Real-time triggering reveals concurrent lapses of attention and working memory

Megan T. deBettencourt [1,2]*, Paul A. Keene[1], Edward Awh[1,2,3] and Edward K. Vogel[1,2,3]

Attention and working memory are clearly intertwined, as shown by co-variations in individual ability and the recruitment of similar neural substrates. Both processes fluctuate over time[1-5], and these fluctuations may be a key determinant of individual variations in ability[6,7]. If these fluctuations are due to the waxing and waning of a common cognitive resource, attention and working memory should co-vary on a moment-to-moment basis. To test this, we developed a hybrid task that interleaved a sustained attention task and a whole-report working memory task. Experiment 1 established that performance fluctuations on these tasks correlated across and within participants: attention lapses led to worse working memory performance. Experiment 2 extended this finding using a real-time triggering procedure that monitored attention fluctuations to probe working memory during optimal (high-attention) or suboptimal (low-attention) moments. In low-attention moments, participants stored fewer items in working memory. Experiment 3 ruled out task-general fluctuations as an explanation for these co-variations by showing that the precision of colour memory was unaffected by variations in attention state. In summary, we demonstrate that attention and working memory lapse together, providing additional evidence for the tight integration of these cognitive processes.

Fluctuations in sustained attention are captured by continuous-performance tasks[1-4]. In these tasks, participants repeatedly press the same response button in the vast majority of trials. Occasionally, participants must inhibit or switch from the prepotent response. Errors on these infrequent trials are prevalent and indicate that attention has lapsed[3,4,8]. Performance on sustained attention to response tasks (SARTs) has revealed substantial and reliable individual differences across the population[6,9]. Studies have used these tasks to track fluctuations in attention through response time (RT): when participants responded more quickly than usual, they were more likely to respond incorrectly to an infrequent trial[3,4]. By repeatedly sampling the RT in a stream of trials, these tasks tracked the intrinsic fluctuations of sustained attention.

A distinct class of tasks has been developed to measure working memory. In these tasks, participants actively maintain information for a brief period of time[7]. Performance is commonly evaluated as working memory capacity (K)—an influential measure that is broadly predictive of intelligence and academic achievement[10,11]. More recently, studies have revealed that trial-to-trial performance variability can, in fact, explain a large proportion of variance in working memory capacity estimates across individuals[5]. These findings suggested that individual differences in working memory capacity may actually reflect differences in the frequency with which people achieve their maximum capacity. In addition, nearly all

computational models of working memory improve when a parameter accounting for trial-to-trial variability in performance is added[12,13]. Recent studies employed whole-report working memory tasks that test memory for each item in a multi-item display to resolve trial-by-trial variations in the number of items stored.

Here, we investigated whether fluctuations in sustained attention and working memory performance are synchronous. Previous work has established that working memory positively correlates with attention control[14,15] and negatively correlates with mind wandering self-reports[16,17]. In addition, participants reported more mind wandering following low-performance working memory trials[18]. However, there remained an historical gulf between sustained attention and memory, as these concepts have been studied with different tasks (SART versus change detection) and discussed with different terminology (vigilance versus capacity). To resolve whether attention and memory indeed fluctuate synchronously, attention state and memory performance needed to be continuously and concurrently assessed. Rather than addressing individual differences across the population, we focused on relating sustained attention and working memory differences within individuals across time.

We hypothesized that fluctuations of attention state would predict performance fluctuations in working memory. This hypothesis was consistent with previous studies that showed a relationship between attention and working memory across individuals. Participants with lower working memory capacities tended to perform worse at tasks that required attention control (for example, antisaccade tasks[10,14]). Across participants, positive latent correlations have been observed between working memory and attention factors[16,18,19]. However, this relationship between individuals does not necessarily presuppose anything about moment-to-moment fluctuations in attention and memory within an individual, as fluctuations of attention could reflect some motoric or visual process unrelated to encoding. In fact, attention and memory could even be anticorrelated in time, as individuals might assign priority to the critical features of one task at the expense of the other.

To test whether attention and memory fluctuate synchronously, we needed a behavioural task that provided an objective measure of both attention and working memory on a moment-to-moment basis. We developed a hybrid design that interleaved two established tasks: a sustained attention task[1,8] and a whole-report working memory task[5,20] (Fig. 1a). On all trials, participants viewed an array of six items—either circles or squares of different colours. The key difference between the sustained attention and working memory tasks was the relevant feature: shape for the sustained attention task and colour for the working memory task. In the sustained attention task, participants pressed a different key if the shapes were squares or circles. We manipulated the probability of each shape (90% circles versus 10% squares), which required participants to make the
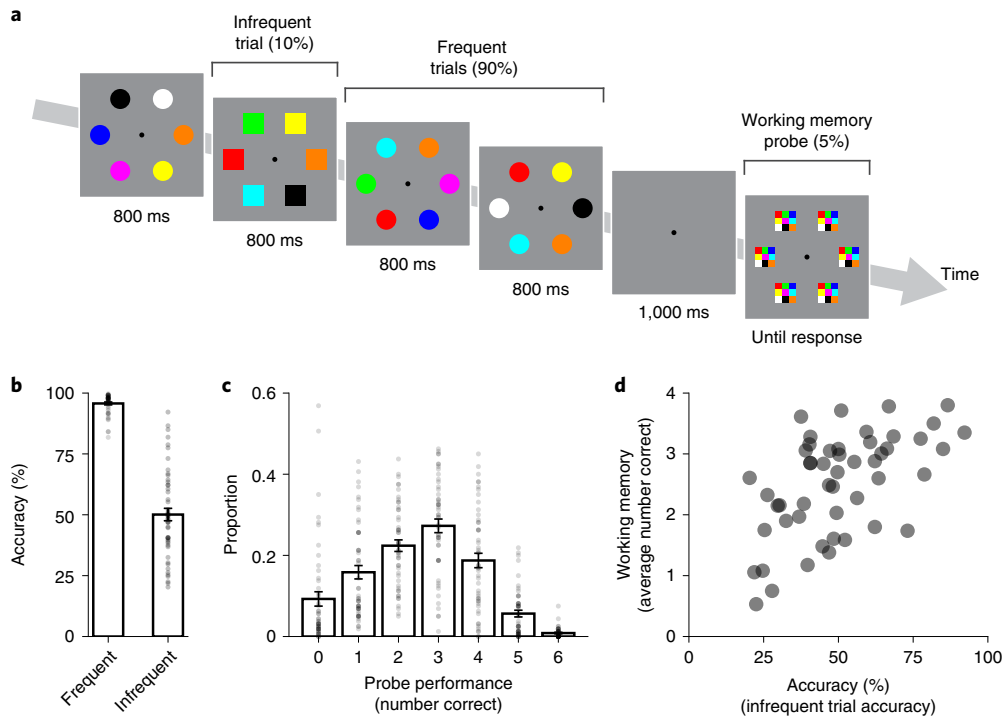
Fig. 1 | Sustained attention relates to working memory performance in an interleaved task. a, Design of the sustained attention and working memory task. On each trial, an array of six items (either circles or squares) of different colours was presented. For the sustained attention task, participants ($n = 50$) responded to the shape. To encourage habitual responding, one of the shapes was much more frequent (90% circles versus 10% squares). For the working memory task, participants clicked the colour of each item after a 1 s delay. In experiment 1a, memory was probed for the infrequent category trials, whereas in experiment 1b, memory was probed for the frequent category trials (as depicted). Participants did not know when working memory probes would appear (5% of total trials). **b**, Sustained attention performance. Accuracy was higher for frequent (96% accuracy) versus infrequent (50% accuracy) trials ($\Delta = 46\%$; $n = 50$; one-tailed $P < 0.001$; $d = 2.69$; 95% CI = 41–50%). The height of each bar depicts the population average. Error bars represent the s.e.m. Each dot depicts the data from each individual in experiments 1a,b. **c**, Working memory performance. In the whole-report working memory task, performance on each trial ranged from 0 (no items correct) to 6 (all items correct). The average number of items correct was above chance ($m = 2.51$; $n = 50$; one-tailed $P < 0.001$; $d = 2.22$; 95% CI = 2.27–2.72). The height of each bar depicts the population average. Error bars represent the s.e.m. for each response. Each dot depicts the proportion of that response for each participant in experiments 1a,b. **d**, Relationship between sustained attention and working memory across participants. Sustained attention accuracy in relation to the infrequent trials was positively correlated with average working memory performance across all participants (Spearman's $r = 0.56$; $n = 50$; $P < 0.001$). Each dot depicts the data from one participant in experiment 1a,b.

same response repeatedly. In the working memory task, participants reported the colour of each item from the most recent array using multicoloured squares at each location. These working memory probes were rare (5%), and participants did not know whether an array would be probed until after it had disappeared.

A key advantage of both tasks is that they provided a relatively continuous assessment of cognitive performance[1,5]. In the sustained attention task, responses were made to the vast majority of trials, including the infrequent category trials. This yielded a more complete assessment of behaviour than if participants had only responded to infrequent targets. In the whole-report working memory task, responses were made for each item on the display. This measured trial-by-trial variability in working memory performance, with enhanced resolution compared with traditional single-probe change detection tasks where the performance on each trial is binary[7]. That is, both of the tasks we used were well suited for tracking moment-to-moment fluctuations in performance while preserving the link to the cognitive constructs of interest. We hypothesized that performance for the shape-based attention task and colour-based working memory task would co-vary across and within individuals, reflecting the shared cognitive resources between attention and memory.

We first examined whether performance on the hybrid task matched characteristic signatures of isolated versions of each task.

In experiment 1a, participants had lower accuracy for infrequent (55% accuracy) than frequent (97% accuracy) trials ($\Delta = 42\%$; $n = 26$; one-tailed $P < 0.001$; Cohen's $d = 2.32$; 95% confidence interval (CI) = 35–49%) in the sustained attention task. We evaluated overall performance in the sustained attention task using a nonparametric measure of sensitivity (A'). Overall performance in the sustained attention task was well above chance ($A' = 0.87$; $n = 26$; one-tailed $P < 0.001$; $d = 15.07$; 95% CI = 0.85–0.89). We also calculated sustained attention performance decrements by calculating the average infrequent trial accuracy for each block. Indeed, the linear slope across the four blocks was reliably negative ($b = -3.99$; $n = 26$; one-tailed $P < 0.001$; $d = 0.84$; 95% CI = −5.87 to −2.28). On half of the infrequent trials, a working memory probe appeared and participants selected the colour of each item from the most recent display. Working memory performance was equal to the number of items correct per display (m), which could range from 0 to 6 (chance = 0.67). On average, participants were well above chance ($m = 3.09$; $n = 26$; one-tailed $P < 0.001$; $d = 5.50$; 95% CI = 2.91–3.24) in the working memory task. These working memory results show—even in this complex hybrid task where colour was most often irrelevant—that participants held colour information in mind.

One potential concern is that the shape still provided some information, as working memory was probed on 50% of the infrequent trials and 0% of the frequent trials. Indeed, infrequent trials themselves
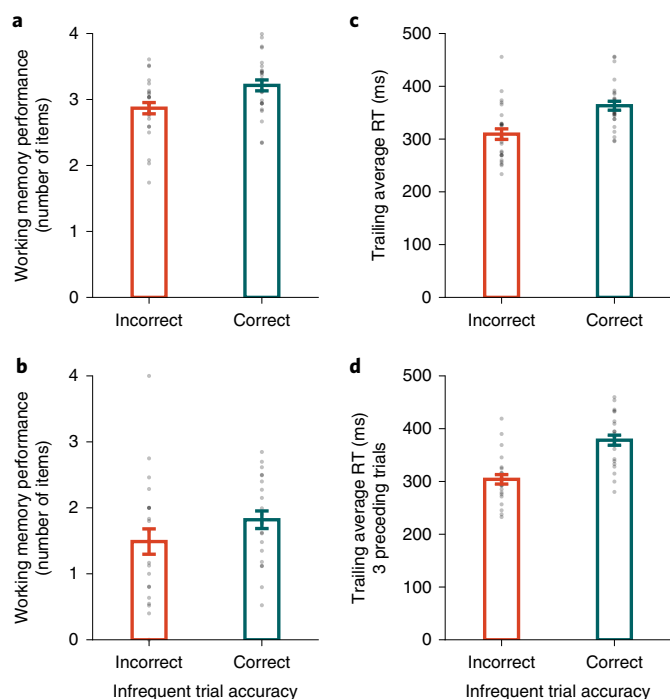
Fig. 2 | Fluctuations of attention predict working memory performance within participants. a, Experiment 1a: attention lapses influence working memory performance. Participants remembered fewer items after an incorrect (orange) versus correct response (teal) in the sustained attention task ($m_{err} = 2.87$; $m_{corr} = 3.21$; $\Delta m = 0.35$; $n = 26$; one-tailed $P < 0.001$; $d = 1.11$; 95% CI = 0.24–0.47). **b**, Experiment 1b: attention lapses influence working memory performance. Participants remembered fewer items after an incorrect (orange) versus correct response (teal) in the sustained attention task ($m_{err} = 1.49$; $m_{corr} = 1.82$; $\Delta m = 0.33$; $n = 22$; one-tailed $P = 0.006$; $d = 0.52$; 95% CI = 0.09–0.61). **c**, Experiment 1a: attention RTs predict attention lapses. Participants made faster responses before an incorrect (orange) versus correct response (teal) in the sustained attention task ($\overline{RT}_{err} = 309$ ms; $\overline{RT}_{corr} = 363$ ms; $\Delta\overline{RT} = 54$ ms; $n = 26$; one-tailed $P < 0.001$; $d = 1.68$; 95% CI = 42–66 ms). **d**, Experiment 1b: attention RTs predict attention lapses. Participants made faster responses before an incorrect response (orange) versus a correct response (teal) in the sustained attention task ($\overline{RT}_{err} = 304$ ms; $\overline{RT}_{corr} = 378$ ms; $\Delta\overline{RT} = 74$ ms; $n = 24$; one-tailed $P < 0.001$; $d = 2.02$; 95% CI = 60–89 ms). In **c** and **d**, the trailing RT was calculated by averaging over the three preceding trials ($i-2$, $i-1$ and i) before each infrequent trial ($i+1$). The height of each bar reflects the population average, and error bars represent the s.e.m. Data from each participant are represented as grey dots.

could influence subsequent attention[21]. To better orthogonalize the sustained attention and working memory tasks, we modified the experimental procedure for experiment 1b and probed memory for frequent trials (that is, circles). That way, participants were less able to anticipate which trials would be probed, as only 5.6% of frequent trials were followed by probes. In experiment 1b, participants still successfully performed both tasks. Overall sustained attention performance was well above chance ($A' = 0.81$; $n = 24$; one-tailed $P < 0.001$; $d = 8.20$; 95% CI = 0.77–0.84), and accuracy was lower for infrequent (45% accuracy) versus frequent (94% accuracy) trials ($\Delta = 50\%$; $n = 24$; one-tailed $P < 0.001$; $d = 3.31$; 95% CI = 43–55%). Also, sustained attention performance declined across the four blocks ($b = -5.12$; $n = 24$; one-tailed $P < 0.001$; $d = 0.83$; 95% CI = −7.59 to −2.73). Average working memory performance was well above chance ($m = 1.88$; $n = 24$; one-tailed $P < 0.001$; $d = 1.80$; 95% CI = 1.60–2.13).

There were substantial differences in performance on each task across individuals (Fig. 1b,c). We hypothesized that greater performance on one task (for example, fewer attention lapses) would correlate with better performance on the other task (for example, higher number of items correct for working memory probes). In contrast, given the dual-task demands, it was also plausible that participants might choose to prioritize one of the tasks, such that performance across tasks was anticorrelated. We found that participants who had fewer attention lapses in the sustained attention task also remembered more items in the working memory task (Spearman's $r = 0.56$; $n = 50$; $P < 0.001$; Fig. 1d). That is, performance was positively correlated across tasks, consistent with our hypothesis. We combined all individuals from experiment 1 to maximize the sample size for the correlation, but correlations were also reliably positive within each experiment (Spearman's $r_{1a} = 0.52$; $n = 26$; $P = 0.007$; Spearman's $r_{1b} = 0.57$; $n = 24$; $P = 0.003$). This relationship was also observed using another index of sustained attention: RT variability. Higher RT variability was anticorrelated with working memory performance (Spearman's $r = -0.51$; $n = 50$; $P < 0.001$; Spearman's $r_{1a} = -0.39$; $n = 26$; $P = 0.047$; Spearman' $r_{1b} = -0.63$; $n = 24$; $P < 0.001$). This relationship between performance on each task corroborates previous work showing that overall attention control correlates with overall working memory capacity[10].

A key feature of the hybrid task is that the attention and working memory tasks are tightly interleaved. This enabled us to explore how moment-to-moment fluctuations of performance on the sustained attention task linked with the fluctuations of performance on the working memory task. Although high-performing participants performed well on both tasks, they could have emphasized colour versus shape at different moments, such that attention and memory would be anticorrelated in time. However, we hypothesized that performance on the two tasks not only co-varied across individuals but also co-varied within individuals across time. Accordingly, we hypothesized that when attention lapsed in the shape task, participants would also remember the colour of fewer items in the working memory task. We examined working memory performance for probe trials separately based on the attention accuracy for that trial. For each participant, we calculated the mean number of items remembered following correct attention responses (m_{corr}) and the mean number of items remembered following incorrect attention responses (m_{err}). Then, we examined whether the difference ($\Delta m = m_{corr} - m_{err}$) was reliably above zero across participants. In experiment 1a, we observed that working memory performance was worse following attention lapses than following non-lapse trials ($m_{err} = 2.87$; $m_{corr} = 3.21$; $\Delta m = 0.35$; $n = 26$; one-tailed $P < 0.001$; $d = 1.11$; 95% CI = 0.24–0.47; Fig. 2a). In experiment 1b, participants made fewer errors responding to the probe trials overall, as probe trials belonged to the frequent category. However, most participants made at least one error responding to a probe trial ($n_{errors} = 10.32$; $n = 22$ of 24; 95% CI = 7.05–14.73). We replicated the finding from experiment 1a that participants remembered fewer items following errors ($m_{err} = 1.49$; $m_{corr} = 1.82$; $\Delta m = 0.33$; $n = 22$; one-tailed $P = 0.006$; $d = 0.52$; 95% CI = 0.09–0.61; Fig. 2b). These results were consistent with the idea that attention and memory co-vary across time together. However, this demonstration did not necessarily allow us to conclude that attention state directly influences memory. An alternative explanation could have been that error-related processing associated with attention lapses deleteriously impacted working memory performance[22,23].

To disentangle attention fluctuations from errors, we turned to RTs. Previous studies have shown that faster RTs directly precede attention lapses (that is, errors in the infrequent trials)[3,4]. We calculated the trailing RT by averaging over the three preceding trials ($\overline{RT}$). Indeed, in experiment 1a, we observed that trailing RTs were faster before an attention lapse versus a non-lapse trial ($\overline{RT}_{err} = 309$ ms; $\overline{RT}_{corr} = 363$ ms; $\Delta\overline{RT} = 54$ ms; $n = 26$; one-tailed $P < 0.001$; $d = 1.68$;

95% CI = 42–66 ms; Fig. 2c). We further replicated this pattern in experiment 1b ($\overline{RT}_{err} = 304$ ms; $\overline{RT}_{corr} = 378$ ms; $\Delta\overline{RT} = 74$ ms; $n = 24$; one-tailed $P < 0.001$; $d = 2.02$; 95% CI = 60–89 ms; Fig. 2d). Moreover, when trailing RTs were sorted and binned within participants, the probability of an attention lapse decreased as the RT increased. This was quantified by a positive linear slope across bins in experiment 1a ($b = 4.36$; $n = 26$; one-tailed $P < 0.001$; $d = 1.99$; 95% CI = 3.49–5.15) and experiment 1b ($b = 5.05$; $n = 24$; one-tailed $P < 0.001$; $d = 2.40$; 95% CI = 4.16–5.81). These analyses suggested that the RT may serve as an index of attention state, to continuously track attention fluctuations irrespective of errors in the sustained attention task.

Although faster RTs preceded attention lapses, and attention lapses correlated with worse working memory performance, decoupling RTs from errors still represented a challenge. For each participant, we computed the Spearman's rank correlation (r) between RTs and working memory performance for correct trials, then examined whether these correlations were reliably positive across participants. In experiment 1a, RTs did not reliably correlate with working memory performance ($r = 0.012$; $n = 26$; one-tailed $P = 0.30$; $d = 0.10$; 95% CI = −0.033–0.055). In experiment 1b, this relationship between trial-by-trial RT and working memory performance was negligible but reliably positive ($r = 0.033$; $n = 24$; $P = 0.021$; $d = 0.42$; 95% CI = 0.001–0.066). We opted to modify the task to incorporate recent advances in real-time triggering methods, which enable the specific targeting of a predefined cognitive or neural state[4]. These techniques use the cognitive state of interest (for example, attention) as the dependent measure to sample rare—but potentially influential—moments. By continuously monitoring attention fluctuations online via the RT, we could trigger an experimental event (for example, the appearance of a working memory probe) at desirable moments. That is, we could probe memory when we detected that the attention state was exceptionally high or low. Triggering working memory probes could potentially increase the power for evaluating the consequences of an extremely low attention state while avoiding an undue influence of errors (and error-related processing) on memory. Notably, the real-time triggering procedure was not designed to maximize working memory performance, but rather to focus on the predictive role of the RT while specifically controlling for other potential explanatory variables (for example, errors).

In experiment 2, we adopted a real-time triggering procedure to more directly link fluctuations of attention and memory (Fig. 3a). We inserted memory probes contingent on attention state (operationalized as moment-to-moment fluctuations of RT). Given the findings from experiments 1a and 1b, we hypothesized that faster responses would index low attention states and therefore predict worse working memory performance. For each participant, we individually tailored and dynamically updated what was considered fast (or slow), based on their cumulative mean RT (μ) and standard deviation (σ). In addition, we calculated a trailing window average of RTs over the three most recent trials ($\overline{RT}$). The moment that this measure exceeded certain predetermined thresholds, we inserted a memory probe for the current trial (i). That is, if participants were responding especially fast, ($\overline{RT} < \mu - \sigma$) or especially slow ($\overline{RT} > \mu + \sigma$), we triggered a memory probe. We have illustrated this procedure using a small selection of trials from a representative participant (Fig. 3b). We hypothesized that memory performance would be worse following probes triggered due to fast (versus slow) RTs.

As an initial validation of this procedure, we wanted to ensure that we successfully targeted low (fast) or high (slow) attentive moments. In experiment 2a, we identified the fast memory probes and slow memory probes. The number of probes did not reliably differ between fast versus slow conditions ($N_{fast} = 69.42$; $N_{slow} = 69.04$; $\Delta = 0.38$; $n = 26$; two-tailed $P = 0.52$; $d = 0.13$; 95% CI = −0.77–1.46). Next, to evaluate whether the procedure was successful at identifying meaningful variation in RT, we examined
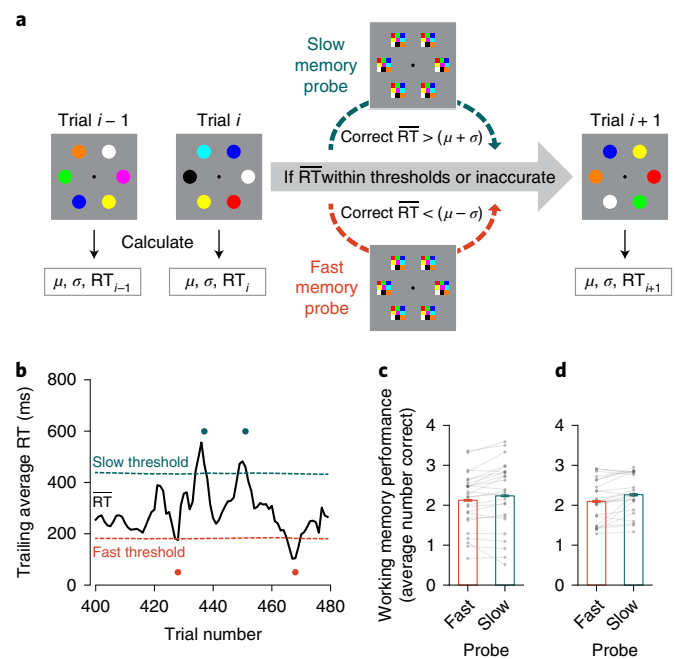
Fig. 3 | Real-time triggering of working memory probes. a, Real-time triggering procedure in experiment 2. For each trial, the attention state was computed using a trailing window of recent RTs ($\overline{RT}$) over the three preceding trials for each participant ($n_{2a} = 26$; $n_{2b} = 23$). Thresholds were individually tailored and dynamically updated, using the cumulative mean (μ) and s.d. (σ) over all preceding trials (1, 2, ..., i). If the trailing RT ($\overline{RT}$) was faster than the fast threshold ($\mu - \sigma$, orange) or slower than the slow threshold ($\mu + \sigma$, teal), a memory probe appeared. For experiment 2b, participants had to respond accurately for all preceding trials, as well as the to-be-probed trial (correct $\overline{RT}$). Otherwise, the trial procedure continued without a memory probe. **b**, Real-time data on a small selection of trials from a representative participant in experiment 2b. The trailing RT ($\overline{RT}$, black line) indicates the average RT of the three preceding trials. The orange dashed line shows the fast threshold ($\mu - \sigma$), and the teal dashed line shows the slow threshold ($\mu + \sigma$). At each trial, if $\overline{RT}$ exceeded either threshold, memory was probed. Dots correspond to memory probes (fast, orange; slow, teal). **c**, In experiment 2a, participants correctly remembered fewer items after fast (orange) versus slow (teal) memory probes ($m_{fast} = 2.12$; $m_{slow} = 2.23$; $\Delta m = 0.11$; $n = 26$; one-tailed $P = 0.007$; $d = 0.49$; 95% CI = 0.02–0.19). **d**, In experiment 2b, participants correctly remembered fewer items after fast (orange) versus slow (teal) memory probes ($m_{fast} = 2.10$; $m_{slow} = 2.26$; $\Delta m = 0.17$; $n = 23$; one-tailed $P < 0.001$; $d = 0.71$; 95% CI = 0.09–0.28). The height of each bar depicts the population average. Error bars represent the within-subject s.e.m. Data from each participant are overlaid as small grey dots, and data from the same participant are connected with lines.

the RTs that triggered a fast memory probe or slow memory probe. As expected, the RTs before fast versus slow memory probes were reliably different ($\overline{RT}_{fast} = 220$ ms; $\overline{RT}_{slow} = 552$ ms; $\Delta\overline{RT} = 333$ ms; $n = 26$; one-tailed $P < 0.001$; $d = 3.85$; 95% CI = 301–365 ms). This analysis validated that real-time triggering was successful at identifying significant deviations in response behaviour.

The critical question was whether this putative attention index, $\overline{RT}$, predicted moment-to-moment working memory performance. We hypothesized that memory performance would be worse for fast memory probes, which appeared when participants were responding more quickly and attention was worse. Indeed, the number of items remembered was lower for fast memory probes versus slow memory probes, and the difference between these conditions was

reliable ($m_{fast} = 2.12$; $m_{slow} = 2.23$; $\Delta m = 0.11$; $n = 26$; one-tailed $P = 0.007$; $d = 0.49$; 95% CI = 0.02–0.19; Fig. 3c). Thus, this real-time triggering procedure established a direct relationship between attention fluctuations and working memory performance. Our triggering design shows that it is possible to predict and influence working memory performance, by carefully selecting the moments when memory is tested.

Although the real-time triggering design was intended to isolate the effect of RT on memory, there are other factors that could have had potentially confounding influences. For example, RTs might have been slower because of an infrequent trial that appeared in the trailing window. Indeed, there were more infrequent trials that appeared before slow versus fast memory probes ($N_{slow} = 27.00$, $N_{fast} = 13.81$; $\Delta = 13.19$; $n = 26$; two-tailed $P < 0.001$; $d = 1.18$; 95% CI = 9.12–17.50). In addition, faster RTs could still have led to errors, which then influenced memory. To test whether these factors accounted for the working memory performance difference, we removed the memory probes with errors and infrequent trials and conducted a post-hoc analysis. The effect of RTs on memory was still observed ($m_{fast} = 2.15$; $m_{slow} = 2.25$; $\Delta m = 0.10$; $n = 26$; one-tailed $P = 0.04$; $d = 0.35$; 95% CI = −0.02–0.19).

However, a key advantage of the real-time triggering procedure is that its predictions are prospective and do not rely on post-hoc analyses. Therefore, in experiment 2b, we made specific improvements to the real-time triggering procedure to more concretely demonstrate that the current attention state predicted future working memory performance. The triggering of a memory probe was based entirely on anticipatory pre-probe RTs. In experiment 2a, the trailing responses (RT_{i-2}, RT_{i-1} and RT_i) that triggered the probe included the to-be-probed array (i). In experiment 2b, we calculated $\overline{RT}$ in the same manner, but then used this attention index to probe memory for the subsequent trial ($i+1$). This modification was intended to eliminate trial-specific encoding signatures of RT. In addition, in experiment 2b, we skipped the triggered memory probe if participants made an error in the sustained attention task. Participants had to make correct responses to all pre-probe trials ($i-2$, $i-1$ and i) as well as the to-be-probed array ($i+1$), which were not requirements for experiment 2a. These modifications to the real-time triggering procedure were intended to further rule out the possibility that basic task compliance could explain the co-variation between attention and memory. We also required all frequent trials in the trailing window so any response-switching demands would not influence the RT. Experiment 2b was intended to provide additional corroboration of the results while more directly targeting an anticipatory, error-free attention state.

In experiment 2b, we were still successful at triggering fast and slow memory probes. Although the overall number of working memory probes decreased, there was no reliable difference ($N_{fast} = 55.74$; $N_{slow} = 53.26$; $\Delta = 2.48$; $n = 23$; two-tailed $P = 0.38$; $d = 0.18$; 95% CI = −2.96–8.09). Next, we calculated the trailing average RT ($\overline{RT}$) for probes triggered due to faster responding or slower responding. As expected based on the triggering procedure, we were successful at capturing a large difference in trailing average RTs ($\overline{RT}_{fast} = 235$ ms; $\overline{RT}_{slow} = 519$ ms; $\Delta\overline{RT} = 284$ ms; $n = 23$; two-tailed $P < 0.001$; $d = 5.51$; 95% CI = 265–306 ms). This analysis validated that the new triggering procedure was still successful at identifying significant deviations in responses.

The critical question was whether this attention index, $\overline{RT}$, predicted upcoming working memory performance. Indeed, participants remembered fewer items for fast memory probes versus slow memory probes ($m_{fast} = 2.10$; $m_{slow} = 2.26$; $\Delta m = 0.17$; $n = 23$; one-tailed $P < 0.001$; $d = 0.71$; 95% CI = 0.09–0.28; Fig. 3d). This real-time triggering procedure in experiment 2b replicated and extended our understanding of the relationship between attention fluctuations and memory performance from experiment 2a, while controlling for other factors that influence performance.

Measurements of the sustained attention state, taken from before the encoding display appeared, were able to predict future working memory performance.

We were interested in whether high- or low-performing individuals in experiment 2 were especially susceptible to the triggering manipulation. We used an independent single-probe change detection task at the end of the session to calculate working memory capacity (K). Average capacity estimates ($K = 2.38$; $n = 45$; 95% CI = 2.04–2.69) were consistent with previous studies[5]. As expected, capacity was indeed correlated with performance in the whole-report task (Spearman's $r = 0.38$; $n = 45$; $P = 0.01$). However, we observed no relationship between capacity and performance difference in response to slow versus fast memory probes (Spearman's $r = 0.10$; $n = 45$; $P = 0.50$). This shows that there is a possible advantage of customizing the triggering manipulation to individuals. Everyone suffers from occasional moments of inattention—even the highest-performing individuals. This real-time procedure exploited those inopportune moments whenever they occurred.

One potential explanation is that attention fluctuations were successful at predicting working memory performance because both rely on a general signature that is important but not specific to these cognitive processes (for example, alertness or task engagement). If so, attention fluctuations would predict most, if not all, measures of performance in this dual-task scenario. Alternatively, attention and memory performance might co-vary only when specific cognitive operations are taxed in both tasks. To test this possibility, we designed experiment 3 to closely match experiment 2b. Memory probes were again triggered based on attention fluctuations. The critical modification was that we measured the precision rather than the number of colours maintained on each trial (Fig. 4a). All items in a given display were the same colour drawn from a continuous colour space, and participants responded by clicking along a colour wheel[24]. An advantage of these continuous report tasks is that small but reliable differences in mouse position were detectable. Previous work has suggested that number and precision reflect distinct aspects of memory ability, which could be attributed to the many differences in memory demands between the two tasks. The number of items is specifically linked to attention control[11]. Therefore, we anticipated that the moment-to-moment variations in attention would no longer co-vary with continuous colour memory in experiment 3. Alternatively, if attention fluctuations reflected a more task-general signature, fluctuations of attention would also co-vary with memory performance.

In experiment 3a, we were first interested in how participants performed this modified hybrid task. Overall performance in the sustained attention task ($A' = 0.85$; $n = 22$; one-tailed $P < 0.001$; $d = 8.95$; 95% CI = 0.83–0.87) was above chance, as in the previous experiments. Performance on the continuous colour memory task was quantified as the absolute response error (err). Participants accurately performed the continuous colour memory task ($err = 17.70°$; 95% CI = 15.59–20.23°), although the hybrid task demands may have contributed to worse performance than in previous demonstrations of continuous colour memory[24]. The distribution of response errors (Fig. 4b) was further characterized as a mixture of a uniform distribution ($g = 7.13\%$; 95% CI = 5.30–9.44%) and a circular normal distribution (s.d. = 17.02°; 95% CI = 15.64–18.44°). Across individuals, we did not observe a correlation between attention lapses and average absolute response error (Spearman's $r_{err} = -0.01$; $n = 22$; $P = 0.95$), guessing (Spearman's $r_g = -0.10$; $n = 22$; $P = 0.66$) or precision (Spearman's $r_{s.d.} = 0.08$; $n = 22$; $P = 0.71$). As in experiments 1 and 2, we captured RT fluctuations ($\overline{RT}_{fast} = 209$ ms; $\overline{RT}_{slow} = 468$ ms; $\Delta\overline{RT} = 259$ ms; $n = 22$; one-tailed $P < 0.001$; $d = 3.67$; 95% CI = 231–288 ms). The critical question was whether better attention (slower RT) predicted better memory (lower absolute response error). There was no reliable co-variation with memory in terms of absolute response error ($err_{fast} = 17.34°$; $err_{slow} = 18.18°$; $\Delta err = 0.84$;
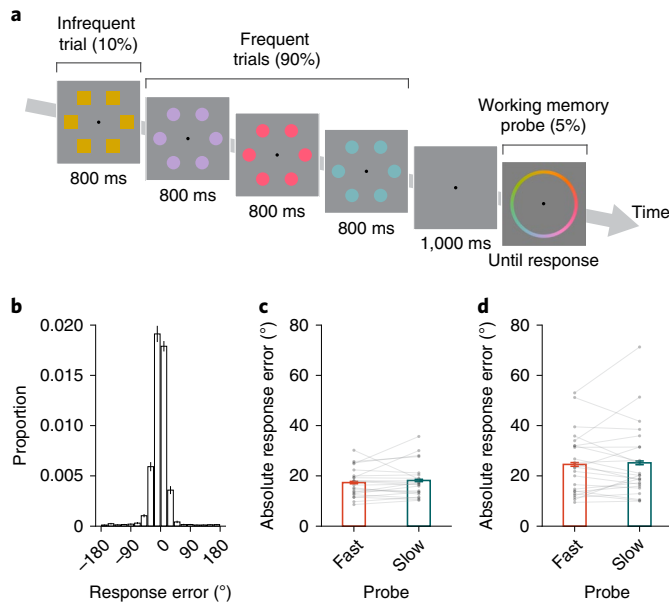
Fig. 4 | Sustained attention and colour memory precision in a continuous report task. a, In experiment 3, participants completed a colour memory precision task. In experiment 3a,b, this colour memory precision was interleaved with a sustained attention task, and the memory probe locations were determined online using a real-time triggering procedure, as in experiment 2b. Each trial was an array of six items (either circles or squares) of the same colour, drawn from a continuous colour space. For working memory probes, participants ($n_{3a} = 22$; $n_{3b} = 24$; $n_{3c} = 23$) selected a colour from a continuous colour wheel. **b**, Distribution of response errors in experiment 3a. The height of each bar in the histogram reflects the mean of the bin across participants (bin width = 20°; $n = 22$). Error bars represent the s.e.m. **c**, In experiment 3a, absolute response error did not reliably differ for fast (orange) versus slow (teal) memory probes ($err_{fast} = 17.34°$; $err_{slow} = 18.18°$; $\Delta err = 0.84°$; $n = 22$; one-tailed $P = 0.82$; $BF_0 = 7.78$; $d = -0.19$; 95% CI = −0.96–2.67°). **d**, In experiment 3b, absolute response error did not reliably differ for fast (orange) versus slow (teal) memory probes ($err_{fast} = 24.57°$; $err_{slow} = 25.19°$; $\Delta err = 0.63°$; $n = 24$; one-tailed $P = 0.66$; $BF_0 = 6.21$; $d = -0.09$; 95% CI = −1.82–3.93°). The height of each bar depicts the population average, and error bars represent the within-subject s.e.m. Data from each participant are overlaid as small grey dots, and data from the same participant are connected with lines.

$n = 22$; one-tailed $P = 0.82$; Bayes factor for the null hypothesis (BF_0) = 7.78; $d = -0.19$; 95% CI = −0.96–2.67°; Fig. 4c). In addition, fast memory probes were not less likely to be remembered ($g_{fast} = 7.07\%$; $g_{slow} = 7.34\%$; $\Delta g = 0.27\%$; $n = 22$; one-tailed $P = 0.60$; $BF_0 = 5.33$; $d = -0.05$; 95% CI = −1.98–2.34%), nor were they remembered less precisely ($s.d._{fast} = 16.98°$; $s.d._{slow} = 17.56°$; $\Delta s.d. = 0.58°$; $n = 22$; one-tailed $P = 0.93$; $BF_0 = 9.96$; $d = -0.31$; 95% CI = −0.20–1.32°).

One possible concern is that this result was due to inadequate task difficulty, perhaps because lingering sensory memory traces could have supported working memory behaviour. In experiment 3b, we therefore reduced the exposure duration and inserted a blank inter-stimulus interval between every trial. Performance on this task was again quantified as average absolute response error ($err = 24.70°$; 20.68–30.46°), guessing ($g = 17.92\%$; 13.19–24.50%) and precision ($s.d. = 15.46°$; 14.27–17.31°). Consistent with the hypothesis that these changes would make the task more difficult, performance was worse in experiment 3b than experiment 3a in terms of response error ($\Delta err = 7.00°$; $n_{3a} = 22$; $n_{3b} = 24$; one-tailed $P = 0.003$; $d = 0.72$; 95% CI = 2.28–13.08) and guessing ($\Delta g = 10.80\%$; $n_{3a} = 22$; $n_{3b} = 24$; one-tailed $P < 0.001$; $d = 1.00$; 95% CI = 5.67–17.69), but not precision ($\Delta s.d. = -1.56°$; $n_{3a} = 22$; $n_{3b} = 24$; one-tailed $P = 0.93$; $d = -0.43$;

95% CI = −3.46–0.68). Higher accuracy in infrequent trials was anticorrelated with absolute response error across participants (Spearman's $r = -0.58$; $n = 24$; $P = 0.003$), unlike what we observed in experiment 3a. However, the critical question was whether attention and memory co-varied across time within participants. There was no reliable memory difference between fast- and slow-triggered memory probes for response error ($err_{fast} = 24.57°$; $err_{slow} = 25.19°$; $\Delta err = 0.63°$; $n = 24$; one-tailed $P = 0.66$; $BF_0 = 6.21$; $d = -0.09$; 95% CI = −1.82–3.93°; Fig. 4d), guessing ($g_{fast} = 17.43\%$; $g_{slow} = 17.71\%$; $\Delta g = 0.28°$; $n = 24$; one-tailed $P = 0.55$; $BF_0 = 5.16$; $d = -0.03$; 95% CI = −3.28–4.50%) or precision ($s.d._{fast} = 15.85°$; $s.d._{slow} = 16.79°$; $\Delta s.d. = 0.94°$; $n = 24$; one-tailed $P = 0.78$; $BF_0 = 7.57$; $d = -0.16$; 95% CI = −0.89–4.24°). The moment-to-moment co-variation between sustained attention and working memory is therefore not present in all tasks, as it would be if attention fluctuations reflected a completely task-general state. Rather, attention and memory co-variations are contingent on the manner in which working memory performance is assessed.

Another way in which we characterized the relationship between the sustained attention and continuous colour memory tasks was through the cognitive impact of the dual tasks. In experiment 3c, participants completed only the continuous report colour memory task without a concurrent sustained attention task. We hypothesized that the dual-task demand of the interleaved sustained attention task deleteriously influenced performance on the colour precision memory task. In contrast, if memory performance was insensitive to the presence of the sustained attention task (perhaps due to inadequate task difficulty), eliminating the dual task would not influence working memory performance. However, consistent with our hypothesis, performance was much better in experiment 3c in terms of the absolute response error ($err = 9.78°$; 9.02–10.80°), guessing ($g = 1.39\%$; 1.09–1.96%) and precision ($s.d. = 11.76°$; 11.04–12.84°). Performance was worse in experiment 3a than experiment 3c in terms of response error ($\Delta err = 7.92°$; $n_{3a} = 22$; $n_{3c} = 23$; one-tailed $P < 0.001$; $d = 1.89$; 95% CI = 5.60–10.55), guessing ($\Delta g = 5.74\%$; $n_{3a} = 22$; $n_{3c} = 23$; one-tailed $P < 0.001$; $d = 1.62$; 95% CI = 3.82–7.90) and precision ($\Delta s.d. = 5.26°$; $n_{3a} = 22$; $n_{3c} = 23$; one-tailed $P < 0.001$; $d = 1.82$; 95% CI = 3.76–7.02). These findings show that the lack of co-variation between tasks was not a consequence of a continuous colour memory task that was insensitive to the presence of the sustained attention task.

Taken together, the results of our three experiments show that attention and memory fluctuate together over time. In experiment 1, we discovered that attention and working memory were related across participants and within participants across time. Experiment 2 more directly linked moment-to-moment fluctuations in attention state with fluctuations in working memory performance using a real-time triggering approach that we developed. By triggering memory probes in real time, we exploited potentially impactful fluctuations of attention state, and decoupled attention fluctuations from other explanatory variables. Experiment 3 was designed to test whether RT reflected a more task-general signal, by examining whether attention fluctuations co-varied with the precision of a single representation in working memory. The same real-time triggering design yielded no co-variation between attention and memory precision fluctuations. In summary, these results support the hypothesis that sustained attention and working memory draw on a common cognitive resource that waxes and wanes.

This online and adaptive triggering design was inspired by behavioural studies that linked attention to long-term recognition memory[4], as well as real-time neuroimaging studies[3]. One advantage of the real-time triggering procedure is that it can prospectively decouple other potential explanatory variables via behaviour. That memory fluctuated even when participants were performing the task correctly is consistent with our previous work, in which we found no difference between early visual responses for trials with

low and high working memory performance[5]. Importantly, we do not mean to suggest that basic task compliance does not influence performance—just that this factor was not key to the observed co-variation between attention and working memory. We propose that the number of items held in working memory depends on a cognitive resource that is also important for over-riding prepotent responses in the sustained attention task[10].

These findings complement previous neural evidence linking attention state to working memory encoding. For example, neural responses to cues[25,26], as well as pre-stimulus oscillatory activity[5,27], are related to working memory performance. Our findings suggest that prestimulus attention fluctuations might be detectable long before a cue or a memory encoding array. Although RT was successful at highlighting good and poor moments, RT by itself is not a perfect measure of attention state. Neural signatures tracking sustained attention fluctuations (for example, multivariate classification[3] or pupillometry[28]) might also similarly or independently co-vary with memory. By combining behavioural and neural indices of attention fluctuations, it might be possible to influence memory performance to a greater extent. Future work could incorporate such neural measures of attention fluctuations to further characterize the links between attention and memory.

Finally, these results provide additional motivation to explore real-time manipulations of attention and memory performance. Here, we have demonstrated through online and adaptive triggers that an entirely anticipatory behavioural measure of attention predicts upcoming working memory. These findings suggest plausible means of enhancing memory performance; for example, by timing memory tests to occur exclusively during optimal times. In addition, triggering approaches could answer neuroscientific questions about the relationship, and potentially decouple other neural and behavioural signals of attention and memory fluctuations (for example, RT variability). The ability to track and predict attention and working memory could be enormously beneficial, especially in situations where lapses are catastrophic. For example, attention could be continuously monitored in educational settings (for example, during a long lecture) or occupational settings (for example, in air traffic control), to select the best moments for memory.

Methods

Participants. A total of 177 people participated across 7 studies for course credit from the University of Chicago or US$20 payment (112 female; mean = 22.1 years). The target sample size a priori was set to approximately 24 for each study, based on previous studies of sustained attention[4]. Because we developed the task specifically to test this hypothesis, no statistical methods were used to predetermine sample sizes before data collection. One participant left the study early without completing it. Eight participants were excluded as their performance in either task exceeded three standard deviations from the population mean in that study. After exclusion, the final sample sizes for each study were: $n = 26$ of 28 for experiment 1a, $n = 24$ of 24 for experiment 1b, $n = 26$ of 28 for experiment 2a, $n = 23$ of 24 for experiment 2b, $n = 22$ of 24 for experiment 3a, $n = 24$ of 24 for experiment 3b and $n = 23$ of 24 for experiment 3c. Participants were allocated to experiments based on when they signed up for the study. We largely conducted within-participant comparisons; therefore, there was no group assignment within experiments. Repeat participation was not prevented, and 17 individuals participated multiple times in experiments 1–3. The individuals who participated in experiments 1a,b were completely non-overlapping. All participants in these experiments reported normal or corrected-to-normal colour vision, and provided informed consent to a protocol approved by the University of Chicago Social and Behavioral Sciences Institutional Review Board.

Apparatus. Participants were seated approximately 88 cm from an LCD monitor (refresh rate = 120 Hz). Stimuli were presented using Python and PsychoPy[29].

Stimuli. Stimuli were shapes—either circles (diameter = 1.5°) or squares (1.5° × 1.5°). Each display comprised 6 shapes at 4° eccentricity (Fig. 1). The shape positions were consistent for all trials to minimize intertrial visual transients[2].

For experiments 1 and 2, each shape was one of nine distinct colours (red, blue, green, yellow, magenta, cyan, white, black or orange) and each display contained six shapes of unique colours. A central black fixation dot (0.1°) appeared at the centre and turned white after a key press. For whole-report

working memory probes, a multicoloured square (1.5° × 1.5°) comprising the 9 colours appeared at each of the 6 locations, and the mouse cursor appeared at the central fixation position.

For experiment 3, the colour of each shape was drawn from a set of 512 colours taken from an International Commission on Illumination L*a*b* space (centred at $L* = 70$, $a* = 20$ and $b* = 38$), and all shapes were the same colour for each display. Each successive display was separated by at least 20°, but the colours were otherwise randomly chosen. For continuous-report working memory probes, a colour wheel (radius 4°) appeared (Fig. 4a), and the mouse cursor appeared at the central fixation position.

Procedure. In the sustained attention task, participants viewed a continuous stream of displays, each of which was an array of squares or circles. Their task was to press keys based on the shapes of the array. If the shapes were squares, they pressed the 's' key, and if the shapes were circles, they pressed the 'd' key. The imbalanced distributions of the shapes across displays were selected from previous SARTs: 90% of the displays were circles and 10% of the displays were squares. The stimuli remained on the screen for 800 ms with no interstimulus interval in all experiments except experiment 3b. In experiment 3b, the exposure duration was shortened (500 ms) and a blank interstimulus interval (300 ms) was introduced between every stimulus. Participants could respond any time before the next stimulus appeared (within 800 ms). Because of the long stimulus exposure durations, it was unlikely that the encoding time seriously constrained performance in the task[30].

In the working memory task, participants were occasionally probed on the colour of all of the items from the most recent display. The length of the retention interval (1 s) eliminated the ability for sensory memory alone to support behaviour. In all experiments, participants used the mouse cursor to select their response. For experiments 1 and 2, participants completed a discrete colour whole-report working memory task. Multicoloured squares that included all nine possible colours appeared at each location. Participants had to select a colour for each item from the previous display before the screen would advance. After making a response, a large black square appeared around the outside of the entire multicoloured square for that item. Participants had to respond to each of the six items. After the last response, the screen went blank again (1 s) before resuming the sustained attention task. For experiment 3, a continuous colour wheel appeared on the screen to probe the precise colour maintained in working memory. After the participants selected a colour along the wheel, the screen went blank again (1 s) before resuming the sustained attention task. The category of the trials selected for memory probes differed across experiments— either the infrequent trials (experiment 1a), the frequent trials (experiment 1b) or both (experiments 2 and 3a,b). In experiment 3c, there was no interleaved sustained attention task, and participants only performed the continuous colour report for the working memory probes.

Participants practised the sustained attention and working memory tasks separately, and then both tasks together, before starting the study. Participants completed multiple blocks with 800 sustained attention trials per block. In most experiments, participants completed four blocks, except in experiment 3b, which had six blocks. The first participant from experiment 1a started but did not complete a fifth block, but only data from the first four blocks were analysed. Each block contained up to 40 working memory probes.

Real-time triggering procedure. In experiments 2a,b and 3a,b, the working memory probes were triggered on the basis of the attention state, operationalized as the speed of responding (Fig. 3). Participants were not informed that their RTs controlled when memory probes would appear.

In experiment 2a, the cumulative mean (μ) and standard deviation (σ) of the RT was calculated for all trials up to that point (1... i). In addition, $\overline{RT}$ was calculated over the three most recent trials ($i - 2$, $i - 1$ and i). A fast memory probe would be triggered when $\overline{RT}$ was faster than the fast threshold: $\overline{RT} < (\mu - \sigma)$. A slow memory probe would be triggered when $\overline{RT}$ was slower than the slow threshold: $\overline{RT} > (\mu + \sigma)$. Memory probes were separated by a minimum distance of three trials.

In experiments 2b and 3a,b, we sought to isolate more anticipatory attention fluctuations, so $\overline{RT}$ was used to trigger the memory probe for the subsequent trial ($i + 1$). That is, the decision on whether a memory probe appeared was unrelated to the RT for that trial (RT_{i+1}). We also required that all three preceding trials ($i - 2$, $i - 1$ and i) were of the frequent category, and that a correct response was made for each of them. Furthermore, to rule out the role of errors, we required that participants respond correctly to the probe trial ($i + 1$). Up to 36 memory probes were inserted on the fly according to fluctuations in response behaviour, split evenly for fast versus slow responding (18 fast and 18 slow). The first 80 trials of a block were used to initalize the estimates of the cumulative mean RT (μ) and standard deviation of RTs (σ) necessary for the real-time triggering procedure. During those trials, we randomly placed four memory probes independent of RT.

Change detection task. After completing the hybrid task, participants in experiments 1a,b, 2a,b and 3a completed 96 trials of a single-probe discrete colour change detection task. We analysed the data from all participants who completed the change detection task (experiment 1a: $n = 16$ of 26; experiment 1b: $n = 22$ of 23;

experiment 2a: $n=24$ of 26; experiment 2b: $n=23$ of 23; experiment 3a: $n=24$ of 24). Some participants did not complete the change detection task (for example, if there was insufficient time) and were therefore not included in this analysis. For each trial, participants viewed an array of 6 coloured squares ($1.5° \times 1.5°$), which appeared anywhere on the screen between 1° and 4° eccentricity, with a minimum distance of 2° between the centroids of the squares. The squares appeared for 500 ms, followed by a 1 s retention interval. Then, one square from the array reappeared as a probe. Participants made an unspeeded response by pressing the '/' key if the colour was the same or the 'z' key if the colour was different. There was an equal probability that the probed square was the same or a different colour. Participants completed 1 block of 96 trials, and working memory capacity was calculated using an established formula for these tasks[31]: $K = N \times (H - FA)$, where N represents the set size, H is the hit rate (probability of correctly identifying a trial where the probe changed colour) and FA is the false alarm rate (probability of incorrectly identifying a trial where the probe did not change colour). In experiment 2, we analysed between-participant correlations with change detection performance. For these analyses, we excluded the two participants from experiment 2a who did not complete the single-probe change detection task and two participants in experiment 2b who had also participated in experiment 2a.

Sustained attention data analysis. Sustained attention performance on each trial was characterized using accuracy (that is, whether the participant made the correct response or an incorrect response). Performance was assessed using hits (correct responses to the frequent category trials) and false alarms (incorrect or no responses to the infrequent category trials). These values were combined into a single non-parametric measure of sensitivity (A′ (for which chance is 0.5). To examine performance decrements over blocks, we calculated the average accuracy in infrequent trials and examined whether there was a reliably negative linear relationship across blocks. For each trial, the RT was calculated. Fluctuations in the trailing RT averaged over the 3 preceding trials ($i-2$, $i-1$ and i) predicted attention accuracy for the subsequent trial ($i+1$). We examined the relationship between trailing RT and accuracy for each participant. First, we examined trailing RT differences between correct versus incorrect trials. We also sorted and binned trailing RTs (11 bins) and calculated the percentage of trials that were responded to correctly within each bin. Then, we examined whether there was a positive linear relationship between bin number (1–11) and accuracy (%). In addition, we calculated another signature of sustained attention abilities: RT variability. RT variability was calculated as the standard deviation of RTs for all correct frequent trials within each block and then averaged across all blocks. Sustained attention performance decrements were calculated by averaging the accuracy to all infrequent trials within each block.

Working memory data analysis. Working memory performance on the whole-report task was summarized as the average number of items per trial for which the participants selected the correct colour. Working memory performance on the continuous report task was assessed using response error—the angular deviation between the selected and original colour. Performance was further quantified by fitting a mixture model to the distribution of response errors for each participant using MemToolbox[32]. We modelled the distribution of response errors as the mixture of a von Mises distribution centred on the correct value and a uniform distribution. We obtained maximum-likelihood estimates for two parameters: (1) the dispersion of the von Mises distribution (s.d.), which reflects response precision; and (2) the height of the uniform distribution (g), which reflects the probability of guessing.

Statistics. Because some of the data violated the assumption of normality, all statistics were computed using a non-parametric, random-effects approach in which participants were resampled with replacement 100,000 times[33]. Null hypothesis testing was performed by calculating the proportion of iterations in which the bootstrapped mean was in the opposite direction. Exact P values are reported (P values smaller than 1 in 1,000 are approximated as $P < 0.001$). Statistical results for directional hypotheses are noted as one tailed and non-directional hypotheses are noted as two tailed. We used an alpha level of 0.05 for statistical tests. The mean and 95% CIs of the bootstrapped distribution are reported as descriptive statistics. Correlations were computed using the non-parametric Spearman's rank-order correlation function included in Scipy. Effect sizes were computed as Cohen's d values in R. For experiment 3, we computed Jeffreys–Zellner–Siow Bayes factors to evaluate the support for the point null (BF_0) versus the directional hypothesis in R using the Bayes Factor library[13]. Parametric statistics are included in the online distribution of data and analyses. Data collection and analysis were not performed blind to the conditions of the experiments.

Reporting Summary. Further information on research design is available in the Nature Research Reporting Summary linked to this article.

Data availability
The data that support the findings of this study are available online in an Open Science Framework repository (https://osf.io/hfeu8/), as well as a GitHub repository (https://github.com/AwhVogelLab/deBettencourt_rtAttnWM).

Code availability
The experimental design was programmed in Python 2.7 using PsychoPy (versions 1.85 and 1.90). All analyses were conducted using custom scripts in Python 3 and R version 3. All codes for running the experiment and regenerating the results are available online in an Open Science Framework repository (https://osf.io/hfeu8/) along with a GitHub repository (https://github.com/AwhVogelLab/deBettencourt_rtAttnWM).

References
1. Robertson, I. H., Manly, T., Andrade, J., Baddeley, B. T. & Yiend, J. 'Oops!': performance correlates of everyday attentional failures in traumatic brain injured and normal subjects. *Neuropsychologia* **35**, 747–758 (1997).
2. Rosenberg, M., Noonan, S., DeGutis, J. & Esterman, M. Sustaining visual attention in the face of distraction: a novel gradual-onset continuous performance task. *Atten. Percept. Psychophys.* **75**, 426–439 (2013).
3. deBettencourt, M. T., Cohen, J. D., Lee, R. F., Norman, K. A. & Turk-Browne, N. B. Closed-loop training of attention with real-time brain imaging. *Nat. Neurosci.* **18**, 470–475 (2015).
4. deBettencourt, M. T., Norman, K. A. & Turk-Browne, N. B. Forgetting from lapses of sustained attention. *Psychon. Bull. Rev.* **25**, 605–611 (2018).
5. Adam, K. C. S., Mance, I., Fukuda, K. & Vogel, E. K. The contribution of attentional lapses to individual differences in visual working memory capacity. *J. Cogn. Neurosci.* **27**, 1601–1616 (2015).
6. Fortenbaugh, F. C. et al. Sustained attention across the life span in a sample of 10,000: dissociating ability and strategy. *Psychol. Sci.* **26**, 1497–1510 (2015).
7. Luck, S. J. & Vogel, E. K. The capacity of visual working memory for features and conjunctions. *Nature* **390**, 279–281 (1997).
8. Esterman, M., Noonan, S. K., Rosenberg, M. & DeGutis, J. In the zone or zoning out? Tracking behavioral and neural fluctuations during sustained attention. *Cereb. Cortex* **23**, 2712–2723 (2013).
9. Esterman, M., Rosenberg, M. D. & Noonan, S. K. Intrinsic fluctuations in sustained attention and distractor processing. *J. Neurosci.* **34**, 1724–1730 (2014).
10. Engle, R. W. Working memory capacity as executive attention. *Curr. Dir. Psychol. Sci.* **11**, 19–23 (2002).
11. Fukuda, K., Vogel, E., Mayr, U. & Awh, E. Quantity, not quality: the relationship between fluid intelligence and working memory capacity. *Psychon. Bull. Rev.* **17**, 673–679 (2010).
12. Van den Berg, R., Awh, E. & Ma, W. J. Factorial comparison of working memory models. *Psychol. Rev.* **121**, 124–149 (2014).
13. Rouder, J. N. et al. An assessment of fixed-capacity models of visual working memory. *Proc. Natl Acad. Sci. USA* **105**, 5975–5979 (2008).
14. Unsworth, N., Schrock, J. C. & Engle, R. W. Working memory capacity and the antisaccade task: individual differences in voluntary saccade control. *J. Exp. Psychol. Learn. Mem. Cogn.* **30**, 1302–1321 (2004).
15. Engle, R. W., Tuholski, S. W., Laughlin, J. E. & Conway, A. R. Working memory, short-term memory, and general fluid intelligence: a latent-variable approach. *J. Exp. Psychol. Gen.* **128**, 309–331 (1999).
16. McVay, J. C. & Kane, M. J. Why does working memory capacity predict variation in reading comprehension? On the influence of mind wandering and executive attention. *J. Exp. Psychol. Gen.* **141**, 302–320 (2012).
17. Kane, M. J. et al. For whom the mind wanders, and when: an experience-sampling study of working memory and executive control in daily life. *Psychol. Sci.* **18**, 614–621 (2007).
18. Unsworth, N. & Robison, M. K. The influence of lapses of attention on working memory capacity. *Mem. Cogn.* **44**, 188–196 (2016).
19. Unsworth, N., Fukuda, K., Awh, E. & Vogel, E. K. Working memory and fluid intelligence: capacity, attention control, and secondary memory retrieval. *Cogn. Psychol.* **71**, 1–26 (2014).
20. Huang, L. Visual working memory is better characterized as a distributed resource rather than discrete slots. *J. Vis.* **10**, 8 (2010).
21. Jonker, T. R., Seli, P., Cheyne, J. A. & Smilek, D. Performance reactivity in a continuous-performance task: implications for understanding post-error behavior. *Conscious. Cogn.* **22**, 1468–1476 (2013).
22. Yeung, N., Botvinick, M. M. & Cohen, J. D. The neural basis of error detection: conflict monitoring and the error-related negativity. *Psychol. Rev.* **111**, 931–959 (2004).
23. Cheyne, J. A., Carriere, J. S. A., Solman, G. J. F. & Smilek, D. Challenge and error: critical events and attention-related errors. *Cognition* **121**, 437–446 (2011).
24. Zhang, W. & Luck, S. J. Discrete fixed-resolution representations in visual working memory. *Nature* **453**, 233–235 (2008).
25. Murray, A. M., Nobre, A. C. & Stokes, M. G. Markers of preparatory attention predict visual short-term memory performance. *Neuropsychologia* **49**, 1458–1465 (2011).

26. Giesbrecht, B., Weissman, D. H., Woldorff, M. G. & Mangun, G. R. Pre-target activity in visual cortex predicts behavioral performance on spatial and feature attention tasks. *Brain Res.* **1080**, 63–72 (2006).

27. Myers, N. E., Stokes, M. G., Walther, L. & Nobre, A. C. Oscillatory brain state predicts variability in working memory. *J. Neurosci.* **34**, 7735–7743 (2014).

28. Robison, M. K. & Unsworth, N. Pupillometry tracks fluctuations in working memory performance. *Atten. Percept. Psychophys.* **81**, 407–419 (2019).

29. Peirce, J. W. PsychoPy—psychophysics software in Python. *J. Neurosci. Methods* **162**, 8–13 (2007).

30. Vogel, E. K., Woodman, G. F. & Luck, S. J. The time course of consolidation in visual working memory. *J. Exp. Psychol. Hum. Percept. Perform.* **32**, 1436–1451 (2006).

31. Cowan, N. The magical number 4 in short-term memory: a reconsideration of mental storage capacity. *Behav. Brain Sci.* **24**, 87–114 (2001).

32. Suchow, J. W., Brady, T. F., Fougnie, D. & Alvarez, G. A. Modeling visual working memory with the MemToolbox. *J. Vis.* **13**, 9–9 (2013).

33. Efron, B. & Tibshirani, R. Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy. *Stat. Sci.* **1**, 54–75 (1986).

Author contributions

M.T.dB., E.A. and E.K.V. conceived of the study and contributed to the study design. M.T.dB. and P.A.K. collected and analysed the data. M.T.dB. wrote an initial draft of the manuscript, which all authors read and edited.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information is available for this paper at https://doi.org/10.1038/s41562-019-0606-6.

Reprints and permissions information is available at www.nature.com/reprints.

Correspondence and requests for materials should be addressed to M.T.dB.

Publisher's note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

nature research

Corresponding author(s): Megan T. deBettencourt

Last updated by author(s): Mar 21, 2019

Reporting Summary

Nature Research wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Research policies, see Authors & Referees and the Editorial Policy Checklist.

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a	Confirmed	
☐	☒	The exact sample size (*n*) for each experimental group/condition, given as a discrete number and unit of measurement
☐	☒	A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
☐	☒	The statistical test(s) used AND whether they are one- or two-sided *Only common tests should be described solely by name; describe more complex techniques in the Methods section.*
☒	☐	A description of all covariates tested
☐	☒	A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
☐	☒	A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
☐	☒	For null hypothesis testing, the test statistic (e.g. *F*, *t*, *r*) with confidence intervals, effect sizes, degrees of freedom and *P* value noted *Give P values as exact values whenever suitable.*
☐	☒	For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
☒	☐	For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
☐	☒	Estimates of effect sizes (e.g. Cohen's *d*, Pearson's *r*), indicating how they were calculated

Our web collection on statistics for biologists contains articles on many of the points above.

Software and code

Policy information about availability of computer code

Data collection	Data was collected in Python 2.7 using Psychopy versions 1.85 and 1.90. Code for running the experiment is posted
Data analysis	Data was analyzed in Python 3.6. All analysis code is made publicly available

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors/reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Research guidelines for submitting code & software for further information.

Data

Policy information about availability of data

All manuscripts must include a data availability statement. This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A list of figures that have associated raw data
- A description of any restrictions on data availability

All code for running the experiment as well as regenerating all of the results from the raw behavioral data files are available online in an Open Science Framework repository [https://osf.io/hfeu8/] as well as a github repository [https://github.com/AwhVogelLab/deBettencourt_attnWM]

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

☐ Life sciences ☒ Behavioural & social sciences ☐ Ecological, evolutionary & environmental sciences

Behavioural & social sciences study design

All studies must disclose on these points even when the disclosure is negative.

Study description	The data are quantitative experimental behavioral data.
Research sample	The research sample was male and female adult humans from Chicago who participated for undergraduate course credit or pay. Those participating for credit were University of Chicago undergraduate students. In total, 177 people participated 112 female, mean 22.1 years).
Sampling strategy	The target sample size a priori of each study was set to approximately 24, based on previous studies of sustained attention (deBettencourt, Norman & Turk-Browne, 2018). Because we developed the paradigm specifically to test this hypothesis, no statistical methods were used to predetermine sample sizes prior to data collection.
Data collection	Data were collected on a PC running Windows 7 using Psychopy running on Python 2.7 version 1.85.2 (Experiments 1&2) and version 1.90.3 (Experiment 3). Participants were seated alone at a computer in a behavioral running room. The researcher was aware of the study hypotheses during data collection.
Timing	Data collection began February 26th, 2018. Data collection terminated January 16th, 2019 Experiment 1a: February 25 – April 12 Experiment 1b: April 12 – May 1 Experiment 2a: May 3 - May 31 Experiment 2b: June 14 – June 22 Experiment 3a: Sept 18 - Sept 28 Experiment 3b: Dec 12 - Jan 16 Experiment 3c: Dec 6 - Dec 10
Data exclusions	One participant left the study early without competing it. Eight participants who completed the study were excluded as their performance in either task exceeded three standard deviations from the population mean in that study.
Non-participation	One participant left the study early without completing it. No participant declined participation.
Randomization	Participants were allocated to experiments based on when they signed up for the study.

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

n/a	Involved in the study
☒	☐ Antibodies
☒	☐ Eukaryotic cell lines
☒	☐ Palaeontology
☒	☐ Animals and other organisms
☐	☒ Human research participants
☒	☐ Clinical data

Methods

n/a	Involved in the study
☒	☐ ChIP-seq
☒	☐ Flow cytometry
☒	☐ MRI-based neuroimaging

Human research participants

Policy information about studies involving human research participants

Population characteristics	See above
Recruitment	Participants were recruited from the University of Chicago and Awh/Vogel Laboratory subject pools for course credit or pay. The credit participants were exclusively University of Chicago students.
Ethics oversight	The study protocol was approved by the University of Chicago Social & Behavioral Sciences Institutional Review Board.

Note that full information on the approval of the study protocol must also be provided in the manuscript.