2026 NDIA MICHIGAN CHAPTER
GROUND VEHICLE SYSTEMS ENGINEERING
AND TECHNOLOGY SYMPOSIUM
MODELING, SIMULATION, PROTOTYPING, AND VALIDATION (MSPV) TECHNICAL SESSION
AUGUST 11-13, 2026 - NOVI, MICHIGAN

THE COST OF DISRUPTION: ANALYZING THE IMPACT OF NEW TASKS ON TIME-EVOLVING COGNITIVE ACTIVITY

S. M. Gordon, PhD[1], J. Touryan, PhD[2]

[1]DCS Corporation, Alexandria VA
[2]US Army DEVCOM Army Research Laboratory

ABSTRACT

Certain aspects of cognitive state, such as attention, are known to oscillate and exhibit time-varying dynamics. This has strong implications for future human-system integration since, across domains, there is often a cost to interrupting ongoing processes. However, to date there is little knowledge about the cost of such disruption when considering cognitive resources. We have constructed a predictive model that captures these oscillations and predicts future cognitive state using only a finite history of prior state. We use this model to investigate what happens when ongoing cognitive dynamics are interrupted. We use data in which participants perform an unrestricted visual search task, while an auditory side task is randomly injected. The results show that when the timing of the side task causes a disruption, the participants' physiology display patterns that have been previously associated with increased cognitive load. This indicates that it is not only the nature of the task that must be considered but also the onset time of the task. When task demands are out of phase with ongoing dynamics there is a cost to this timing error in addition to the challenge of the task itself.

Citation: S. M. Gordon and J. Touryan, "The Cost of Disruption: Analyzing the impact of new tasks on time-evolving cognitive activity," In *Proceedings of the Ground Vehicle Systems Engineering and Technology Symposium* (GVSETS), NDIA, Novi, MI, Aug. 11-13, 2026.

1. INTRODUCTION

There are many challenges to developing the next generation of manned and unmanned ground vehicles. With respect to the human dimension, as crew sizes are reduced the burden on individual operators is going to increase. The increased burden will, in turn, increase the likelihood of committing an error. Mitigating these errors through optimized system design requires knowledge of the fundamental limitations imposed by the human cognitive system [1-3] – specifically, *when* and *why* errors occur, and ultimately *what* actions can be taken to improve the situation?

Most studies on human performance tend to view errors as the outcome of instantaneous environmental and operator states. For example, changes in the task lead to changes in cognitive state that, in turn, yields increased error rate. Since the environment cannot be easily predicted the assumption is that errors also cannot be predicted. Therefore, while such an approach has a certain amount of explanatory power, it has little-to-no predictive power.

To investigate this issue, recent work published by our team [4] showed that ongoing cognitive dynamics (e.g., naturally occurring oscillations in attentional state) significantly predicted future periods of high error probability in a ground vehicle driving task. For this analysis, we used data collected at the DEVCOM Ground Vehicle Systems Center (GVSC). Participants were required to operate a vehicle in naturalistic conditions while obeying numerous traffic rules. In four of the five conditions, participants had access to autonomous driving aids that provided speed, or speed and lane, control. In half of the driving aid conditions the autonomous driving aid was rigged to be faulty [5].

We applied a predictive model, termed the Generative Cognitive Modeling Tool (GCMT), that used prior (i.e., historical) cognitive state estimates to estimate future cognitive state. The GCMT did not consider information from the environment. The model leveraged generative AI principles, "listened to" ongoing attentional oscillations, and then predicted how those oscillations would unfold in the future with a maximum horizon of 25 seconds. These predicted states explained roughly two-thirds of the errors committed by the participants when operating the vehicle without a driving aid.

The GCMT was inspired by prior work in the cognitive neuroscience community that has reliably shown that certain cognitive states, such as attention, wax and wane on multiple timescales [6-9]. By showing that an AI system could learn these variations and apply that knowledge to predicting activity in novel domains (i.e., new tasks and new participants), our prior work implied the existence of a participant- and task-agnostic transition law governing the moment-to-moment changes in cognitive state.

While the GCMT helps address *when* errors occur, there are still questions about *why* they occur and *what* can be done to prevent them? We believe that additional findings from that prior work, though, offer a clue for understanding these two issues. In our analysis of the GVSC driving data, we also found that vehicle operators would adapt their state to meet current task demands – i.e., the observed state did not always match the predicted state (output of the GCMT) and this mismatch was often in response to unpredictable changes in task demands. We then found that when this adaptation was required, the operators were more likely to offload tasking to an autonomous driving aid even when that aid was known to be faulty [4]. In other words, disrupting their ongoing dynamics led the participants to offload tasks. We hypothesize that this was due to a hidden cost (perceived by the operators) of disrupting their ongoing cognitive dynamics.

The idea of incurring costs for disrupting ongoing physiological processes is not new. For example, it is known that disrupting one's ongoing circadian rhythm can impact both physical and cognitive function [10-13]. While there is a large body of knowledge describing both the acute and long-term impacts of disrupting circadian rhythm, there is much less knowledge on the impact of disrupting moment-to-moment cognitive dynamics. In the current work, we investigate the physiological cost of such disruption. We begin by adapting our original predictive

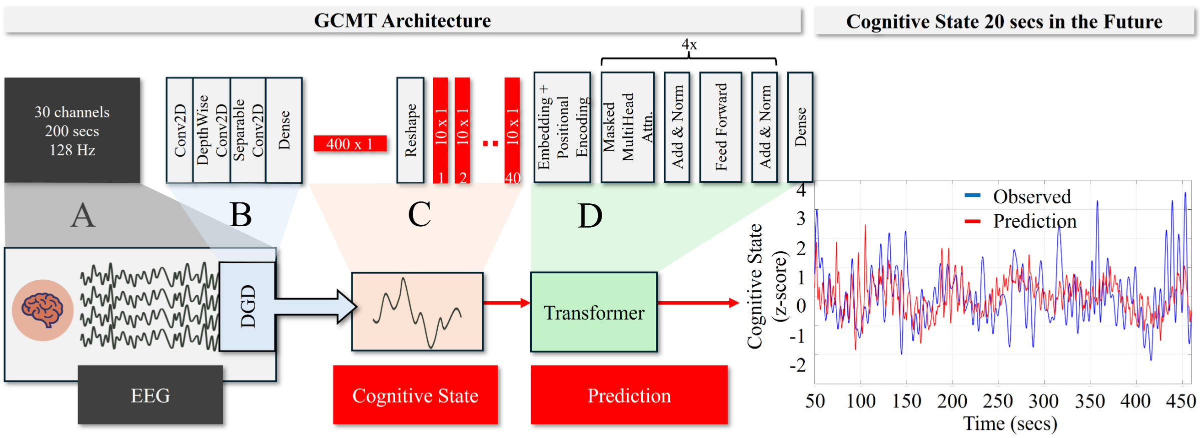

Figure 1: Physiological data (A) is decoded using the DGD (B) into low-dimensional probability estimates associated with the words in our lexicon (C). The probability estimates are passed to predictive layers (D) and used to estimate how the low-dimensional state will unfold.

approach by switching from a Long-term Short-Term Memory (LSTM) model to a Transformer model. We then apply this model to data from a visual foraging task in which randomized events have been injected, while we record multiple concurrent physiological signals including electroencephalogram (EEG), pupillometry, and eye movements. Finally, we compare the physiological responses when the injected events disrupt ongoing dynamics versus coincide with ongoing dynamics. The results show a significant relationship between the phase of the ongoing cognitive state at the time of task onset and subsequent pupillary and fixation-based activity. When the injected task causes a disruption, the physiological changes resemble those associated with increased cognitive load. This indicates that time, with respect to ongoing cognitive dynamics, is a critical factor impacting cognitive load and internal perception of task difficulty.

2. METHODS

2.1. Generative Cognitive Modeling Tool

The GCMT is a multilayered approach originally developed to model EEG-based cognitive activity. Figure 1 shows an overview of the GCMT architecture.

Decoding Layer – The first component in the GCMT is a domain-generalized decoder (DGD). The DGD extracts a predefined state from the neurophysiological (EEG) data. The DGD is trained independently of the rest of the GCMT and performs multiple critical functions. First, it reduces the dimensionality of the original signal, which is necessary for the subsequent predictive layers when learning long-range patterns. Second, the DGD abstracts the data away from the details of the given physiological response and projects that data onto a participant- and domain-agnostic representation. This is important so that data can be concatenated across experiments, so that a larger amount of diverse data can be supplied to the predictive layers when those layers are trained.

We used the EEGNet architecture [14] to implement the DGD. EEGNet has successfully been used to develop participant- and task-agnostic models of EEG activity [15-17]. A full description of how to use EEGNet for decoding phasic cognitive processes or tonic states is available in [18] and [19], respectively. Here, we provide a brief overview.

To train the DGD to decode a specific cognitive state we must first identify a source training set in which the desired cognitive state is known to be present. We refer to this as the Lexicon Source (LS) data set – as it defines the lexicon used by the GCMT. For instance, the LS data set could be a laboratory-based study on visual versus auditory task load, visual evoked detection, affective state, or attentional oscillations. Depending on the nature of the LS data set the "words" in the GCMT's lexicon would be "high task load" versus "low task load", "visual target detected" versus "no target detected", or "high attentional state" versus "low attentional state". To provide the lexicon, the LS data set must simply contain labeled time slices when the cognitive state or process of interest is present (or not), paired with the desired physiological signal to be measured.

All data passed to a given instance of the DGD must be preprocessed in a similar fashion. This includes the LS data used to train the DGD, the data used to train the predictive layers, and all subsequent test data. For EEG data, common preprocessing typically involves reference signal, bandpass filter settings, and channel montage (or interpolated montage) selection. The LS data is then partitioned based on the known labels. Data is balanced via randomized downsampling of the majority class and concatenated across participants. In practice, since many studies on cognitive state have imbalanced conditions, we repeat the randomized downsampling 5x to produce an ensemble of 5x DGDs. This ensures better usage of the available LS data and guards against a single model encountering local minimum during training.

To interrogate novel test data (Figure 1 A) the ensemble of trained DGDs (Figure 1 B) is stepped across the new data at a predefined sampling rate. Note that the new data has already been preprocessed using the same settings applied to the LS data. A group average is then obtained. This produces a single continuous probability estimate that the physiological data at that moment is conceptually more similar to one (or the other) "words" in our lexicon (Figure 1 C).

We used the EEGNet architecture with the settings (8-4-2), as described in [14]. The EEG data was sampled at 128 Hz and we used a DGD input window size of 385 samples (3 seconds of data). The EEG montage was set to the 30-channel intersection of the Neuroscan NuAmps system (LS data set) and BioSemi Active II system (all other data sets). The full details for training the exact DGD used in this work are provided in [19].

Predictive Layers – The predictive layers in the GCMT are designed to accept a finite history of cognitive state estimates and predict how that signal will unfold in the next several seconds. The original GCMT detailed in [4] used a recurrent network, implemented with an LSTM block. Here, we adapt the model to a decoder-only transformer with stacked blocks for multihead attention and feedforward feature extraction [20]. We made the change from LSTM to transformer to balance the needs of phase accuracy and prediction horizon. From our ongoing work with the GCMT we have empirically identified tradeoffs from using different predictive layers. The LSTM has shown better accuracy at predicting lower-frequency oscillations, giving it a longer predictive horizon. The transformer approach has, so far, shown better accuracy at predicting higher frequency oscillations, but with a slightly reduced horizon.

Within the predictive layers, we first reshape the data into chunks. This is referred to as input patching [21]. In our prior work, we performed this step to reduce the impact on vanishing gradients when using the

LSTM. We retained this step with the transformer to reduce computational cost and improve overall accuracy [22]. Next, we use a stack of 4 single-headed transformer blocks with embedding dimension = 128 and feed-forward dimension = 256. Each block first applies a MultiHeadAttention layer, then LayerNormalization with epsilon = 1e-6, then a Dense layer with 'gelu' activation function, and a final LayerNormalization with epsilon = 1e-6. Within each MultiHeadAttention layer we apply a causal mask to force the model to learn causal relationships. Many language-based models do not always use such a mask; however, this is because, at the level of individual sentences, language is not explicitly causal and sequential [23] – modifiers, for example, can appear before or after the object being modified. By applying a causal mask we are asserting that learned relationships between successive cognitive state estimates could only go in one direction. Our transformer architecture is shown in detail in Figure 1 D.

2.2. Data

Lexicon Source – To define our lexicon and train the DGD, we consider a single cognitive state, defined empirically by a previously conducted, controlled research study. The original study is described in [24] and required participants to drive a simulated vehicle along an empty closed-circuit at night. Experimental manipulations were used to induce drowsiness. During the driving task, a lateral perturbation (i.e., simulated wind) would push the vehicle out of the lane and the time required for the participants to respond and correct the vehicle's trajectory was recorded as reaction time. From this behavioral recording, periods of drowsy versus alert driving were identified, yielding two words in our lexicon: "drowsy" and "alert". Once trained, we applied that DGD to each of the data sets in Table 1 to produce the training data for the predictive layers.

Training and Model Selection – The predictive layers were designed to accept a finite sequence of 400 samples at 2 Hz, which was patched into 40 blocks of 10 samples each. We used the same data sets previously described in [4] to train the predictive model.

Table 1. Training Data.

Data Set	*# Participants / Hours of EEG*	*OpenNeuro ID*
1	10 (10.78 hrs)	ds004853
2	14 (16.32 hrs)	ds004842
3	13 (27.61 hrs)	ds004843
4	158 (52.40 hrs)	ds004123
5	109 (116.71 hrs)	ds004122
6	28 (22.37 hrs)	ds004121
7	24 (33.60 hrs)	ds004120
8	13 (16.53 hrs)	ds004855
9	66 (18.51 hrs)	ds004851
SUM	**435 (314.83 hrs)**	

These data sets encompass a total of 435 participants, and approximately 314.83 hours of recorded EEG data. This is in addition to the LS data set used to train the DGD. All training data sets can be accessed publicly via the OpenNeuro repository. Data set 9 was used as a hold-out validation set. The predictive model was trained for 10,000 epochs using a logcosh error function and a learning rate of 2e-4.

Testing Data – Data was collected from 45 participants from the greater Los Angeles area (17 female, mean age 36.8 ± 12.3, 28 male mean age 41.6 ± 14.4) and, later, 21 participants from San Antonio, TX (3 female, mean age 43.7 ± 6.7, 18 male mean age 32.6 ± 7.9). Data from nine participants were subsequently excluded due to overall data quality issues in either the EEG or eye tracking record yielding a total of 57 participants. Data collection was performed at two locations due to the COVID pandemic and regional restrictions on in-person data

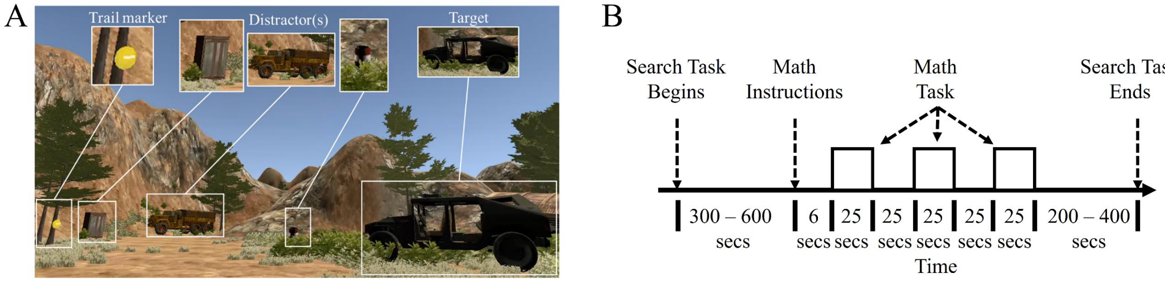

Figure 2: (A) Sample scene with trail marker, distractor objects, and target object highlighted, and (B) temporal structure for the primary search task with the auditory side math task.

collection. All participants were at least 18 years of age and signed an Institutional Review Board approved informed consent form prior to participation [ARL 19-122]. All participants had normal, or corrected-to-normal, vision. EEG data was collected using a 64 channel BioSemi Active II using the 10-10 montage, with an original sampling rate of 2048 Hz, and two mastoid channels for reference. Electrooculogram signals (EOG) were collected using four bipolar external leads – one placed on each outer canthus, one on the left infraorbital ridge, and one on the left supraorbital ridge. Eye tracking data was collected using a Tobii Pro Spectrum (300 Hz). Participants were seated at approximately 70 cm from the Tobii Pro Spectrum monitor (EIZO FlexSCan EV2451) in a quiet room with sound and lighting control. Experimental data was presented on a 24” monitor with a resolution of 1,920 x 1,080 pixels. The eye tracking data were synchronized with the game state and EEG data using Lab Streaming Layer [25].

During the task, participants were asked to navigate through a virtual environment using a mouse and keyboard to control their position and orientation. Participants were encouraged to follow a designated course through the environment (Figure 2 A) by searching for yellow trail markers but could deviate from the course at their discretion. The primary task was to visually search for and count the occurrence of target objects that had been randomly distributed in the environment amongst an array of distractor objects. This allowed the participants to operate in a closed-loop, or coupled, fashion with the environment where their behavior at one time point directly impacted their experience of the environment at the current and next time points.

During the primary task, an auditory side task was injected at a random time (Figure 2 B). This side task involved computing the sum of a sequence of numbers and verbally reporting the answer to the experimenter. The onset of the task could not be predicted and had onset times randomly distributed between 300-600 seconds after the start of the primary task. The side task involved the presentation of three different sequences of numbers. Each sequence was presented for 25 seconds with a 25 second gap between sequences. Prior to the first sequence a brief set of instructions indicating the side task was about to begin was provided for the participants. The duration of this instruction was approximately 6 seconds long.

2.3. Analysis

In our analysis of the test data, we assess the impact of the side task as a function of the ongoing dynamics. We use the GCMT to obtain estimates of cognitive state 20 seconds into the future. We examine the impact the side task had on the EEG, eye movement (fixation duration and blink rate), and

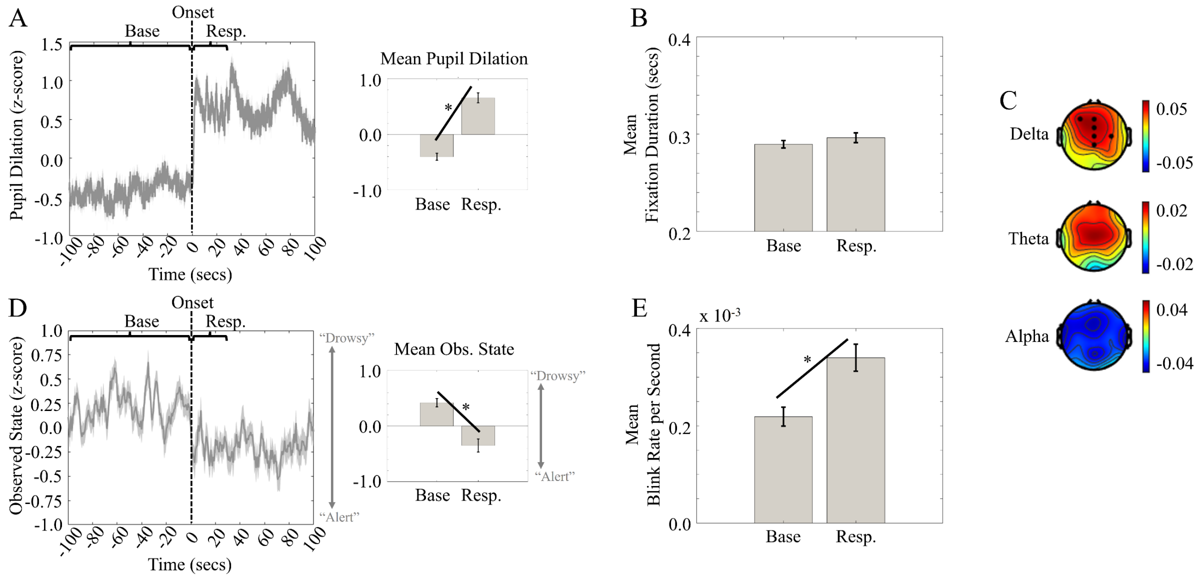

Figure 3: (A) Mean pupil dilation, (B) fixation duration, (C) EEG PSD changes, (D) observed cognitive state, and (E) blink rate at the onset of the first side task.

pupillary responses. We also examine the impact on the observed cognitive state (i.e., output of the DGD) and how this differs from the predicted cognitive state. To compute fixations, we use the velocity algorithm previously applied to this data set and detailed in [26]. To compute blink rate, we use a threshold detection method applied directly to the vertical EOG signal and tailored to each participant.

We also assess individual power-spectral density (PSD) in the EEG signal. We focus on the EEG frequency bands Delta (0.1-3.0) Hz, Theta (4.0-7.0) Hz, and Alpha (8.0-12.0) Hz. To compute spectral power in each of the EEG frequency bands we first filtered the data to the frequency band of interest using finite-impulse response filters implemented in Matlab and the EEGLab toolbox [27]. Next, we Hilbert transformed the data to provide a moment-to-moment estimate of the power in the selected frequency band. We rejected noisy epochs of the filtered signal using a threshold detection method that removed any data point more than 3 deviations away from the median.

To compute the phase of the predicted cognitive state we first z-score the predictions across an entire recording for each participant. We then apply the Hilbert transform to separate the magnitude and phase components. We compute an average phase around the onset time of the side task using a window of $\pm$ 3 seconds. Finally, we use the Circular Statistics Toolbox [28] to regress changes in physiological response against the phase of the predicted state.

3. RESULTS

Figure 3 shows the grand average response across participants for pupil dilation, observed cognitive state, mean fixation duration, mean blink rate, and EEG activity at the onset of the first side task. In each case, the response (*Resp*) is averaged over the first full math sequence (25 seconds of data) and compared to a baseline (*Base*) period of 100 seconds before the start of the task. In Figure 3 A we see significant changes in pupil dilation ($p < 0.0001$) using a paired t-test implemented in Matlab. We also see significant changes ($p < 0.0001$) in the

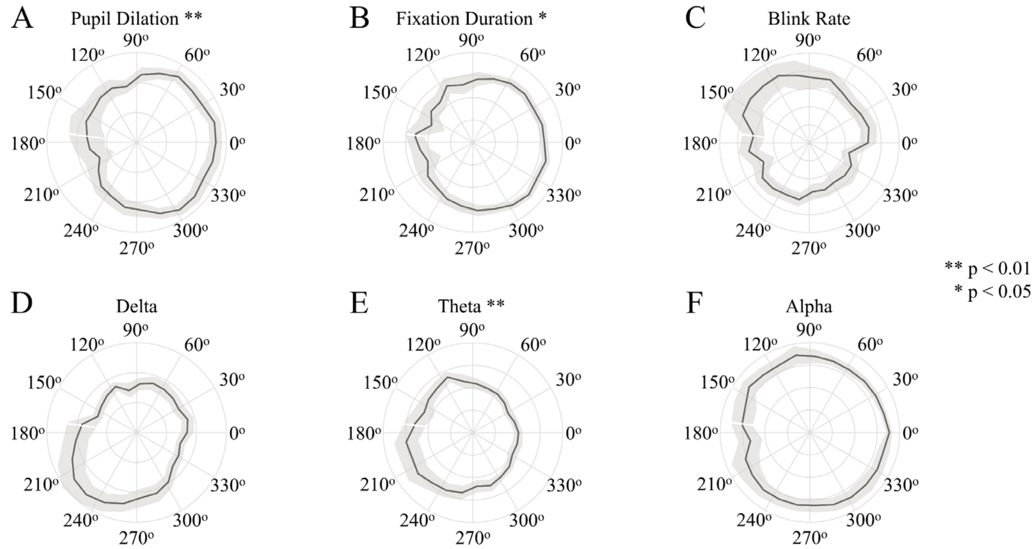

Figure 4: Circular relationship between changes in (A) pupil dilation, (B) fixation duration, (C) blink rate, (D) Delta EEG power, (E) Theta EEG power, and (F) Alpha EEG power and phase of the predicted cognitive state.

observed cognitive state (Figure 3 B). Note, the observed cognitive state is the output of the DGD and is driven by broad spectrum (i.e., frequencies [0.1-50] Hz) EEG activity across the scalp. Since the data for both Figure 3 A and B are derived from continuous time series signals, we present both the time series activity and the bar plots used to determine significance.

Figure 3 C and D show the mean fixation duration and blink rate averaged over the same response and baseline periods from Figure 3 A and B. Here we see there is no change in fixation duration, but there is a significant increase in blink rate ($p < 0.0005$). Figure 3 E shows the changes in EEG PSD for the Delta, Theta, and Alpha bands over the scalp. To produce this figure, we subtracted the Base PSD from the Resp PSD. Channels with significant differences are marked with a black dot. We only see significant differences in EEG activity in mid-frontal regions for the Delta band.

Next, we extracted the phase of the predicted state (i.e., output of the GCMT) at the onset of the side task and compared that phase to the magnitude of the physiological responses, i.e., *Resp* minus *Base* in Figure 3. This produced one circular variable (phase) and one linear variable (*Resp* minus *Base*) for each physiological signal. As previously stated, we used the Circular Statistics Toolbox [28] to test this circular-linear correlation for significance. The results are presented in Figure 4. We used the individual data points across participants to compute significance, but to produce these plots we used a sliding window to visually provide an estimate of the average magnitude of the physiological change (dark line) and standard errors (shaded region). For display purposes, we selected a sliding window size of 90^o. To produce the Delta, Theta, and Alpha plots we took the response from the Cz electrode.

From Figure 4 we see that the change in the physiological response for pupil dilation ($p < 0.01$), fixation duration ($p < 0.05$), and central Theta ($p < 0.01$) all significantly correlate with the phase of the predicted cognitive state. Similar results are obtained for Theta at electrodes FCz, C1, and C2 indicating the relationship is to be found primarily in central regions. Central Delta ($p < 0.07$) approached significance, and blink rate was not significantly correlated.

To better understand what we are seeing, Figure 5 plots the predicted cognitive state after dividing the participants into two groups. To do this, we identified those participants whose predicted state is in phase with the observed change induced by the side task (shown in Figure 3 B). We call this group the "in-phase" group. Participants were placed in the "in-phase" group if the phase of their predicted cognitive state was within $\pm 90^o$ of the induced change ($\sim 0^o$) The remaining participants were placed in the "out-of-phase" group. From Figure 5 we can visually verify that the "in-phase" group's cognitive state (blue line) *was predicted to already be in the state that is then induced by the task*. The "out-of-phase" group was predicted to be on the other side of the spectrum away from that state. Note that the prediction was made 20 seconds before the start of the side task and 14 seconds before the verbal instructions.

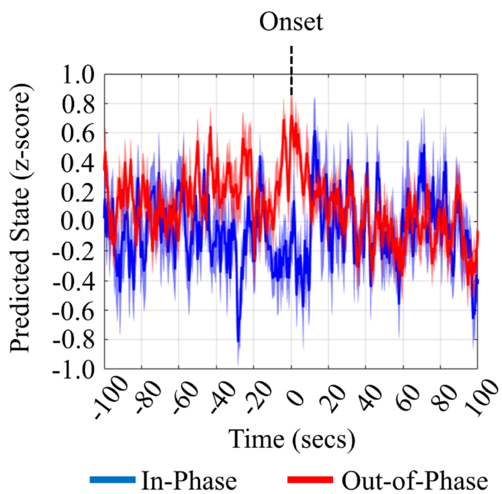

Figure 5: Predicted cognitive state, time-locked to the onset of the first side task.

Finally, we are prepared to test our hypothesis of a hidden cost for disrupting the "out-of-phase" group. To do this, we compare the magnitude of change in physiological response (i.e., *Resp* minus *Base* from Figure 3) across the two groups. These results are presented in Figure 6 and show that the "out-of-phase" group has a significantly stronger pupillary response ($p < 0.05$) and an increase in Fixation Duration ($p < 0.05$). This group also has a significantly smaller increase in Blink Rate ($p < 0.05$). Figure 6 also shows the difference in EEG activity, with the "out-of-phase" group showing less Theta activity ($p < 01$).

4. DISCUSSION

A review of the literature reveals that, in general, as cognitive load increases: 1) pupil dilation increases, 2) fixation duration *increases*, 3) mid-frontal Theta increases, and 4) blink rate *decreases* [29-33]. Mental arousal, such as that associated with stress or anxiety, on the other hand, is more likely to produce 1) increased pupil dilation, 2) *decreased* fixation duration, and 3) *increased* blink rate [34-37]. When considering the aggregate response across all participants (Figure 3), the physiological response appears more similar to a stress response than cognitive load, which could be expected since the primary visual search task was not designed to be difficult and the side task required on-the-spot mental arithmetic. However, when we factored in the participants' ongoing cognitive dynamics, we found that for those participants in the "out-of-phase" group the physiological changes, relative to the "in-phase" group, more closely matched that associated with cognitive load, i.e., increased pupil dilation, increased fixation duration, and decreased blink rate. This, of course, assumes that the task required participants to be in a more "alert" state, but this assumption is confirmed by the group response in Figure 3 b.

The only physiological variables that did not exhibit this pattern of increased cognitive load were the EEG variables Delta, Theta,

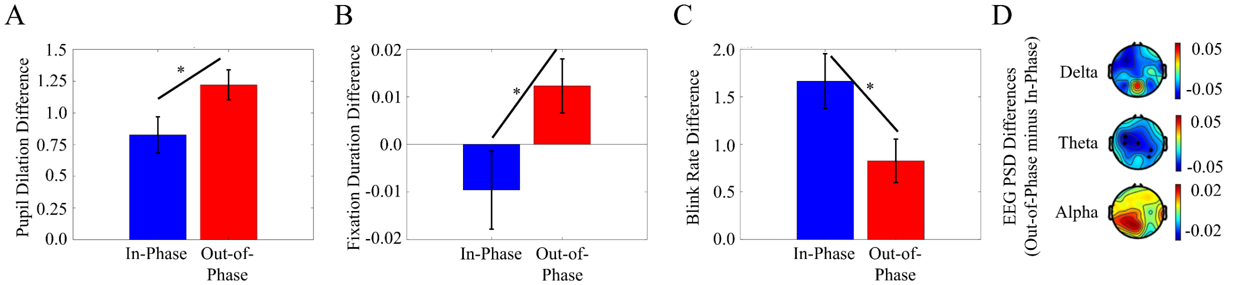

Figure 6: Differences in (A) pupil dilation, (B) fixation duration, (C) blink rate, and (D) EEG PSD between the "out-of-phase" and "in-phase" groups.

and Alpha. Contrary to prior work, Theta showed less activity in the "out-of-phase" group versus the "in-phase" group. We believe, however, that this explains our results rather than contradicts them. Naturally, Theta activity was contributing to the observed cognitive state estimate, thus oscillations in Theta were helping to form the oscillations in the observed cognitive state. In this way, the "in-phase" group got increased Theta activity "for free" as part of their ongoing dynamics. Prior work has shown that frontal-central Theta activity is an indicator of cognitive resource recruitment and deployment [38-39], which would provide a benefit to the "in-phase" group. The "out-of-phase" group received no such benefit. This group had fewer resources available based on ongoing dynamics and, thus, had to disrupt those dynamics to address the task requirements. This, in turn, produced a net effect of increased cognitive load when compared to the "in-phase" group.

We believe these results help shed light on *why* errors occur. If cognitive resources are oscillating due to their natural design, when those oscillations do not align with task demands the individual must disrupt them, but as with any process, disruption comes at a cost. Whether that cost is derived directly from the disruption or the inability to pull-in enough resources for the task during the disruption is beyond the scope of the current work. What we can say is that, from the physiological perspective, not having to disrupt ongoing dynamics appears preferable. This helps explain our result [4] in which errors between observed and predicted state led drivers to offload tasks to driving aids.

As for *what* can be done to mitigate errors, we argue the answer is clear: when possible align task requirements with ongoing dynamics. In practice, this can be done in two ways. The first is through learning. In [4], we used the GCMT to show that as participants learned a task their ongoing dynamics began to align with task demands, through knowledge of the environment, anticipation of outcomes, and behavioral freedom. Our current work was designed to remove such anticipatory effects. Since the real world, and events within it, can be both predictable and unpredictable, learning can only take the individual so far. Therefore, we believe the second way to mitigate errors is through improved human-technology design. In this case, adaptive technologies would understand the time-varying limitations of the human to help schedule tasks in a manner consistent with human capabilities. When such scheduling is not possible, the system would then need to be prepared for error states and have in-built remediation strategies (e.g., take the wheel).

5. CONCLUSION

As long as there is a role for the human within modern (or future) military domains there will be a need for the military to understand human capabilities and how those capabilities can be leveraged to produce optimal performance. While the human sciences community has long understood that there is a temporal aspect to task performance and cognitive function, due to a lack of tools to measure and predict these variations there has been little effort to leverage this knowledge to improve system design. Only by understanding when and why cognitive function can lead to task errors, though, can system designers develop solutions to prevent them. The current work was just another step in addressing these questions, but firmly demonstrates that ongoing, cognitive dynamics provide critical information that cannot be overlooked.

6. REFERENCES

[1] Shiffrin, R.M., 1976. Capacity limitations in information processing, attention, and memory. *Handbook of learning and cognitive processes*, *4*, pp.177-236.

[2] Fougnie, D. and Marois, R., 2006. Distinct capacity limits for attention and working memory: Evidence from attentive tracking and visual working memory paradigms. Psychological science, 17(6), pp.526-534.

[3] Sharpe, B.T. and Tyndall, I., 2025. The sustained attention paradox: a critical commentary on the theoretical impossibility of perfect vigilance. Cognitive Science, 49(4), p.e70061.

[4] Gordon, S.M., Lawhern, V.J. and Touryan, J., 2025. *Human Performance Modeling Using Generative Ai Principles* (No. 2025-01-0486). SAE Tech. Paper.

[5] Metcalfe, J.S., Marathe, A.R., Haynes, B., Paul, V.J., Gremillion, G.M., Drnec, K., Atwater, C., Estepp, J.R., Lukos, J.R., Carter, E.C. and Nothwang, W.D., 2017, May. Building a framework to manage trust in automation. In Micro-and nanotechnology sensors, systems, and applications IX (Vol. 10194).

[6] Niebur, E., Koch, C. and Rosin, C., 1993. An oscillation-based model for the neuronal basis of attention. Vision research, 33(18), pp.2789-2802.

[7] Fiebelkorn, I.C. and Kastner, S., 2019. A rhythmic theory of attention. Trends in cognitive sciences, 23(2), pp.87-101.

[8] Frey, J.N., Ruhnau, P. and Weisz, N., 2015. Not so different after all: The same oscillatory processes support different types of attention. Brain research, 1626

[8] Valdez, P., 2019. Circadian rhythms in attention. The Yale journal of biology and medicine, 92(1), p.81.

[10] Potter, G.D., Skene, D.J., Arendt, J., Cade, J.E., Grant, P.J. and Hardie, L.J., 2016. Circadian rhythm and sleep disruption: causes, metabolic consequences, and countermeasures. Endocrine reviews, 37(6), pp.584-608.

[11] West, A.C. and Bechtold, D.A., 2015. The cost of circadian desynchrony: evidence, insights and open questions. Bioessays, 37(7), pp.777-788.

[12] Sletten, T.L., Cappuccio, F.P., Davidson, A.J., Van Cauter, E., Rajaratnam, S.M. and Scheer, F.A., 2020. Health consequences of circadian disruption. Sleep, 43(1).

[13] Xu, S., Akioma, M. and Yuan, Z., 2021. Relationship between circadian rhythm and brain cognitive functions. Frontiers of optoelectronics, 14(3), pp.278-287.

[14] Lawhern, V.J., Solon, A.J., Waytowich, N.R., Gordon, S.M., Hung, C.P. and Lance, B.J., 2018. EEGNet: a compact convolutional neural network for EEG-based brain–computer interfaces. Journal of neural engineering, 15(5), p.056013.

[15] Pandey, V., Panwar, N., Kumbhar, A., Roy, P.P. and Iwamura, M., 2025. Enhanced cross-task EEG classification: Domain adaptation with EEGNet. In International Conference on Pattern Recognition (pp. 354-369). Spring. Cham.

[16] Song, H., She, Q., Fang, F., Liu, S., Chen, Y. and Zhang, Y., 2025. Domain generalization through latent distribution exploration for motor imagery EEG classification. Neurocomputing, 614,

[17] Li, M., Li, J., Song, Z., Deng, H., Xu, J., Xu, G. and Liao, W., 2024. EEGNet-based multi-source domain filter for BCI transfer learning. Medical & Biological Engineering & Computing, 62(3)

[18] Gordon, S.M., Dalangin, B. and Touryan, J., 2024. Saccade size predicts onset time of object processing during visual search of an open world virtual environment. NeuroImage, 298,

[19] Gordon, S.M., McDaniel, J.R., King, K.W., Lawhern, V.J. and Touryan, J., 2023. Decoding neural activity to assess individual latent state in ecologically valid contexts. Journal of Neural Engineering, 20(4), p.046033.

[20] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł. and Polosukhin, I., 2017. Attention is all you need. Advances in neural information processing systems

[21] Anh, L.H., Vu, D.T., Oh, S., Yu, G.H., Han, N.B.N., Kim, H.G., Kim, J.S. and Kim, J.Y., 2024. Partial transfer learning from patch transformer to variate-based linear forecasting model. Energies, 17(24), p.6452.

[22] Han, K., Xiao, A., Wu, E., Guo, J., Xu, C. and Wang, Y., 2021. Transformer in transformer. Advances in neural information processing systems, 34

[23] Chen, X., O'Donnell, T. and Reddy, S., 2024. When does word order matter and when doesn't it?. arXiv preprint arXiv:2402.18838.

[24] Lin, C.T., Wu, R.C., Liang, S.F., Chao, W.H., Chen, Y.J. and Jung, T.P., 2005. EEG-based drowsiness estimation for safety driving using independent component analysis. IEEE Transactions on Circuits and Systems I: Regular Papers, 52(12), pp.2726-2738.

[25] Kothe, C., Shirazi, S.Y., Stenner, T., Medine, D., Boulay, C., Grivich, M.I., Artoni, F., Mullen, T., Delorme, A. and Makeig, S., 2025. The lab streaming layer for synchronized multimodal recording. Imaging Neuroscience, 3, pp.IMAG-a.

[26] Enders, L.R., Smith, R.J., Gordon, S.M., Ries, A.J. and Touryan, J., 2021. Gaze behavior during navigation and visual search of an open-world virtual environment. Frontiers in Psychology, 12,

[27] Delorme, A. and Makeig, S., 2004. EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of neuroscience methods*, *134*(1), pp.9-21.

[28] Berens, P., 2009. CircStat: a MATLAB toolbox for circular statistics. Journal of statistical software, 31, pp.1-21.

[29] Pomplun, M. and Sunkara, S., 2019. Pupil dilation as an indicator of cognitive workload in human-computer interaction. In Human-Centered Comp.. CRC Press.

[30] Othman, N. and Romli, F.I., 2016. Mental workload evaluation of pilots using pupil dilation. International Review of Aerospace Engineering, 9(3), pp.80-84.

[31] Marquart, G., Cabrall, C. and De Winter, J., 2015. Review of eye-related measures of drivers' mental workload. Procedia Manufacturing, 3, pp.2854-2861.

[32] Biondi, F.N., Saberi, B., Graf, F., Cort, J., Pillai, P. and Balasingam, B., 2023. Distracted worker: Using pupil size and blink rate to detect cognitive load during manufacturing tasks. Applied ergonomics

[33] Diaz-Piedra, C., Sebastián, M.V. and Di Stasi, L.L., 2020. EEG theta power activity reflects workload among army combat drivers: an experimental study. Brain sciences, 10(4), p.199.

[34] Bradley, M.M., Miccoli, L., Escrig, M.A. and Lang, P.J., 2008. The pupil as a measure of emotional arousal and autonomic activation. Psychophysiology, 45(4), pp.602-607.

[35] Kuniecki, M., Wołoszyn, K.B., Domagalik, A. and Pilarczyk, J., 2017. Effects of scene properties and emotional valence on brain activations: A fixation-related fMRI study. Frontiers in human neuroscience, 11, p.429.

[36] Mott, R.O., Hawthorne, S.J. and McBride, S.D., 2020. Blink rate as a measure of stress and attention in the domestic horse (Equus caballus). Scientific reports, 10(1), p.21409.

[37] Calvo, M.G. and Avero, P., 2005. Time course of attentional bias to emotional scenes in anxiety: Gaze direction and duration. Cognition & Emotion, 19(3),

[38] Cavanagh, J.F. and Frank, M.J., 2014. Frontal theta as a mechanism for cognitive control. Trends in cognitive sciences, 18(8), pp.414-421.

[39] Nigbur, R., Ivanova, G. and Stürmer, B., 2011. Theta power as a marker for cognitive interference. Clinical Neurophysiology, 122(11),

[40] Gordon, S.M., Lawhern, V.J. and Touryan, J., 2025b. *Assessing Cognitive State Adaptations using Predictive Models*. Interservice/Industry Training, Simulation, and Education Conference

7. CONTACT INFORMATION

Stephen M. Gordon
DCS Corporation
6909 Metro Park Dr., Suite 500
Alexandria, VA 22310
sgordon@dcscorp.com
571-227-6167

8. ACKNOWLEDGMENT

This project was sponsored by the DEVCOM Army Research Laboratory under contract W911QX-21-D-0004. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the US Government.

9. ACRONYMS

GCMT	Generative Cognitive Modeling Tool
GVSC	Ground Vehicle System Center
EEG	Electroencephalogram
DGD	Domain Generalized Decoder
LS	Lexicon Source