacted. Following a microlapse, the utility scalar in Eq. 1 is decremented by an adjustable scalar, FP_{dec} , according to $U_s = U_s \cdot FP_{dec}$. This increases the likelihood of microlapses in subsequent production cycles. Across such a series of cycles, the probability of responding decreases progressively, causing behavioral lapses. The final adjustable parameter, *cycle time*, controls the duration of conflict resolution at the start of each production cycle. In total, the ACT-R model contains four free parameters: U_s , U_T , FP_{dec} , and *cycle time*.

Our model harnessed two sources of temporal variability. The first related to the variable sequence of productions selected in a trial, and the second related to the stochastic duration of production and cycle times. Each trial's RT, then, was determined by the summed durations of the productions and their associated cognitive and motor processes. In this way, the ACT-R model can produce a full distribution of RTs, rather than an approximation of an aggregate mean RT (Walsh, et al., 2014).

Simulation Method

We simulated an idealized selective influence experiment (Donkin, et al., 2011) in which the parameters of each model were systematically varied one at a time while all others were set to default values. This approach allowed us to examine (1) our ability to accurately recover parameters of each model, (2) the extent to which the models mimicked each other and (3) how the parameters were correlated between models. Parameter ranges were drawn from the published model fits of PVT performance by 13 well-rested individuals in the control condition of a sleep deprivation experiment (Doran, Van Dongen, & Dinges, 2001; see also Walsh et al., 2014). We set the default value of each parameter to the median estimate from the individual model fits, and the range of each parameter to the complete range of estimates from the individual fits (Table 1). We varied parameters at ten equally spaced intervals over their ranges, resulting in 40 ACT-R parameter sets (10 levels per parameter by 4 parameters) and 50 LBA parameter sets (10 levels per parameter by 5 parameters). We simulated 50,000

PVT trials for each model and parameter set to minimize the role of sampling error and bias in our analyses.

Table 1. Default parameters and ranges in the simulation.

LBA	b	A_{SI}	V_{SI}	t_0	V_{ISI}
Default	0.68	0.44	3.42	0.15	-2.34
Min	0.54	0.1	3	0.15	-2.95
Max	0.98	0.56	3.9	0.18	-2.01
ACT-R	U_s	U_T	FP_{dec}	<i>Cycle Time</i>	$U_s - U_T$
Default	4.85	4.39	0.98	0.04	0.46
Min	4.01	4.07	0.91	0.029	-0.38
Max	5.6	5.02	0.99	0.057	1.21

Each model was fit to the 90 simulated datasets using quantile maximum likelihood estimation (Heathcote, Brown & Mewhort, 2002). RTs that occurred prior to stimulus onset or within 150 ms of stimulus onset were combined into a false start bin (Lim & Dinges, 2008). The remaining portion of the distribution was further divided into 20 quantile bins. Likelihood estimates were calculated from the observed and expected proportions of RTs within each quantile bin. A simplex algorithm embedded within a grid search was used to find the model parameters that maximized the likelihood of each simulated dataset. Large-scale computing resources (Harris, 2008) were leveraged for ACT-R, as it is computationally intensive.

Results

Model RT Distributions

Figure 2 shows four of the most distinctive RT distributions produced by ACT-R and the LBA. The distributions, which vary in terms of numbers of false starts and lapses as well as median RTs (Table 2), are within the ranges of those produced by well-rested and sleep deprived individuals (cf., Walsh et al., 2014). In the 90 simulated datasets, the models produced similar proportions of false starts and lapses and similar median RTs. However, the LBA model consistently yielded distributions with more pronounced skew.

Table 2. Proportions of false starts and lapses, and median RTs from the simulated distributions in Fig. 2.

Model	Curve	False Starts	Lapses	Median RT (ms)
ACT-R	Blue	.006	.000	245
	Red	.008	.005	272
	Black	.010	.083	305
	Green	.101	.222	381
LBA	Blue	.006	.000	242
	Red	.008	.010	271
	Black	.011	.085	306
	Green	.106	.210	381

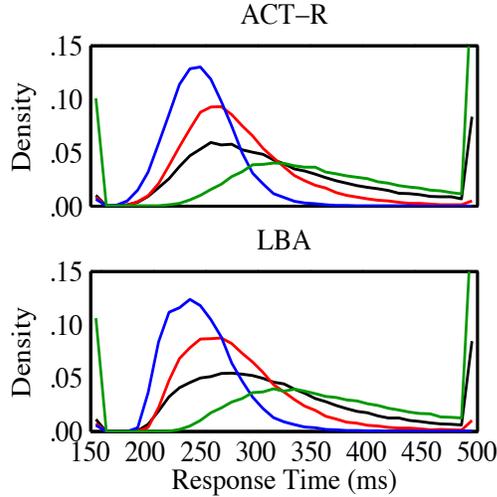


Figure 2. Proportion of RTs in 10 ms bins ranging from 150 ms to 500 ms. The first bin contains all RTs before 150 ms, and the last bin contains all RTs after 500 ms. Blue, red, black, and green lines show fast, medium, slow and sleep deprived RT distributions.

Parameter Recovery

The parameter recovery model fits address how accurately the parameters can be estimated from PVT data. In these analyses, the models were fit to their self-generated data. Two metrics were used to assess the quality of the parameter recovery: correlation to measure the linear association between the true and recovered parameters, and relative bias to measure the precision of the estimates.

Table 3 (upper) shows the parameter recovery results for ACT-R. The high correlation for *cycle time* indicates that this parameter is recoverable. Correlations for U_s and U_T were moderate, but the correlation for the difference between U_s and U_T was high. This indicates that the utility scalar and threshold jointly influence performance dynamics in the ACT-R model. The low correlation for FP_{dec} is due to the relatively infrequent occurrence of lapses in well-rested individuals. Relative bias was low across all parameters, indicating the high precision of the estimates.

Table 3 (lower) displays the parameter recovery results for the LBA. The high correlations and low relative bias indicate that the parameter recovery was successful. Collectively, these results show that parameters from both models can be reliably estimated from their own simulations of PVT data.

Table 3. Parameter recovery results for ACT-R and LBA.

ACT-R	U_s	U_T	FP_{dec}	<i>Cycle Time</i>	$U_s - U_T$
Correlation	0.85	0.77	0.56	0.99	0.99
Relative Bias	1%	1%	0%	0%	4%
LBA	b	A_{SI}	V_{SI}	t_0	V_{ISI}
Correlation	0.93	0.97	0.85	0.85	0.98
Relative Bias	-3%	2%	-1%	3%	0%

Model Mimicry

The model mimicry analyses address whether ACT-R and the LBA produce different predictions on the PVT. In these simulations, the ACT-R and LBA models were cross-fit to data generated by each other. The Bayesian Information Criterion (BIC) was used to determine whether the data-generating model provided a better fit to the RT distributions than the alternate model while adjusting for parametric sources of model complexity. Smaller values denote better fit.

Figure 3 shows the BICs averaged across datasets for each model. In all 90 simulated data sets, both models provided better fits to their own data than the alternate model. This shows that although the models make very similar predictions they are identifiable in simulations with very large sample sizes.

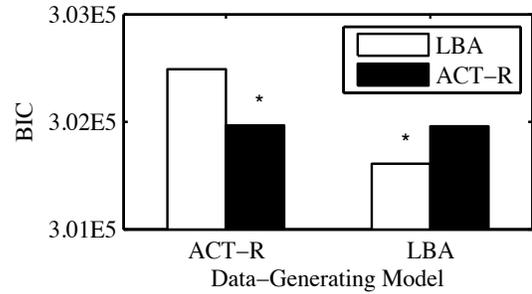


Figure 3. BIC averaged across datasets. Stars denote fit of data-generating model to itself.

Parameter Correspondence

We examined the manner in which parameters in the two models corresponded to one another. In our simulations, parameters were varied one at a time while the other parameters were fixed. In the simplest case, a change in one parameter would be captured by variation in a single, analogous parameter in the alternate model. For simplicity, we considered three core parameters in the ACT-R model ($U_s - U_T$, FP_{dec} , and *cycle time*), and four in the LBA (V_{SI} , V_{ISI} , t_0 , and $b - A_{SI}/2$). The composite parameter $b - A_{SI}/2$, called *response caution*, is derived from the threshold and the center of the start point distribution, and measures the average amount of information that is needed to reach the decision threshold (Donkin, et al., 2009).

We first examined how ACT-R responded to manipulations of the LBA parameters (Table 4). No

Table 4. Correlations between LBA (data generating) and ACT-R (best fitting) parameter values. * $p < .05$

LBA	ACT-R		
	FP_{dec}	<i>Cycle Time</i>	$U_s - U_T$
V_{ISI}	-0.06	0.08	0.04
V_{SI}	0.10	-0.09	0.16
t_0	0.04	0.20	0.22
<i>Response Caution</i>	-0.63*	0.91*	0.68*

Table 5. Correlations between ACT-R (data generating) and LBA (best fitting) parameter values. *p < .05

ACT-R	LBA			
	V_{ISI}	V_{SI}	t_0	Response Caution
FP_{dec}	-0.22	0.16	0.08	0.30
Cycle Time	-0.18	-0.01	-0.08	0.89*
$U_S - U_T$	0.07	0.96*	-0.41*	0.08

parameters in the ACT-R model were selectively influenced by changes to V_{SI} , V_{ISI} and t_0 , but all parameters were affected by changes to *response caution*. Next, we examined how the LBA responded to manipulations of ACT-R parameters (Table 5). Changes to *cycle time* were captured by *response caution*, and changes to $U_S - U_T$ were captured by V_{SI} . No parameter in the LBA was selectively influenced by changes to FP_{dec} . In sum, there was a direct mapping between individual ACT-R parameter manipulations and LBA parameters, but not between individual LBA parameter manipulations and ACT-R parameters.

Discussion

The detection of a single stimulus is among the most-widely studied topics in cognitive science. Yet, despite the simplicity of one-choice RT tasks, the RT distributions they produce are complex and difficult to account for in detail. Here, we compared two computational cognitive models of the PVT. One model was based on ACT-R and consists of a sequence of discrete cognitive events while the other was based on the LBA, which involves continuous evidence accumulation. The results of our simulations support three findings. First, both models produced the qualitative shapes of RT distributions found in the PVT, including the long right tail of RT distribution, and occasional false starts and lapses (Fig. 2). Second, most model parameters were recoverable and the PVT was capable of distinguishing between the models. Third, isolated parameters in the LBA model captured the effects of varying ACT-R parameters, but the reverse was not always true. The correspondence between ACT-R parameters and LBA parameters suggests similarity between these differing modeling formalisms.

Model Comparison

The correspondence between parameters in the LBA and ACT-R models was complex. In some cases, parameters in one model were affected by parametric variations in the other in intuitive ways. For example, drift rate (V_{SI}) in the LBA captured changes in the difference between the utility scalar and threshold ($U_S - U_T$) in ACT-R. This makes sense because both fundamentally control the signal-to-noise ratio in the decision process.

In other cases, unexpected model parameters corresponded to one another. For example, changes in *response caution* in the LBA were captured by *cycle time* in ACT-R and vice versa. *Response caution* is thought to be

sensitive to instructions designed to prioritize speed or accuracy, whereas *cycle time* is conceptualized as a stable property of the cognitive architecture that only varies among individuals. ACT-R posits that production selection is instantiated in the basal ganglia, which receives input from multiple excitatory and inhibitory pathways. It is conceivable that the duration of production selection, represented by *cycle time*, varies with dynamic activity from these pathways. In other words, the relationship between *response caution* and *cycle time* may be real, despite the current standard of fixing *cycle time* within ACT-R models of individuals.

In a third set of cases, we found little correspondence between model parameters. For example, ACT-R failed to capture manipulations of non-decision time in the LBA. This relationship was relatively symmetrical in that non-decision time showed little or no systematic relationship to the manipulation of any ACT-R parameters. Such a lack of correspondence suggests that an experimental manipulation of non-decision time could potentially discriminate between ACT-R and the LBA. Moreover, this finding indicates that conclusions will depend critically upon which model is used to evaluate data.

Effects of Fatigue on Psychomotor Vigilance

We demonstrated that the ACT-R and LBA models produce a range of response profiles that are similar to each other, and similar to those observed in well-rested individuals. The models rarely responded before 150 ms of stimulus presentation (false starts), and they rarely responded more than 500 ms after the stimulus appeared (lapses). False starts and lapses, though present in baseline RT distributions, are greatly exacerbated by fatigue from sleep loss. As shown by Walsh et al. (2014), ACT-R can be integrated with a biomathematical model of fatigue to predict the effects of time awake and time of day on PVT performance. The LBA model has not been expanded to account for the effects of fatigue on PVT performance, yet it should be conceptually straightforward to do so.

Evaluating the models under conditions of fatigue might also enhance model discriminability. More confidence can be placed in a model that captures normal as well as impaired cognitive functioning. Certain parameters that are essential to capturing the effects of fatigue minimally affect alert performance on the PVT (FP_{dec} and U_T in ACT-R, and V_{ISI} in the LBA). In this sense, sleep deprivation protocols provide a unique opportunity to distinguish among models of the PVT (Walsh et al., 2014) and could be leveraged as a general strategy for model comparison.

Towards an Integration of ACT-R and the LBA

Sequential sampling models and ACT-R explain cognition using different modeling formalisms. Sequential sampling models provide detailed accounts of empirical RT distributions. This emphasis comes at the cost of limited generalizability beyond well-constrained decision-making tasks utilizing fixed trial structures. Cognitive architectures,

by contrast, focus on the unification and generalization necessary to model complex tasks. Because of this focus, cognitive architectures neglect certain details of low-level decision processes.

Efforts to capitalize on the complimentary strengths of sequential sampling models and cognitive architectures have been made recently. Van Maanen, van Rijn, and Taatgen (2012) combined the DDM and ACT-R to form RACE/A, which accounts for the dynamics of declarative memory in a picture-word interference task. A DDM with multiple accumulators governs how the activation values of information in declarative memory change over time and determine retrieval latencies. ACT-R, in turn, provides the control structure necessary for coordinating the multitude of decision and non-decision processes evoked by the task.

Within the context of the PVT, sequential sampling models could be used as a mechanism for production selection. Presently, the duration of production selection in ACT-R is treated as a uniform random variable with a mean of about 40 ms (Table 1). Each production could instead be represented as an accumulator with a drift rate determined by the match between the state of the world and the production's conditions. Integrating these approaches would provide a theory of production selection (implemented as a sequential sampling model) along with a theory of task control (implemented as production rules). The LBA would be a natural choice for the sequential sampling model for three reasons: (1) it is applicable to selection among two or more alternatives, (2) it is more parsimonious than other sequential sampling models, and (3) parameter estimation is efficient and mathematically tractable.

Incorporating a sequential sampling model into a cognitive architecture would provide a more detailed, formal account of the time course of production selection. Such an account would provide a rationale for changes in the stochastic duration of cycle time. Although such an account may be unnecessary for modeling the PVT, incorporating both representational levels would be useful for capturing complete performance dynamics in more complex tasks. Factors in multi-alternative choice tasks such as decision conflict and value influence decision times (Ratcliff & Frank, 2012). Likewise, factors in single-alternative choice tasks such as stimulus contrast and luminosity influence decision times. Presently, these effects are difficult to explain in ACT-R. Implementing production selection as a sequential sampling process could overcome these challenges.

Acknowledgments

The views expressed in this paper are those of the authors and do not reflect the official policy or position of the Department of Defense or the U.S. Government. M.M.W. held a National Research Council Research Associateship Award with the AFRL while conducting this research. This research was supported by an AFOSR grant to L.M.B. Distribution A: Approved for public release; distribution unlimited. 88ABW Cleared 03/09/2015; 88ABW-2015-0914.

References

- Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* NY, NY: Oxford University Press.
- Brown, S. D., & Heathcote, A. (2008). The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive Psychology*, *57*, 153-178.
- 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.
- Donkin, C., Brown, S., Heathcote, A., & Wagenmakers, E.-J. (2011). Diffusion versus linear ballistic accumulation: different models but the same conclusions about psychological processes? *Psychonomic Bulletin & Review*, *18*(1), 61-69.
- 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*(1), 106-115.
- Harris, J. (2008). MindModeling@Home: A large-scale computational cognitive modeling infrastructure. In *Proceedings of the Sixth Annual Conference on Systems Engineering Research 2008* (pp. 246-252). Los Angeles, CA, USA.
- Heathcote, A., Brown, S., & Mewhort, D. J. K. (2002). Quantile maximum likelihood estimation of response time distributions. *Psychonomic Bulletin & Review*, *9*, 349-401.
- Lim, J., & Dinges, D. F. (2008). Sleep deprivation and vigilant attention. *Annals of the New York Academy of Sciences*, *1129*, 305-322.
- Luce, R. D. (1986). *Response times*. New York, NY: Oxford University Press.
- Ratcliff, R., & Frank, M. (2012). Reinforcement-based decision making in corticostriatal circuits: Mutual constraints by neurocomputational and diffusion models. *Neural Computation*, *24*, 1186-1229.
- 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.
- 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 Maanen, L., van Rijn, H., & Taatgen, N. (2012). RACE/A: An architectural account of the interactions between learning, task control, and retrieval dynamics. *Cognitive Science*, *36*, 62-101.
- Walsh, M. M., Gunzelmann, G., & Van Dongen, H. P. A. (2014). Comparing accounts of psychomotor vigilance impairment due to sleep loss. In *36th Annual Conference of the Cognitive Science Society*, Quebec City, Canada.