Research Article

Psychological Science
2019, Vol. 30(4) 587–595

DOI: 10.1177/0956797619830398
www.psychologicalscience.org/PS

Reactivation of Previous Experiences in a Working Memory Task

Gi-Yeul Bae and Steven J. Luck
Department of Psychology and Center for Mind and Brain, University of California, Davis

Abstract

Recent experiences influence the processing of new information even when those experiences are irrelevant to the current task. Does this reflect the indirect effects of a passively maintained representation of the previous experience, or is this representation reactivated when a new event occurs? To answer this question, we attempted to decode the orientation of the stimulus on the previous trial from the electroencephalogram on the current trial in a working memory task. Behavioral data confirmed that the previous-trial stimulus orientation influenced the reported orientation on the current trial, even though the previous-trial orientation was now task irrelevant. In two independent experiments, we found that the previous-trial orientation could be decoded from the current-trial electroencephalogram, indicating that the current-trial stimulus reactivated or boosted the representation of the previous-trial orientation. These results suggest that the effects of recent experiences on behavior are driven, in part, by a reactivation of those experiences and not solely by the indirect effects of passive memory traces.

Keywords

serial dependence, working memory, ERP decoding, previous trial decoding, open data

Received 7/22/18; Revision accepted 11/10/18

Recent experiences can have a large and automatic impact on our current perceptions, thoughts, and actions. Studies of priming and serial dependency show that we more readily perceive recently encountered objects and words, that recent experiences unconsciously influence how we categorize people and objects, and that we tend to repeat recent actions (Banaji & Hardin, 1996; Bertelson, 1965; Fischer & Whitney, 2014; Neely, 1991; Tipper, 1985). Despite the ubiquity of such effects, little is known about the mechanisms by which recent experiences that are no longer relevant can nonetheless influence the processing of current information.

Information from a given event leaves a trace in the synaptic connections among neurons (e.g., short-term synaptic facilitation; Mongillo, Barak, & Tsodyks, 2008), and there are two ways that these activity-silent traces could impact subsequent behavior. First, changes in synaptic weights could impact behavior indirectly by influencing the flow of information for subsequent events (Grill-Spector, Henson, & Martin, 2006; Stokes, 2015). Although these changes in synaptic weights are maintained without continued electrophysiological activity, and they are therefore invisible in neural recordings, they could nonetheless impact behavior by altering the neural response to new events. A second (but not incompatible) possibility is that activity-silent synaptic traces are also used to generate an active representation when the next stimulus is presented. This active representation could then directly impact processing, much like processing is impacted by active representations in working memory. Unlike changes in synaptic weights, an active representation should produce recordable neural signals, and the identity of the previous stimulus should be decodable from these signals.

Prior decoding studies have already shown that information can be moved back and forth between active and activity-silent states within a single trial of a

Corresponding Author:
Gi-Yeul Bae, University of California, Davis, Center for Mind and Brain, 267 Cousteau Pl., Davis, CA 95618
E-mail: gybae@ucdavis.edu

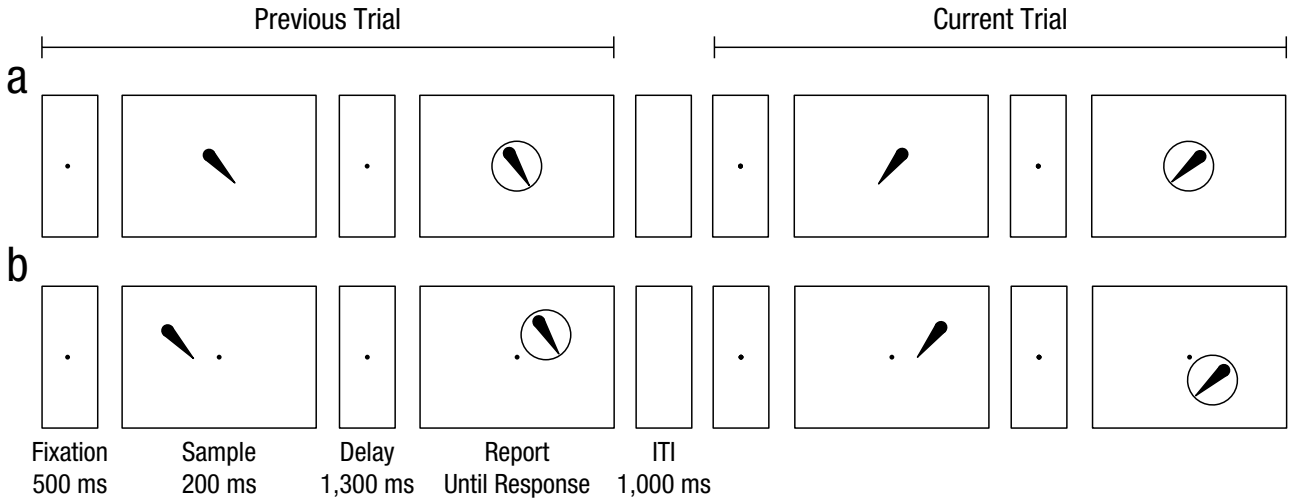

Fig. 1. Example trial sequences from (a) Experiment 1 and (b) Experiment 2. In both experiments, participants first saw a sample teardrop drawn in 1 of 16 orientations. After a delay, participants reported the remembered orientation by adjusting the orientation of a test teardrop until it matched the remembered orientation. The Experiment 2 task was identical to the Experiment 1 task except that the location and the orientation of the sample teardrop were independently manipulated, and the location of the test teardrop was independent of the location of the sample teardrop. ITI = intertrial interval.

working memory task (Rose et al., 2016; Wolff, Jochim, Akyürek, & Stokes, 2017). However, reactivation from activity-silent representations is thought to be possible only for task-relevant information from the current trial (Stokes, 2015). Here, we asked whether reactivation from an activity-silent representation is also possible for information from previous trials that is no longer task relevant, which would provide a new potential mechanism to explain how priming from previous trials impacts behavioral performance on a current trial.

We conducted a decoding study in which we recorded event-related potentials (ERPs) while participants performed a delayed estimation task. On each trial, they saw a single teardrop-shaped sample stimulus in 1 of 16 possible orientations and then adjusted the orientation of a test teardrop to match the remembered orientation after a short delay (see Fig. 1a). Because the orientation on one trial was independent of the orientation on the next trial, there would be no reason to intentionally maintain orientation information across trials. However, previous research on serial dependency has shown that the feature value presented on one trial can influence the reported value on the next trial (Fischer & Whitney, 2014; Fritsche, Mostert, & de Lange, 2017), and we expected to replicate this behavioral effect. In the ERP analyses, we attempted to decode the orientation of the stimulus from the scalp distribution of the ERP signal, using the high temporal resolution of ERPs to precisely assess the timing of the signal being decoded. We have previously shown that we can decode the current-trial orientation during the delay period in this data set (Bae & Luck, 2018). Here, we tested whether we could also decode the previous-trial orientation from the current-trial ERP scalp distribution. To establish the replicability and generality of the results, we also analyzed the data from a second experiment in which the location and orientation of the sample teardrop were independently manipulated (see Fig. 1b).

If the onset of a new trial automatically triggers a reactivation of the previous trial from an activity-silent state, then we should be able to decode the previous-trial orientation from the ERP signals following the current-trial sample stimulus (just as we have previously shown that we can decode the current-trial orientation from these signals). However, if information from the previous trial is maintained solely by means of activity-silent mechanisms, then the previous-trial orientation should not be decodable from the current-trial ERP signals (or should be only briefly decodable as the sample stimulus passes through the brain, as in the study by Wolff et al., 2017).

In both experiments, we found sustained decoding of the previous-trial stimulus orientation that began shortly after the onset of the current-trial sample stimulus, indicating that the previous-trial orientation information was reactivated. This finding demonstrates that the processing of a new input can trigger the reactivation of a previous experience even when that experience is no longer relevant to the task, providing a potential mechanism by which recent experiences can automatically influence current processing.

General Method

This article reports new analyses of previously published data (Bae & Luck, 2018). Additional details of

the stimuli, task, and analyses can be found in that article.

Participants

Two groups of 16 college students between the ages of 18 and 30 years participated in each experiment (Experiment 1: 10 female, 6 male; Experiment 2: 9 female, 7 male). The sample size was determined a priori on the basis of similar decoding studies (Fahrenfort, Grubert, Olivers, & Eimer, 2017; Foster, Sutterer, Serences, Vogel, & Awh, 2016). The study was approved by the University of California, Davis, Institutional Review Board.

Stimuli and tasks

The stimuli were presented on an LCD monitor at a viewing distance of 100 cm. A black fixation dot was continuously present in the center of the display except during the intertrial interval. In Experiment 1 (see Fig. 1a), each trial began with a fixation dot (500 ms) followed by a black, teardrop-shaped sample stimulus (200 ms; 2.17° long, 0.8° maximum width) that was centered on the fixation dot. The teardrop was presented in 1 of 16 equally spaced (22.5°) orientations. After a 1,300-ms delay period, a response ring became visible, and the participant attempted to reproduce the sample orientation. After the participant started moving the mouse, a test teardrop appeared at the center of the response ring with an orientation that depended on the current mouse position. After adjusting the test teardrop to match the remembered sample teardrop, the participant clicked a mouse button to finalize the report and initiate a 1,000-ms intertrial interval. Each participant received 40 trials for each of the 16 orientations, in random order.

The task in Experiment 2 (see Fig. 1b) was identical to that in Experiment 1, except that the locations of the sample and test teardrops varied independently from trial to trial. The teardrop was presented at 1 of 16 equally spaced locations on a notional circle (radius = 2.17°) centered on the fixation dot. The location and orientation of the sample teardrop on a given trial were chosen independently and randomly from the 16 orientations and 16 locations (256 combinations). The response ring for the test teardrop was positioned independently of the location of the sample teardrop.

Analyses of behavioral data

The behavioral response on each trial was transformed into a response error (i.e., reported orientation – sample orientation). Positive values indicated clockwise errors, and negative values indicated counterclockwise errors. The mean response error was computed for each of the 16 stimulus-orientation differences between the previous-trial orientation and the current-trial orientation (i.e., previous-trial orientation – current-trial orientation; see Fig. 1d). The behavioral analyses excluded trials on which the response error was greater than 60° (0.29% of the total trials in Experiment 1, and 0.18% of the total trials in Experiment 2), which were likely to reflect lapses of attention.

Following previous research (Bae & Luck, 2017), we expected the reported orientation on the current trial to be repelled away from the previous-trial orientation. In other words, when the difference in orientation between the previous and current trials was negative (counterclockwise), we expected the reported orientation to be biased in a positive (clockwise) direction, and when the orientation difference was positive (clockwise), we expected a negative (counterclockwise) bias. We used an approach developed by Fischer and Whitney (2014) to quantify the magnitude of this serial-dependence effect. Specifically, we fitted the single-trial response errors over the 16 orientation differences with a function defined by the first derivative of a Gaussian. This function is given by the following equation:

$$E = \alpha \times wx \times ce^{-(wx)^2} + \beta,$$

where E is the response error for the current-trial orientation, x is the difference between the previous-trial and current-trial orientations, α is the amplitude of the peak of the curve ($\alpha = 0$ means no serial dependence), w is the scale of the curve width, and β is an intercept that compensates for any overall bias away from zero. The constant c is $t\sqrt{2} / e^{-0.5}$. This function was fitted to each individual participant's data, and a set of parameters that minimized the root-mean-square error was estimated. We then used the alpha parameter to represent the magnitude of the serial-dependence effect in our statistical analyses. The presence of a serial-dependence effect was assessed by determining whether this parameter was significantly different from zero using a one-sample, two-tailed t test.

Electroencephalogram (EEG) recording and preprocessing

Using a BrainVision actiCHamp system (Brain Products, Gilching, Germany), we recorded from 27 broadly distributed scalp sites, from the left and right mastoids, and from electrodes lateral to the external canthi and below the right eye (for additional details, see Bae & Luck, 2018). Electrode impedances were maintained below 50 KΩ (and were typically below 10 KΩ). The signals were digitized at 500 Hz with a 130-Hz antialiasing filter.

The EEG signals were referenced off-line to the average of the left and right mastoids. The horizontal

electrooculogram was computed as the difference between the two lateral canthi, and the vertical electrooculogram was computed as the difference between Fp2 and the electrode below the right eye. The data were band-pass filtered from 0.1 to 80 Hz and resampled at 250 Hz, and independent component analysis was used to remove voltage deflections associated with blinks and eye movements. We have previously shown that this eliminates any ocular signals that could be used to decode the stimuli (Bae & Luck, 2018). The EEG was then segmented for each trial from −500 ms to 1,500 ms relative to the onset of the sample teardrop. The segmented EEG was low-pass filtered at 6 Hz and resampled at 50 Hz (one data point per 20 ms) prior to decoding.

Decoding analysis

A more extensive description of the decoding approach is provided in our previous article (Bae & Luck, 2018), in which we demonstrated that the current-trial orientation could be decoded during the delay interval. The only difference was that a more appropriate and less liberal statistical approach was used in the present analyses (described below).

We used support vector machines (SVMs) combined with error-correcting output codes (ECOCs; Dietterich & Bakiri, 1995) to classify the previous-trial teardrop orientation on the basis of the distribution of the ERP signal over the 27 scalp electrodes (collapsed across the current-trial teardrop orientation). To take advantage of the temporal resolution of the EEG method, we decoded the data independently at each of the 100 time points from −500 ms to 1,480 ms (relative to sample array onset).

To decode the previous-trial orientation from the scalp distribution of the sustained ERP response[1] on the current trial, we organized the data with respect to the previous-trial orientation, irrespective of the current-trial orientation. Because the previous-trial orientation was undefined on the first trial, we removed the first trial from the analysis, leaving 39 trials per orientation. The decoding used 20 iterations of a threefold cross-validation procedure. The data set was divided into three equal-size groups of trials (three groups of 13 trials for each of the 16 orientations), yielding a scalp distribution at each time point being analyzed (a matrix of 3 groups × 16 orientations × 27 electrodes for each time point). To increase the signal-to-noise ratio, we averaged the 13 trials in a given group, yielding three averaged ERP waveforms for each orientation, and the averaged ERP voltage at a given time point was fed into the SVM-ECOC classifier. Two of the three groups of trials for a given orientation were used for training, and the third was used for testing. Each of the 16 SVMs was trained using a one-versus-all approach, in which each SVM was trained to distinguish between 1 orientation and the other 15 orientations.

The set of 16 trained SVM-ECOC models was then used to predict the previous-trial orientation for each of the averaged ERPs that were reserved for testing (1 for each orientation). A single predicted orientation was chosen for each tested ERP (at each time point) by minimizing the average binary loss over the 16 SVMs. A prediction was considered correct only if it exactly matched the true previous-trial orientation. Because there were 16 equiprobable orientations, chance was 1 in 16 (.0625).

This procedure was repeated 3 times at each time point, once with each of the three groups of data for a given orientation serving as the testing data set. This procedure was then iterated 20 times with new random assignments of trials to the three groups to minimize any idiosyncrasies associated with the assignment of trials to groups. After all iterations of the cross-validation procedure were completed, decoding accuracy was collapsed across the 16 orientations, across the three cross-validations, and across the 20 iterations, producing a decoding percentage for a given time point that was based on 960 decoding attempts (16 orientations × 3 cross validations × 20 iterations). Because this approach involved randomly subdividing trials into training and testing sets, there was no way that decoding accuracy could consistently fall below the chance level of 1 in 16.

As noted above, this decoding procedure was applied separately at each time point for a given participant, producing one decoding accuracy value for each time point (aggregated across cross-validations and iterations). To minimize noise in this function, we smoothed the decoding accuracy values across time points using a 5-point moving window (equivalent to a time window of ±40 ms). The smoothed functions were then submitted to statistical analyses (which took into account the autocorrelation produced by the smoothing). The temporal precision resulting from the entire data-processing pipeline was ±50 ms.

Statistical analysis of decoding accuracy

If the pattern of voltage over the 27 electrodes at a given time point contained information about the previous-trial orientation, then decoding accuracy would be greater than chance (.0625). To compare decoding accuracy with chance at each time point while controlling for multiple comparisons, we used a nonparametric cluster-based permutation technique (Groppe, Urbach, & Kutas, 2011). This method provides an intelligent

correction for multiple comparisons and does not require normally distributed data. Note that this approach was slightly different from and more conservative than that used in our previous study (Bae & Luck, 2018), in which we used a Monte Carlo approach rather than a permutation approach and did not fully account for temporal autocorrelation in the EEG data.

We first used one-sample *t* tests to determine whether the decoding accuracy at each individual time point during the entire 2,000-ms epoch was greater than chance. One-tailed tests were used at this step because below-chance decoding accuracy was not meaningful in our decoding analysis. We then found clusters of contiguous time points for which the single-point *t* tests were significant ($p < .05$), and the *t* scores within each cluster were then summed to produce a cluster-level *t* mass.

We then determined whether the observed *t* mass was greater than the 95th percentile of a null distribution that was created by randomly permuting the true target labels when testing the accuracy of the decoder. For each iteration of the permutation procedure, we randomly shuffled the labels indicating the previous-trial orientation and then computed the accuracy of the decoder at predicting these shuffled labels. This procedure simulated the accuracy values that would be obtained by chance if the decoder had no information about the orientation. Importantly, we used the same shuffled labels for all the time points in a given trial, instead of using different shuffled target labels for each time point independently, to reflect the temporal autocorrelation of the continuous EEG data. As in our main decoding procedure, accuracy was computed 960 times (16 directions × 3 cross-validations × 20 iterations) for a given time point, and these values were then aggregated to compute the decoding accuracy at that time point for that permutation iteration. The time course of the permutation-based decoding accuracy was then smoothed with the same 5-point running average filter that was applied to the real decoding accuracy values. For each iteration, this procedure was performed separately for each of the 16 participants. We then performed one-sample *t* tests against chance for each time point and computed the sum of the *t* values (the *t* mass) of the largest cluster (with a mass of 0 if there were no significant *t* values). If there was more than one cluster of individually significant *t* values, we took the mass with the largest summed *t* values.

This procedure was iterated 1,000 times, yielding one maximum cluster mass per iteration, to produce a null distribution of *t* mass values (yielding a resolution of $p = 10^{-3}$). The *p* value for a *t* mass from the actual data set was then computed from the percentile of this value within the null distribution (using linear interpolation; we report $p < 10^{-3}$ if the observed mass was greater than all masses in the null distribution). We concluded that the decoding was significantly above chance for a given cluster if the *t* mass for that cluster was in the top 95th percentile of the null distribution.

The decoding procedure and statistical analysis were identical for Experiments 1 and 2. In Experiment 2, we collapsed across sample locations when we decoded the previous-trial orientation. Because orientation and location were completely counterbalanced, the decoding of orientation could not have been influenced by information about location.

Results

The behavioral data from Experiment 1 are summarized in Figure 2a, which shows the direction of the response error as a function of the difference in orientation between the current trial and the previous trial. A serial-dependence effect was clearly present: The orientation reported on the current trial was biased away from the orientation presented on the previous trial. In other words, when the current-trial orientation was counterclockwise to the previous-trial orientation, the reported orientation was shifted further counterclockwise, and the converse was found for clockwise differences. This effect was tested statistically by fitting the data with the first derivative-of-Gaussian function (Fischer & Whitney, 2014). A one-sample *t* test indicated that the amplitude of this function was significantly greater than zero, $t(15) = -9.731$, $p = 7.143^{e-08}$.

We applied a machine-learning approach to the ERP data to decode the previous-trial sample orientation from the current-trial scalp distribution, independent of the current-trial orientation. Because there were 16 possible orientations, and we required an exact match for a classification to be considered correct, chance was 1 in 16 (.0625). We found that the decoding accuracy—assessed on data not used for training—was above this chance level (three significant clusters of time points according to a cluster mass permutation test, with $p < .001$, $p = .030$, and $p < .001$ for the individual clusters, respectively) starting approximately 100 ms after the onset of the current-trial sample stimulus (see Fig. 2cF This shows that new information about the previous-trial orientation became available in the current-trial EEG shortly after the onset of the current-trial sample stimulus and that this information was maintained during the delay interval of the current trial. Moreover, the decoding was not limited to the first few hundred milliseconds after the onset of the current-trial sample stimulus, as would be expected if an activity-silent representation of the previous-trial orientation modulated the neural response to the sample stimulus (as in Wolff

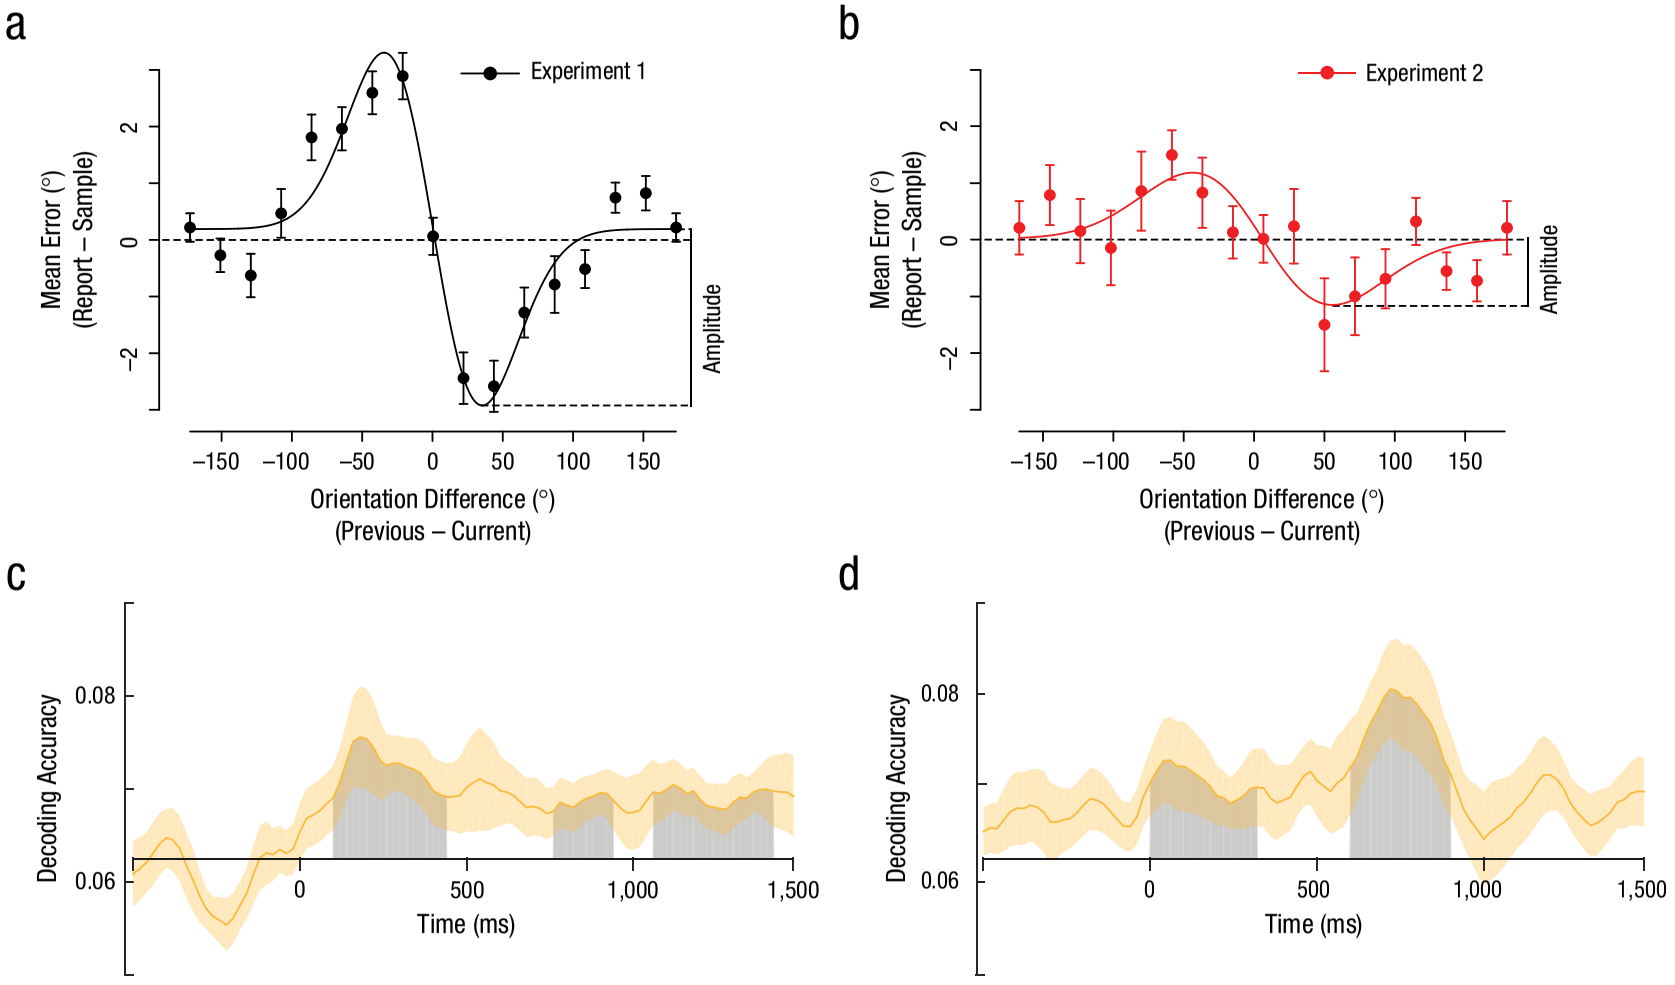

Fig. 2. Results. In the top row, mean response error (reported orientation – sample orientation) is shown as a function of the difference between the current-trial and previous-trial orientations in (a) Experiment 1 and (b) Experiment 2. Error bars represent ±1 *SE*. The curves show the best-fitting first derivative-of-Gaussian function, which was used to quantify the amplitude of the serial-dependence effect (indicated by the area labeled "Amplitude"). Mean accuracy for decoding the previous-trial orientation from the current-trial scalp topography is shown for (c) Experiment 1 and (d) Experiment 2. Orange shading represents ±1 *SE*. Gray areas indicate clusters of time points that produced above-chance decoding after correction for multiple comparisons.

et al., 2017). Instead, statistically significant decoding was present throughout much of the 1,500-ms period following the current-trial sample stimulus.

It is possible that this above-chance decoding reflected location information (i.e., the location of the tip of the teardrop) rather than bona fide orientation information. To address this possibility and assess the replicability of the Experiment 1 finding, we performed a similar decoding analysis on the data from Experiment 2. In this experiment (see Fig. 1b) the location and the orientation of the sample teardrop varied independently, and the location of the sample and test teardrops also varied independently, eliminating the possibility that location information could be used to remember the sample orientation (see extensive discussion in Bae & Luck, 2018). Again, we found a serial-dependence effect in the behavioral responses, $t(15) = -2.326$, $p = .034$ (one-sample t test on the amplitude of the first derivative-of-Gaussian function; see Fig. 2b) As has been shown previously (Fischer & Whitney, 2014), this effect was weaker when the stimulus location varied across trials (Experiment 2) than when it remained fixed (Experiment 1), $t(30) = -3.224$, $p = .003$ (two-sample t test comparing the amplitude of the first derivative-of-Gaussian function across the two experiments).

We attempted to decode the previous-trial orientation from the ERP data in Experiment 2, independently of the current-trial location and orientation and independent of the previous-trial location. As can be seen in Figure 2d, decoding accuracy was again significantly above chance in two clusters of time points ($p = .002$ and $p < .001$), extending more than 500 ms after the sample stimulus. This demonstrates both the replicability and the location independence of the previous-trial orientation decoding. However, there was some hint that decoding was weaker when location information was removed (e.g., smaller clusters of significant decoding accuracy in Experiment 2 than in Experiment 1). This may reflect the fact that many orientation-specific cells in the visual cortex are also spatially specific, decreasing the consistency of the orientation-specific scalp topography when the location varies across trials.

It is important to note that the ERP data in both experiments were baseline corrected by subtracting the mean voltage during the prestimulus baseline of the current-trial sample stimulus, eliminating any information about the previous-trial orientation that was present in the EEG prior to the sample stimulus. Thus, the finding that previous-trial decoding accuracy rose above chance shortly after the onset of the current-trial sample stimulus indicates that new information about the previous-trial orientation became available at this time. Thus, the previous-trial orientation representation was either reactivated from a completely activity-silent representation or boosted above the pretrial activity level. Consistent with this conclusion, results also showed no significant decoding of the previous-trial orientation during the intertrial interval when we used the prestimulus period of the previous trial as the baseline (see the Supplemental Material available online).

Given that the scalp EEG consists of small neural signals that are filtered through the skull and mixed with both biological and nonbiological noise, it is remarkable that we could decode not only the working memory representation of the current-trial orientation but also the task-irrelevant orientation from the previous trial. To push this even further, we attempted to decode the orientation of the $N - 2$ trial, but we found that decoding accuracy was near chance (see the Supplemental Material). This may indicate that more distant experiences are not reactivated or may instead reflect the limitations of scalp EEG decoding.

Discussion

In Experiment 1, we found that the orientation of the stimulus on the previous trial could be decoded from the ERP scalp topography on the following trial after the onset of the sample stimulus, demonstrating that the presentation of the current-trial sample stimulus reactivated or boosted a representation of the previous-trial stimulus. In Experiment 2, we replicated this effect in a task in which the location and orientation of the stimulus were independently manipulated, demonstrating that the reactivation of the previous-trial orientation was not location specific.

These findings are inconsistent with the hypothesis that task-irrelevant information from the previous trial is maintained solely in a passive, activity-silent manner. Such synaptic storage would not produce decodable ERP signals and could not directly underlie the sustained decoding of the previous-trial orientation that was observed after the onset of the current trial. Because ERPs reflect neural activity (primarily from postsynaptic potentials; see Buzsáki, Anastassiou, & Koch, 2012), the observed decoding necessarily reflects an active representation of the previous-trial orientation. Reactivation of previously silent representations has been observed in working memory tasks (Rose et al., 2016; Wolff et al., 2017), but reactivation was thought to be impossible when the trial is over and the information is no longer relevant (Wolff et al., 2017). The present study demonstrated that it is indeed possible to reactivate information that is no longer relevant, providing a potential mechanism for serial dependence and other priming-like effects. Note, however, that these results do not argue against the possibility that activity-silent representations of previous experiences also impact current-trial behavior.

In addition to producing a decodable signal in the ERP data, the previous-trial orientation also impacted the behavioral report of the current-trial orientation. The present study could not establish a causal relationship between the decoding effect and the behavioral effect, but it did demonstrate that previous experiences can be automatically reactivated by new stimuli, providing a potential mechanism by which previous experiences can automatically impact current behavior (as previously hypothesized by Logan, 1990). That is, if a given stimulus reactivates a representation of previous episodes involving stimuli of the same general class, including the behavioral response that was produced during the previous episode, this may be an efficient means of responding quickly to a stimulus. Additional research will be necessary to provide a firm link between the reactivation observed in the present study and the broad range of priming-related behavioral effects that have been reported in the literature.

We previously showed that the current-trial orientation could also be decoded during the current-trial delay period in this same data set (Bae & Luck, 2018). Thus, the EEG contained concurrent multiplexed signals about both the current-trial and previous-trial orientations. The EEG also contained concurrent information about the location of the current-trial stimulus, even when this information was task irrelevant (Bae & Luck, 2018; Foster, Bsales, Jaffe, & Awh, 2017). These findings are consistent with the hypothesis that multiple representations can be in an active state simultaneously (Sutterer, Foster, Adam, Vogel, & Awh, 2018), which contrasts with the hypothesis that only a single representation can be active at a given moment (McElree, 2006; Olivers, Peters, Houtkamp, & Roelfsema, 2011). However, additional research is necessary to rule out alternative explanations for the concurrent decoding, such as rapid switching between representations.

Action Editor

Edward S. Awh served as action editor for this article.

Author Contributions

Both authors designed the study. G.-Y. Bae recorded and analyzed the data. Both authors wrote the manuscript and approved the final manuscript for submission.

ORCID iD

Gi-Yeul Bae https://orcid.org/0000-0001-8598-0869

Acknowledgments

We thank Aaron Simmons for assistance with data collection.

Declaration of Conflicting Interests

The author(s) declared that there were no conflicts of interest with respect to the authorship or the publication of this article.

Funding

This research was made possible by Grant No. R01MH 076226 from the National Institute of Mental Health.

Supplemental Material

Additional supporting information can be found at http://journals.sagepub.com/doi/suppl/10.1177/0956797619830398

Open Practices

All data have been made publicly available via the Open Science Framework and can be accessed at osf.io/dbgh6. The design and analysis plans were not preregistered. The complete Open Practices Disclosure for this article can be found at http://journals.sagepub.com/doi/suppl/10.1177/0956797619830398. This article has received the badge for Open Data. More information about the Open Practices badges can be found at http://www.psychologicalscience.org/publications/badges.

Note

1. We also conducted the same decoding analyses using alpha-band oscillatory activity but found no significant decoding for either experiment (see the Supplemental Material).

References

Bae, G.-Y., & Luck, S. J. (2017). Interactions between visual working memory representations. *Attention, Perception, & Psychophysics*, *79*, 2376–2395. doi:10.3758/s13414-017-1404-8

Bae, G.-Y., & Luck, S. J. (2018). Dissociable decoding of spatial attention and working memory from EEG oscillations and sustained potentials. *The Journal of Neuroscience*, *38*, 2860–2817. doi:10.1523/JNEUROSCI.2860-17.2017

Banaji, M. R., & Hardin, C. D. (1996). Automatic stereotyping. *Psychological Science*, *7*, 136–141. doi:10.1111/j.1467-9280.1996.tb00346.x

Bertelson, P. (1965). Serial choice reaction-time as a function of response versus signal-and-response repetition. *Nature*, *206*, 217–218. doi:10.1038/206217a0

Buzsáki, G., Anastassiou, C. A., & Koch, C. (2012). The origin of extracellular fields and currents — EEG, ECoG, LFP and spikes. *Nature Reviews Neuroscience*, *13*, 407–420. doi:10.1038/nrn3241

Dietterich, T. G., & Bakiri, G. (1995). Solving multiclass learning problems via error-correcting output codes. *Journal of Artificial Intelligence Research*, *2*, 263–286. doi:10.1.1.72.7289

Fahrenfort, J. J., Grubert, A., Olivers, C. N. L., & Eimer, M. (2017). Multivariate EEG analyses support high-resolution

tracking of feature-based attentional selection. *Scientific Reports*, *7*(1), Article 1886. doi:10.1038/s41598-017-01911-0

Fischer, J., & Whitney, D. (2014). Serial dependence in visual perception. *Nature Neuroscience*, *17*, 738–743. doi:10.1038/nn.3689

Foster, J. J., Bsales, E. M., Jaffe, R. J., & Awh, E. (2017). Alpha-band activity reveals spontaneous representations of spatial position in visual working memory. *Current Biology*, *27*, 3216–3223.e6. doi:10.1016/j.cub.2017.09.031

Foster, J. J., Sutterer, D. W., Serences, J. T., Vogel, E. K., & Awh, E. (2016). The topography of alpha-band activity tracks the content of spatial working memory. *Journal of Neurophysiology*, *115*, 168–177. doi:10.1152/jn.00860.2015

Fritsche, M., Mostert, P., & de Lange, F. P. (2017). Opposite effects of recent history on perception and decision. *Current Biology*, *27*, 590–595. doi:10.1016/j.cub.2017.01.006

Grill-Spector, K., Henson, R., & Martin, A. (2006). Repetition and the brain: Neural models of stimulus-specific effects. *Trends in Cognitive Sciences*, *10*, 14–23. doi:10.1016/j.tics.2005.11.006

Groppe, D. M., Urbach, T. P., & Kutas, M. (2011). Mass univariate analysis of event-related brain potentials/fields I: A critical tutorial review. *Psychophysiology*, *48*, 1711–1725. doi:10.1111/j.1469-8986.2011.01273.x

Logan, G. D. (1990). Repetition priming and automaticity: Common underlying mechanisms? *Cognitive Psychology*, *22*, 1–35. doi:10.1016/0010-0285(90)90002-L

McElree, B. (2006). Accessing recent events. In B. H. Ross (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 46, pp. 155–200). San Diego, CA: Elsevier.

Mongillo, G., Barak, O., & Tsodyks, M. (2008). Synaptic theory of working memory. *Science*, *319*, 1543–1546. doi:10.1126/science.1150769

Neely, J. H. (1991). Semantic priming effects in visual word recognition: A selective review of current findings and theories. In D. Besner & G. W. Humphreys (Eds.), *Basic processes in reading: Visual word recognition* (pp. 264–336). Hillsdale, NJ: Erlbaum.

Olivers, C. N. L., Peters, J., Houtkamp, R., & Roelfsema, P. R. (2011). Different states in visual working memory: When it guides attention and when it does not. *Trends in Cognitive Sciences*, *15*, 327–334. doi:10.1016/j.tics.2011.05.004

Rose, N. S., LaRocque, J. J., Riggall, A. C., Gosseries, O., Starrett, M. J., Meyering, E. E., & Postle, B. R. (2016). Reactivation of latent working memories with transcranial magnetic stimulation. *Science*, *354*, 1136–1139. doi:10.1126/science.aah7011

Stokes, M. G. (2015). 'Activity-silent' working memory in prefrontal cortex: A dynamic coding framework. *Trends in Cognitive Sciences*, *19*, 394–405. doi:10.1016/j.tics.2015.05.004

Sutterer, D. W., Foster, J. J., Adam, K. C. S., Vogel, E. K., & Awh, E. (2018). Item-specific delay activity demonstrates concurrent storage of multiple items in working memory. *Biorxiv*. doi:10.1101/382879

Tipper, S. P. (1985). The negative priming effect: Inhibitory priming by ignored objects. *The Quarterly Journal of Experimental Psychology A*, *37*, 571–590. doi:10.1080/14640748508400920

Wolff, M. J., Jochim, J., Akyürek, E. G., & Stokes, M. G. (2017). Dynamic hidden states underlying working-memory-guided behavior. *Nature Neuroscience*, *20*, 864–871. doi:10.1038/nn.4546