THE UNIVERSITY OF CHICAGO

# Identifying Predictors of Lapses in Sustained Attention

By

Alfred F. Chao

August 2022

A paper submitted in partial fulfillment of the requirements for the Master of Arts degree in the Master of Arts in Computational Social Science

Faculty Advisor: Dr. Monica Rosenberg
Preceptor: Elizabeth Huppert

## Abstract

Fluctuations in sustained attention occur at fine time scales. In many cognitive tasks, behavioral measures such as response time can predict the onset of these fine-scale fluctuations. Contemporaneous work has indicated a narrowing trend in the time scales on which predicting sustained attention fluctuations on the basis of fluctuations in functional connectivity among brain regions is possible. Uniting these ideas, we apply connectome-based predictive modeling to derive novel brain networks whose degree of activation may contribute to forecasting upcoming lapses in sustained attention on the order of 2-3 seconds prior to those lapses. We describe these networks' predictive power relative to canonical and predefined network models of task-relevant cognitive processes as well as models relating response time behavior to upcoming lapses. We find our novel network models offer the best single-feature prediction of upcoming lapses and are essential to their best overall prediction at this scale. More broadly, we find that both neural and behavioral measures may enable the prediction of upcoming attention lapses, but that the signatures of lapsing attention also differ between tasks.

## Introduction

Sustained attention is the preservation of directed focus over time (Timmers, 2013). In daily life we often rely on sustained attention, be it while conversing among friends or driving on the interstate. A consequence of sustained attention's ubiquity is the commonality of experiencing its lapses. If sustained attention is the preserved direction of focus over time, a lapse in sustained attention is the temporary loss of that direction (Cheyne, 2010). In conversation, such lapses are trivial; we may ask our friends to repeat what they just said and apologize for zoning out. On the interstate, consequences for such episodes may be more severe. For this reason—that we would often prefer to detect lapses in sustained attention before their behavioral consequences occur—the question emerges whether and how such predictions may be possible.

Prior work has identified some generally consistent behavioral predictors of lapses in sustained attention across a variety of task contexts. One such example is the sustained attention to response task (SART) (Robertson, 2001), in which participants are presented with streams of stimuli and asked to respond for a frequently-presented stimulus category while withholding responses to infrequently-presented stimuli. In this context, it is typical for participants to perform a trivial task (e.g. "press the spacebar if the stimulus is a scene image, withhold a response if the stimulus is a face), therefore we assume erroneous responding is a consequence of lapsing sustained attention. In these SART paradigms, it has been demonstrated that averaging reaction times (RTs) over intervals as short as ~4 seconds preceding infrequent stimulus trials can meaningfully predict the outcomes of those trials (Robertson, 2001; Sakai, 2013). In other words, information about a participant's attentional state on upcoming trials may be present in their behavior during preceding ones. This finding has been supported in conceptually similar tasks across a variety of time scales ranging from 40s-3s (deBettencourt, 2018; Rosenberg, 2015; Vaurio, 2009)

The SART's validity as a measure of sustained attention has been criticized on the basis of its trials' abrupt onsets. Potentially attention-reorienting effects of trial-based designs are thought to complicate the SART's role as a measure of sustained attention. To address this issue, the gradual-onset continuous performance task (gradCPT) was developed and applied to explore how patterns in both brain and behavior may associate with sustained attention performance (Esterman, 2013). As in the SART, behavioral data collected from participants performing the gradCPT indicated that pretrial RT (namely RT variability) could be applied to predict attentional state during future trials. By boiling traditional vigilance tasks into their essential character with gradCPTs, a concrete behavioral predictor of attentional state—pretrial RT—was obtained. However, the gradCPT is exactly that: a distillation of task features thought to best-apprehend sustained attention specifically. Real life tasks demanding sustained attention are quite unlike this ideal laboratory setup. One question accompanying the behavioral gradCPT findings, which this present work seeks to address, is the extent to which those findings generalize beyond the context of gradCPTs.

Implicit to these RT findings is the idea that sustained attention fluctuations may occur at high frequencies. That is, if RT data collected during the seconds preceding task trials contain enough information to make reliable predictions about the outcomes of those trials (i.e. about the likelihood of an attentional lapse), then we may conclude that whatever process underlies those outcomes might be evident at those time scales. This conclusion is supported by converging evidence from other tasks, such as the Stroop task (Jong, 1999) and a global-local letter recognition task (Weissman, 2006), both of which are argued to at least tangentially measure sustained attention and for which inter-trial attentional lapses on the order of seconds are thought to have influenced task performance. Further evidence supporting the existence of fine-scale fluctuations in sustained attention may be observed in clinical literature (Vaurio, 2006) and animal models (Cohen, 2011).

Although high-frequency measurements are relatively straightforward in the collection of behavioral data, if we intend to pursue the goal of predicting lapses in sustained attention *before* they are evident in behavior, our job grows complex. Some prior work has applied a variety of univariate and multi-voxel pattern analysis (MVPA) methods to analyzing neural precursors of lapses in sustained attention.

In the gradCPT, elevated dorsal attention network (DAN) activity preceded correct responses whereas default-mode network (DMN) activity preceded incorrect responses (i.e., lapses) when analyzing trial-level network antecedents of trial-level task performance; however, this result was complicated by a closer analysis which revealed that these effects were largely predicated on larger-scale dynamics (Esterman, 2013). That is, the precise relationship between DAN/DMN activity and task performance on the scale of ~1.6 seconds was, in part, a function of the participant's attentional state as indexed by their performance over a much broader temporal window spanning ~16 seconds. Another study found that MVPA classifiers could predict behavioral accuracy for a SART-like task on current trials using the voxel-wise activation data gathered between 6 and 12 seconds prior (deBettencourt, 2015). Further work in this vein confirmed that, when focusing on trial-level DMN/DAN activity, univariate analysis could not distinguish between attentional states while MVPA could (Rosenberg, 2015). In the 10 second regime, it was found that DMN activity alone was insufficient to predict lapsed attention during CPT performance whereas on the scale of 40 seconds robust within-network connectivity and between-network anticorrelations were observed (Kucyi, 2018).

Collectively, these results suggest that the notion of a fine-scale brain-based prediction of lapses in sustained attention may hold promise. However, they are all somewhat tangent to the question. For example: the behavioral results indicate that attention may fluctuate on the order of seconds, so a tentative brain-based metric of these fluctuations will require similar granularity. Another complication is the prominence held by so-called canonical networks like the DAN and DMN. Recent work has shown that these canonical networks may not offer the best description of what functional architecture underlies cognitive processes such as attention (Rosenberg, 2016). By narrowing the scope of their investigation to activity observed in these canonical networks, the prior work may be missing critical information required for brain-based predictions. Finally, although univariate and MVP analyses have been historically successful means of translating functional magnetic resonance imaging (fMRI) data into insights on the nature of cognition, these approaches are based on the interpretation of brain activity. It has been shown that sustained attention performance is better-characterized by analyses of functional connectivity than regional activation (Rosenberg, 2016).

Functional connectivity is a description of the statistical relationships among discrete brain regions' activity. For example, two regions that are spatially distant but which consistently activate in tandem are considered likely to be involved in the same sorts of processes and so exhibit a connection in function despite their physical separation. Functional connectivity has been applied in conjunction with connectome-based predictive modeling (CPM), which is a data-driven method of asking whether edges in a set of brain activity graphs are meaningfully related to a corresponding set of outcomes (Shen, 2017). In this application, functional connectivity measurements provide the graphs while some operationalization of success or failure at preserving sustained attention provides the outcomes.

Functional connectivity and CPM are therefore united to ask whether patterns in a brain's organization reliability emerge together with particular sustained attention outcomes. This line of work first showed that an individual's general ability to sustain attention is predictable on the basis of their brain's resting-state functional architecture (Rosenberg, 2016), and a novel sustained attention connectome-predictive model (saCPM) was derived. Later elaboration found that fluctuations in attentional state as indexed by changes in participants' task performance were predictable on the basis of fluctuations in those participants' saCPM network strengths (Rosenberg, 2020). This suggests functional connectivity can be applied to predict not only a participant's sustained attention performance on average, but also variation about that mean associated with changes in the brain's network organization. These findings encourage us to apply cutting-edge functional connectivity methods in pursuing a deeper understanding of sustained attention's neural bases.

However, there is an incongruence between the time scale on which fluctuations in sustained attention have been predicted using the CPM approach and the previously-outlined time scale on which sustained attention fluctuations have been observed in behavioral experiments. The highest-frequency fluctuations investigated during task performance in prior work have been on the order of minutes, whereas the behavioral literature suggests these fluctuations may occur on the order of seconds, if not faster. Therefore, another goal of the present work is pushing the limits of these brain-based predictions and asking whether the functional connectivity/CPM approach may be applied profitably to predicting sustained attention performance on the scale of seconds.

To quantify functional connectivity with this level of temporal granularity, we apply the measurement of connectome cofluctuation described in (Esfahlani, 2020). These cofluctuation measurements provide a way to profile the functional connectome at a given instant. Specifically, by quantifying the extent to which each brain region's activity is coordinated with each other region's activity at a given moment, we develop a graph representation of the entire brain wherein its parcellated regions form the nodes and functional connectivity among them form its edges. We then use these fine-scale graphs in the CPM framework to develop predictive models of the transient neural precursors to lapses in sustained attention. Prior work has identified magnitudinous episodes of cofluctuation as a key drivers of functional connectivity (Esfahlani, 2020). That is, high-amplitude bursts of cofluctuation have been shown disproportionately predictive of the functional connectome's structure. Other work has shown that patterns of cofluctuation may associate with transitions between brain states (Song, 2022). Cofluctuations among brain regions involved in connectome-based models of sustained attention, measured at critical moments, may therefore offer information about transitions into or out from attentive states, i.e. of fluctuations and therefore of lapses in sustained attention, prior to target trials.

We revisit the behavioral prediction of lapses in sustained attention as described by prior research, ask whether those findings extend beyond their original context, and apply our novel cofluctuation-based CPM strategy to a large fMRI dataset. Furthermore, we ask how each of these sources of information—brain and behavior—contribute to the prediction of upcoming lapses. In many cases, such as on the interstate, it would be ideal to have information about the likelihood of a lapse before its consequences are behaviorally evident. This research broadly seeks to advance that ideal.

## Methods

### ***Dataset and Participants***

We analyzed data from the publicly available Adolescent Brain Cognitive Development (ABCD) baseline year release, collected when participants were 9-10 years old. The ABCD Study involves recruitment and longitudinal measurement of 11,875 participants' behavior, brain structure, and brain function with MRI throughout their adolescence, with coordinated data acquisition protocols across 21 collection sites in the United States.

Our analysis specifically focused on the behavioral and fMRI data collected during task performance. We analyzed a subset of ABCD Study participants described in previous work (Kardan, 2021). This subset excludes participants whose task data couldn't be time-aligned with their scan data due to uncertainties in measurement as well as participants whose scans were missing entirely or taken with Philips scanners (due to a systematic data preprocessing issue), reducing our sample to 9155. We performed visual quality control on these participants' MRI data and applied an additional frame displacement exclusion threshold of mean < 0.2mm/max < 2mm, after which 1548 participants remained whose scan data were of reasonable quality. Finally, we excluded one of the data collection sites which was reduced to three usable participants after the previous exclusions. The final sample was 1545 participants (851 female) ranging from 9 to 11 years old ($\mu$ = 10.03) collected at 18 sites.

### ***Behavioral Task***

The ABCD Study protocol includes fMRI during three cognitive tasks: a stop-signal task, a monetary incentive delay task, and an emotional *n*-back task. In lieu of a direct sustained attention task, we analyzed the data collected during performance of the emotional *n*-back (EN-back) task. The EN-back task involves two difficulty conditions (0-back and 2-back) intended to differentially load working memory.

In 0-back blocks, participants are presented first with a target image for 2.5 seconds, then with a stream of 10 category-matched images, each shown for 2 seconds with an inter-stimulus interval of 1 additional second. Participants are instructed to respond to every image with a button press, submitting a "match" response when the stream image matches the target and a "no match" response otherwise.

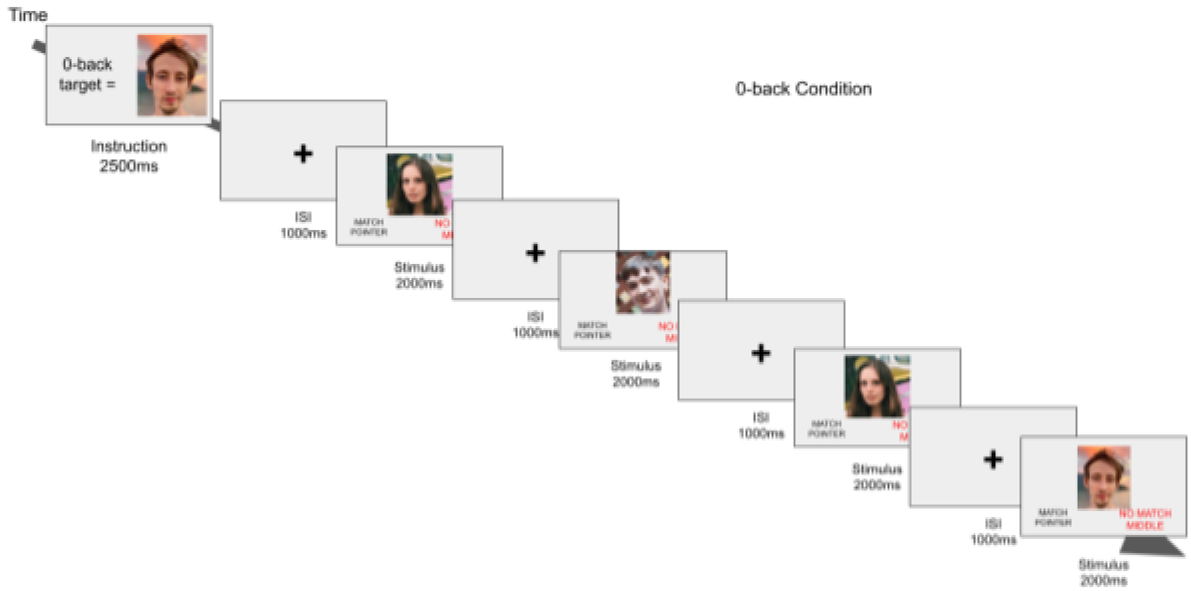


*Fig. 1: schematic diagram of EN-back 0-back task*

In 2-back blocks, participants are presented with a similar stream of images, but are tasked with responding "match" when the current image matches the one presented two images ago and "no match" otherwise. The images themselves consist of either face or scene stimuli, with the face stimuli being further subdivided into faces depicting positive, negative, or neutral emotional states.

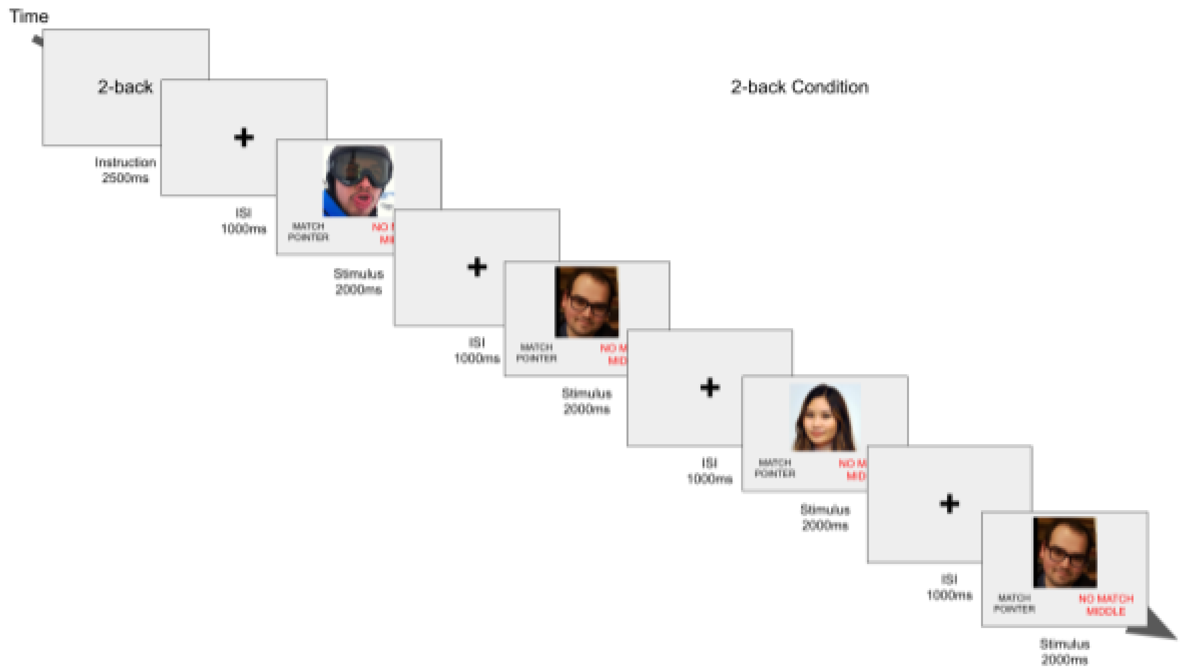


*Fig 2: schematic diagram of EN-back 2-back task*

Over the course of one scan session, a participant completed 16 blocks evenly divided between 0-back and 2-back tasks and assigned in a random order. We take erroneous responses to target trials as indicating lapses in sustained attention.

***Predicting Lapses from Pre-trial Reaction Times***

RT data were collected for each EN-back task trial. Trends in RT such as variability or short-term slowing have been shown predictive of upcoming sustained attention lapses in previous work (Esterman, 2013; deBettencourt, 2018). To ask whether RT information predicts upcoming lapses in the 0-back and 2-back tasks, we analyzed RT data for the three trials preceding each EN-back target trial. First, we omitted a preceding trial's RT if no response was given, no RT was collected, or if that preceding trial happened to contain another target. Target trials for which more than one preceding trial was omitted in this manner were then themselves omitted. The remaining target trials were split according to response accuracy, and a mean pre-trial RT was separately calculated by averaging RTs on the three trials prior to these targets.

In a separate analysis, RTs from only the single trials directly preceding target trials were z-scored within-participants and used to fit general linear mixed-effects models (GLMM) predicting lapses. These models apply the simple formula *Trial Accuracy ~ Previous Trial RT + (1|Participant)*, and we apply mixed-effects modeling in part to account for the well-documented extent of individual variation in sustained attention and its propensity to fail (Welhaf, 2020; Broadbent, 1982; Unsworth, 2010; Vaurio, 2009).

### ***MRI Data and Preprocessing***

ABCD Study baseline year data (ABCD Release 2.0.1) were acquired from the National Institutes of Mental Health data archive. Reanalysis of these de-identified data was approved by the University of Chicago's Institutional Review Board. Prior to their release, MRI data were subjected to initial preprocessing steps (Hagler, 2019), including motion and distortion correction, tissue segmentation and co-registration, normalization to the MNI152 nonlinear 6th generation template, and the regression of 36 nuisance regressors including global signal, mean signal from cerebro-spinal fluid, mean white matter signal, and head motion (Power, 2012). Data were bandpass filtered (0.008Hz-0.12Hz) and voxelwise time courses were averaged into 268 parcels as defined by the Shen functional atlas (Shen, 2013).

### ***Cofluctuation***

By taking the product of two given parcels' averaged BOLD measurements, we obtain a signed representation of their activity's coordination. For example, if two regions both exhibit relatively low activity during one pretrial window then their BOLD measurements will both be negative, which naturally gives a positive result when multiplied. On the other hand, if one of these regions is relatively active while the other is not, then the product of their signal values will be negative. In both cases, the product's sign describes whether our regions are generally synchronized while its magnitude indicates the extent to which that's true, and we apply this method to every combination of pairs of brain regions. The result is a graph which describes the state of a participant's functional connectome (i.e. the presence or absence of connectivity between each parceled brain region) at a given instant. We derive cofluctuation from standardized BOLD signal measurements associated with brain activity underway over the course of three TRs (2400ms) preceding target trials during EN-back task, though in practice these are the fifth through seventh TRs *following* target presentation on account of the hemodynamic delay .

### ***Predicting Lapses from Pre-trial Cofluctuation***

#### ***<u>Canonical Attention Networks</u>***

We tested whether pre-trial functional connectivity as indexed by cofluctuation within canonical networks associated with attentive processes such as the DMN and DAN could be used to predict EN-back target trial outcomes, which we take to proxy lapses in sustained attention. These networks are traditionally associated with off- and on-task performance, respectively. However, recent work has shown this relationship may not be so straightforward (Rosenberg, 2015; Esterman, 2013), and to our knowledge their relevance on fine time scales has yet to be tested at all. We measured each of these networks' strengths via cofluctuations preceding target trials, z-scored these measurements within-participants, and applied those standardized scores to fit separate GLMMs predicting lapses for each canonical attention network with the formula *Accuracy ~ DAN or DMN Activity + (1|Participant)*.

***Predefined Connectome-based Predictive Models***

We also tested whether two predefined connectome-based predictive models could be used to predict target trial outcomes on the basis of pretrial cofluctuation. These networks—the previously-described saCPM and its conceptual sibling in working memory (the wmCPM)—have been shown to predict cognitive abilities relevant to the EN-back task such as sustained attention (Rosenberg, 2016) and working memory (Avery, 2020), albeit on larger time scales. As above, we applied the within-participant z-scored cofluctuation observed preceding target trials in these networks to fit separate GLMMs predicting lapses with the formula *Accuracy ~ saCPM or wmCPM Activity + (1|Participant)*

***Training and Testing a Lapse Network***

To examine whether the cofluctuation graphs described above were related with performance on their associated upcoming trials in a data-driven manner, we calculated the correlation between each graph's edge (indicative of momentary functional connectivity) and its trial's outcome. Under the umbrella of a split-half cross-validation strategy, in a randomly assigned model-training half of participants, edges correlated with outcomes ($p < 0.01$) were partitioned according to whether their association was positive or negative. This process allows us to obtain task-positive and task-negative networks. Cofluctuation measured in those task-negative networks may then be subtracted from the cofluctuation measured in the task-positive networks to give a summary metric of the whole network's relative activation.

In addition to the authentic lapse network model, we use the same model-training data to generate a null network model for each cross-validation iteration. First, we shuffle the task outcomes associated with our pretrial cofluctuation measurements to preserve the structure of our functional connectivity data but remove its actual relationship with lapses in sustained attention. Then, as above, we examine the correlation of edges in these graphs with the randomized outcomes. By chance, some of these edges will survive feature selection, but we know *a priori* that this relationship is spurious. These spurious edges collectively form a null network model, also comprising task-positive and task-negative subgraphs, against which our novel lapse network may be compared.

In a randomly held-out half of participants, both lapse and null network strengths were gauged, z-scored within-participants, then applied to fit separate GLMMs relating EN-back target trial accuracy and network strength with the formula *Accuracy ~ Lapse or Null Network Activity + (1|Participant)*. This procedure is repeated over 50 randomized split-halves of our participants. If there is a statistically significant difference between the resulting distributions of log-odds coefficients made by the lapse network and null network GLMMs, then we have reason to believe the former are not a result of chance associations between cofluctuation and sustained attention lapse.

We generated separate GLMMs describing previous-trial RT, DAN, DMN, saCPM, wmCPM, the novel attention-lapse connectome-based predictive model (alCPM), and the null network model's fits to target trial outcomes in 50 random split-halves of our 1545 participants. From this generation we obtain 50 GLMMs for each feature, each describing what relationship exists between a given feature and sustained attention lapses in the held-out test set of participants. The GLMMs are summarized by $\beta$ coefficients assigned to their features, and the extent to which these summary $\beta$ distributions differ from a zero-mean normal distribution is calculated in all cases save the alCPM/null models, for which the above-described nonparametric significance test was performed instead.

## Results

### ***Predicting lapses in sustained attention with behavioral measurements***

***<u>Pretrial RT predicts lapses in sustained attention</u>***

To ask whether this developmental population's pretrial RT predicts upcoming lapses during 0- and 2-back task performance, we examined RT data collected during the EN-back task.

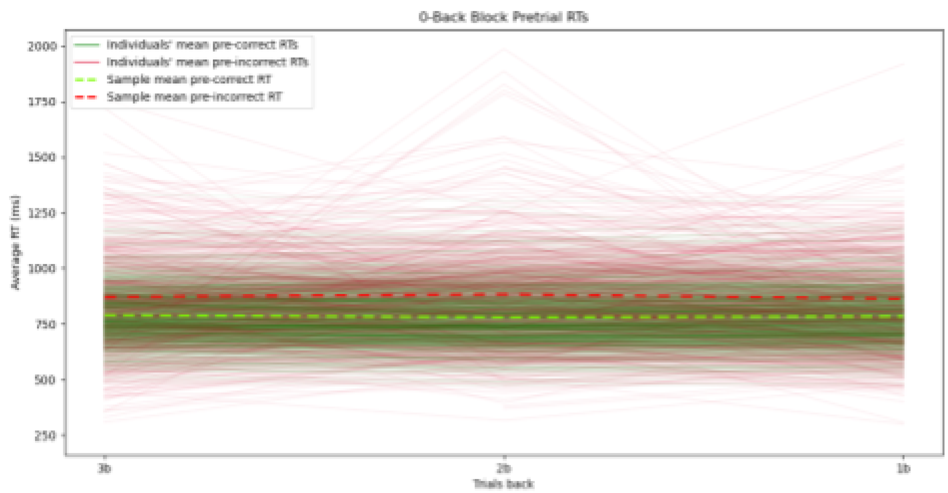


*Fig 3: within- and between-participant mean pre-target RTs according to 0-back target trial outcome. Faster RTs consistently precede correct responses ($p < 10^{-134}$)*

As shown in figure 3, for 0-back task blocks, we found a systematic difference in the RTs preceding correct and incorrect responses to targets. Combined, the three preceding responses were generally faster before a correct response than before an incorrect response (785ms +/- 270ms pre-correct & 871ms +/- 321ms pre-incorrect; $t = -24.8697$ $p < 10^{-134}$). This pattern is preserved when examining each individual trial among these preceding three. That is, for 0-back task blocks, RTs are consistently faster before correct responses than incorrect responses even when they are only examined three trials pre-target (789ms +/- 272 ms pre-correct vs. 870ms +/- 324ms pre-incorrect; $t = -12.988$, $p < 10^{-36}$), two trials pre-target (780ms +/- 268ms pre-correct vs. 881ms +/- 321ms; $t = -17.0479$, $p < 10^{-63}$), and one trial pre-target (784ms +/- 269ms pre-correct vs. 861ms +/- 316ms pre-incorrect; $t = -13.0700$, $p < 0.01x10^{-37}$).

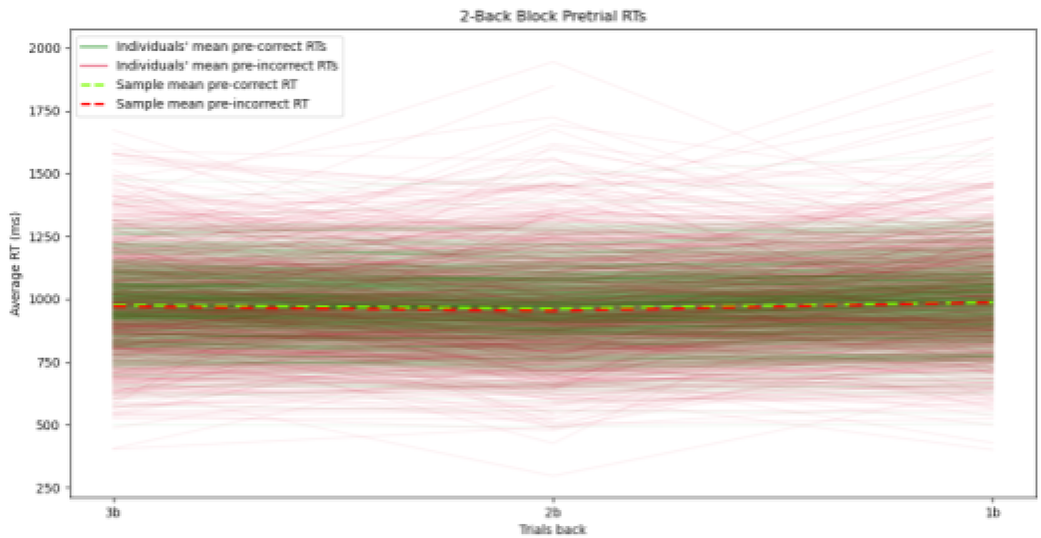


*Fig 4: within- and between-participant mean pre-target RTs according to 2-back trial outcome. No significant difference was observed either in aggregate or at any individual point*

Conversely, as figure 4 shows, for 2-back blocks we found no systematic difference in the three preceding trials' RTs, either combined or at any individual trial. Mean RT did not differ 3 trials before correct vs. incorrect target trials (975ms +/- 338ms vs. 970ms +/- 339ms; $t = 0.0579$, $p = 0.953$). RT also did not differ 2 trials before correct vs. incorrect targets (961ms +/- 327ms vs. 952ms +/- 334 ms; $t = 1.725$, $p = 0.085$) or 1 trial before correct vs. incorrect targets (986ms +/- 320ms vs. 986ms +/- 325ms; $t = 1.005$, $p = 0.315$). Combining all three preceding trials' RTs, the pre-correct mean was 974ms +/- 329ms and the pre-incorrect mean was 969ms +/- 334ms ($t = 1.6151$, $p = 0.106$).

Although RTs before correct vs. incorrect target trials systematically differed in 0-back blocks, we observed substantial mean RT differences between individuals. To account for this individual variability, and in accordance with prior work emphasizing the level of individual variation in these measures, we reanalyzed RT data with mixed-effects models incorporating pre-target RT as a fixed effect and participant as a random effect predicting target trial outcomes. When we model target trial outcomes as a function of immediately-preceding RT and include participants as a random effect, we find that RT contributes significant predictive power on similar orders of magnitude in both cases (0-back RT mean GLMM $\beta$: -0.051 +/- 0.018, $p < 10^{-24}$ vs. 2-back GLMM mean RT $\beta$: -0.049 +/- 0.019, $p < 10^{-23}$). To accommodate later comparisons among identified predictors of sustained attention lapses, this GLMM approach was applied within the context of our analysis-wide split-half cross validation, repeated for 50 iterations.

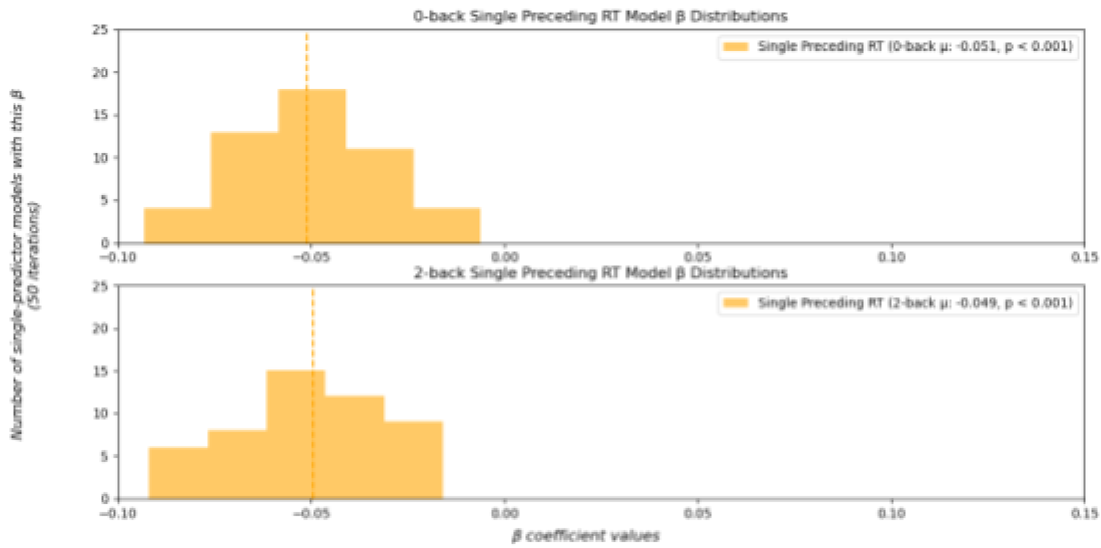


*Fig. 5: β coefficient distributions associated with mixed-effects models of target trial outcomes as a function of single preceding trial RT in 0- and 2-back blocks over 50 random split-halves*

***Predicting lapses in sustained attention with cofluctuation***

To ask whether high-frequency changes in brain activity observed with fMRI predicted upcoming lapses in sustained attention, we analyzed cofluctuation in the 2.4s leading up to target trial presentation. This is conceptually most similar to immediately-preceding RT behavior analysis given the EN-back task's 3s trial duration.

**Canonical networks associated with attention and mind-wandering are inconsistent predictors of lapses in sustained attention**

First we investigated the relationship between cofluctuations within canonical attention networks and upcoming lapses. Specifically, we tested the extent to which activity in DAN and DMN could contribute to predicting target trial outcomes. We applied the measured cofluctuations observed in these canonical networks during 2.4s preceding target trials to model those trials' outcomes in separate GLMMs with participants as a random effect. Repeating this process across 50 split-halves of our sample, we derive a distribution of $\beta$ coefficients, each of which summarizes a given feature's predictive power in that iteration's model. We find that only DMN activity makes significant contributions to lapse prediction in 0-back blocks (DMN mean $\beta = 0.0306$ +/- 0.02; $p < 10^{-14}$), while the distribution of DAN $\beta$s is statistically inseparable from a zero-mean distribution (DAN mean $\beta = 0.0051$ +/- 0.02, $p = 0.1$). For 2-back blocks, we find that both networks make significant contributions to lapse prediction (DMN mean $\beta = 0.0158$ +/- 0.02, $p < 10^{-7}$; DAN mean $\beta = 0.0239$ +/- 0.02, $p < 10^{-11}$).

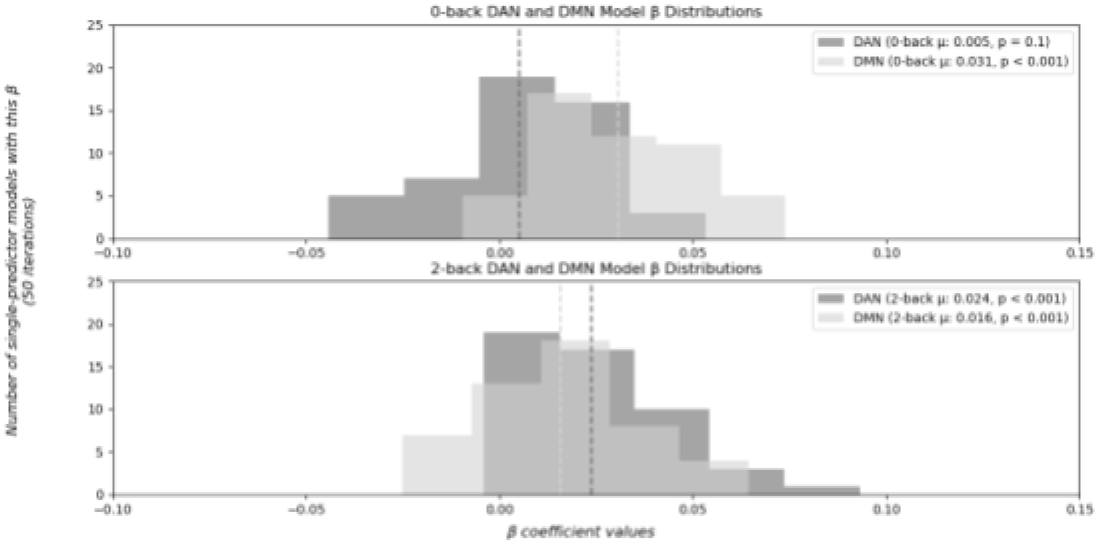


*Fig. 6: β coefficient distributions associated with mixed-effects models of target trial outcomes as functions of pre-target DAN and DMN connectivity in 0- and 2-back blocks over 50 random split-halves*

**<u>Predefined connectome models of sustained attention and working memory are inconsistent predictors of lapses in sustained attention</u>**

Next we applied predefined network models of cognition to ask whether activity in the previously-described saCPM and wmCPM held information relevant for the prediction of target trial outcomes. As above, we modeled those trial outcomes as a function of the cofluctuations observed in these predefined networks during 2.4s preceding target trials with a random intercept for each unique participant, and repeated this process over 50 random split-halves. We find that both networks make significant and roughly similarly magnitudious contributions to lapse prediction in 0-back blocks (saCPM mean $\beta$ = 0.0413 +/- 0.02, $p < 10^{-19}$; wmCPM mean $\beta$ = 0.0449 +/- 0.02, $p < 10^{-20}$). However, for 2-back blocks, only the saCPM $\beta$ distribution approaches a significant difference from zero (saCPM mean $\beta$ = 0.005 +/- 0.02, $p < 0.05$; wmCPM mean $\beta$ = -0.007 +/- 0.0251, $p = 0.13$).

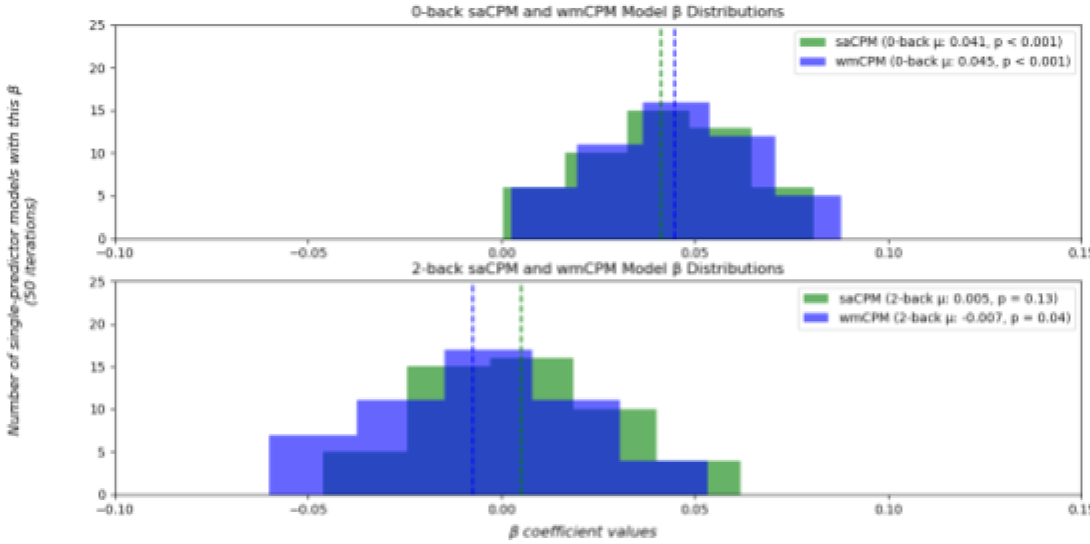


*Fig. 7: β coefficient distributions associated with mixed-effects models of target trial outcomes as functions of pre-target saCPM and wmCPM connectivity in 0- and 2-back blocks over 50 random split-halves*

**Attention-lapse network models predict lapses in sustained attention**

We built a novel attention-lapse connectome-based predictive model (alCPM) from half of our participants' data and applied it to predicting target trial outcomes on the basis of cofluctuations observed in the remaining half prior to EN-back target trials for 50 random split-half iterations. We additionally derived and applied null network models built to predict upcoming trial outcomes which had been shuffled within-subject to compare our alCPMs against in contrast with the canonical and predefined network analyses, for which their predictive model $\beta$ distributions are tested against the null hypothesis of a zero-mean normal distribution (i.e. no consistent, general effect of predictor on trial outcomes). The alCPM and null models were trained and tested within-block, meaning that a 0-back alCPM does not necessarily consist of the same edges as a 2-back alCPM.

For 0-back blocks we find alCPM models significantly predict upcoming lapses (alCPM mean $\beta$: 0.0745 +/- 0.02, null model mean $\beta$: -0.005 +/- 0.03; nonparametric $p < 10^{-13}$). For 2-back blocks we also find the alCPM significantly predicts lapses in sustained attention (alCPM mean $\beta$: 0.0581 +/- 0.03, null model mean $\beta$: -0.006 +/- 0.03; nonparametric $p < 10^{-20}$)

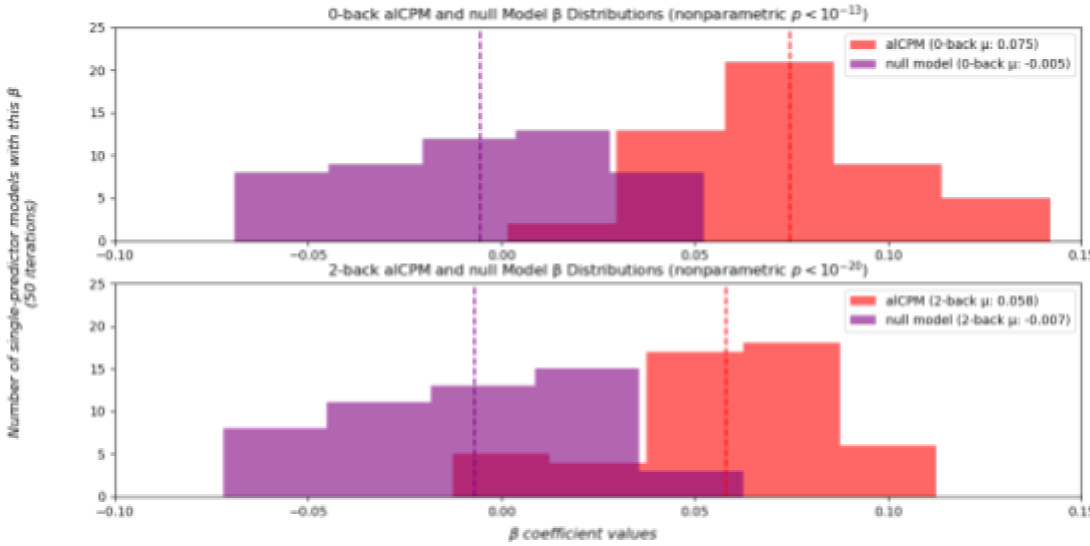


*Fig. 8: β coefficient distributions associated with mixed-effects models of target trial outcomes as functions of pre-target alCPM and null model connectivity in 0- and 2-back blocks over 50 random split-halves*

## <u>Comparing predictors of lapses in sustained attention</u>

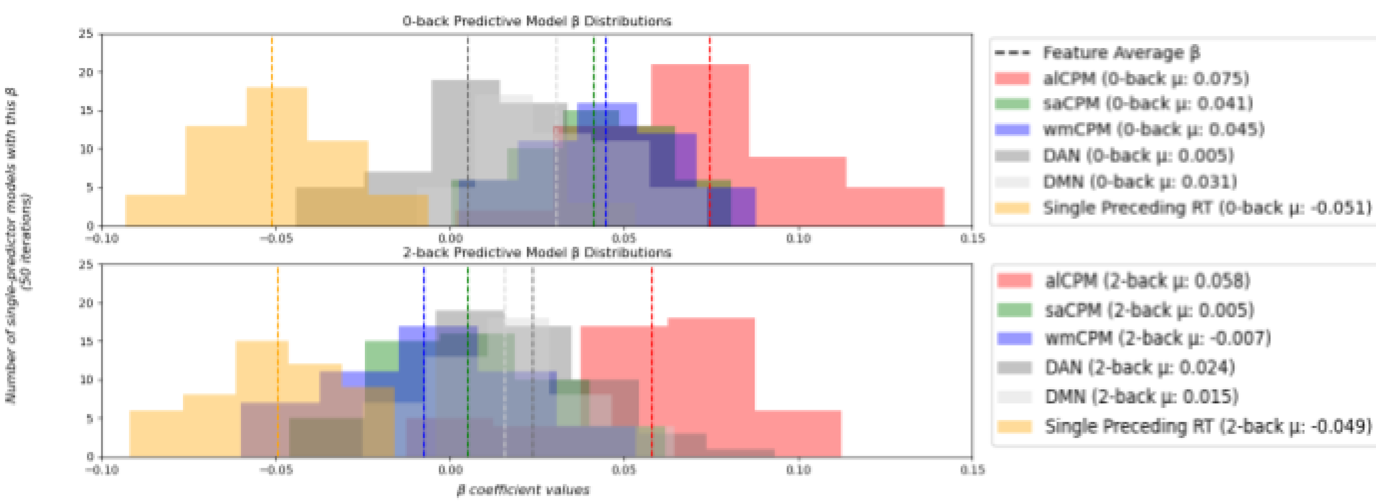


*Fig. 9: joint comparison of predictive model β distributions*

As shown in figure 9, we observe the strongest contributions to lapse prediction from the novel alCPM and the immediately-preceding trial RT. We fit models combining these data in 50 random split-halves, reasoning that these feature coefficient distributions could indicate independent sources of information possibly jointly contributing to prediction.

We then compared the difference in deviance between each predictive model and an empty model. This “empty model” has the formula *Accuracy ~ (1|Participant),* and describes a case where we know there is zero contribution of fixed effects to the outcome variable. This represents the worst-possible fit available under the mixed-effects paradigm (as our random effect of participants is preserved); distance from this empty model as quantified by the difference between empty model deviance and predictive model deviance is therefore a metric by which predictive models may be compared. That is, a larger difference between some given predictive model and the empty model, which we know to be maximally poor *a priori*, indicates a better fit between the predictive model’s predictor and outcome. These differences in deviance may then be averaged over split-half iterations to summarize the relative predictive power of each model tested across our cross-validation scheme.

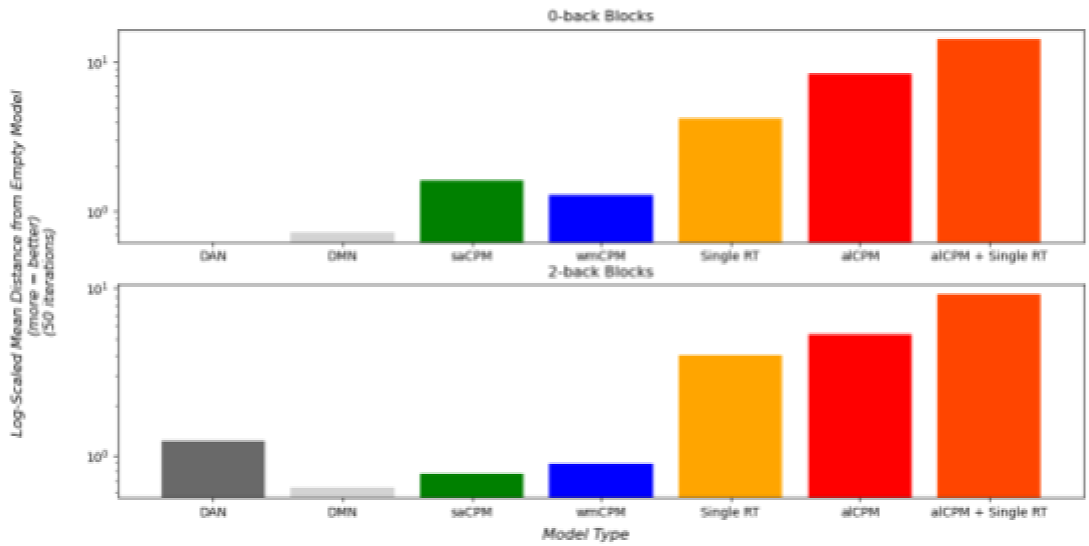


*Fig. 10: predictive model comparison*

As shown in figure 10, the highest-quality models incorporate information from both aICPM cofluctuations and single-preceding-trial RT. Additionally, when evaluating single-predictor statistical models of upcoming lapses, aICPM models consistently outperform the alternatives, including canonical network models, predefined CPMs, and behavior-only RT models. It is worth noting that this superiority is preserved between task difficulty contexts despite the clear reorganization in predictive power exhibited by the alternative models between 0-back and 2-back lapse models.

## Discussion

In this work we have identified some predictors of lapses in sustained attention. Specifically, we applied various models of behavioral and neural predictors of sustained attention lapses described in previous work to novel task, population, and temporal contexts. Additionally, we developed a new attention-lapse connectome-based predictive model (alCPM) shown to capture meaningful variance associated with upcoming EN-back target trial performance on granular timescales which previously-described models either do not fully notice or miss entirely.

***Behavioral Predictors***

We found that RT confers predictive power on statistical models of lapsing sustained attention in the low-memory-load task condition; however, the direction of association between RT and lapses is contrary to the expectation furnished by prior research (Robertson, 2001). We found that, for the EN-back task's 0-back difficulty condition, faster sample mean pre-target RT on the scale of 10s prior to target trials in fact tended to precede correct responses. This directly contradicts findings from other sustained attention tasks such as the SART or gradCPT, wherein faster pre-target RT tends to precede incorrect responses (Robertson, 2001; deBettencourt, 2018). Additionally, we found no association between sample-wide mean pre-trial RT on the 10s scale and task performance for the high-memory-load task condition. This finding has several possible interpretations which could motivate future work. For example, the original SART and gradCPT research was conducted on adults whereas the ABCD Study participants are still maturing. The ABCD Study's longitudinal character allows researchers to track a consistent population over time, and possibly identify whether the relationship observed here between pretrial RT and lapses in sustained attention converges to the expectation over the course of development. Another question involves task differences. The ABCD EN-back task features relatively many breaks, few trials per block, and attention-orienting instruction screens and inter-stimulus intervals. This is quite unlike SARTs and gradCPTs, which are designed to demand more consistent performance. Investigating whether task characteristics such as reorienting stimuli or frequent breaks influence the attentional dynamics underlying task performance could shed light on the occasionally-inconsistent results reported in sustained attention literature.

We applied general linear mixed-effects modeling to look for a relationship between RTs gathered during trials immediately preceding target trials and those target trials' outcomes, assuming the individual variation associated with sustained attention performance may be complicating our analysis while also restricting our time interval of interest to better resemble our connectivity analyses. From these GLMMs we observed that there is indeed sufficient information contained in these single preceding trial RTs to make statistical predictions about the outcomes of their subsequent trials. This may be taken as evidence in support of RT variability as a superior behavioral predictor than mean RT in the moments leading up to lapses in sustained attention (Esterman, 2013). That is, despite the statistical inseparability of sample-wide mean RTs preceding correct vs. incorrect responses to target trials for 2-back blocks, the mixed-effects approach demonstrated that there was still some relevant signal on the same order as that present in 0-back blocks for single pretrial RTs predicting lapses in sustained attention. Although our modeling approach has offered evidence that some information consistently relating single pretrial RTs to lapses exists, future work should disambiguate its source and clarify the relationship between both slowing and variability in RT and lapses in sustained attention. This matter is naturally complicated by the well-documented extent of individual variation in all three constructs, as well as possible associations among them such as the observation that some high-variability response periods may be driven by particularly slow responding (Unsworth, 2021).

***Brain Predictors***

We applied a variety of established network models associated with cognitive processes putatively involved in the ABCD EN-back task to predicting task outcomes on the basis of brain activity in those networks, indexed via functional cofluctuation, during the moments leading up to target trials.

***Canonical Networks***

Applying models of canonical networks such as the DAN and DMN, which have been previously implicated in attentive processes such as effortful attention (Vossel, 2013) and mind-wandering (Kucyi, 2018) respectively, we find that only DMN activity predicts upcoming lapses in 0-back blocks whereas both networks hold significant predictive power for upcoming lapses in 2-back blocks, with the DAN exhibiting numerically higher average contributions to these statistical predictions. The 0-back result roughly coheres with a traditional view of the DAN's involvement in effortful tasks and the DMN's involvement in practiced, less-difficult processes; however, the 2-back results challenge conceptions of these networks as strictly task-positive or task-negative. That is, the average predictive model incorporating information about activity in both networks describes a positive relationship between network activation and task performance. If the DAN and DMN were strictly oppositional, this should not be possible. As such, this finding contributes to our evolving understanding of how these canonical networks are involved in task performance and attentive processes.

***Predefined CPMs***

We applied predefined connectome-based predictive models—the saCPM and wmCPM—previously shown predictive of individual differences in cognitive performance on sustained attention (Rosenberg, 2016) and working memory tasks (Avery, 2020). Additionally, fluctuations within individual saCPMs have been shown to relate with fluctuations in sustained attention on a variety of time scales (Rosenberg, 2020). Extending that effort, this work continues to explore the limits of these models' predictive power by applying them on the finest scales yet. We found that both saCPM and wmCPM activity are similarly predictive of lapses in sustained attention for 0-back blocks; however, only the saCPM extends that power into 2-back blocks, and does so with relative weakness. This finding is somewhat surprising, as the wmCPM should theoretically be well-attuned to a task which engages working memory as extensively as the EN-back 2-back blocks do. Although the task used to develop wmCPM and the EN-back task are quite similar, one critical difference between their contexts is the developmental population studied here. Future work, much in the same vein as the behavioral predictors, should ask whether the inconsistent findings reported here with regard to previously-established models are a consequence of developmental processes. Additionally, the relative relevance saCPM continued to hold for the 2-back task predictions may indicate an underlying cognitive priority; that is, future work could ask whether there is a directionality in the relationship between working memory and sustained attention.

***Novel attention-lapse CPMs***

In addition to the canonical and predefined network models, we applied a data-driven method of deriving new CPMs in to develop models better-predictive of the second-scale relationship between brain activity and upcoming lapses in sustained attention. We found that cofluctuation activity in the resulting attention-lapse connectome predictive models (alCPMs) prior to the presentation of target trials best-predicted those trials' outcomes. In other words, of all the network models tested, information contained in the alCPMs was most relevant to the prediction of lapses in sustained attention. This remained the case for both 0-back and 2-back alCPMs. Future work could investigate whether any commonalities exist between alCPMs derived from similar tasks with differential load conditioning to ask if these fine-scale lapse precursors are task-specific (as the prior network model results suggest) or share an underlying mechanism. For example, we could attempt applying the 0-back alCPM to predicting upcoming 2-back block lapses and vice versa. Another interesting question here deals with the topic of individual differences; namely, one alternative to the present approach which addresses the individual variation literature could involve developing personal alCPMs and then asking whether and the extent to which they generalize—a step impossible for the present data due to their relative poverty of target trials. The slightly larger standard deviations associated with alCPM coefficients across cross-validation iterations as compared with other network model coefficients may indicate that overall alCPM model performance is more affected by the vicissitudes of randomized half-splitting than its siblings, which in turn may point to a larger role for individual variation in alCPM activity's relationship with momentarily lapsing attention than that for e.g. saCPM activity. Another future direction could continue pushing the temporal boundaries of lapse prediction. The present decision to average cofluctuation over 2.4 seconds preceding target trials reflects a conservative preference for higher signal-to-noise ratio at the expense of finer precision. In principle, there is no reason the same basic approach could not be applied to data on scales as small as a single fMRI frame; indeed, the fact that novel alCPMs offered a more robust prediction of upcoming lapses on the basis of preceding brain activity than comparable models trained on much larger time intervals suggests that the mechanisms underlying lapses in sustained attention may involve processes underway at a variety of frequencies. As foreshadowed by our final model comparison, it is likely that the best predictions will result from data combined across many time scales and sources.

***Conclusion***

The present study's limitations offer further directions for elaboration in future work. Some obvious candidates include its population and specific features of the EN-back task. Opportunities to follow up on the present findings in longitudinal research have been discussed and broadly encourage us to ask whether the discrepancies with prior work reported here may be a consequence of systematic differences resulting from developmental processes. More accurately, without that follow-up we can not safely eliminate the confounding influence of participant age on our findings and therefore their generalizability is unclear. Another potential confound is the question of whether our novel alCPM is measuring intrinsic or task-dependent attentional dynamics. Although the EN-back task provides a relatively more natural scenario than SART or gradCPT alternatives, that naturalness involves a variety of attentionally relevant events such as abrupt stimulus onsets and inter-block task breaks. Future work could revisit these classic sustained attention tasks and ask whether lapses in those more-dedicated contexts may be predicted using the analytic approach outlined in the present work. That being said, we may also argue that the rapid reorientation of attention within task blocks may stave off lower-frequency attentional dynamics and allows us to better-investigate high-frequency predictors of lapses in sustained attention.

Despite these minor inconveniences, we have made some contributions to the ongoing program of characterizing sustained attention and its fluctuations. Returning to our initial motivation: if we would like the ability to predict lapses before their behavioral consequences manifest, an important first step would consist of demonstrating whether and the temporal regimes at which such lapses could be predicted on the basis of brain activity. Modifying the functional connectivity/connectome-based predictive modeling approach by incorporating cofluctuation for traction on fine time scales, we have shown that a novel attention-lapse network model can make more meaningful contributions to the statistical prediction of upcoming lapses than canonical and predefined network models of task-relevant cognitive functions such as sustained attention and working memory. These novel alCPMs provide predictive power in this task and population context greater than even its most successful siblings in behavior, and is an essential part of the combined brain/behavior models which best-account for imminent, momentary lapses in sustained attention.

In addition to many directions outlined above, other future work could examine whether and how the network structure and anatomy of alCPMs may offer theoretical insight into the functional mechanisms underlying lapses. Still further work could develop our motivating ideal, replicating the present analyses with more-portable neuroimaging methods such as functional near-infrared spectroscopy to ask whether these neural antecedents observed in-lab are general to the world at large. This last direction may be of particular importance, as MRI scanners are simply not a practical means of sustaining progress to the realization of brain-based forecasting implemented in the daily life activities during which we would most want to predict lapses in sustained attention.

## References

Timmers, D. (2013, October 25). Treating attention deficits and impulse control. Clinical Neurotherapy. Retrieved from https://www.sciencedirect.com/science/article/pii/B9780123969880000064

Cheyne, A. (2010). Attention lapses . The Corsini Encyclopedia of Psychology. Retrieved from https://onlinelibrary.wiley.com/doi/full/10.1002/9780470479216.corpsy0095

Robertson, I. H., Manly, T., Andrade, J., Baddeley, B. T., & Yiend, J. (2001, August 23). `oops!': Performance correlates of everyday attentional failures in traumatic brain injured and normal subjects. Neuropsychologia. Retrieved from https://www.sciencedirect.com/science/article/abs/pii/S0028393297000158
Sakai H, Uchiyama Y, Shin D, Hayashi MJ, Sadato N. Neural activity changes associated with impulsive responding in the sustained attention to response task. PLoS One. 2013 Jun 25;8(6):e67391. doi: 10.1371/journal.pone.0067391. PMID: 23825657; PMCID: PMC3692459.

deBettencourt, M.T., Norman, K.A. & Turk-Browne, N.B. Forgetting from lapses of sustained attention. Psychon Bull Rev 25, 605–611 (2018). https://doi.org/10.3758/s13423-017-1309-5

Rosenberg, M. D., Finn, E. S., Constable, R. T., & Chun, M. M. (2015, March 20). Predicting moment-to-moment attentional state. NeuroImage. Retrieved from https://www.sciencedirect.com/science/article/abs/pii/S1053811915002141?via%3Dihub

Vaurio, R. G., Simmonds, D. J., & Mostofsky, S. H. (2009, January 22). Increased intra-individual reaction time variability in attention-deficit/hyperactivity disorder across response inhibition tasks with different cognitive demands. Neuropsychologia. Retrieved from https://www.sciencedirect.com/science/article/abs/pii/S0028393209000050

Michael Esterman, Sarah K. Noonan, Monica Rosenberg, Joseph DeGutis, In the Zone or Zoning Out? Tracking Behavioral and Neural Fluctuations During Sustained Attention, Cerebral Cortex, Volume 23, Issue 11, November 2013, Pages 2712–2723, https://doi.org/10.1093/cercor/bhs261

Jong, R. D., Berendsen, E., & Cools, R. (1999, October 19). Goal neglect and inhibitory limitations: Dissociable causes of interference effects in conflict situations. Acta Psychologica. Retrieved from https://www.sciencedirect.com/science/article/abs/pii/S0001691899000128?via%3Dihub
Weissman DH, Roberts KC, Visscher KM, Woldorff MG. The neural bases of momentary lapses in attention. Nat Neurosci. 2006 Jul;9(7):971-8. doi: 10.1038/nn1727. Epub 2006 Jun 11. PMID: 16767087.
Cohen MR, Maunsell JH. When attention wanders: how uncontrolled fluctuations in attention affect performance. J Neurosci. 2011 Nov 2;31(44):15802-6. doi: 10.1523/JNEUROSCI.3063-11.2011. PMID: 22049423; PMCID: PMC3579494.

deBettencourt MT, Cohen JD, Lee RF, Norman KA, Turk-Browne NB. Closed-loop training of attention with real-time brain imaging. Nat Neurosci. 2015 Mar;18(3):470-5. doi: 10.1038/nn.3940. Epub 2015 Feb 9. PMID: 25664913; PMCID: PMC4503600.
Kucyi A. Just a thought: How mind-wandering is represented in dynamic brain connectivity. Neuroimage. 2018 Oct 15;180(Pt B):505-514. doi: 10.1016/j.neuroimage.2017.07.001. Epub 2017 Jul 3. PMID: 28684334.
Rosenberg, M., Finn, E., Scheinost, D. et al. A neuromarker of sustained attention from whole-brain functional connectivity. Nat Neurosci 19, 165–171 (2016). https://doi.org/10.1038/nn.4179

Shen, X., Finn, E., Scheinost, D. et al. Using connectome-based predictive modeling to predict individual behavior from brain connectivity. Nat Protoc 12, 506–518 (2017). https://doi.org/10.1038/nprot.2016.178

Rosenberg MD, Scheinost D, Greene AS, Avery EW, Kwon YH, Finn ES, Ramani R, Qiu M, Constable RT, Chun MM. Functional connectivity predicts changes in attention observed across minutes, days, and months. Proc Natl Acad Sci U S A. 2020 Feb 18;117(7):3797-3807. doi: 10.1073/pnas.1912226117. Epub 2020 Feb 4. PMID: 32019892; PMCID: PMC7035597.
Zamani Esfahlani F, Jo Y, Faskowitz J, Byrge L, Kennedy DP, Sporns O, Betzel RF. High-amplitude cofluctuations in cortical activity drive functional connectivity. Proc Natl Acad Sci U S A. 2020 Nov 10;117(45):28393-28401. doi: 10.1073/pnas.2005531117. Epub 2020 Oct 22. PMID: 33093200; PMCID: PMC7668041.
Omid Kardan, Andrew J. Stier, Carlos Cardenas-Iniguez, Julia C. Pruin, Kathryn E. Schertz, Yuting Deng, Taylor Chamberlain, Wesley J. Meredith, Xihan Zhang, Jillian E. Bowman, Tanvi Lakhtakia, Lucy Tindel, Emily W. Avery, Qi Lin, Kwangsun Yoo, Marvin M. Chun, Marc G. Berman, Monica D. Rosenberg. Connectome-based predictions reveal developmental change in the functional architecture of sustained attention and working memory. bioRxiv 2021.08.01.454530; doi: https://doi.org/10.1101/2021.08.01.454530

Song H., Shim W. M., Rosenberg, M.D. (2022) Neural dynamics in a low-dimensional state space reflect cognitive and attentional dynamics. Presented at the Organization for Human Brain Mapping (OHBM) 2022 Annual Meeting, Glasgow, Scotland, June 22

Welhaf MS, Smeekens BA, Meier ME, Silvia PJ, Kwapil TR, Kane MJ. The Worst Performance Rule, or the Not-Best Performance Rule? Latent-Variable Analyses of Working Memory Capacity, Mind-Wandering Propensity, and Reaction Time. J Intell. 2020 Jun 2;8(2):25. doi: 10.3390/jintelligence8020025. PMID: 32498311; PMCID: PMC7713012.
Broadbent DE, Cooper PF, FitzGerald P, Parkes KR. The Cognitive Failures Questionnaire (CFQ) and its correlates. Br J Clin Psychol. 1982 Feb;21(1):1-16. doi: 10.1111/j.2044-8260.1982.tb01421.x. PMID: 7126941.
Unsworth, N., Redick, T. S., Lakey, C. E., & Young, D. L. (2010). Lapses in sustained attention and their relation to executive control and fluid abilities: An individual differences investigation. Intelligence, 38(1), 111–122. https://doi.org/10.1016/j.intell.2009.08.002

Hagler DJ Jr, Hatton S, Cornejo MD, Makowski C, Fair DA, Dick AS, Sutherland MT, Casey BJ, Barch DM, Harms MP, Watts R, Bjork JM, Garavan HP, Hilmer L, Pung CJ, Sicat CS, Kuperman J, Bartsch H, Xue F, Heitzeg MM, Laird AR, Trinh TT, Gonzalez R, Tapert SF, Riedel MC, Squeglia LM, Hyde LW, Rosenberg MD, Earl EA, Howlett KD, Baker FC, Soules M, Diaz J, de Leon OR, Thompson WK, Neale MC, Herting M, Sowell ER, Alvarez RP, Hawes SW, Sanchez M, Bodurka J, Breslin FJ, Morris AS, Paulus MP, Simmons WK, Polimeni JR, van der Kouwe A, Nencka AS, Gray KM, Pierpaoli C, Matochik JA, Noronha A, Aklin WM, Conway K, Glantz M, Hoffman E, Little R, Lopez M, Pariyadath V, Weiss SR, Wolff-Hughes DL, DelCarmen-Wiggins R, Feldstein Ewing SW, Miranda-Dominguez O, Nagel BJ, Perrone AJ, Sturgeon DT, Goldstone A, Pfefferbaum A, Pohl KM, Prouty D, Uban K, Bookheimer SY, Dapretto M, Galvan A, Bagot K, Giedd J, Infante MA, Jacobus J, Patrick K, Shilling PD, Desikan R, Li Y, Sugrue L, Banich MT, Friedman N, Hewitt JK, Hopfer C, Sakai J, Tanabe J, Cottler LB, Nixon SJ, Chang L, Cloak C, Ernst T, Reeves G, Kennedy DN, Heeringa S, Peltier S, Schulenberg J, Sripada C, Zucker RA, Iacono WG, Luciana M, Calabro FJ, Clark DB, Lewis DA, Luna B, Schirda C, Brima T, Foxe JJ, Freedman EG, Mruzek DW, Mason MJ, Huber R, McGlade E, Prescot A, Renshaw PF, Yurgelun-Todd DA, Allgaier NA, Dumas JA, Ivanova M, Potter A, Florsheim P, Larson C, Lisdahl K, Charness ME, Fuemmeler B, Hettema JM, Maes HH, Steinberg J, Anokhin AP, Glaser P, Heath AC, Madden PA, Baskin-Sommers A, Constable RT, Grant SJ, Dowling GJ, Brown SA, Jernigan TL, Dale AM. Image processing and analysis methods for the Adolescent Brain Cognitive Development Study. Neuroimage. 2019 Nov 15;202:116091. doi: 10.1016/j.neuroimage.2019.116091. Epub 2019 Aug 12. PMID: 31415884; PMCID: PMC6981278.

Power JD, Barnes KA, Snyder AZ, Schlaggar BL, Petersen SE. Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. Neuroimage. 2012 Feb 1;59(3):2142-54. doi: 10.1016/j.neuroimage.2011.10.018. Epub 2011 Oct 14. Erratum in: Neuroimage. 2012 Nov 1;63(2):999. PMID: 22019881; PMCID: PMC3254728.

Shen X, Tokoglu F, Papademetris X, Constable RT. Groupwise whole-brain parcellation from resting-state fMRI data for network node identification. Neuroimage. 2013 Nov 15;82:403-15. doi: 10.1016/j.neuroimage.2013.05.081. Epub 2013 Jun 4. PMID: 23747961; PMCID: PMC3759540.

Avery EW, Yoo K, Rosenberg MD, Greene AS, Gao S, Na DL, Scheinost D, Constable TR, Chun MM. Distributed Patterns of Functional Connectivity Predict Working Memory Performance in Novel Healthy and Memory-impaired Individuals. J Cogn Neurosci. 2020 Feb;32(2):241-255. doi: 10.1162/jocn_a_01487. Epub 2019 Oct 29. PMID: 31659926; PMCID: PMC8004893.

Unsworth, N., Robison, M. K., & Miller, A. L. (2021). Individual differences in lapses of attention: A latent variable analysis. Journal of Experimental Psychology: General, 150(7), 1303–1331. https://doi.org/10.1037/xge0000998

Vossel S, Geng JJ, Fink GR. Dorsal and ventral attention systems: distinct neural circuits but collaborative roles. Neuroscientist. 2014 Apr;20(2):150-9. doi: 10.1177/1073858413494269. Epub 2013 Jul 8. PMID: 23835449; PMCID: PMC4107817.