

Contralateral Delay Activity Tracks Fluctuations in Working Memory Performance

Kirsten C. S. Adam¹, Matthew K. Robison², and Edward K. Vogel¹

Abstract

■ Neural measures of working memory storage, such as the contralateral delay activity (CDA), are powerful tools in working memory research. CDA amplitude is sensitive to working memory load, reaches an asymptote at known behavioral limits, and predicts individual differences in capacity. An open question, however, is whether neural measures of load also track trial-by-trial fluctuations in performance. Here, we used a whole-report working memory task to test the relationship between CDA amplitude and working memory performance. If working memory failures are due to decision-based errors and retrieval failures, CDA amplitude would not differentiate good and poor performance trials when load is held constant. If failures arise during storage, then CDA amplitude should track both working memory load and trial-by-trial performance. As expected, CDA amplitude

tracked load (Experiment 1), reaching an asymptote at three items. In Experiment 2, we tracked fluctuations in trial-by-trial performance. CDA amplitude was larger (more negative) for high-performance trials compared with low-performance trials, suggesting that fluctuations in performance were related to the successful storage of items. During working memory failures, participants oriented their attention to the correct side of the screen (lateralized P1) and maintained covert attention to the correct side during the delay period (lateralized alpha power suppression). Despite the preservation of attentional orienting, we found impairments consistent with an executive attention theory of individual differences in working memory capacity; fluctuations in executive control (indexed by pretrial frontal theta power) may be to blame for storage failures. ■

INTRODUCTION

Our attention fluctuates from moment to moment, both in laboratory tasks and in our daily lives (Kane et al., 2017; Reason, 1984). During periods of relative inattention, participants have more erratic RTs and are likely to miss targets (e.g., Esterman, Noonan, Rosenberg, & DeGutis, 2013). Although fluctuations of attention are most commonly studied with simple RT paradigms, recent work has revealed that these fluctuations of attention also impact more complex processes, such as memory. For example, participants are less likely to remember items when they were encoded during a period of suboptimal attention (Aly & Turk-Browne, 2016a, 2016b), and causally presenting items during periods of optimal attention increases memory performance (deBettencourt, Norman, & Turk-Browne, 2017). In other words, attentional state has a strong influence on the fate of memories. Many synonymous terms have been used to describe fluctuations of attentional state in the literature (e.g., attentional control, executive attention). Throughout the article, we use the term “executive control” to refer to the allocation of central attentional resources to the task at hand.

In this work, we examined how fluctuations of attentional state may influence working memory performance. Most individuals are capable of storing around three to four simple items in visual working memory (Cowan, 2001; Luck & Vogel, 1997), but they frequently underperform this potential capacity (Adam, Mance, Fukuda, & Vogel, 2015). Critically, the rate of these working memory “failures” is strongly predictive of working memory capacity as a whole. A key question, then, is when and why working memory failures occur. During each trial of a working memory task, there are many aspects of task performance that could go awry. Participants must attend to the task at hand, encode and individuate items, store them, protect them from interference, and decide on a response. To test which aspects of task performance are disrupted during working memory failures, we took advantage of previously established ERPs and oscillatory markers of processes thought to be critical for successful working memory performance. Using these markers, we sought to identify the most critical aspects of task performance that are disrupted during working memory failures.

To perform poorly on a working memory test, it seems obvious that participants would fail to maintain working memory representations throughout the entire delay period. However, this is not necessarily the case. For example, participants could successfully maintain items during the retention period, but then experience interference or fail to retrieve this information at test (Souza, Rerko, &

This paper is part of a Special Focus deriving from a symposium at the 2017 annual meeting of Cognitive Neuroscience Society, entitled, “Fluctuations in Attention and Cognition.”

¹University of Chicago, ²University of Oregon

Oberauer, 2016; Harlow & Donaldson, 2013). To test whether items are dropped from working memory during maintenance on failure trials, we looked at the amplitude of the contralateral delay activity (CDA). The CDA is a measure of working memory storage, and it is measured in lateralized working memory tasks in which participants are asked to remember items in one visual hemifield and ignore items in the other visual hemifield. During the maintenance period, there is a sustained negativity in contralateral electrodes relative to ipsilateral electrodes. This negative difference is the CDA. The CDA tracks working memory load, becoming more negative in amplitude until hitting an asymptote around typical capacity estimates (Vogel & Machizawa, 2004). The CDA also correlates with individual differences in working memory performance (Luria, Balaban, Awh, & Vogel, 2016; Vogel, McCollough, & Machizawa, 2005), supporting its role as a relevant marker of both between- and within-subject variation in working memory maintenance. Thus, the CDA is widely believed to index the amount of information held in working memory on each trial. However, previous studies have lacked trial-by-trial resolution to see whether CDA is disrupted during momentary failures of working memory performance.

Maintenance is one critical aspect of working memory performance, but other stages of processing are also important. In addition to using the CDA to track whether working memory maintenance failures explain trial-by-trial fluctuations in working memory performance, we can also use established ERP markers to assess other aspects of task performance. For example, in lateralized working memory task designs, participants must first use a cue to orient their attention to one side of the display. Then, they must select the relevant items from the cued side of the display. These two processes are reflected in the lateralized P1 component (e.g., Van Voorhis & Hillyard, 1977; for a review, see Mangun, 1995) and the N2PC (e.g., Luck & Hillyard, 1990, 1994), respectively. Other processes of interest can be tracked with oscillatory markers. For example, theta power (4–7 Hz) at frontal electrodes is thought to covary with executive control processes (Cavanagh & Frank, 2014) and is inversely correlated with activity in the default mode network (e.g., Scheeringa et al., 2008). In previous work, we found that low frontal theta power predicted poor working memory performance (Adam et al., 2015). Here, we sought to replicate this finding. Finally, alpha power (8–12 Hz) suppression in contralateral electrodes has previously been shown to track the locus of covert attention (e.g., Thut, Nietzel, Brandt, & Pascual-Leone, 2006; Sauseng et al., 2005; Worden, Foxe, Wang, & Simpson, 2000). By examining alpha power suppression, we can test whether sustained covert attention before or during maintenance predicts fluctuations in working memory performance.

To preview results, we found that the amplitude of the CDA tracked working memory performance. When par-

ticipants successfully recalled more items at test, the CDA amplitude was more negative. We also replicated previous findings that theta power (4–7 Hz) at frontal electrodes predicted working memory performance even before the memory array had appeared (Adam et al., 2015), implicating suboptimal executive control during working memory failures. Interestingly, however, some aspects of task performance were preserved during working memory failures. For example, participants still correctly oriented their attention to the cued side of the screen (lateralized P1) and strongly maintained covert attention to the attended side throughout the entire trial (lateralized alpha power suppression). Together, our results suggest a locus for working memory failures after encoding has occurred, during working memory maintenance. Our findings also support models that propose a tight link between the consistency of executive attention and working memory ability (e.g., Souza & Oberauer, 2017; Kane, Conway, Hambrick, & Engle, 2008).

Overview of Experiments

In two experiments, participants completed a lateralized whole-report task (e.g., Adam et al., 2015; Huang, 2010) while EEG data were collected. In Experiment 1, we measured changes in performance and the CDA amplitude with set size. In Experiment 2, we held the memory set size constant at six items and measured trial-by-trial fluctuations in performance, lateralized ERP components, and oscillatory signals. In both experiments, participants also completed a separate behavioral color change detection task (e.g., Luck & Vogel, 1997) at the beginning of the experimental session.

These experiments replicate several findings from Adam et al. (2015) but also make novel contributions to our understanding of working memory failures. Here, we introduced balanced, lateralized displays that allowed us to exploit well-characterized ERP and oscillatory signals (lateralized P1, N2PC, CDA, and lateralized alpha). This updated task design allowed us to test whether working memory failures disrupted attentional orienting, item selection, item maintenance, and sustained spatial attention. We also assessed whether individual differences in the CDA correlated with behavior, and we put forth a novel theoretical account of individual differences in CDA amplitude. Finally, we replicated analyses of global signals (frontal theta, global alpha, global P1) measured in Adam et al. (2015).

METHODS

Participants

Participants were recruited from the University of Oregon and surrounding community. All participants were between the ages of 18 and 35 years and had self-reported normal or corrected-to-normal visual acuity and normal

color vision. All participants gave informed consent and completed the 3-hr session for \$30 in compensation. A total of 31 participants (12 women, $M = 22.0$ years, $SD = 3.7$) participated in Experiment 1, and 48 participants (24 women, $M = 21.6$ years, $SD = 3.97$) participated in Experiment 2. Three participants were excluded from Experiment 2 before artifact rejection (one had a missing behavior file, two left the session after a few blocks), leaving 45 participants with usable EEG data. Participants were excluded from Experiment 1 analyses if they had fewer than 75 trials in any set size condition after artifact rejection (remaining $n = 29$). Participants were excluded from Experiment 2 analyses if they had fewer than 40 trials per condition after artifact rejection (remaining $n = 38$).¹ In Experiment 2, the two conditions of interest were “high performance” trials (>3 correct) and “low performance” trials (<3 correct). For combined experiment correlation analyses, participants were excluded if they were missing change detection data or if they had fewer than 75 trials per set size after artifact rejection (total $n = 72$).

Stimuli

Stimuli were rendered using the Psychophysics toolbox (Brainard, 1997; Pelli, 1997) and presented on a 17-in. cathode ray tube monitor. Participants were seated 100 cm from the screen, though a chin rest was not used so all visual angle calculations are approximate. In all experiments, participants remembered colored squares presented on a medium gray background (RGB = 127.5 127.5 127.5). Participants maintained fixation on a small black dot (0.12°). Colors of the squares were chosen from a pool of nine distinct colors: red (RGB = 255 0 0), green (0 255 0), blue (0 0 255), yellow (255 255 0), magenta (255 0 255), cyan (0 255 255), orange (255 128 0), white (255 255 255), and black (1 1 1). Each square subtended 1.2° , and there was a minimum distance requirement of at least 1.5 squares between the centroids of any two squares. Squares could appear anywhere within a portion of the display subtending 7.0° to the left or right of fixation and 5.2° above or below fixation. For the lateralized whole-report task, participants were cued to attend either the left- or right-half of the display before the onset

of the memory array with a small pink and green diamond (inset of Figure 1). The diamond was approximately 0.2° tall by 0.4° wide and was presented 0.4° above the fixation cross.

Tasks

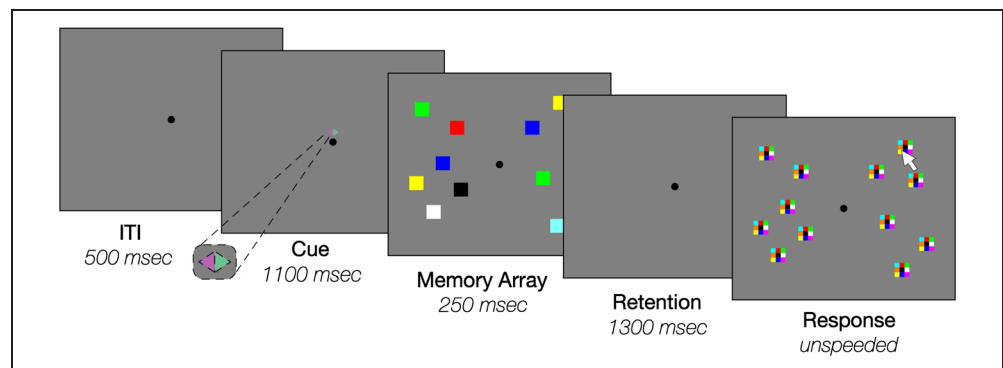
Discrete Whole-Report Task

The whole-report task was the primary task of interest while EEG data were recorded. Each trial began with a blank intertrial interval (500 msec) followed by a small diamond-shaped cue (1100 msec). Participants were instructed to direct their attention to the side of the display indicated by the green side of the diamond and to remember the items on that side. The cue stayed on the screen for the remainder of the trial. After the cue period ended, a memory array was presented (250 msec). The memory array contained an equal number of items on the cued side and the uncued side. Colors were chosen without replacement within each side (i.e., all cued colors were unique but might be repeated on the uncued side of the display). After encoding, participants remembered the items across a blank delay (1300 msec). At test, a 3×3 matrix of the nine possible colors was presented at the location of each item on both the attended and unattended side. Participants were instructed to click the color in each matrix corresponding to the color presented at the location. The response period ended after participants made a response for all items on the attended side. Participants clicked the mouse to initiate the beginning of the next trial.

Color Change Detection Task

We used a separate behavioral change detection task to measure participants’ working memory capacity. Each trial began with an intertrial interval of 1000 msec. Next, the memory array (three, six, or eight colored squares) appeared for 250 msec; participants remembered the colors and locations of the squares for a blank delay of 1000 msec. For set sizes 3 and 6, colors were chosen without replacement from the pool of nine colors. For

Figure 1. Trial events in Experiment 1 and Experiment 2. Trial events are depicted from left to right. In Experiment 1, the memory array could contain one, three, or six items. In Experiment 2, the memory array always contained six items.



set size 8, colors were chosen randomly from a doubled list of the colors (i.e., each color could be repeated up to one time in the array). After the blank delay, participants were presented with a probe at one of the remembered locations. On 50% of trials, the probe was the same color as the remembered item at that location (“no change” trials). On the other 50% of trials, the probe was a different color from the remembered item (“change” trials). Participants gave an unspeeded response; they were instructed to press the “Z” key for no-change trials and the “?” key for change trials. The next trial began immediately after participants responded.

Procedures

Session length was ~2.5 hr in Experiment 1 and ~3 hr in Experiment 2. Participants first completed 144 trials of the change detection task (48 trials each of set sizes 3, 6, and 8). Change detection data were not collected for one participant in Experiment 1. Partial change detection data were obtained for one participant in Experiment 2 because of a computer crash (108/144 trials). After beginning the EEG recording, participants did the whole-report task for the remainder of the session. Trials were self-paced and were collected in blocks of 30 trials (10 trials each of set sizes 1, 3, and 6 in Experiment 1, 30 trials of set size 6 in Experiment 2). After each block, participants received a short break (~30 sec) before continuing. Participants completed an average of 21.1 blocks ($SD = 5.3$) in Experiment 1 and 16.7 blocks ($SD = 3.3$) in Experiment 2.

EEG Acquisition

Before completing the tasks, participants were fitted with an elastic cap with 20 electrodes (ElectroCap International, Eaton, OH). We recorded from International 10/20 sites F3, Fz, F4, T3, C3, Cz, C4, T4, P3, Pz, P4, T5, O1, and O2 along with five nonstandard sites: OL midway between T5 and O1, OR midway between T6 and O2, PO3 midway between P3 and OL, PO4 midway between P4 and OR, and POz midway between PO3 and PO4. All sites were recorded with a right mastoid reference, and the data were rereferenced offline to the algebraic average of the left and right mastoids. Horizontal EOG (HEOG) was recorded from electrodes placed about 1 cm from the left and right of the external canthi of each eye to measure horizontal eye movements. To detect blinks, vertical EOG was recorded from an electrode mounted beneath the right eye. The EEG and EOG signals were amplified with an SA Instrumentation amplifier (Fife, Scotland) with a band-pass of 0.01–80 Hz and were digitized at 250 Hz in Labview 6.1 (Fife, Scotland) running on a PC. EEG activity was collected during the discrete whole-report task only.

EEG Analysis

Artifact Rejection

Participants were instructed not to move their eyes or blink during the trial until the test array appeared on the screen. Trials including horizontal eye movements, blinks, blocking (amplifier saturation after drift), or excessive noise were rejected. For horizontal eye movement rejection, we used a split-half sliding window approach (window size = 200 msec, step size = 10 msec, threshold = 20 μ V) on the HEOG signal. We slid a 200-msec time window in steps of 10 msec from the beginning to the end of the trial. If the change in voltage from the first half to the second half of the window was greater than 20 μ V, it was marked as an eye movement and rejected. We also used a sliding window step function to check for blinks in the vertical EOG (window size = 200 msec, step size = 10 msec, threshold = 50 μ V). For blocking rejection, we slid a 200-msec time window in steps of 50 msec and excluded trials for blocking if any EEG electrode had at least 15 consecutive time points (i.e., 60 msec) that were within 1 μ V of each other. We excluded trials for excessive noise if any electrode had peak-to-peak amplitude greater than 200 μ V within a 15-msec time window. Finally, we visually inspected the data to confirm automatic rejection criteria.

ERPs

For ERP analyses, we baselined the signal over the 200 msec before the time-locking event (onset of the memory array). Lateralized waveforms were built by subtracting the average of the ipsilateral electrodes from the average of the contralateral electrodes. Lateral-occipital and posterior-parietal electrodes used for lateralized waveforms were O1, O2, OL, OR, P3, P4, PO3, PO4, T5, and T6. Statistics were performed on the baselined, unfiltered data. For visualization purposes, trials were low-pass filtered with a two-way least squares finite impulse response filter (*eegfilt.m*; Delorme & Makeig, 2004) with a cutoff of 30 Hz.

Time–Frequency Analyses

For time–frequency analyses, we bandpass-filtered the raw EEG using a two-way, least squares finite impulse response filter using the *eegfilt.m* function from the EEGLAB Toolbox (Delorme & Makeig, 2004) and applied the MATLAB Hilbert transform (*hilbert.m*) to extract the instantaneous power values for the theta band (4–7 Hz) and the alpha band (8–12 Hz). Percent change in power was calculated relative to a baseline period before the onset of the cue (–1500 to –1100 msec relative to memory array onset). Electrodes for frontal theta and posterior alpha were chosen a priori from the literature (Adam et al., 2015; Fukuda, Mance, & Vogel, 2015). Frontal theta was calculated as the average of theta power in the electrodes F3, F4, and Fz. Lateralized alpha power was calculated by subtracting the percent change waveform for

ipsilateral electrodes from the percent change waveform for contralateral electrodes. Lateral occipital and posterior parietal electrodes used for lateralized alpha power were the same as for the CDA: O1, O2, OL, OR, P3, P4, PO3, PO4, T5, and T6.

Change Detection Performance

Change detection performance was converted into a capacity estimate (K) for each set size in the change detection task, following the formula $K = N \times (H - FA)$, where N represents the set size, H is the hit rate (proportion of correct change trials), and FA is the false alarm rate (proportion of incorrect no-change trials). This formula (Cowan, 2001) is most appropriate for single-probe displays like the ones used here (Rouder, Morey, Morey, & Cowan, 2011). Average change detection performance (mean K) was calculated as the average of performance for all set sizes (three, six, and eight items).

RESULTS

Behavioral Results

Mean performance on the change detection and discrete whole-report tasks was similar to prior studies (e.g., Adam et al., 2015; Luck & Vogel, 1997) and is shown in Figure 2. Average change detection capacity (K) was 2.62 ($SD = 1.00$, range = 0.83–4.54) in Experiment 1 and 2.64 ($SD = .70$, range = 1.13–4.40) in Experiment 2. In Experiment 1, participants reported 0.95 ($SD = 0.04$) items correct for set size 1, 2.41 ($SD = 0.33$) items correct for set size 3, and 2.53 ($SD = 0.53$) items correct for set size 6. In Experiment 2, participants correctly reported on average 2.64 items ($SD = 0.35$, range = 1.95–3.55) out of 6. In both experiments, average whole-report performance was positively correlated with average change detection K ($r = .59, p = .001$ in Experiment 1; $r = .34, p = .04$ in Experiment 2).

Electrophysiological Results

Experiment 1: Replication of Neural Correlates of Set Size

First, we examined typical neural correlates of set size in our whole-report working memory task. We found that markers of working memory storage as well as covert attention tracked set size in the whole-report task (Figure 3). We examined CDA amplitude as a marker of storage and N2PC amplitude as a measure of selection. The N2PC (200–300 msec) became more negative with set size, $F(2, 56) = 18.89, p < .001, \eta_p^2 = .40$, and reached an asymptote between three and six items ($p = .23$). Likewise, CDA amplitude (400–1500 msec) became more negative with set size, $F(1.67, 46.51) = 25.10, p < .001, \eta_p^2 = .47$, and reached an asymptote between three and six items (difference between set size 3 and 6, $p = .79$). Thus, our novel whole-report demands did not alter the typical pattern of results observed in change detection studies.

Next, we examined changes in alpha power suppression as a function of set size. Lateralized alpha power suppression (contralateral–ipsilateral) has been shown to index the deployment of covert spatial attention to one side of the display (e.g., Worden et al., 2000). Independent of this signal, global alpha power suppression (across all posterior electrodes) has been shown to covary with working memory load (Fukuda, Kang, & Woodman, 2016). We tested whether one or both of these signals would covary with set size in Experiment 1. To do so, we ran a repeated-measures ANOVA with the factors Hemifield (contralateral vs. ipsilateral) and Set size (1, 3, or 6) for both pretrial lateralization of alpha power (–1100 to 0 msec) and delay period lateralization of alpha power (400–1500 msec). During the pretrial cue period, alpha power was significantly lateralized as shown by a significant effect of Hemifield, $F(1, 28) = 26.94, p < .001, \eta_p^2 = .49$. As expected, there was no pretrial effect of Set size on alpha power ($p = .73$) or an interaction between Set size and Hemifield ($p = .48$). During the delay

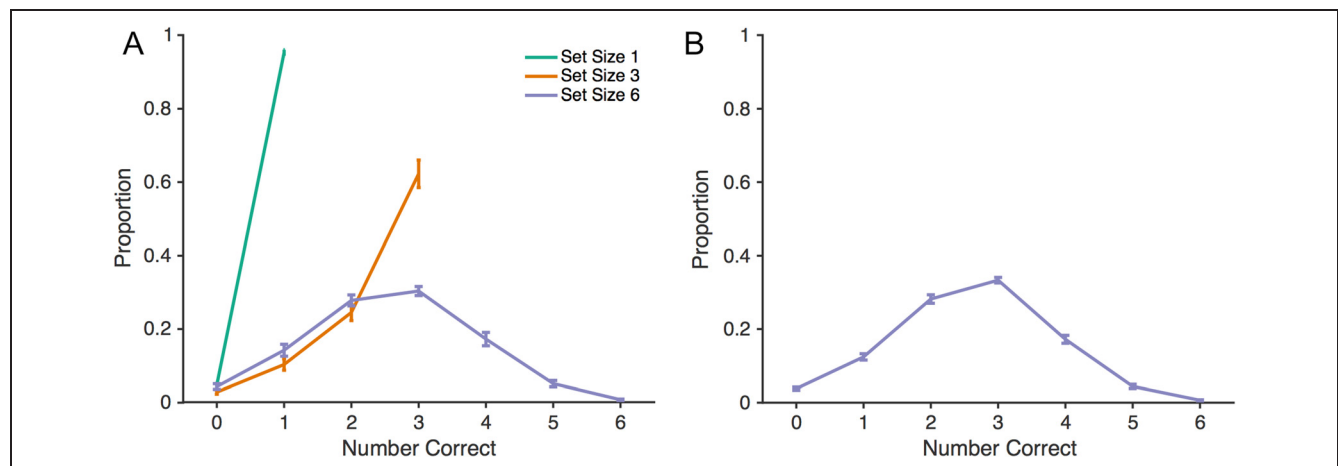


Figure 2. Behavioral performance in Experiment 1 (A) and Experiment 2 (B). Error bars represent 1 SEM.

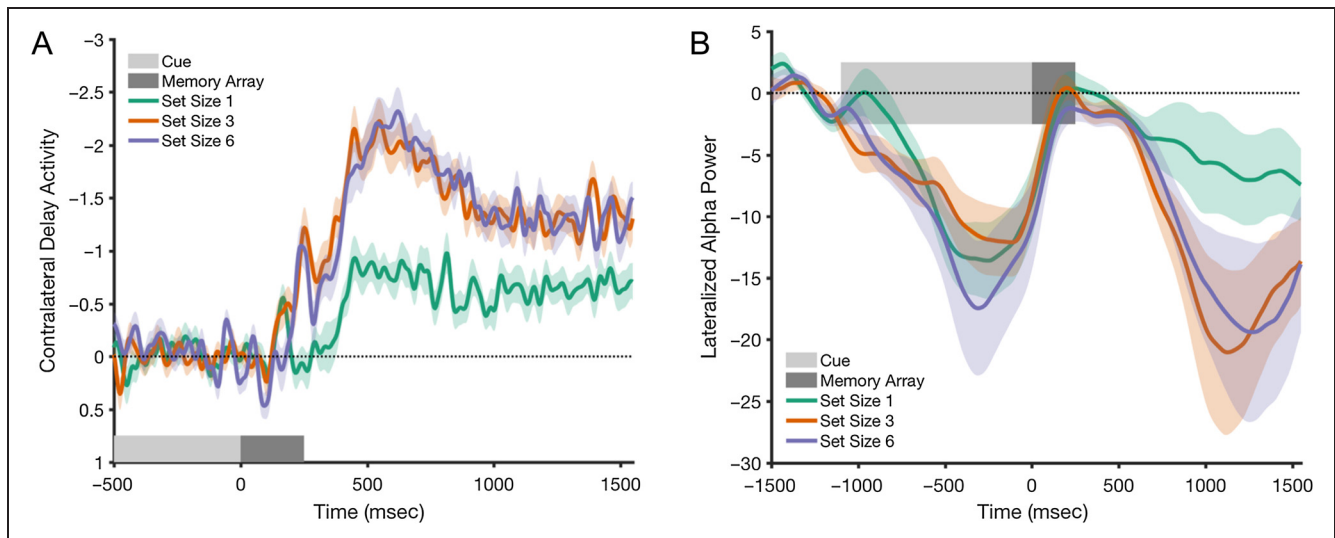


Figure 3. CDA (A) and lateralized alpha power (B) as a function of set size in Experiment 1. Shaded error bars represent 1 *SEM*. During the delay period (400–1500 msec), CDA and lateralized alpha power both tracked memory load.

period, there was a main effect of Hemifield, indicating systematic lateralization of alpha power, $F(1, 28) = 9.68, p = .004, \eta_p^2 = .26$. We also found a main effect of Set size on alpha power across all electrodes, $F(1.4, 39.2) = 7.36, p = .005, \eta_p^2 = .21$, consistent with previous work demonstrating that global alpha power is more suppressed for higher set sizes during visual working memory tasks (Fukuda et al., 2015). Finally, we also observed an interaction between Set size and Hemisphere, $F(1.34, 37.48) = 5.87, p = .013, \eta_p^2 = .17$, indicating that alpha power was more lateralized for higher set sizes. Post hoc comparisons revealed that the difference between contralateral and ipsilateral alpha power was significantly smaller for set size 1 compared with set size 3 ($p = .01$), but not for set size 3 to set size 6 ($p = .51$). In summary, we replicated the finding that global alpha power suppression tracks memory load (Fukuda et al., 2015, 2016). In addition, we found that participants sustained their covert attention (lateralized alpha power suppression) throughout the memory delay, but did so less strongly for sub-capacity set size 1 arrays. This finding is consistent with Sauseng et al. (2009) but inconsistent with Fukuda et al. (2016), who found no effect of memory load on lateralization of alpha power.

Experiment 2: Neural Correlates of Trial-by-trial Fluctuations in Performance

In Experiment 2, we examined predictors of performance rather than of set size. Critically, if poor performance effectively modulates working memory load (e.g., fewer items were stored), then we would predict that correlates of performance fluctuations should be similar to correlates of set size. If instead working memory failures are primarily caused by errors at the retrieval or decision stage, then poor performance should not covary with

markers of storage. To examine predictors of trial-by-trial performance, we analyzed the difference between “good” trials (four or more items correct) and “poor” trials (two or fewer items correct). This specific behavioral threshold was chosen for a couple of reasons. First, because participants are required to report all of the items, they will sometimes get additional items correct by chance; eliminating the middle category (three correct) minimizes the overlap between the “good” and “poor” categories. Second, we wanted the current results to be directly comparable to earlier work on this topic (Adam et al., 2015).

First, we examined whether participants successfully stored items throughout the working memory delay period, using CDA amplitude as a proxy for successful storage. The CDA (400–1500 msec) discriminated between good and poor performance trials, $t(37) = 2.81, p = .008, 95\% \text{ CI } [.08, .48]$, indicating that participants successfully maintained fewer items when working memory performance was poor (Figure 4A). The change in CDA amplitude across good and poor trials could not be explained by a decrease in eye movement artifacts. There was no main effect of behavioral performance on HEOG amplitude during the delay period ($p = .42$). In addition, there was no difference in artifact rejection rates for good performance trials relative to poor performance trials ($p = .84$). Consistent with previous work (Adam et al., 2015), this suggests that participants were no more likely to be task-noncompliant (e.g., blinking during the memory array) during poor performance trials.

One potential explanation for smaller CDA amplitude during poor performance trials is that participants were completely disengaged from the task at hand. If disengaged from the task, participants may have failed to use the spatial cue altogether or mistakenly attended the wrong side of the display. To test whether participants were wholly disengaged, we looked at markers of

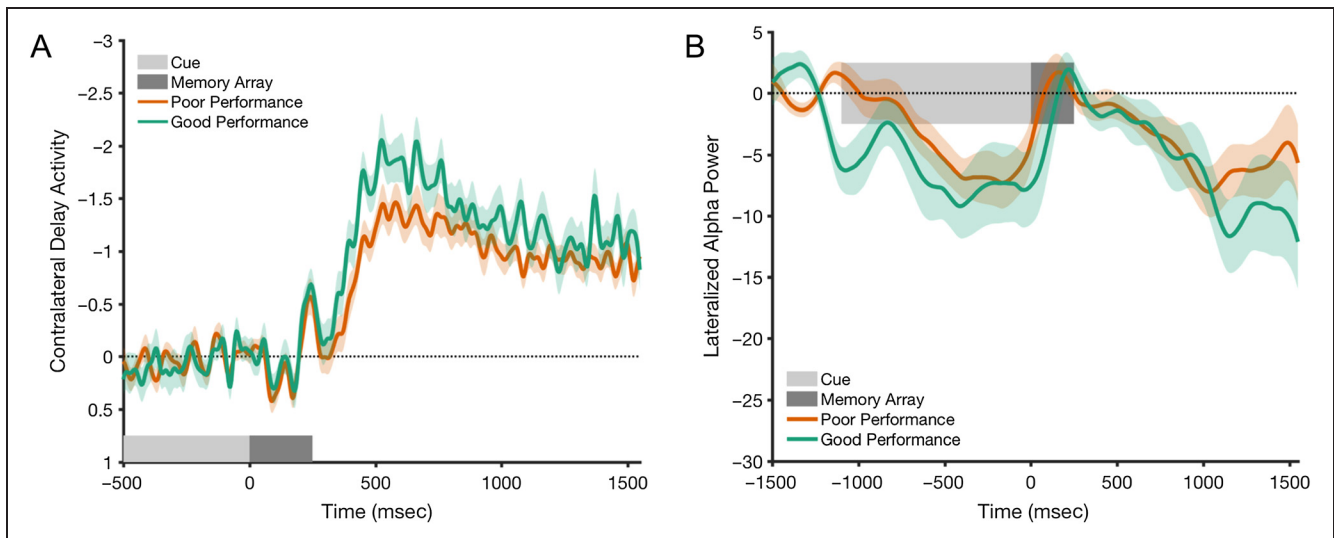


Figure 4. CDA (A) and lateralized alpha power (B) as a function of accuracy in Experiment 2. Shaded error bars represent 1 *SEM*. During the delay period (400–1500 msec), CDA tracked variability in memory performance but lateralized alpha power did not.

early attentional selection and individuation, the lateralized P1 component and the N2PC component. If participants did not selectively attend to the cued side during working memory failures, we would expect to see a diminished lateralized P1 component (70–120 msec). If participants also selected fewer items on the correct side, we should also see a diminished N2PC response (200–300 msec). Somewhat surprisingly, we found that attentional orienting and selection were preserved during poor performance trials (Figure 5). A repeated-measures ANOVA with the factors Hemifield (contralateral vs. ipsilateral) and Memory performance (good vs. poor) revealed that the P1 component was significantly lateralized overall, as indicated by a main effect of Hemifield, $F(1, 37) = 15.8$, $p < .001$, $\eta_p^2 = .30$. However, there was no interaction between Hemifield and Performance, indicating no difference in lateralized P1 amplitude for good versus poor performance trials, $F(1, 37) = 1.04$, $p = .32$, $\eta_p^2 = .027$.

In addition, there was no main effect of performance on global PI amplitude, $F(1, 37) = .014$, $p = .91$, $\eta_p^2 < .001$. Likewise, the N2PC time window was significantly lateralized overall, $F(1, 37) = 9.07$, $p = .005$, $\eta_p^2 = .20$, but there was no interaction between Lateralization and Performance, $F(1, 37) = 2.03$, $p = .16$, $\eta_p^2 = .05$. In summary, early attentional selection did not predict fluctuations in working memory performance.

Next, we checked whether a sustained measure of covert attention might be more sensitive to fluctuations in spatial attention. The P1 component is relatively transient; it only briefly measures the allocation of attention at the moment a stimulus is presented. To look at sustained spatial attention, we used lateralized alpha power. Lateralized alpha power suppression tracks the location of covert attention in a sustained, fine-grained fashion (e.g., Foster, Sutterer, Serences, Vogel, & Awh, 2017). Using lateralized alpha power suppression, we found that

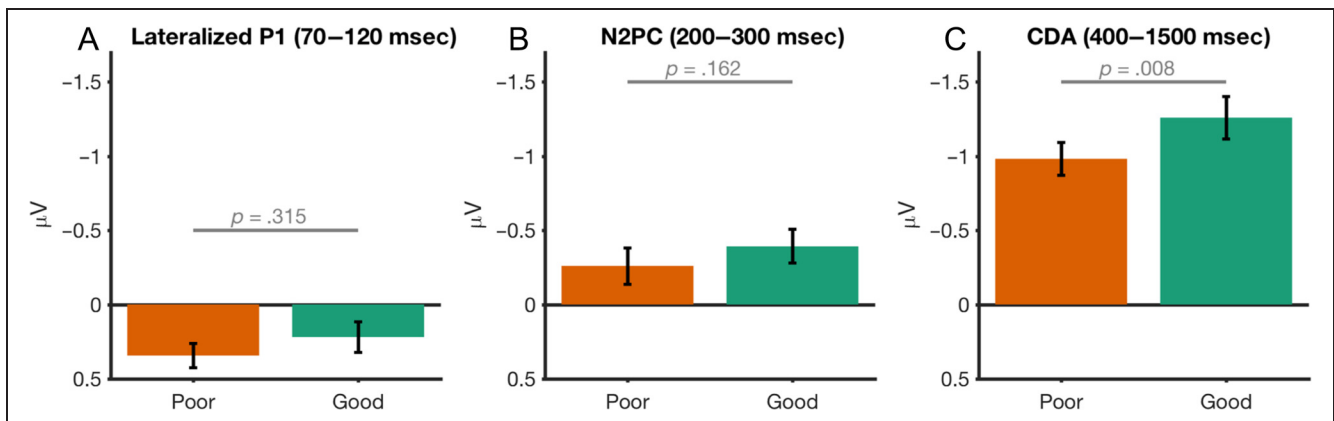


Figure 5. Bar graph of the lateralized ERP components in Experiment 2. The lateralized P1 (A) and N2PC (B) components were significantly larger than 0 but did not differ as a function of trial accuracy, indicating that participants attended the correct side of the memory array even during poor performance trials. On the other hand, CDA amplitude (C) was smaller for poor performance trials, indicating that participants successfully maintained fewer items. Error bars represent 1 *SEM*.

sustained spatial attention did not track trial-by-trial fluctuations in working memory performance (Figure 4B). Although participants successfully maintained spatial attention to the cued side, as indicated by significant overall lateralization of alpha power, this lateralization was not different for poor and good performance trials in either the pretrial period or the delay period. To determine this, we again ran a repeated-measures ANOVA with the factors Hemifield (contralateral vs. ipsilateral) and Performance (high or low) for both pretrial lateralization of alpha power (−1100 to 0 msec) and delay period lateralization of alpha power (400–1500 msec). We found a significant main effect of Hemifield in both the pretrial period, $F(1, 37) = 20.2, p < .001, \eta_p^2 = .35$, and the delay period, $F(1, 37) = 13.8, p = .001, \eta_p^2 = .27$. However, there was no interaction between Hemifield and Performance, indicating a similar degree of lateralization for both good and poor performance trials during both the cue period, $F(1, 37) = 1.73, p = .20, \eta_p^2 = .05$, and the delay period, $F(1, 37) = .64, p = .43, \eta_p^2 = .02$.

We next examined whether global alpha suppression, another proposed marker of working memory storage, tracked working memory performance. We looked at both pretrial cue period activity and retention interval activity, all baselined to before the cue (−1500 msec to −1100 msec). There was no main effect of Performance on global alpha power during the pretrial cue period, $F(1, 37) \leq .001, p = .99, \eta_p^2 \leq .001$, but a trending effect during the delay period, $F(1, 37) = 3.71, p = .06, \eta_p^2 = .09$, in the predicted direction (greater global alpha suppression for greater number of items remembered). Because there were no significant pretrial effects, we rebaselined to eliminate noise during the long cue period; by baselining far in advance of the memory array, we may have introduced more noise to estimates of memory array-related activity. With a clean baseline closer to the memory period (−1500 to −100 msec), we found a significant effect of performance on global alpha power during the retention interval, $F(1, 37) = 9.22, p = .004, \eta_p^2 = .84$, but no interaction between lateralization and performance, $p = .85$. These results are consistent with prior work emphasizing the dissociation between global and lateralized measures of working memory performance (Fukuda et al., 2016). Global alpha power during the delay period is thought to track working memory load, with lower alpha power corresponding to higher memory load. Consistent with our CDA measure of decreased storage during poor performance trials, global alpha power was higher for poor performance trials. On the other hand, lateralized alpha power suppression is thought to track the allocation of sustained spatial attention, and this separate aspect of alpha power did not track fluctuations in working memory success.

Finally, we examined whether fluctuations in executive control may underlie fluctuations in working memory storage. Previous work (Adam et al., 2015) found that pretrial frontal theta power predicted trial-by-trial

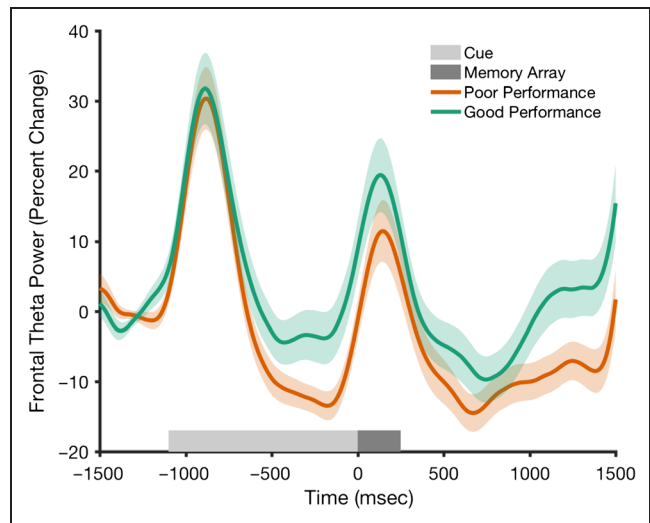


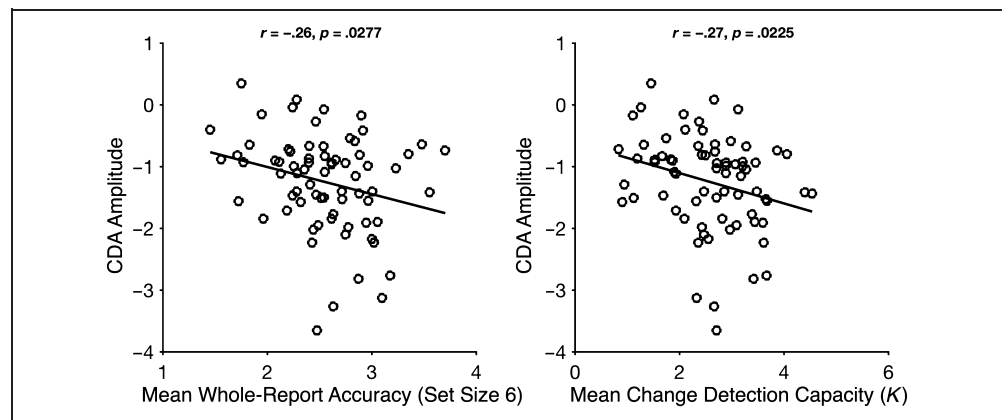
Figure 6. Frontal theta power as a function of trial accuracy in Experiment 2. Frontal theta power tracked trial-by-trial fluctuations in working memory performance, both during the pretrial period (−500 to 0 msec) and during the retention interval (400–1500 msec). Shaded error bars represent 1 SEM.

working memory performance. Here, we replicated the finding that frontal theta power 500 to 100 msec before stimulus onset predicted working memory performance, $t(37) = -2.94, p = .006, 95\% \text{ CI} [-14.28, -2.63]$, and this difference persisted during the memory delay period (400–1500 msec), $t(37) = -3.07, p = .004, 95\% \text{ CI} [-12.88, -2.64]$ (Figure 6). Thus, even before the memoranda had been presented for encoding, this frontal theta power differentiated poor trials from good trials.

Across Experiments: CDA Predicts Individual Differences in Working Memory Performance

In addition to between-subject effects, we replicated the finding that overall CDA amplitude for set size 6 predicted individual differences in working memory performance (Figure 7). CDA amplitude predicted whole-report performance during EEG acquisition ($r = -.26, p = .028$) as well as performance on a separate color change detection task ($r = -.27, p = .023$). The magnitude of these effects is relatively small, but consistent with previously observed effects in the literature (Unsworth, Fukuda, Awh, & Vogel, 2015). Also note that the correlation between behavior and a CDA amplitude for a single set size appears to be smaller than the difference in CDA amplitude between set sizes (Luria et al., 2016). On the basis of previous findings by Unsworth and colleagues (2015), the expected correlation strength between color change detection performance and set size 6 CDA amplitude is $-.33$. With 72 participants, we should have had been able to detect this expected effect with power ($1 - \beta$) of .90 (calculated using G*Power 3.1; Faul, Erdfelder, Buchner, & Lang, 2009).

Figure 7. Individual differences in set size 6 CDA amplitude indexed behavior. Set size 6 CDA amplitude correlated both with performance on set size 6 trials during the EEG recording (left) as well as for a separate change detection estimate of capacity (right). The correlation includes all participants from Experiments 1 and 2 with at least 100 set size 6 trials ($n = 72$).



DISCUSSION

Failures of attention are ubiquitous and can have profound consequences on nearly every aspect of cognition (Unsworth & Robison, 2016; Esterman, Rosenberg, & Noonan, 2014; Unsworth & McMillan, 2014a, 2014b; Reason, 1984). Although it is clear that attentional fluctuations impact many behavioral outcomes, we still have relatively poor understanding of the specific cognitive processes they disrupt. Here, we examined which subprocesses of working memory performance were disrupted during performance failures (i.e., trials where the participant performs poorly on the working memory task).

First, we found that the CDA, a neural measure of ongoing working memory maintenance, was sensitive to fluctuations in performance. In addition to tracking within-subject fluctuations in performance, the CDA also significantly predicted individual differences in working memory capacity. This is consistent with previous work (Unsworth et al., 2015) and also hints at an underlying explanation for the correlation between raw CDA amplitude and working memory capacity. Namely, individual differences in CDA amplitude may be related to individual differences in the consistency of storage. In this view, individuals with larger CDA amplitude more consistently fill their capacity, whereas individuals with smaller CDA amplitude more frequently have storage failures.

Similar to our CDA results, McCollough, Machizawa, and Vogel (2007) found that change detection error trials had smaller amplitude CDA than correct change detection trials, though they lacked a fine-grained behavioral measure of precisely how much information the participant could recall on incorrect trials. For example, because change detection probes only one location, an “incorrect” trial could represent many different levels of task performance: from trials where participants stored nothing at all to trials where participants performed well (e.g., four items correct) but were probed on an item they did not store. By instead having participants report all items in the array, we had better resolution to distinguish between these very different cognitive states.

Together, our key CDA results suggest that participants experience storage failures during poor performance trials and that failures cannot be fully explained by retrieval failures or interference at test. Because we did not examine a proposed neural measure of retrieval failures, we cannot dismiss the possibility that retrieval failures and misbinding (Bays, Catalao, & Husain, 2009) may explain an additional portion of working memory failures. Future work will be needed to determine the relative contribution of each of these processes.

We found that some aspects of task performance were impaired (e.g., executive control), but others were preserved (e.g., covert spatial attention). Before the onset of the memoranda, decreased frontal theta power predicted poor performance. In the context of our whole-report task, elevated frontal theta power before stimulus onset suggests that participants proactively increased executive control to better deal with the challenging task demands of individuating and storing items. This finding replicated our own previous work (Adam et al., 2015) and is also in line with other areas of the literature suggesting that frontal theta power is a marker of executive control (Cavanagh & Frank, 2014; Scheeringa et al., 2008) and is implicated in the success of both working and episodic memory (Hsieh & Ranganath, 2014; Itthipuripat, Wessel, & Aron, 2013). However, other key aspects of task performance did not differentiate between poor and good performance trials. Participants continued to correctly orient to the cued side of the display (lateralized P1), and they sustained covert attention to the remembered side throughout the entire memory delay period (lateralized alpha power suppression). Similarly, previous work has found that separable aspects of attentional control predict working memory performance. For example, Unsworth and Robison (2016) found that mind-wandering frequency and filtering ability both predicted individual differences in working memory capacity, yet mind-wandering and filtering were dissociable predictors.

Lateralized alpha power produced seemingly inconsistent results across the two experiments, and we would like to briefly discuss this particular result. In Experiment 1, we found less lateralization of alpha power for set size

1 relative to the other set sizes (3 and 6), suggesting that this signal might covary with working memory load. However in Experiment 2, there was no difference in alpha power lateralization for high versus low working memory performance trials, despite differences in two separate markers of working memory load (CDA and global alpha power). Thus, the effects of working memory load on alpha power lateralization were ambiguous. Indeed, load-dependent effects on lateralization of alpha power have only been inconsistently observed in the literature (Sauseng et al., 2009, vs. Fukuda et al., 2016). Furthermore, in this study we cannot rule out a confounding factor. In the only condition where we observed decreased alpha power lateralization (set size 1), there was no need to bind the color information to the space information. As such, a decrease in alpha power lateralization may be limited to this particular case rather than representing a true load-dependent signal. Future work is needed to establish the consistency and reliability of load-dependent alpha lateralization effects.

The present results replicate key features of previous work (Adam et al., 2015) but also provide novel insights into the mechanisms underlying failures of working memory performance. We replicated the findings that pretrial theta power was significantly lower during failures trials, that early visual processing (global P1) did not predict working memory failures, and that stimulus-locked global alpha power was less suppressed during failures. In addition to directly replicating past work, the current work sheds new light on cognitive processes contributing to working memory failures. By employing a balanced, lateralized design, we were able to take advantage of well-characterized ERP and oscillatory signals. With this design, we found that participants maintained fewer items during failure trials (indexed by a smaller CDA) despite successfully orienting attention to the cued side (lateralized P1), individuating items (N2PC), and sustaining attention to the cued side (lateralized alpha power). The finding that some cognitive processes are preserved whereas others are impaired suggests potential avenues for behavioral interventions and real-time neural feedback.

Recently, the importance of fluctuations in executive control for working memory performance has also been corroborated by work using pupillometry. It has long been known that pupil dilation tracks working memory load, but recent work has additionally shown that pretrial pupil dilation predicts fluctuations in working memory success (Unsworth & Robison, 2015). Unsworth and Robison found that error trials were preceded by smaller pupil dilation relative to accurate trials. Individual differences also covaried with pupil dilation; individuals with lower working memory capacity had more variable pupil dilation during the pretrial period, indicating that they less consistently maintained high levels of executive control throughout the task. In summary, these pupillometry results corroborate an account whereby shifts in atten-

tional state (and perhaps general arousal) impact working memory success. Our findings are consistent with these pupillometry results but offer better temporal resolution and insight into specific processes that are disrupted during working memory failures.

Individual differences in working memory capacity are reliable (Xu, Adam, Fang, & Vogel, 2017; Beckmann, Holling, & Kuhn, 2007; Klein & Fiss, 1999) and predict important higher-order cognitive abilities like fluid intelligence. As such, better understanding individual differences in capacity has been a long-standing goal of working memory research. Our work makes a key advance toward this goal. Individual differences in working memory capacity are typically conceptualized as individual differences in the ceiling of working memory performance (i.e., the largest array that may be perfectly stored), but our work suggests that the consistency of working memory performance is the key defining feature of individual differences. This view is corroborated by previous work and models of individual differences in working memory capacity. First, participants with low working memory capacity have particular deficits in excluding irrelevant information from working memory (Awh & Vogel, 2008; Vogel et al., 2005). Second, there is a strong relationship between working memory and attentional control (e.g., Unsworth, Fukuda, Awh, & Vogel, 2014). Indeed, previous models of working memory have proposed that variation in working memory performance is largely due to variation in executive control (e.g., Kane et al., 2008; Engle & Kane, 2004; Kane & Engle, 2002). Our findings support these models of individual differences and suggest that the consistency of executive control is key for working memory success.

Acknowledgments

We thank Richard Matullo, Will McGuirk, and Zhilong (Joshua) Wu for assistance with data collection. Research was supported by grants awarded to E. V. (National Institutes of Health grant 5R01-MH087214-08 and Office of Naval Research grant N00014-12-1-0972). Data sets for all experiments are available online on Open Science Framework at <https://osf.io/8xuk3/>. K. A. performed analyses and drafted the manuscript. K. A. and M. R. collected data. All authors planned experiments and revised the manuscript.

Reprint requests should be sent to Kirsten C. S. Adam, Department of Psychology, University of Chicago, 940 E 57th St, Chicago, IL 60637, or via e-mail: kadam1@uchicago.edu.

Notes

1. We also checked that the Experiment 2 accuracy effects survived when only participants with at least 75 trials per accuracy condition were included (the threshold for Experiment 1 exclusion). Although this resulted in a smaller number of participants (remaining $n = 21$), the pattern of results was the same for all reported effects.
2. Greenhouse–Geisser corrected values are reported whenever the assumption of sphericity is violated.

REFERENCES

- Adam, K. C. S., Mance, I., Fukuda, K., & Vogel, E. K. (2015). The contribution of attentional lapses to individual differences in visual working memory capacity. *Journal of Cognitive Neuroscience*, *27*, 1601–1616.
- Aly, M., & Turk-Browne, N. B. (2016a). Attention promotes episodic encoding by stabilizing hippocampal representations. *Proceedings of the National Academy of Sciences, U.S.A.*, *113*, E420–E429.
- Aly, M., & Turk-Browne, N. B. (2016b). Attention stabilizes representations in the human hippocampus. *Cerebral Cortex*, *26*, 783–796.
- Awh, E., & Vogel, E. K. (2008). The bouncer in the brain. *Nature Neuroscience*, *11*, 5–6.
- Bays, P. M., Catalao, R. F. G., & Husain, M. (2009). The precision of visual working memory is set by allocation of a shared resource. *Journal of Vision*, *9*, 7.
- Beckmann, B., Holling, H., & Kuhn, J.-T. (2007). Reliability of verbal–numerical working memory tasks. *Personality and Individual Differences*, *43*, 703–714.
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*, 433–436.
- Cavanagh, J. F., & Frank, M. J. (2014). Frontal theta as a mechanism for cognitive control. *Trends in Cognitive Sciences*, *18*, 414–421.
- Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, *24*, 87–114.
- deBettencourt, M. T., Norman, K. A., & Turk-Browne, N. B. (2017). Forgetting from lapses of sustained attention. *Psychonomic Bulletin & Review*. doi:10.3758/s13423-017-1309-5.
- Delorme, A., & Makeig, S. (2004). EEGLAB: An open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of Neuroscience Methods*, *134*, 9–21.
- Engle, R. W., & Kane, M. J. (2004). Executive attention, working memory capacity, and a two-factor theory of cognitive control. In B. Ross (Ed.), *Psychology of learning and motivation: Advances in research and theory* (Vol. 44, pp. 145–199). New York: Elsevier.
- Esterman, M., Noonan, S. K., Rosenberg, M., & DeGutis, J. (2013). In the zone or zoning out? Tracking behavioral and neural fluctuations during sustained attention. *Cerebral Cortex*, *23*, 2712–2723.
- Esterman, M., Rosenberg, M. D., & Noonan, S. K. (2014). Intrinsic fluctuations in sustained attention and distractor processing. *Journal of Neuroscience*, *34*, 1724–1730.
- Faul, F., Erdfelder, E., Buchner, A., & Lang, A.-G. (2009). Statistical power analyses using G*Power 3.1: Tests for correlation and regression analyses. *Behavior Research Methods*, *41*, 1149–1160.
- Foster, J. J., Sutterer, D. W., Serences, J. T., Vogel, E. K., & Awh, E. (2017). Alpha-band oscillations enable spatially and temporally resolved tracking of covert spatial attention. *Psychological Science*, *28*, 929–941.
- Fukuda, K., Kang, M.-S., & Woodman, G. F. (2016). Distinct neural mechanisms for spatially lateralized and spatially global visual working memory representations. *Journal of Neurophysiology*, *116*, 1715–1727.
- Fukuda, K., Mance, I., & Vogel, E. K. (2015). Alpha power modulation and event-related slow wave provide dissociable correlates of visual working memory. *Journal of Neuroscience*, *35*, 14009–14016.
- Harlow, I. M., & Donaldson, D. I. (2013). Source accuracy data reveal the thresholded nature of human episodic memory. *Psychonomic Bulletin & Review*, *20*, 318–325.
- Hsieh, L.-T., & Ranganath, C. (2014). Frontal midline theta oscillations during working memory maintenance and episodic encoding and retrieval. *Neuroimage*, *85 Pt 2*, 721–729.
- Huang, L. (2010). Visual working memory is better characterized as a distributed resource rather than discrete slots. *Journal of Vision*, *10*, 8.
- Itthipuripat, S., Wessel, J. R., & Aron, A. R. (2013). Frontal theta is a signature of successful working memory manipulation. *Experimental Brain Research*, *224*, 255–262.
- Kane, M. J., Conway, A. R. A., Hambrick, D. Z., & Engle, R. W. (2008). Variation in working memory capacity as variation in executive attention and control. In A. Conway, C. Jarrold, M. Kane, A. Miyake, & J. Towse (Eds.), *Variation in working memory* (pp. 21–48). New York: Oxford University Press.
- Kane, M. J., & Engle, R. W. (2002). The role of prefrontal cortex in working-memory capacity, executive attention, and general fluid intelligence: An individual-differences perspective. *Psychonomic Bulletin & Review*, *9*, 637–671.
- Kane, M. J., Gross, G. M., Chun, C. A., Smekens, B. A., Meier, M. E., Silvia, P. J., et al. (2017). For whom the mind wanders, and when, varies across laboratory and daily-life settings. *Psychological Science*, *28*, 1271–1289.
- Klein, K., & Fiss, W. H. (1999). The reliability and stability of the Turner and Engle working memory task. *Behavior Research Methods, Instruments, & Computers*, *31*, 429–432.
- Luck, S. J., & Hillyard, S. A. (1990). Electrophysiological evidence for parallel and serial processing during visual search. *Perception & Psychophysics*, *48*, 603–617.
- Luck, S. J., & Hillyard, S. A. (1994). Spatial filtering during visual search: Evidence from human electrophysiology. *Journal of Experimental Psychology: Human Perception and Performance*, *20*, 1000–1014.
- Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, *390*, 279–281.
- Luria, R., Balaban, H., Awh, E., & Vogel, E. K. (2016). The contralateral delay activity as a neural measure of visual working memory. *Neuroscience and Biobehavioral Reviews*, *62*, 100–108.
- Mangun, G. R. (1995). Neural mechanisms of visual selective attention. *Psychophysiology*, *32*, 4–18.
- McCollough, A. W., Machizawa, M. G., & Vogel, E. K. (2007). Electrophysiological measures of maintaining representations in visual working memory. *Cortex*, *43*, 77–94.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*, 437–442.
- Reason, J. T. (1984). Lapses of attention in everyday life. In R. Parasuraman & D. R. Davies (Eds.), *Varieties of attention*. Orlando, FL: Academic Press.
- Rouder, J. N., Morey, R. D., Morey, C. C., & Cowan, N. (2011). How to measure working memory capacity in the change detection paradigm. *Psychonomic Bulletin & Review*, *18*, 324–330.
- Sauseng, P., Klimesch, W., Heise, K. F., Gruber, W. R., Holz, E., Karim, A. A., et al. (2009). Brain oscillatory substrates of visual short-term memory capacity. *Current Biology*, *19*, 1846–1852.
- Sauseng, P., Klimesch, W., Stadler, W., Schabus, M., Doppelmayr, M., Hanslmayr, S., et al. (2005). A shift of visual spatial attention is selectively associated with human EEG alpha activity. *European Journal of Neuroscience*, *22*, 2917–2926.
- Scheeringa, R., Bastiaansen, M. C. M., Petersson, K. M., Oostenveld, R., Norris, D. G., & Hagoort, P. (2008). Frontal theta EEG activity correlates negatively with the default mode network in resting state. *International Journal of Psychophysiology*, *67*, 242–251.

- Souza, A. S., & Oberauer, K. (2017). The contributions of visual and central attention to visual working memory. *Attention, Perception, & Psychophysics*, *79*, 1897–1916.
- Souza, A. S., Rerko, L., & Oberauer, K. (2016). Getting more from visual working memory: Retro-cues enhance retrieval and protect from visual interference. *Journal of Experimental Psychology: Human Perception and Performance*, *42*, 890–910.
- Thut, G., Nietzel, A., Brandt, S. A., & Pascual-Leone, A. (2006). Alpha-band electroencephalographic activity over occipital cortex indexes visuospatial attention bias and predicts visual target detection. *Journal of Neuroscience*, *26*, 9494–9502.
- Unsworth, N., Fukuda, K., Awh, E., & Vogel, E. K. (2014). Working memory and fluid intelligence: Capacity, attention control, and secondary memory retrieval. *Cognitive Psychology*, *71*, 1–26.
- Unsworth, N., Fukuda, K., Awh, E., & Vogel, E. K. (2015). Working memory delay activity predicts individual differences in cognitive abilities. *Journal of Cognitive Neuroscience*, *27*, 853–865.
- Unsworth, N., & McMillan, B. D. (2014a). Fluctuations in pre-trial attentional state and their influence on goal neglect. *Consciousness and Cognition*, *26*, 90–96.
- Unsworth, N., & McMillan, B. D. (2014b). Trial-to-trial fluctuations in attentional state and their relation to intelligence. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *40*, 882–891.
- Unsworth, N., & Robison, M. K. (2015). Individual differences in the allocation of attention to items in working memory: Evidence from pupillometry. *Psychonomic Bulletin & Review*, *22*, 757–765.
- Unsworth, N., & Robison, M. K. (2016). The influence of lapses of attention on working memory capacity. *Memory & Cognition*, *44*, 188–196.
- Van Voorhis, S., & Hillyard, S. A. (1977). Visual evoked potentials and selective attention to points in space. *Perception & Psychophysics*, *22*, 54–62.
- Vogel, E. K., & Machizawa, M. G. (2004). Neural activity predicts individual differences in visual working memory capacity. *Nature*, *428*, 748–751.
- Vogel, E. K., McCollough, A. W., & Machizawa, M. G. (2005). Neural measures reveal individual differences in controlling access to working memory. *Nature*, *438*, 500–503.
- Worden, M. S., Foxe, J. J., Wang, N., & Simpson, G. V. (2000). Anticipatory biasing of visuospatial attention indexed by retinotopically specific alpha-band electroencephalography increases over occipital cortex. *Journal of Neuroscience*, *20*, RC63.
- Xu, Z., Adam, K. C. S., Fang, X., & Vogel, E. K. (2017). The reliability and stability of visual working memory capacity. *Behavior Research Methods*. doi:10.3758/s13428-017-0886-6.