

Forgetting from lapses of sustained attention

Megan T. deBettencourt^{1,2,3} · Kenneth A. Norman^{1,4} · Nicholas B. Turk-Browne^{1,4}

Published online: 5 June 2017
© Psychonomic Society, Inc. 2017

Abstract When performing any task for an extended period of time, attention fluctuates between good and bad states. These fluctuations affect performance in the moment, but may also have lasting consequences for what gets encoded into memory. Experiment 1 establishes this relationship between attentional states and memory, by showing that subsequent memory for an item was predicted by a response time index of sustained attention (average response time during the three trials prior to stimulus onset). Experiment 2 strengthens the causal interpretation of this predictive relationship by treating the sustained attention index as an independent variable to trigger the appearance of an encoding trial. Subsequent memory was better when items were triggered from good versus bad attentional states. Together, these findings suggest that sustained attention can have downstream consequences for what we remember, and they highlight the inferential utility of adaptive experimental designs. By continuously monitoring attention, we can influence what will later be remembered.

Electronic supplementary material The online version of this article (doi:10.3758/s13423-017-1309-5) contains supplementary material, which is available to authorized users.

✉ Megan T. deBettencourt
debetten@uchicago.edu

- ¹ Princeton Neuroscience Institute, Princeton University, Princeton, NJ, USA
- ² Institute for Mind and Biology, University of Chicago, Chicago, IL, USA
- ³ Department of Psychology, University of Chicago, Chicago, IL, USA
- ⁴ Department of Psychology, Princeton University, Princeton, NJ, USA

Keywords Goal-directed attention · Distraction · Episodic memory · Real time

Introduction

It is often necessary to sustain attention for long periods of time, such as when listening to a lecture or driving a car. Despite this ubiquitous need, and the potential consequences of losing focus, humans are notoriously bad at sustaining attention. Lapses occur frequently in a wide range of laboratory studies, such as continuous performance and response inhibition tasks (e.g., deBettencourt, Cohen, Lee, Norman, & Turk-Browne, 2015; Robertson, Manly, Andrade, Baddeley, & Yiend, 1997) and low-prevalence visual search (e.g., Wolfe, Horowitz, & Kenner, 2005; c.f. Fleck & Mitroff, 2007). Indeed, sustained attention is characterized by continuous fluctuations between good and bad attentional states (Esterman, Noonan, Rosenberg, & DeGutis, 2013; Esterman, Rosenberg, & Noonan, 2014; Rosenberg, Noonan, DeGutis, & Esterman, 2013). The purpose of this study was to investigate the longer-term consequences of such fluctuations for memory.

Attention and memory are intricately related in general (see Aly & Turk-Browne, 2017). However, attention is a broad term that encompasses a wide range of tasks (see Chun, Golomb, & Turk-Browne, 2011), and the relationship between attention and memory has been primarily examined for divided attention (see Craik, 2001) or selective attention (e.g., Moray, 1959; Turk-Browne, Golomb, & Chun, 2013; Uncapher, Hutchinson, & Wagner, 2011; Yi & Chun, 2005). Here, we explore the relationship between another attentional construct—sustained attention—and subsequent recognition memory. There has been extensive work investigating the consequences of fluctuations of sustained attention on

immediate task performance (Esterman et al., 2013, 2014; Robertson et al., 1997; Rosenberg et al., 2013), but relatively less work on the mnemonic consequences. One study found that mindfulness correlated with sustained attention in a response inhibition task and memory for background distractor scenes (Rosenberg et al., 2013). Other studies have reported memory differences based on the response demands of the task (Chiu & Egner, 2015a, 2015b; Makovski, Jiang, & Swallow, 2013).

Here, we aim to more directly link sustained attention to memory encoding. That is, we test the hypothesis that the state of sustained attention leading into a trial determines the mnemonic fate of that trial. Various response signatures, such as increased response variability (Esterman et al., 2013, 2014; Rosenberg et al., 2013) or faster responding (deBettencourt et al., 2015; Robertson et al., 1997), have been associated with attentional lapses in continuous performance tasks. Faster responses are associated with habitual responding as opposed to carefully attending to stimulus properties. Infrequent trials, where the habitual response and the correct response differ, thus provide a critical interrogation of sustained attentional state. Therefore, in this study, we operationalize the state of sustained attention as response times (RTs) on preceding trials and use that to predict behavior on infrequent trials. In Experiment 1, we show that this measure of attention predicts whether upcoming information will be encoded successfully. In Experiment 2, we seek to establish the causal nature of this relationship using a real-time design in which the attentional measure is treated as an independent variable for triggering encoding opportunities.

Experiment 1

The goal of this experiment is to test whether the state of sustained attention going into a trial predicts memory for that trial. Specifically, we hypothesize that faster preceding RTs (indicating poor attention) will be associated with worse encoding of upcoming information.

Methods

Participants Thirty-two undergraduates (22 female; mean age = 20.5 years) from Princeton University participated for course credit or US\$10 payment. This sample size was chosen to be twice as large as previous studies of sustained attention (e.g., deBettencourt et al., 2015), because the effects on memory were anticipated to be smaller. One additional participant was excluded for task performance more than 3 SDs below the mean. All participants in both experiments reported normal or corrected-to-normal color vision, and provided informed consent to a protocol approved by the Princeton University IRB.

Stimuli Color scene images (550 indoor, 550 outdoor) were selected from the SUN database (Xiao, Hays, Ehinger, Oliva, & Torralba, 2010). They subtended approximately 7° in the center of a gray background, with a central black fixation dot (0.1°). The dot turned white after each response.

Apparatus Participants were seated approximately 70 cm from a CRT monitor (100-Hz refresh rate). Stimuli were presented using MATLAB (MathWorks, Natick, MA, USA) and the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997).

Procedure The design consisted of two phases: sustained attention and surprise memory test. In the sustained attention phase (Fig. 1a), each participant viewed 500 unique images for 1000 ms each with no interstimulus interval. Overall, 90% of the images (450) were from the “frequent” category (e.g., outdoor) and 10% (50) were from the “infrequent” category (e.g., indoor); categories were counterbalanced across participants. Participants were instructed to press “h” with their right index finger for the frequent category and “j” with their right middle finger for the infrequent category. They completed a short practice block until achieving 80% accuracy.

In the surprise memory test phase (Fig. 1b), each participant viewed 200 unique images (100 per category). Half of the images had appeared in the sustained attention task (“old”) and the other half were novel (“new”). All 50 old images from the infrequent category were included, along with 50 of the 450 old images from the frequent category (10 images per 100 trials to balance onset time). Participants were instructed to press buttons 1–4 with their right index finger to indicate their confidence that the image had appeared previously (response mapping presented below each image). The image and mapping remained on the screen until response, after which the response was also shown for 500 ms, followed by a blank 500-ms interstimulus interval. Participants were instructed to balance their responses across the buttons.

Analysis

The state of sustained attention for a trial was operationalized as the average RT over the three preceding trials (deBettencourt et al., 2015; cf. Robertson et al., 1997). Before relating this measure to memory on a trial-by-trial basis, we removed the linear drift in RTs throughout the sustained attention phase to control for generic time-dependent effects (e.g., practice, fatigue). For subsequent memory analyses, high-confidence old responses were treated as remembered and all other responses as forgotten (e.g., Kim, Lewis-Peacock, Norman, & Turk-Browne, 2014; Wagner et al., 1998). Across infrequent items within participant, logistic regression was used to predict this binary memory variable from the RT index of sustained attention at encoding. We

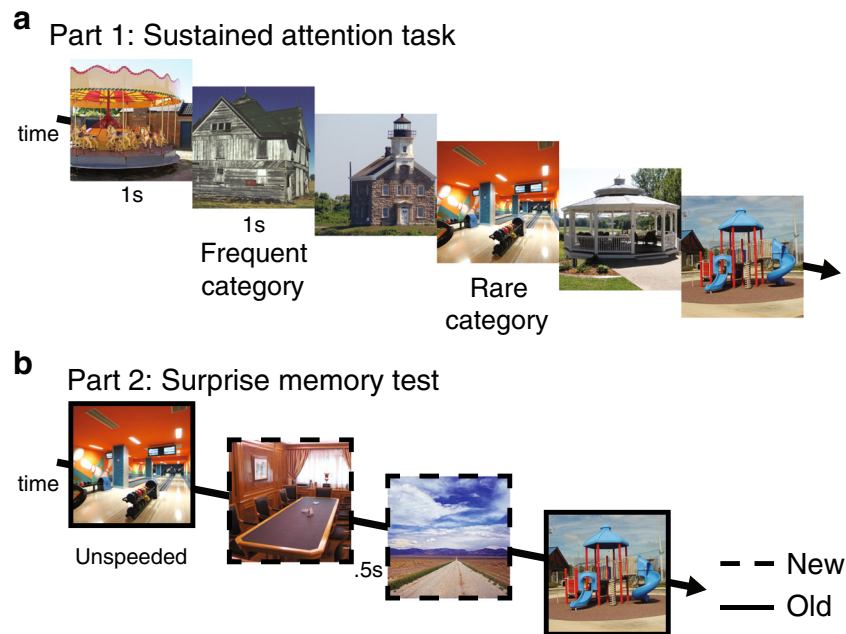


Fig. 1 Experimental design. **a** Participants first completed a sustained attention task in which they viewed trial-unique scene images and made an indoor/outdoor judgment. Of the scenes, 90% were from one of these categories (e.g., outdoor) and 10% were from the other category (e.g., indoor). Because of this imbalance, responding correctly to the infrequent category required inhibiting the prepotent response to the frequent

category. **b** Participants then completed a surprise memory test in which they reported their confidence that each scene had appeared in the first part of the experiment. Of the images, 50% were from the sustained attention task (*old*) and 50% were novel to the experiment (*new*). Among these, 50% were from the frequent category and 50% were from the infrequent category

verified that the effects were consistent across a range of sizes of the trailing window (Supplemental Fig. 1).

Because some of the data violated the assumption of normality, all statistics were computed using a nonparametric random-effects approach in which participants were resampled with replacement 100,000 times (Efron & Tibshirani, 1986). Null hypothesis testing was performed by calculating the proportion of the iterations in which the bootstrapped mean was in the opposite direction. The mean and 95% confidence interval (CI) of the bootstrapped distribution are reported as descriptive statistics. All results that were significant at $p < 0.05$ with nonparametric tests remained significant with parametric tests, except where explicitly noted. All data and analyses are available online in a jupyter notebook with this publication (http://github.com/PrincetonCompMemLab/deBettencourt_realtimeBehav).

Results

During the sustained attention task, the overall sensitivity of the category judgments was above chance ($A' = 0.92$, 0.90 – 0.93 ; vs. chance = 0.5 , $p < 0.00001$). The error rate was higher for trials from the infrequent category (0.30 , 0.26 – 0.34) versus frequent category (0.02 , 0.02 – 0.03 ; $p < 0.00001$), reflecting the need to inhibit a prepotent response. Furthermore, our RT measure of sustained attention—average RT over preceding

three trials—was predictive of the accuracy on infrequent trials: responses were slower before a correct response (520 ms, 502 – 537) than an incorrect response (446 ms, 422 – 471 ; $p < 0.00001$). This effect remained significant when each trial was analyzed separately ($\Delta RT_{i-1} = 98$ ms, $\Delta RT_{i-2} = 70$ ms, $\Delta RT_{i-3} = 58$ ms; $ps < 0.00001$; Supplemental Fig. 1). The timecourse of this effect over preceding timepoints is depicted in Fig. 2a.

During the surprise memory test, sensitivity was above chance for items from both the frequent ($A' = 0.66$, 0.61 – 0.69 ; $p < 0.00001$) and infrequent category (0.81 , 0.79 – 0.83 ; $p < 0.00001$), although higher for the infrequent category ($p < 0.00001$). Among the old infrequent items—the basis of the subsequent memory analysis—38% (33–44) were classified as remembered and 62% (56–67) as forgotten. The critical test of our hypothesis concerned the relationship between this measure of memory at test and the measure of sustained attentional state at encoding. As illustrated in Fig. 2b, this was evaluated with logistic regression across items within participant. The resulting slopes were positive across participants ($\beta = 1.07$, -0.17 – 2.18 ; $p = 0.038$ nonparametric, $p = 0.087$ parametric). The distribution of the slopes is shown in Fig. 2c. This effect was present when using the RT from only the immediately preceding trial ($\beta_{i-1} = 0.88$; $p = 0.038$ nonparametric, $p = 0.092$ parametric; Supplemental Fig. 1). We also replicated this relationship using the frequent trials ($\beta = 1.69$, 0.33 – 2.95 ; $p = 0.0065$).




Fig. 2 Relating sustained attention and memory. **a** RTs were slower before a correct response (*blue*) vs. incorrect response (*pink*) to an infrequent trial (all time points, $p_s < 0.00001$). Individual participants are plotted in *thinner lines* and the average in *thicker lines*. Raw RTs are reported in the text for this analysis but normalized RTs are depicted here because they were the input to the subsequent memory analysis shown in the other panels (statistics were unaffected). **b** Illustration of approach for quantifying relationship between sustained attention and

memory in one representative participant. For every item from the infrequent category, the average RT over the three preceding trials in the sustained attention task was fit to the binarized recognition judgment from the surprise memory test using logistic regression. Each *dot* is one item and the *line* is the fitted logistic function. **c** The logistic functions for all participants are plotted, revealing a reliably positive slope on average ($p = 0.038$). That is, slower preceding RTs correlated with better memory

Discussion

This experiment demonstrates that it is possible to predict whether an image will be remembered even before it appears, based on the state of sustained attention. This temporal precedence is consistent with attentional state being part of the causal chain that determines which memories are formed. However, this analysis is based on a correlation of idiosyncratic fluctuations in attention not under experimental control. We seek to strengthen the causal interpretation of this relationship in Experiment 2 using a real-time, adaptive design in which attentional state serves as an independent variable.

Experiment 2

The goal of this experiment is to test whether triggering encoding trials based on the current attentional state will influence what they remember. Specifically, we hypothesize that images presented while participants are responding faster than usual (indicating a bad sustained attentional state) will be remembered worse.

Methods

Participants Twenty-four undergraduates (15 female; mean age = 19.2 years) from Princeton University participated for course credit. This sample size was chosen to be smaller than Experiment 1 because we anticipated that the triggering design would lead to larger effects. One additional participant was excluded for task performance more than 3 SDs below the mean.

Stimuli and Apparatus Same as Experiment 1.

Procedure As in Experiment 1, participants completed a sustained attention task and a surprise memory test, with the same display and response procedure. During the sustained attention phase, each participant viewed 500 unique images. The first and last 50 trials had the same distribution of categories as Experiment 1, with 10% being from the infrequent category. The middle 400 trials (trials 51–450) began with 100% of trials from the frequent category. Up to 40 of those 400 trials could be replaced with images from the infrequent category, depending on real-time measurements of behavior. Specifically, infrequent trials were inserted when our measure of sustained attention deviated above an upper bound or below a lower bound. These bounds were recomputed for each trial i in several steps: analogous to Experiment 1, linear drift in the RTs was removed from trials 1 to $i-1$. The overall mean and standard deviation (SD) of these trials was then calculated from the residuals. The average RT of the preceding three trials ($i-3$, $i-2$, $i-1$) was calculated to index momentary sustained attention. If the moving-window RT was slower than the mean +1 SD or faster than the mean -1 SD, then the image for trial i was drawn from the infrequent category (Fig. 3). Otherwise, the frequent image on trial i was not replaced. We required a minimum of three frequent trials between infrequent trials (to not contaminate the moving-window RT measure). The average number of infrequent trials during the real-time period was 27.8 (25.3–30.4). The average list position for trials triggered due to slower versus faster RTs was not reliably different ($p = 0.28$). Participants were not informed that their RTs controlled when they would be shown images from the infrequent category.




Fig. 3 Online adaptive experimental design. The real-time triggering design is depicted for a period of trials from a representative participant. The *dashed black line* depicts the RTs to individual trials. The *solid black line* depicts the trailing window average RT over the 3 preceding trials. The bounds for triggering an infrequent trial, ± 1 SD, are plotted in a *solid gray line*. Infrequent trials were triggered if the average RT was slower than the upper bound (*blue dots*) or faster than the lower bound (*pink dots*). Otherwise, frequent trials were shown

In the surprise memory test, participants were shown all old images from the infrequent category (including triggered images), 50 old images from the frequent category, and a matched number of new images from each category (mean = 177, range = 154–204).

Analysis Infrequent trials from the real-time period were sorted based on whether they were triggered by faster or slower RTs. Otherwise, the analyses and statistics were similar to Experiment 1.

Results

During the sustained attention task, overall sensitivity was above chance (mean $A' = 0.89$, 95% CI = 0.88–0.91; vs. chance = 0.5; $p < 0.00001$). For the infrequent category, the error rate was higher for images triggered by faster (0.66, 0.55–0.77) than slower RTs (0.26, 0.21–0.32; $p < 0.00001$; Fig. 4a). However, trials triggered by faster (vs. slower) RTs appeared after shorter sequences of frequent trials (13.5 vs. 19.8; $p = 0.013$ nonparametric, $p = 0.064$ parametric). To investigate whether this accounted for differences in performance, we ran a logistic regression analysis to relate the length of the frequent trial sequence to infrequent trial accuracy. There was no reliable relationship in either the faster ($\beta = -0.01$, -0.19–0.34; $p = 0.56$) or slower RT conditions ($\beta = 0.03$, -0.14–0.35; $p = 0.43$) for participants who made correct and incorrect responses in each condition ($n = 17$).

During the surprise memory test, sensitivity was above chance for items from both the frequent category ($A' = 0.71$, 0.67–0.74) and the infrequent category (0.80, 0.80–0.82), although higher for the infrequent category ($p < 0.00001$). The critical test of our hypothesis concerned the infrequent trials during the real-time period, and specifically whether memory was worse for trials triggered in a worse sustained attentional state. Indeed, participants remembered fewer images from

trials triggered by faster (0.24, 0.18–0.32) versus slower RTs (0.38, 0.31–0.44; $p = 0.00023$; Fig. 4b).

One possible explanation for why RT triggering affected memory is that the state of sustained attention determined whether a behavioral error would be made, and such errors prevented encoding into memory. Although this kind of indirect effect would be consistent with our claims about the causal role of sustained attention in memory encoding, we evaluated whether RT per se was additionally responsible by restricting the subsequent memory analysis to infrequent trials with correct responses. We included all participants who made at least one correct response in both conditions ($n = 17$). Even in this restricted sample, participants remembered fewer images triggered by faster (0.26, 0.17–0.40) versus slower RTs (0.40, 0.32–0.49; $p = 0.0059$; Fig. 4c).

Discussion

This experiment employed a triggering design to more directly assess the relationship between attentional state and subsequent memory. Attentional state, operationalized as fluctuations in RT, was used as an independent variable. The design of each participant's session was customized adaptively by monitoring attentional state until it was particularly good or bad and then triggering critical probe trials. By directly manipulating whether items were presented in a good or bad attentional state, this experiment provides stronger support for the claim that sustained attention can control memory encoding.

General Discussion

The goal of this study was to test the nature of the relationship between moment-to-moment fluctuations in sustained attentional state and episodic memory. In Experiment 1, we confirmed that quicker responses predicted upcoming errors in a sustained attention task and then used this measure of attentional state to predict how well stimuli from the task had been encoded incidentally. In Experiment 2, we treated RT as an independent variable in determining the timing of the experimental design, providing an endogenous manipulation of attentional state that allowed us to more precisely probe encoding during good and bad states. Differences in RT have been linked to many psychological phenomena. Studies employing tasks very similar to ours have interpreted RTs as reflecting cognitive control (Chiu & Egner, 2015a, 2015b) or motor processes (Makovski et al., 2013). The relative contribution of such processes—above and beyond sustained attention—to explaining variance in memory will need to be examined in future studies that better isolate the component processes involved.

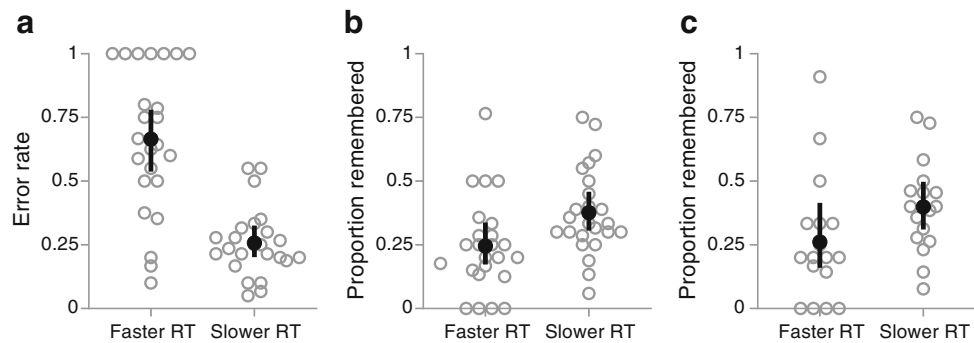


Fig. 4 The fate of triggered trials. **a** During the sustained attention task, participants made more errors on infrequent trials that were triggered by faster vs. slower RTs. **b** During the surprise memory test, participants remembered fewer images that were triggered by faster vs. slower RTs. **c** Participants still remembered fewer images that were triggered by faster

vs. slower RTs when the analysis was restricted to correct responses in the sustained attention task. Individual participants are depicted in *gray circles*. The average is the *solid black circle* and the *error bars* depict 95% CIs of the mean

The online and adaptive design employed in Experiment 2 represents a methodological approach inspired by real-time neuroimaging studies (Stoekel et al., 2014; Sulzer et al., 2013). This approach could have broad applications for many other cognitive domains in which performance can be predicted in advance from neural or behavioral measures, including attentional capture (Leber, 2010), visual detection (Salari, Büchel, & Rose, 2012), cognitive flexibility (Leber, Turk-Browne, & Chun, 2008), memory-based decision making (Duncan, Sadanand, & Davachi, 2012), and memory recollection (Otten, Quayle, Akram, Ditlewicz, & Rugg, 2006). Once predictive measures have been identified in a correlative manner, they can then be monitored in real time to trigger the onset of trials and control performance experimentally. Although this has been done with fMRI (Yoo et al., 2012) and EEG (Salari & Rose, 2016), we have shown that behavioral studies also stand to gain. For example, when re-examining the data from Experiment 1, relatively few trials (25%) exceeded our bounds for fast and slow RTs from Experiment 2 (beyond ± 1 SD). This is precisely the benefit of real-time adaptive designs, as we were able to ensure in Experiment 2 that all infrequent trials occurred at moments of clearly good and bad attentional states. Triggering designs can thereby support strengthened causal inference by leading to a cleaner separation of cognitive states of interest. By comparing different real-time measures of sustained attention, such as accuracy, mean preceding RT, and RT variability (Esterman et al., 2013, 2014; Rosenberg et al., 2013), it may be possible in future studies to explain different facets of memory performance and assess these ways of operationalizing attention. This may require adapting these attentional measures for real-time use: For example, RT variability is often calculated over a large time window, extending both before and after the trial of interest, but would need to be fully indexed in advance such that trials could be triggered in a leading-edge manner.

Beyond the predictive RT relationship for infrequent item memory, we additionally found better overall memory for

infrequent items, where participants had to deviate from a prepotent response, compared to frequent items, where participants did not. This may seem discrepant with some prior studies (Chiu & Egner, 2015a, 2015b), which found worse memory for items where participants had to inhibit their response. However, other results have demonstrated memory improvements for trials that require deviating from the default response (Makovski et al., 2013). In contrast to these studies, we used a frequency-based manipulation, such that trials that required inhibition of a prepotent response and execution of a different response were exceedingly rare (10%). The greater distinctiveness of these infrequent items in our study likely accounts for why they were better remembered.

In conclusion, our study demonstrates that attentional state can be monitored behaviorally using RT and that such states can have profound consequences for memory. By tracking attention in real time, it may be possible to avoid bad attentional states during learning and to reduce the likelihood of forgetting.

Acknowledgments This work was supported by NSF fellowship DGE1148900, NIH grant R01 EY021755, and Intel Corporation. The opinions expressed in this publication are those of the authors and do not necessarily reflect the views of these funding agencies.

Author contributions All authors helped conceive the study and contributed to the study design. M.T. deBettencourt collected and analyzed the data under the supervision of K.A. Norman and N.B. Turk-Browne. M.T. deBettencourt wrote an initial draft of the manuscript, which all authors read and edited.

References

- Aly, M., & Turk-Browne, N. B. (2017). How hippocampal memory shapes, and is shaped by, attention. In D. E. Hannula & M. C. Duff (Eds.), *The Hippocampus from cells to systems: Structure, connectivity, and functional contributions to memory and flexible cognition* (pp. 369–403). Switzerland: Springer International Publishing AG.

- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, *10*(4), 433–436.
- Chiu, Y.-C., & Egner, T. (2015a). Inhibition-induced forgetting results from resource competition between response inhibition and memory encoding processes. *The Journal of Neuroscience*, *35*(34), 11936–11945.
- Chiu, Y.-C., & Egner, T. (2015b). Inhibition-induced forgetting when more control leads to less memory. *Psychological Science*, *26*(1), 27–38.
- Chun, M. M., Golomb, J. D., & Turk-Browne, N. B. (2011). A taxonomy of external and internal attention. *Annual Review of Psychology*, *62*(1), 73–101.
- Craik, F. I. M. (2001). Effects of dividing attention on encoding and retrieval processes. In H. L. Roediger, J. S. Nairne, I. Neath, & A. M. Surprenant (Eds.), *The nature of remembering: Essays in honor of Robert G. Crowder* (pp. 55–68). Washington, DC, US: American Psychological Association.
- deBettencourt, M. T., Cohen, J. D., Lee, R. F., Norman, K. A., & Turk-Browne, N. B. (2015). Closed-loop training of attention with real-time brain imaging. *Nature Neuroscience*, *18*(3), 470–475.
- Duncan, K., Sadanand, A., & Davachi, L. (2012). Memory's Penumbra: Episodic memory decisions induce lingering mnemonic biases. *Science*, *337*(6093), 485–487.
- Efron, B., & Tibshirani, R. (1986). Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy. *Statistical Science*, *1*(1), 54–75.
- 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*(11), 2712–2723.
- Esterman, M., Rosenberg, M. D., & Noonan, S. K. (2014). Intrinsic fluctuations in sustained attention and distractor processing. *The Journal of Neuroscience*, *34*(5), 1724–1730.
- Fleck, M. S., & Mitroff, S. R. (2007). Rare targets are rarely missed in correctable search. *Psychological Science*, *18*(11), 943–947.
- Kim, G., Lewis-Peacock, J. A., Norman, K. A., & Turk-Browne, N. B. (2014). Pruning of memories by context-based prediction error. *Proceedings of the National Academy of Sciences of the United States of America*, *111*(24), 8997–9002.
- Leber, A. B. (2010). Neural predictors of within-subject fluctuations in attentional control. *The Journal of Neuroscience*, *30*(34), 11458–11465.
- Leber, A. B., Turk-Browne, N. B., & Chun, M. M. (2008). Neural predictors of moment-to-moment fluctuations in cognitive flexibility. *Proceedings of the National Academy of Sciences of the United States of America*, *105*(36), 13592–13597.
- Makovski, T., Jiang, Y. V., & Swallow, K. M. (2013). How do observer's responses affect visual long-term memory? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *39*(4), 1097–1105.
- Moray, N. (1959). Attention in dichotic listening: Affective cues and the influence of instructions. *Quarterly Journal of Experimental Psychology*, *11*(1), 56–60.
- Otten, L. J., Quayle, A. H., Akram, S., Ditewig, T. A., & Rugg, M. D. (2006). Brain activity before an event predicts later recollection. *Nature Neuroscience*, *9*(4), 489–491.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*(4), 437–442.
- Robertson, I. H., Manly, T., Andrade, J., Baddeley, B. T., & Yiend, J. (1997). 'Oops!': Performance correlates of everyday attentional failures in traumatic brain injured and normal subjects. *Neuropsychologia*, *35*(6), 747–758.
- Rosenberg, M., Noonan, S., DeGutis, J., & Esterman, M. (2013). Sustaining visual attention in the face of distraction: A novel gradual-onset continuous performance task. *Attention, Perception, & Psychophysics*, *75*(3), 426–439.
- Salari, N., Büchel, C., & Rose, M. (2012). Functional dissociation of ongoing oscillatory brain states. *PLoS One*, *7*(5), e38090.
- Salari, N., & Rose, M. (2016). Dissociation of the functional relevance of different pre-stimulus oscillatory activity for memory formation. *NeuroImage*, *125*, 1013–1021.
- Stoeckel, L. E., Garrison, K. A., Ghosh, S., Wightton, P., Hanlon, C. A., Gilman, J. M., ... Evins, A. E. (2014). Optimizing real time fMRI neurofeedback for therapeutic discovery and development. *NeuroImage: Clinical*, *5*, 245–255.
- Sulzer, J., Haller, S., Scharnowski, F., Weiskopf, N., Birbaumer, N., Blefari, M. L., ... Sitaram, R. (2013). Real-time fMRI neurofeedback: Progress and challenges. *NeuroImage*, *76*, 386–399.
- Turk-Browne, N. B., Golomb, J. D., & Chun, M. M. (2013). Complementary attentional components of successful memory encoding. *NeuroImage*, *66*, 553–562.
- Uncapher, M. R., Hutchinson, J. B., & Wagner, A. D. (2011). Dissociable effects of top-down and bottom-up attention during episodic encoding. *The Journal of Neuroscience*, *31*(35), 12613–12628.
- Wagner, A. D., Schacter, D. L., Rotte, M., Koutstaal, W., Maril, A., Dale, A. M., ... Buckner, R. L. (1998). Building memories: Remembering and forgetting of verbal experiences as predicted by brain activity. *Science*, *281*(5380), 1188–1191.
- Wolfe, J. M., Horowitz, T. S., & Kenner, N. M. (2005). Cognitive psychology: Rare items often missed in visual searches. *Nature*, *435*(7041), 439–440.
- Xiao, J., Hays, J., Ehinger, K. A., Oliva, A., & Torralba, A. (2010). SUN database: Large-scale scene recognition from abbey to zoo. In *2010 I.E.E.E. Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 3485–3492).
- Yi, D.-J., & Chun, M. M. (2005). Attentional modulation of learning-related repetition attenuation effects in human parahippocampal cortex. *The Journal of Neuroscience*, *25*(14), 3593–3600.
- Yoo, J. J., Hinds, O., Ofen, N., Thompson, T. W., Whitfield-Gabrieli, S., Triantafyllou, C., & Gabrieli, J. D. E. (2012). When the brain is prepared to learn: Enhancing human learning using real-time fMRI. *NeuroImage*, *59*(1), 846–852.