Real-Time Prediction of Fluctuations in Cognitive Workload.

Udo Boehm^a

Dora Matzke^b

Matthew Gretton^c

Spencer Castro^d

Joel Cooper^e

Michael Skinner^f

David Strayer^g

Andrew Heathcote^h

Author Note

^a Department of Psychology, University of Amsterdam, Amsterdam, The Netherlands,
email: u.bohm@uva.nl

^b Department of Psychology, University of Amsterdam, Amsterdam, The Netherlands,
email: d.matzke@uva.nl

^c Department of Psychology, University of Tasmania, Sandy Bay, Australia, email:
matthew.gretton@utas.edu.au

^d Department of Psychology, University of Utah, Utah, United States, email:
spencer.castro@utah.edu

^e Department of Psychology, University of Utah, Utah, United States, email:
joel.cooper@psych.utah.edu

^f Aerospace Division, Defence Science and Technology Group, Melbourne, Australia,
email: michael.skinner@dst.defence.gov.au

^g Department of Psychology, University of Utah, Utah, United States, email:
david.strayer@utah.edu

^hDepartment of Psychology, University of Tasmania, Sandy, Bay, Australia, email:
andrew.heathcote@utas.edu.au

The authors declare no competing financial interests. This research was supported by Competitive Evaluation Research Agreement (CERA) grant to AH, DM, DS and MS. DM is supported by a Veni grant (451-15-010) from the Netherlands Organization of Scientific Research (NWO).

Correspondence concerning this article should be addressed to Udo Boehm, Department of Psychology, University of Amsterdam, PO Box 15906, 1001 NK Amsterdam, The Netherlands, Email: u.bohm@uva.nl

Abstract

Human operators often experience large fluctuations in cognitive workload that can lead to sub-optimal performance, ranging from overload to neglect. Help from automated support systems could potentially address this issue, but to do so the system would ideally need to be aware of real-time changes in operators' cognitive workload, so it can provide help in times of peak demand and take advantage of troughs to elicit operator engagement. We used the ISO standard Detection Response Task (DRT) to measure cognitive workload approximately every 4 seconds in a demanding task requiring monitoring and refuelling of a fleet of unmanned aerial vehicles (UAVs). We showed that the DRT provided a valid measure that can detect changes in workload due to changes in the number of UAVs. We used a cross-validation analysis to assess whether measures related to task performance immediately preceding the DRT could be used to predict detection performance as a proxy for cognitive workload. Although the simple occurrence of task events had weak predictive ability, composite measures that tapped operators' situational awareness with respect to fuel levels were much more effective. We conclude that real-time prediction of operators' cognitive workload shows promise as an avenue for improved human-automation teaming.

keywords: Cognitive workload; Detection response task; Cross validation

Real-Time Prediction of Fluctuations in Cognitive Workload.

Introduction

Modern computer interfaces provide human operators with an unprecedented level of information and control over local and remote systems. When teamed with supporting automation that can take over routine functions or provide recommendations for the operator to action, these systems support tremendous increases in productivity. However, they also come with challenges, one of the most prominent being variations in workload. Fluctuations in workload can be associated with sub-optimal performance, which can lead to poor outcomes or complete failures of control. Underload during prolonged periods where automation takes care of all operations can cause mind wandering (Hawkins, Mittner, Boekel, Heathcote, & Forstmann, 2015), sometimes called “automation neglect”, leaving the operator ill prepared for emergencies (Quandt, 20017). Although automation can help to avoid overload, it lacks the situational awareness of human assistants (Strayer, Biondi, & Cooper, 2017), and it can make recommendations at times that strain the operator’s cognitive capacity, and can even create unnecessary failures by intruding at critical times with low-priority tasks. Thus, methods to monitor operators’ cognitive engagement and workload in real time could enhance human-autonomy teaming through “adaptive automation” systems (e.g., Aricò et al., 2016) that better manage workload fluctuations. Our aim here is to explore the possibility of an even greater enhancement through being able to *predict* cognitive workload in real time.

Traditional cognitive-workload measurement via retrospective questionnaires, although reliable (e.g., Hart, 2006), suffers from the twin disadvantages of subjectivity and a lack of temporal specificity. Psychophysiological measures of workload are available in real time (e.g., Wilson & Russell, 2003), but the International Standards Organization has opted for a behavioural approach, the detection-response task (DRT, ISO 17488, 2016). The DRT has been shown to provide a valid measure of cognitive workload that is more sensitive than psychophysiological measures and suitable for use in active tasks such as driving (Strayer et al., 2015). At random 3-5 seconds (s) intervals operators are required to press a button in response

to the occurrence of a simple stimulus (e.g., a buzzer on the skin or a light on a head-mounted wand). Primary task performance (e.g., driving) is not meaningfully impaired (Strayer et al., 2015), whereas slowed DRT response times (RTs) and response omissions provide a sensitive index of increased cognitive workload.

We use the DRT, where responses were made with a foot pedal to tactile stimulation from a buzzer, to measure workload in a multiple-asset monitoring task, simulating an operator managing the refuelling of a team of unmanned aerial vehicles (UAVs). Difficulty was manipulated over blocks where operators had to supervise either 3, 5 or 7 UAVs moving randomly over an ocean background and reflecting at screen boundaries. Pilot testing confirmed that this difficulty manipulation caused the task to vary from engaging but manageable to extremely demanding. Fuel levels, which could be monitored by hovering over a UAV with a mouse cursor, drained at a constant rate, and UAVs could only be refuelled below a critical level. When a UAV ran out of fuel it crashed and exploded and was replaced by a new UAV that appeared from the side of the screen with a high but variable fuel load. Points were awarded for checking fuel levels during the critical period, and for successfully refuelling, and points were subtracted for monitoring or refuelling at the wrong time, and for UAVs exploding. Although demanding, the task was relatively simple, so that participants could develop some expertise over the course of an initial one-hour session. Measurement over a second one-hour session performed on a subsequent day allowed us to examine the moderating effects of skill acquisition.

In an initial analysis we determined participants' sensitivity to the payoff scheme by examining how their scores, and the frequency of the five possible task events (e.g., successful or failed monitoring, successful or failed refuelling, and a UAV explosion) changed from the first to the second session. We then examined the effect of the difficulty manipulation on DRT performance, with the hypothesis that clear increases in RT and omissions with increased difficulty would validate its sensitivity to cognitive workload. These analyses set the stage for our main aim, to assess whether and how we could predict real-time fluctuations in cognitive workload as measured through the proxy of DRT performance.

We first examined the predictive ability of the occurrence of each of the five task events and the time scale over which they had influence within the resolution of DRT monitoring (3-5 s). The large differences in DRT performance that we observed as a function of the difficulty manipulation provided a benchmark against which to judge the effect of these predictors, and our analyses also took account of potential learning and fatigue effects over the course of a session. Although the initial analysis did not reveal any strong relationships, it did provide guidance which allowed us to develop a set of more refined predictors, most of which aimed to index the operator's situational awareness with respect to the fuel levels of the UAVs. These refined predictors were then used in a second analysis. Because of our focus on prediction, the primary method of inference in both cases was a five-fold cross-validation analysis. Linear mixed-effect models of DRT RT and omission rates were fit to data from mutually exclusive groups each consisting of around 20% of the participants and the models were assessed by their ability to predict corresponding DRT performance for the left-out groups of 80% of participants. The stability of these predictions over participants was assessed by comparisons among results for the five "folds".

Methods

Participants

Forty six participants were recruited from a participant pool database maintained by the University of Utah. Participants were compensated \$40 for completing the full study, which consisted of two one-hour sessions on different days. Ethical approval for the study was granted by the University of Tasmania's Social Sciences Human Research Ethics Committee (Ref No: H0016488).

Equipment and Task Software

Four Windows 7 desktop computers, equipped with second generation i7 processors with integrated graphics processors and 21 inch high definition screens, presented the primary tracking task. In pilot testing, these systems were verified as presenting the UAV task at a

steady 35 frame per second. Each computer interfaced with DRT hardware, which included a millisecond-accurate controller box running an Advanced Reduced-Instruction-Set-Computer (RISC) Machine (ARM) embedded processor and a rubberised vibrotactile motor for tactile alerting. Participants used their preferred hand to interact with the UAV task by operating a Microsoft Basic Optical Mouse connected via USB. The other hand was held at rest on the desk and used to collect Galvanic Skin Response (GSR) data. GSR readings are not analysed in the present report.

DRT units followed the performance specifications outlined in ISO 17488, with the exception that a foot pedal switch was used to collect responses rather than the finger mounted micro-switch. Responses were recorded on a TEMCo Heavy Duty cast aluminium foot switch. A small vibration motor provided the stimuli to be detected. The motor was fastened to participants' left shoulder using medical tape. When the motors are active they provide a gentle vibration for up to 1 second, every 3-5 s. Participants were instructed to press a foot pedal switch as soon as they detect the vibration stimulus. The DRT task was active throughout the experiment. Reaction times collected from the DRT units were gathered on a micro-controller that provided sub-millisecond accuracy. Each reaction-time event was timestamped using the host computer's clock. These timestamps were used to synchronize with data generated by the UAV task.

The UAV task was developed using the Unity® game engine (Unity Technologies, 2017). The core objective of the task is to refuel small UAV targets that randomly move around the computer screen. The task is presented from a birds-eye perspective with a static ocean image shown in the background. UAVs appear as small grey planes seen from above, randomly moving about the full space of the screen (see Figure 1). Fuel levels depleted at a rate of 5% per second. Clicking on the UAV when fuel entered the lower 25% of the gauge (a "Hit refuel") refuelled the plane and resulted in a score increase. Clicking on a UAV when fuel levels were above 25% (a "False-Alarm Refuel") resulted in a score deduction. If the fuel fully depleted (a "Miss Refuel"), the UAV crashed, the score was reduced, and a new UAV immediately entered the screen.

Figure 1: Task display with seven UAVs. Hovering the mouse over a UAV displays the fuel gauge which shows that the fuel level (yellow bar) is not yet in the green region where it can be refuelled. The time remaining in the block is displayed at the top, the difficulty at the bottom left, and the score at the bottom right of the screen.

Table 1

Payoff scheme.

Task event	Points
False-Alarm Hover Over (FAHO)	-50
False-Alarm Refuel (FAR)	-500
Hit Hover Over (HHO)	50
Hit Refuel (HR)	2000
Miss Refuel (MR)	-1000

When a new UAV entered the screen it began with 70-100% fuel level. Each time a refuel event occurred the refuelled plane gained between 50% and 100% of their full tank. When hovered over with the mouse cursor for at least 0.3s a circular fuel-gauge appeared around the UAV depicting the level of remaining fuel in that vehicle and its optimal refuelling zone. Participants were allowed one hover-over event to gauge initial fuel levels. Subsequent hover-over events rewarded points (a “Hit Hover Over”) if the remaining fuel level was under 25% or penalised points (a “False-Alarm Hover Over”) if it was over. Each scoring event was briefly presented as a signed points number above the relevant UAV while the total score (shown in the bottom right) was updated. Table 1 shows the points associated with the five different game events. As can be seen, payoffs were maximised by increasing the number of well-timed refuels and decreasing the number of exploding UAVs.

Design and Procedure

A 3 (task difficulty) x 2 (session) repeated-measures design was used. Task difficulty was varied by adjusting the number of UAV targets that are simultaneously present on the screen. Sessions consisted of 21 blocks. Task difficulty was varied in pseudo-random order across blocks, with each consecutive subset of 3 blocks containing an easy (3 UAV), medium (5 UAV), and hard (7 UAVs) block. Blocks lasted 2 minutes, allowing for a short break of 15 s between blocks, during which the score for different task events was tallied. Participants were given a

countdown 5 s to the start of the next block. Three practice blocks were given at the beginning of each session in a fixed easy, medium, hard order, and these were not analysed. We define a “trial” as the time period between two DRT prompts; hence a trial ranges randomly from 3-5 s according to a uniform distribution.

Upon arrival at the first session of the study, participants were provided with a consent form to sign, acknowledging the purpose of the study and their level of compensation. They were then introduced to the GSR sensors and the DRT task. Participants were read a script that informed them of their goal to refuel each unmanned aerial vehicle (UAV) by hovering to check the fuel level and then left-clicking to refuel, the difficulty levels (easy, medium, hard), and the number of UAVs that pertained to each level (3,5, and 7 respectively). Participants were also informed that the experiment would consist of 3 practice blocks, during which time they would be allowed to ask questions and receive additional instruction or clarification, followed by 21 test blocks, during which time they would not be allowed to communicate with the experimenter. Finally, they were informed that the blocks would persist for two minutes at a time, that the experiment would last approximately 1 hour, and that they were free to stop for any reason at any time and still receive full compensation of \$20 per hour. Following this, participants completed the practice block and the seven experimental sets. The second session occurred between one and three days after the first session and consisted only of the practice and experimental blocks. When complete, participants were compensated for their time and thanked for participation.

Results

Points and Task Events

The goal of our first analysis was to establish whether participants were sensitive to the payoff structure of the UAV task. If they are sensitive to payoffs, over the course of the two experimental sessions participants should learn to avoid task events that are associated with a large loss in points and to seek strategies that maximise the points they earn. This analysis was based on the complete data of 46 participants performing 21 experimental blocks on each of two testing days. Note that the number of task events naturally tends to increase with difficulty

simply because there are more UAVs on the screen, hence the effect of difficulty is not of primary interest in this analysis except to the degree that it interacts with day. The same is true of points, beyond the observation that participants were on average able to make use of the greater number of opportunities to earn points as difficulty increased and so the average score increased with difficulty.

Figure 2 summarises different aspects of participants' performance in the UAV task. The top left plot shows the average number of points participants earned in each experimental block on day 1 and day 2 in each of the three difficulty conditions. As can be seen, the number of points earned increased over days in all three difficulty conditions, with average increases being largest in the most difficult condition. We assessed these qualitative observations statistically using linear mixed effects models (LMEs) (Pinheiro & Bates, 2000). To this end, we created three models with the average number of points earned per block as the dependent variable and a random intercept for each participant. The first model, $\mathcal{M}_{diff}$, only included difficulty as a predictor, the second model, $\mathcal{M}_{diff+day}$, additionally included an additive effect of day, and the third model, $\mathcal{M}_{diff \times day}$, furthermore allowed for an interaction between difficulty and day. The results of the formal model comparison are shown in the first row of Table 2. As can be seen, *AIC* as well as *BIC* were smallest for $\mathcal{M}_{diff \times day}$, indicating that the average number of points earned increased from day 1 to day 2, with gains being larger in the most difficult condition.

The second plot in the top row of Figure 2 shows the average number of "False-Alarm Hover Over" (FAHO) events per experimental block on day 1 and day 2 in each of the three difficulty conditions. This plot suggests that the number of FAHOs did not change across days. We used generalised linear mixed effects models (GLMs) to assess this observation statistically. Specifically, we fitted the average number of FAHOs using Poisson family models with a logarithmic linking function. Similar to the analysis for the average number of points earned, we created three models that allowed for a random intercept for each participant. The results of the formal model comparison are shown in the second row of Table 2. The Akaike Information Criterion (*AIC*) (Akaike, 1973) indicated a slight advantage of $\mathcal{M}_{diff+day}$ over the

Figure 2: Overview of points earned and frequency of different task events (see Table 1 for definitions) per block for both testing days and the three difficulty conditions. Thick error bars without caps show mean ± 1.96 SEM, thin error bars with caps show 95% CIs based on mixed effects models (see Baguley, 2012).

Table 2

Results of Model Comparison for Effect of Testing Day and Difficulty on Points and Task Events (see Table 1 for definitions).

Dependent Variable	AIC			BIC		
	$\mathcal{M}_{diff}$	$\mathcal{M}_{diff+day}$	$\mathcal{M}_{diff \times day}$	$\mathcal{M}_{diff}$	$\mathcal{M}_{diff+day}$	$\mathcal{M}_{diff \times day}$
Average Points	41398	41306	41272	41426	41339	41316
FAHO	13916	13913	13914	13938	13941	13953
FAR	12167	12169	12137	12189	12197	12176
HHO	11924	11923	11915	11946	11950	11954
MR	8579.6	8446.3	8431.3	8601.8	8474.2	8470.3
HR	11076	11038	11036	11098	11066	11075

Note. Values in bold indicate the best-fitting model.

other two models, while the Bayesian Information Criterion (*BIC*) (Schwarz, 1978) indicated strong support for the simplest model, $\mathcal{M}_{diff}$. These formal results align with the qualitative observations in that they suggest that there was not substantial evidence for a change in the number of FAHOs across days.

The third plot in the top row of Figure 2 shows the average number of “False-Alarm Refuel” (FAR) events per experimental block across difficulty conditions and days. The plot suggests that the number of FARs did not change systematically across difficulty conditions, but rather increased in the easiest condition and decreased in the hardest condition. We again assessed these observations using Poisson family GLMs that were specified as for the analysis of the FAHOs. The results of formal model comparison, shown in the third row of Table 2, indicate a strong preference for $\mathcal{M}_{diff \times day}$. However, as $\mathcal{M}_{diff+day}$ performed worse than $\mathcal{M}_{diff}$ in terms of *AIC* as well as *BIC*, the latter effect is solely driven by the interaction of day and difficulty. Due to the absence of a main effect of difficulty this interaction effect should not be interpreted. These results suggest that there was no statistically reliable change in the number of FAR events

between testing days.

The first plot in the bottom row of Figure 2 shows the average number of “Hit Hover Over” (HHO) events per experimental block across difficulty conditions and days. There only appears to be a difference between day 1 and day 2 in the easy condition but not in the remaining difficulty conditions. Our statistical analysis used three GLMs that were specified as the models for the FAHO events. The results of the model comparison are shown in the fourth row of Table 2. *AIC* indicates equal performance of the $\mathcal{M}_{diff}$ model and the $\mathcal{M}_{diff+day}$ model, which means that there was no evidence for a main effect of day. *BIC*, on the other hand, prefers $\mathcal{M}_{diff}$ over both alternative models. These results suggest that there was no change in the number of HHO events between testing days.

The second plot in the bottom row of Figure 2 shows the average number of “Miss Refuel” (MR) events per experimental block across difficulty conditions and days. The average number of MRs decreased from day 1 to day 2 in all difficulty conditions, with decreases being largest in the most difficult condition. We again tested these observations statistically by comparing three GLMs. This model comparison confirmed the qualitative observations, with both *AIC* and *BIC* indicating a substantial advantage of $\mathcal{M}_{diff \times day}$ over the other models. These results thus suggest that the number of MR events decreased substantially between testing days.

Finally, the third plot in the bottom row of Figure 2 shows the average number of “Hit Refuel” (HR) events per experimental block across difficulty conditions and days. The average number of HRs seems to increase from day 1 to day 2, with gains being largest in the most difficult condition. A formal model comparison confirmed these observations. *AIC* indicated a strong advantage of $\mathcal{M}_{diff+day}$ over $\mathcal{M}_{diff}$ but only a negligible advantage of $\mathcal{M}_{diff \times day}$ over $\mathcal{M}_{diff+day}$. *BIC* indicated a clear advantage of $\mathcal{M}_{diff+day}$ over both competitor models. These results suggest that there was a substantial increase in the number of HR events from day 1 to day 2.

Taken together, the analysis of the task data indicates that participants learnt to improve the number of points earned. This increase from day 1 to day 2 was principally driven by those

UAV-task events that were associated with considerable penalties and rewards, indicating that participants adjusted their behaviour according to the payoff manipulation. Subsidiary analyses of arcsine-transformed mean number of events per participant using Bayesian ANOVAs confirmed these conclusions (see Supplementary Materials).

Detection-Response Task

We carried out our analysis of the DRT data in two steps. In a first step we evaluated whether DRT behaviour was sensitive to the workload manipulation in the UAV task. In a second step we explored the relationship between UAV-task events and DRT behaviour. We were interested in how different types of task events affect RTs and omissions on the DRT, and how these relationships evolve as participants gain more experience with the UAV task.

Difficulty Manipulation. Figure 3 shows the effect of the number of UAVs on mean RTs and omission rates. As can be seen, both mean RTs and omission rates increased as the number of UAVs increased, which suggests that the increased difficulty in the UAV task reduced participants' DRT performance. Moreover, RTs and omission rates decreased from day 1 to day 2 in all difficulty conditions, suggesting that cognitive workload decreased as participants became more skilled at the task. RTs, for instance, decreased on average by 5.4% in easy blocks and by 3.7% in medium and high difficulty blocks. We tested these qualitative observations statistically using LMEs for the log-RTs and binomial family GLMs for the omission rates. We specified the same three models of each type as in earlier analyses that included random intercepts for participants and used log-RTs or omission rates as dependent variable. The results of the model comparison are shown in Table 3. For log-RT (top row), both *AIC* and *BIC* indicated $\mathcal{M}_{diff+day}$ as the best model, which supports the observation that log-RT decreased across days and was longer for more difficult conditions but the change in log-RT between days was comparable for all difficulty conditions in a proportional sense (as additivity on a log scale indicates a constant proportional difference).

For omissions, (bottom row in Table 3), *AIC* preferred $\mathcal{M}_{diff+day}$ over $\mathcal{M}_{diff}$ and $\mathcal{M}_{diff \times day}$ over $\mathcal{M}_{diff+day}$, which suggests that omission rate increased with increasing task

Figure 3: Mean response time (RT) and omission rate for the three task difficulty levels, separately for the two testing days. Thick error bars without caps show mean ± 1.96 SEM, thin error bars with caps show 95% CIs based on mixed effects models (see Baguley, 2012).¹

difficulty and decreased across days, but the decrease was differently affected by different difficulty conditions. *BIC*, on the other hand, suggests $\mathcal{M}_{diff+day}$ as the best fitting model. Nevertheless, both model selection criteria agree that task difficulty and testing day both affected omission rates.

Taken together, these results confirm that participants' DRT performance was sensitive to the difficulty manipulation, with higher workload in the more difficult condition resulting in decreased DRT performance. Moreover, there appear to be clear practice effects reflected in improved DRT performance on day 2 compared to day 1.

Predicting Workload Fluctuations

In the second step of our analysis we investigated how the UAV task affects cognitive workload on a real-time basis, as reflected in fluctuations in DRT performance. In particular,

¹Model-based CIs are mean-adjusted and appropriately transformed to account for the fact that RTs and omissions are modeled on the log and logit-scale in our linear mixed models, respectively.

Table 3

Results of Model Comparison for Effect of Testing Day and Difficulty on DRT Performance.

Dependent Variable	AIC			BIC		
	$\mathcal{M}_{diff}$	$\mathcal{M}_{diff+day}$	$\mathcal{M}_{diff \times day}$	$\mathcal{M}_{diff}$	$\mathcal{M}_{diff+day}$	$\mathcal{M}_{diff \times day}$
Log-RT	61493	61399	61403	61536	61451	61472
Omissions	43725	43577	43569	43761	43621	43632

we were interested in predicting cognitive workload based on recent events in the UAV task.

To this end, we took a cross-validation approach to our statistical analysis. We first identified a number of observable task events that might potentially affect workload as measured by the DRT. Subsequently, we refined and complemented these predictors with other variables that should reflect psychological influences on participants' cognitive workload. To guarantee that sufficient data were available for all these analyses, we removed the data of all participants who failed to respond to more than 45% of DRT prompts in two or more difficulty conditions. This resulted in the removal of seven participants on day 1 and five participants on day 2.

Time on Task. Before performing the predictive analysis, we investigated time on task, as we expected that fatigue might affect DRT performance independent of specific UAV task events and cognitive workload. Figure 4 shows the development of DRT performance across blocks. The top row shows how mean RT changes across experimental blocks on day 1 (left) and day 2 (right). There does not appear to be a systematic relationship between experimental block and RT on day 1, but on day 2 mean RT appears to increase across blocks, consistent with a fatigue effect. The apparent lack of a relationship between RT and block on day 1 might be due to any fatigue effect being masked by learning in the UAV task that decreased its cognitive demands.

The bottom row shows how mean omission rate changes across experimental blocks. There appears to be a slight increase in omission rate across blocks on both days. However, this tendency is difficult to interpret due to the high standard errors in the observed mean omission

Figure 4: Development of DRT performance across blocks. Top row: mean RT across blocks. Bottom row: mean omission rate across blocks. Error bars show mean ± 1.96 SEM.

rates.

To assess these observations statistically, we used LMEs models to test the relationship between RT and block, and binomial family GLMs to test the relationship between omission rate and block. In light of the results for the relationship between difficulty and RT/ omission rate above, we considered three different models with RT/ omission rate as dependent variables and random intercepts for each participant. The first model, $\mathcal{M}_{diff+day}$ only used day and difficulty as predictors. The second model, $\mathcal{M}_{diff+day+block}$, additionally allowed for an effect of block. The third model, $\mathcal{M}_{diff+day \times block}$ furthermore allowed for an interaction effect between block and day. The results of the model comparisons are shown in Table 4. For the models

Table 4

Results of Model Comparison for Effect of Block on DRT Performance.

Dependent Variable	AIC			BIC		
	$\mathcal{M}_{diff+day}$	$\mathcal{M}_{diff+day+block}$	$\mathcal{M}_{diff+day \times block}$	$\mathcal{M}_{diff}$	$\mathcal{M}_{diff+day}$	$\mathcal{M}_{diff \times day}$
Log-RT	57652	57641	57620	57704	57701	57689
Omissions	34573	34541	34541	34617	34594	34602

with RT as dependent variable, shown in the top row, *AIC* indicated $\mathcal{M}_{diff+day \times block}$ as the best model, supporting an effect of block on RT, but only on day 2. *BIC*, on the other hand, preferred $\mathcal{M}_{diff+day+block}$, which suggests that block affected RT equally on both days. The results for the models with omissions as dependent variable are shown in the bottom row of Table 4; both *AIC* and *BIC* suggest that $\mathcal{M}_{diff+day+block}$ accounts best for the observed data, which means that block affected omission rate equally on both days.

Taken together, these results suggest the presence of fatigue effects that led to a decline in DRT performance across blocks. The effect of block on RT might have been mitigated by learning effects on day 1, leading to an interaction between block and testing day.

Prediction from UAV-task events occurring at different time scales. We first considered the occurrence of different UAV-task events as potential predictors. The UAV-task events FAHO, FAR, MR, and HR were defined as in our previous analysis. However, instead of HHOs, we considered a new type of game event, “Hover Over” (HO), which indicates whether a participant requested information about any UAV, irrespective of whether that UAV had reached a critical fuel level. FAHO and HHO events were combined into HO to assess whether HHO events convey additional information beyond FAHO events about participants’ mental workload.²

There is no a priori way of determining the time scale on which task events affect workload. That is, we do not know whether the occurrence of a task event will increase mental workload, and hence affect DRT performance, 3, 4, or even 5 seconds later. Therefore, we

²We also repeated many of the analyses reported in this section with separate FAHO and HHO events with substantially similar results. The results of these analyses are reported in the Supplementary Materials.

repeated our analysis three times, once for RT and omissions if an event of a specific type occurred up to 3 s prior to the DRT prompt, once for the case that an event occurred up to 4 s prior to the DRT prompt, and once for the case that an event occurred up to 5 s prior to the DRT prompt. We focused on the 3-5 s intervals because they spanned the range of intervals between DRT stimuli.

Figure 5 shows how different UAV events affect RT on different time scales. Each plot shows the difference in mean RT between trials on which the specific event occurred and trials on which the event did not occur, for day 1 and day 2. The top row shows the results for events occurring up to 3 s prior to the DRT prompt, the middle row shows the results for a 4 s time window, and the bottom row shows the results for a 5 s time window. As can be seen, on both days the occurrence of any event during any of the three time windows appears to delay the detection response. Differences between different classes of events appear to be relatively small and differences between day 1 and day 2 for the same type of event seem negligible. However, it should be noted that different types of events might be considerably correlated, either because some types of game events logically precede others, or because several different events occur within the same time window. We will, therefore, rely on the following multivariate statistical analysis to disentangle these correlations.

As we found clear learning effects over days, which may change what events, if any, are predictive, we analysed the data for each testing day separately. We split the available data for each day into five non-overlapping subsets, the cross-validation folds, with approximately equal numbers of participants. Each fold served in turn as the training set and the remaining four subsets served as cross-validation sets.

Similar to the analyses already reported, our cross-validation analysis used LMEs as statistical descriptions, as these can flexibly account for random effects on the participant level, and thus minimise the influence of individual differences. In the first set of analyses, all models used log-RT as the dependent variable, and included a random intercept for each participant. We considered five types of UAV-task events as predictors (FAHO, FAR, MR, HR, and HO) that

Figure 5: Effect of different UAV-task events on response time (RT) for different time windows prior to the DRT prompt. Top row: 3s time window; middle row: 4s time window; bottom row: 5s time window. Error bars show mean difference ± 1.96 SEM.

might influence mental workload and thus affect log-RT. We considered experimental block as a potential proxy for learning and fatigue effects that might influence log-RT. Moreover, based on the results of our previous analyses, we also included difficulty as a potential predictor of log-RT. These seven potential predictors can be combined into $2^8 = 128$ candidate models. However, using 128 models in 5-fold cross-validation would result in formidable computational costs. To keep computational costs at an acceptable level, we only considered all possible combinations of the five UAV-task events as predictors, with difficulty and experimental block always included as predictors. In addition, we considered four baseline models that included only block and difficulty, only block, only difficulty, or only an overall intercept as predictors. For reasons of computational efficiency, we also only allowed for additive effects, with no interactions between predictors.

We fitted each model to the training set and used the resulting parameter estimates to predict the data of the remaining four cross-validation sets. To measure predictive performance, we computed the mean squared deviation (MSD) between observed and predicted log-RT for each cross-validation fold, and normalised it by subtracting the average MSD across all five folds (NMSD). This analysis was repeated for time windows of 3 s, 4 s, and 5 s prior to the DRT prompt.

Figure 6 shows the results of the cross-validation analysis for different time windows and testing days. Results are displayed in order of increasing mean NMSD across folds. A ‘+’ in the label below each data point indicates that the corresponding variable was included as a predictor in the model. NMSD showed similar patterns on both testing days and across time windows, which suggests that UAV-task events can affect DRT performance over a time span of several seconds. NMSD was considerably lower for models that included difficulty and experimental block as predictors than for the baseline models that did not include these predictors. This means that both predictors convey essential predictive information. Among the five UAV events, only HO appears consistently in all models with good predictive performance, for all time windows and testing days. HO should therefore be considered an essential predictor. HR is also included in

a large number of models with low NMSD, but also appears in some models with higher NMSD. On day 2, and to a lesser extent on day 1, Fold 4 produced results that were inconsistent with the other folds.

Figure 7 shows how different UAV events affect omission rates on different time scales. Each plot shows the difference in mean omission rate between trials on which the specific event occurred and trials on which the event did not occur, for day 1 and day 2. The top row shows the results for events occurring within a time window of 3 s prior to the DRT prompt, the middle row shows the results for a 4 s time window, and the bottom row shows the results for a 5 s time window. Similar to the results for the effect on RT, all event types except for FAHO and HO seem to affect omissions to the same extent on both days, and differences between these event types seem negligible. FAHO and HO, on the other hand, only appear to influence omissions on day 2 but not on day 1. However, we again note that the occurrence of different event types is highly correlated and further analysis is needed to disentangle these correlations.

Our further statistical analysis again used a five-fold cross-validation approach, now using binomial family GLMs to predict omissions. We considered the same set of predictors as in the analysis of log-RT and specified the models in the same way. Predictive performance was measured in terms of NMSD between the observed and predicted probability of an omission.

The results of the cross-validation analysis for different time windows and testing days are shown in the same way as for RT in Figure 8. NMSD showed similar patterns on both testing days and across time windows. However, there are considerable differences visible between different cross-validation folds, with some folds showing the reverse pattern of the remaining folds. This trend is visible but somewhat weaker on day 2 than on day 1. Nevertheless, these patterns suggest that there are no universally good predictors of omissions. Similar to the results for RT, models that included difficulty and experimental block as predictors, in addition to UAV-task events, had lower NMSD in three of the five folds. In the remaining two folds the best performance was for models that only had difficulty and experimental block as predictors.

Among the UAV-task events only HO was consistently included in all models with low

Figure 6: Cross-validation results for different UAV-task event predictors of log-RT. Normalised mean squared deviations for different LMEs for different time windows prior to the DRT prompt are shown. Top row: 3s time window; middle row: 4s time window; bottom row: 5s time window. Left column: results for day 1; right column: results for day 2. Red squares show the mean across folds.

Figure 7: Effect of different UAV events on omissions for different time windows prior to the DRT prompt. Top row: 3 s time window; middle row: 4 s time window; bottom row: 5 s time window. Error bars show mean difference ± 1.96 SEM.

mean NMSD across folds. However, this pattern again only held for three of the five folds, and differences in NMSD between models with different subsets of the UAV-task events as predictors were negligible compared to the differences in NMSD for the same model across folds.

Taken together, the results of our cross-validation analysis so far suggest that participants' DRT performance is strongly influenced by global levels of difficulty and fatigue. Of the UAV events considered, only HO seems to have a moderate influence on DRT performance. A possible interpretation of this effect is that, as their workload increases, participants tend to lose track of the state of the different UAVs. To compensate for this loss of situational awareness, they need to hover over the UAVs to update and refresh their mental representation of the situation. In the next section we investigate this suggestion by investigating the predictive ability of statistics constructed to measure participants' situational awareness with respect to fuel load.

Refined Predictors. We considered five new predictor variables in our second cross-validation analysis, in addition to difficulty and block. The first two new predictors were based on requesting fuel information as the previous cross-validation analysis suggested that this was indicative of increased cognitive workload. As the results of the previous cross-validation analysis were similar for different time windows, and were slightly more stable across folds for shorter time windows, we considered only events within a 3 s time window.

As more fine-grained information might be provided by the frequency with which participants request fuel information within a certain time period, our first new predictor was the number of HO events within the last 3 s, “#HO3s”. Secondly, the number of FAHO events within a certain time period might convey more specific information about participants' current workload than the number of HO events. The total number of UAVs a participant can request fuel information for is limited by the speed at which they can move the mouse between UAVs. If participants hover over UAVs to compensate for a loss of situational awareness, they should be more likely to request fuel information for UAVs that are at a high fuel level. Consequently, although the total number of HO events might remain fairly stable with cognitive workload, the relative number of FAHO events might increase as cognitive workload increases. We therefore

Figure 8: Cross-validation results for different UAV-task event predictors of omissions.

Normalised mean squared deviations for different binomial family generalised linear mixed effect models for different time windows prior to the DRT prompt are shown. Top row: 3 s time window; middle row: 4 s time window; bottom row: 5 s time window. Left column: results for day 1; right column: results for day 2. Red squares show the mean across folds.

considered the number of FAHO events within a 3 s time window, “#FAHO3s”, as an additional potential predictor.

The third and fourth potential predictors were also related to participants’ knowledge of fuel levels, but focused on the fact that as participants received a particularly high penalty for exploding UAVs, so their workload should mostly be driven by the UAVs with the lowest fuel levels. If they maintain an accurate mental representation of the current situation, participants should mostly request fuel information for the UAVs with the lowest fuel levels. Hence, our third potential predictor was the ordinal position of the UAV participants hovered over most recently, “OrdFuel”. That is, did the last UAV checked in the 3 s before a DRT response have the lowest fuel (OrdFuel=1), the second lowest (OrdFuel=2), and so on.

Multiple UAVs reaching a critical fuel level in a short period of time should increase participants’ cognitive workload considerably more than a single UAV reaching a critical fuel level, as that might leave insufficient time to refuel all of them before one or more crashed. Our fourth new predictor combines the fuel level of all UAVs on screen, but gives more weight to lower fuel levels, “WeightFuel”. This computation was based on the Minkowski distance of the vector of fuel levels from 0, normalised by the number of UAVs:

$$\text{WeightFuel} = \left(\sum_{k=1}^K \frac{f_k^p}{K} \right)^{1/p}, \quad (1)$$

where K is the total number of UAVs on screen and f_k is the fuel level of the k -th UAV. Here we set $p = 1/4$ (we discuss this choice further below).

We added one final, and somewhat different, potential predictor, whether there was a omission on the previous DRT trial, “MissPrevDRT”, in response to the overall poorer prediction of DRT omissions in the previous cross-validation analysis. An omission on the previous trial might predict an omission or a slower RT on the present trial, due to a period of sustained higher workload. However, an omission on the present trial might also be due to reasons not directly related to workload, such as the occurrence of periods where participants neglected instructions to respond to the DRT and instead responded only to the UAV task. To the degree that the latter process was in play we might expect MissPrevDRT to predict omissions but not RT.

As information about the preceding trial was not available for the first trial of each block, we excluded these trials from all further analysis steps. Moreover, the data recording system used for the experiment only logged information about UAVs when participants interacted with a UAV. This means that we could only compute the new predictor variables for trials on which participants had interacted with a UAV within a period before being presented the DRT prompt. We therefore removed all trials on which participants had not interacted with a UAV or on which the last interaction with a UAV occurred more than 100s before the DRT prompt. This resulted in the removal of only a small number of trials, a total of 280 from the data set for day 1 and a total of 253 from the data set for day 2.

We used the same cross-validation setup as in the previous analysis to test whether any of the five predictors conveys additional information about participants' DRT performance. We again considered all possible combinations of the five predictors together with difficulty and experimental block. In addition, we considered four baseline models that included only difficulty and experimental block as predictors, only difficulty, only experimental block, or only an overall intercept.

Figure 9 shows the LME results for predicting log-RT. Models are ordered by increasing average NMSD. As can be seen, models that included block and difficulty in addition to the refined predictors performed better than the baseline models that did not include any refined predictors. Among the refined predictors, #HO3s and WeightFuel were consistently included in the models with the lowest average NMSD on both testing days. These results hold across all cross-validation folds on day 1. On day 2, however, this pattern only holds for four of the five folds. In Fold 4 more parsimonious models that included either OrdFuel or #FAHO3s, but not both, performed best. Nevertheless, by and large, #HO3s and WeightFuel were most consistently included in the models with the lowest NMSD across folds.

The results for predicting omissions are shown in Figure 10. Models that included block and difficulty in addition to the refined predictors performed better than the baseline models that did not include any refined predictors. Among the refined predictors, MissPrevDRT was

Figure 9: Cross-validation results for predicting log-RT based on refined predictors.

Figure 10: Cross-validation results for predicting omissions based on refined predictors.

consistently included in all models with low NMSD. WeightFuel was also consistently included in many of the models with low NMSD. However, WeightFuel was also included in a group of models with higher NMSD, although it did not appear in any of the models with the highest average NMSD. All other refined predictors were also included in models with high NMSD and thus most likely do not convey critical predictive information. These patterns were consistent across testing days and folds, and noticeably much more consistent across folds than in the previous cross-validation analysis.

Figure 11 shows approximate effect sizes for the three refined predictors that performed best for RT and/or omissions, and for Block and Difficulty. The left panel shows the unstandardised beta weights from an LME with RT as dependent variable and WeightFuel, #HO3s, MissPrevDRT, Block, and Difficulty as predictors, and a random intercept for each participant. WeightFuel and #HO3s were z-standardised within each participant so that beta

Figure 11: Approximate effect size of different predictors on RT and log odds of omissions. Bars show the unstandardised beta weights from a linear mixed model for response time (RT) and a generalized linear mixed model for log odds of omissions. Predictors WeightFuel and #HO3s were z-standardised within each participant.

weights indicate the change in predicted RT if the predictor value increases by 1 SD. Difficulty was entered as a factor with the case of 3 UAVs as the reference level, which is absorbed in the intercept term of the model, and levels +2UAVs and +4UAVs representing the cases of a total of 5 and 7 UAVs, respectively. The bar for Block shows the total effect size for completing 21 blocks. Black bars show the beta weights for day 1, grey bars show the beta weights for day 2.

As can be seen, effects were similar in size on both testing days for all predictors except for Block and MissPrevDRT. WeightFuel had a negative effect on RT, with a one standard deviation increase in weighted fuel levels decreasing RT by about 40ms. #HO3s had a slightly positive effect, with each additional HO event increasing RT by about 20ms. MissPrevDRT had a negligible negative effect on RT on day 1 and a small positive effect on day 2. Block had a

negligible effect on RT on day 1 and a considerable positive effect on RT on day 2, resulting in an increase in RT by about 60ms after 21 blocks. Increasing the number of on screen UAVs by 2 to a total of 5 increased RT by 50ms, and increasing the number of on-screen UAVs by 4 to a total of 7 increased RT by 90ms.

The right panel of Figure 11 shows the unstandardised beta weights from an binomial family GLM with log odds of omission as the dependent variable.³ Predictors were specified as for the RT model. Results were similar to those for RT except for MissPrevDRT. WeightFuel, +2UAVs, and +4UAVs considerably increased the log odds of an omission. The log odds for an omission also increased considerably over the course of 21 blocks, and this effect was more pronounced on day 2. Increases in the number of HO events had a comparatively small effect on the log odds of an omission. Finally, in contrast to the results for RT, MissPrevDRT had a large effect on both days, leading to a considerable increase in the log odds of an omission.

One question that remains is the influence our choice of the exponent of the Minkowski distance in Equation (1) has on the predictive performance of WeightFuel. Our choice of $p = 1/4$ was based on the assumption that participants attach higher importance to UAVs with lower fuel levels. According to this reasoning, computing the predictor with a different exponent $p < 1$ should yield similar results, but using an exponent $p > 1$ should yield worse predictive performance. To test this hypothesis, we repeated the cross-validation analysis for log-RTs with four different values of p : $p = 1/8$ (with corresponding WeightFuel predictor WF $p=1/8$), $p = 1/4$ (WF $p=1/4$), $p = 2$ (WF $p=2$), and $p = 4$ (WF $p=4$). The results of this analysis are shown in Figure 12.

Models are again presented in order of increasing average NMSD. As can be seen, for the two values of the exponent smaller than 1, the corresponding WeightFuel predictor is consistently included in the models with the lowest average NMSD. For the two values of the exponent

³The log odds of an omission are computed as $\log(p/(1 - p))$, where p is the probability of an omission.

Although the mapping is not linear in p , it is strictly increasing, which means that larger beta weights correspond to a higher probability of an omission.

Figure 12: Cross-validation results for predicting log-RT based on refined predictors with different exponents for the weighted average fuel-load predictor. Note that in order to make comparison easy the results in the top right panel repeat those in Figure 9.

greater than 1, on the other hand, the corresponding WeightFuel predictor is included in several of the worst models.

Discussion

In this paper we attempted, for the first time to our knowledge, to explore the possibility of *predicting* real-time fluctuations in cognitive workload. To do so, we examined performance in a task where an operator had to manage the refuelling of a fleet of Unmanned Ariel Vehicles (UAVs) in two sessions, broken up into 24 two-minute blocks, occurring on different days. We used the Detection-Response Task (DRT), a new ISO standard (ISO 17488, 2016) methodology,

to quantify cognitive workload in terms of responses to tactile stimulation occurring every 3-5 s. DRT performance is defined along two related dimensions, response time (RT) and omissions. We therefore analysed RT and omission data separately to identify potential predictors of cognitive workload. Our aim was to determine if DRT RT and/or omission rates, both of which increase with cognitive workload, could be predicted from the recent history of the UAV task and the operator's interactions with it.

Preliminary analyses showed that over the two days of practice operators improved their performance in response to a payoff scheme, increasing occurrences of the two types of positive task events—checking for and replenishing fuel when it was low—and reducing occurrences of the three types of negative events—checking fuel and trying to refuel when the level was high, or allowing a UAV to run out of fuel so that it crashed. Both DRT omissions and RTs increased strongly with task difficulty, which was manipulated by varying the number of UAVs that had to be managed. In contrast, these measures decreased from day 1 to day 2, indicating that the cognitive demands of the task decreased with practice, although the difficulty effect remained clearly present on the second day. There was also a clear effect of blocks within sessions (i.e., time-on-task) with omissions increasing over the course of the sessions on both days. The same effect was seen on RT during day 2, but not on day 1, which may have been due to an underlying increase due to fatigue being mitigated by a decrease due to learning.

Our initial analysis assessed whether the occurrence of each of five task events in the last 3 s, 4 s, or 5 s predicted cognitive-workload fluctuations. We were not merely interested in assessing whether DRT performance differed significantly as a function of whether an event occurred or not. Instead, because of our focus on prediction, we used a much more demanding cross-validation test, which assessed whether including a task-event predictor in a model fit to DRT performance from one (training) subset of participants would enable that model to better predict DRT performance for another (test) set of participants. In all models, we included difficulty and time-on-task as predictors and then compared the effects of adding all possible combinations of the five events. This was repeated for five different folds—divisions of our

participants into training and test sets—to assess the stability of predictions over individual differences.

The initial analysis confirmed that difficulty and time-on-task (block) needed to be included, but among the task events only checking the fuel level, either successfully (i.e., when the UAV was ready for fueling) or unsuccessfully, had predictive ability, with little variation in its performance across the three time intervals. Similar effects were found for RT and omissions, but for omissions the effect of fuel checking was weaker and there was much greater variation over the folds. We took the effect of checking fuel as suggesting that, as workload increases, participants lost track of the UAV's fuel level, and in order to compensate for this loss of situational awareness they needed to more frequently check their fuel levels.

We used the results of the initial analysis to create a refined set of five real-time predictors, and focused on only the last 3 s as that provided slightly more consistent results in the initial analysis. Four predictors sought to better measure awareness about fuel load. The first simply refined the most successful measures from the initial analysis, by counting the number of times fuel levels were checked, either successfully or unsuccessfully. The second counted the number of times they were unsuccessfully checked, on the assumption that doing so was more indicative of a failure of situational awareness. The third, the rank of the fuel load (lowest to highest) of the last-checked UAV, was created with a similar rationale (i.e., that higher ranks corresponded to reduced situational awareness). The fourth was the average fuel levels of all UAVs with a higher weight given to lower loads because they were more indicative of an incipient UAV crash. Finally, we used a fifth predictor in the hope of specifically improving the poorer and more variable prediction performance for omissions that we found in the initial analysis, the occurrence of an omission for the previous DRT stimulus. We speculated that omissions might occur in periods where participants forgot to respond to the DRT task, so an omission for the present DRT stimulus might be indexed by a previous omission. A previous omission might also predict a current omission occur due to a prolonged period of high workload, in which case it would also predict RT, whereas forgetting would seem to predict a selective

influence on omission rate.

The results of the second analysis were much more encouraging for the prospect of developing real-time predictions. Once again difficulty and time-on-task (block) were included in the best models for both RT and omissions. For RT there was also clear support for the refined version of the best predictor in the initial analysis, the number of fuel checks in the last 3 s and also for the weighed average fuel load. The latter predictor was also supported for omissions, as was the occurrence of a miss on the previous trial, and in contrast to the first analysis, the pattern of results was stable over folds. This was also true for RT on day 1, but on day 2 there was some inconsistency for one of the five folds, with the ordinal position of the fuel load for the last check and the number of unsuccessful fuel checks receiving support.

Estimates of effect sizes suggested that a one standard deviation decrease in the weighted fuel level, and a one standard deviation increase in the number of fuel checks had approximately the same effect on RT as an increase of 2 UAVs. For omissions, there was a similar effect of the weighted fuel level. Hence, these real-time predictors appear to be able to account for the same sort of substantial change in cognitive workload as an increase in the level of difficulty. For omissions, the effect of failing to respond to a previous DRT stimulus had an even greater effect than the additions of two UAVs. The fact that this effect was selective to omissions suggests that it might index a somewhat different mechanism to the other predictors, perhaps related to goal neglect with respect to DRT responding. This might, in turn, be related to cognitive workload, but might also have other potential causes. The effect of the number of fuel checks was also selective, but for RT, also suggesting that omissions and RT may have partially different drivers.

Although the second analysis was clearly exploratory, so that we cannot rule out that other better predictors exist, it did at least confirm that real-time prediction appears viable. It was also salutatory, as it indicated that at least one route to successful prediction relied on quantifying factors related to the mental states that operators required to enable successful control, which in the case of our task was situational awareness about UAV fuel levels. Given that there was some divergence between RT and omission effects it would be desirable in future work to attempt

to develop an analysis that simultaneously accounted for both measures. Statistical models are typically not equipped to account for such correlated, multidimensional behavioural measures. In a future investigations we plan to use the reduced set of candidate predictors in combination with evidence-accumulation models (e.g., Tillman, Strayer, Eidels, & Heathcote, 2017) extended to simultaneously model RT and omission data, in order to assess whether this yields better real-time prediction.

References

- Akaike, H. (1973). Information theory and an extension of the maximum likelihood principle. In B. N. Petrov & S. Caski (Eds.), *Proceedings of the second international symposium on information theory*. (pp. 267–281). Akademiai Kiado, Budapest, Hungary.
- Aricò, P., Borghini, G., Di Flumeri, G., Colosimo, A., Bonelli, S., Golfetti, A., . . . Babiloni, F. (2016, October). Adaptive Automation Triggered by EEG-Based Mental Workload Index: A Passive Brain-Computer Interface Application in Realistic Air Traffic Control Environment. *Frontiers in Human Neuroscience*, 10(198), 230–13.
- Baguley, T. (2012). Calculating and graphing within-subject confidence intervals for ANOVA. *Behavior Research Methods*, 44(1), 158–175. doi:10.3758/s13428-011-0123-7
- Hart, S. G. (2006). Nasa-Task Load Index (NASA-TLX); 20 Years Later. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 50(9), 904–908.
- Hawkins, G. E., Mittner, M., Boekel, W., Heathcote, A., & Forstmann, B. U. (2015, December). Toward a model-based cognitive neuroscience of mind wandering. *Neuroscience*, 310, 290–305.
- ISO 17488. (2016). Road Vehicles -Transport information and control systems - Detection Response Task (DRT) for assessing selective attention in driving.
- Pinheiro, C. J. & Bates, D. M. (2000). *Mixed-effects models in S and S-Plus*. New York: Springer.
- Quandt, J. (20017). *Automatic vehicle control systems [Investigation No. PE 16–007]*. US Department of Transportation: National Highway Traffic Safety Administration.
- Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, 6(2), 461–464.
- Strayer, D. L., Biondi, F., & Cooper, J. M. (2017). Dynamic Workload Fluctuations in Driver/Non-Driver Conversational DYADS. In *Driving assessment conference* (pp. 362–367). Iowa City, Iowa: University of Iowa.

- Strayer, D. L., Turrill, J., Cooper, J. M., Coleman, J. R., Medeiros-Ward, N., & Biondi, F. (2015). Assessing Cognitive Distraction in the Automobile. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 57(8), 1300–1324.
- Tillman, G., Strayer, D., Eidels, A., & Heathcote, A. (2017). Modeling cognitive load effects of conversation between a passenger and driver. *Attention, Perception, & Psychophysics*, 70(6), 1795–1803.
- Unity Technologies. (2017). Unity® (version 2017.3.1f1) [Computer software]. Retrieved from <http://unity3d.com>
- Wilson, G. F. & Russell, C. A. (2003). Real-time assessment of mental workload using psychophysiological measures and artificial neural networks. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 45(4), 635–643.