

The Role of Meaning in Visual Memory: Face-Selective Brain Activity Predicts Memory for Ambiguous Face Stimuli

Timothy F. Brady,¹ George A. Alvarez,² and Viola S. Störmer¹

¹Department of Psychology, University of California, San Diego, La Jolla, California 92093, and ²Department of Psychology, Harvard University, Cambridge, Massachusetts 02138

How people process images is known to affect memory for those images, but these effects have typically been studied using explicit task instructions to vary encoding. Here, we investigate the effects of intrinsic variation in processing on subsequent memory, testing whether recognizing an ambiguous stimulus as meaningful (as a face vs as shape blobs) predicts subsequent visual memory even when matching the perceptual features and the encoding strategy between subsequently remembered and subsequently forgotten items. We show in adult humans of either sex that single trial EEG activity can predict whether participants will subsequently remember an ambiguous Mooney face image (e.g., an image that will sometimes be seen as a face and sometimes not be seen as a face). In addition, we show that a classifier trained only to discriminate between whether participants perceive a face versus non-face can generalize to predict whether an ambiguous image is subsequently remembered. Furthermore, when we examine the N170, an event-related potential index of face processing, we find that images that elicit larger N170s are more likely to be remembered than those that elicit smaller N170s, even when the exact same image elicited larger or smaller N170s across participants. Thus, images processed as meaningful, in this case as a face, during encoding are better remembered than identical images that are not processed as a face. This provides strong evidence that understanding the meaning of a stimulus during encoding plays a critical role in visual memory.

Key words: EEG; face memory; memory; N170; visual memory

Significance Statement

Is visual memory inherently visual or does meaning and other conceptual information necessarily play a role even in memory for detailed visual information? Here we show that it is easier to remember an image when it is processed in a meaningful way, as indexed by the amount of category-specific brain activity it elicits. In particular, we use single-trial EEG activity to predict whether an image will be subsequently remembered, and show that the main driver of this prediction ability is whether or not an image is seen as meaningful or non-meaningful. This shows that the extent to which an image is processed as meaningful can be used to predict subsequent memory even when controlling for perceptual factors and encoding strategies that typically differ across images.

Introduction

If you are asked to remember the visual details of an image—say, which exact face you saw—do you encode this information in a purely perceptual format, like a camera would store a photograph? Or does human memory always, even for visual details, depend on how meaningful that stimulus is to a given person (Bartlett, 1932)? Memory for visual details is often quite impres-

sive (Hollingworth, 2004; Brady et al., 2008, 2013), suggesting that detailed visual information is frequently stored in visual long-term memory. At the same time, it has long been argued that memory is semantically organized (Collins and Loftus, 1975), and that both elaborative encoding and strong retrieval cues are critical to the success of long-term memory (Baddeley et al., 2009). A meaningful interpretation of an image may allow for a more elaborative initial encoding, creating more or stronger paths to access the memory (Bower et al., 1975; Bradshaw and Anderson, 1982).

Consistent with the idea that how a stimulus is initially processed is critical for memory performance, a large body of research has shown that neural activity patterns during encoding can predict subsequent memory (Sanquist et al., 1980; Wagner et al., 1998; Paller and Wagner, 2002; Daselaar et al., 2004; Kuhl et

Received July 5, 2018; revised Nov. 20, 2018; accepted Nov. 27, 2018.

Author contributions: T.F.B., G.A.A., and V.S.S. edited the paper; T.F.B., G.A.A., and V.S.S. designed research; T.F.B. and V.S.S. performed research; T.F.B. and V.S.S. analyzed data; T.F.B. and V.S.S. wrote the paper.

This work was partially supported by an NSF Grant (BCS-1829434) to T.F.B. and V.S.S.

The authors declare no competing financial interests.

Correspondence should be addressed to Timothy F. Brady at timbrady@ucsd.edu.

<https://doi.org/10.1523/JNEUROSCI.1693-18.2018>

Copyright © 2019 the authors 0270-6474/19/391100-09\$15.00/0

Example stimuli




Figure 1. Left, Unambiguous Mooney face stimuli. Center, Ambiguous stimuli that some participants see as faces whereas others do not see as faces (when briefly presented). Right, Non-faces designed to match the face stimuli in low-level characteristics; made from scrambling and inverting parts of the unambiguous and ambiguous faces.

al., 2012). However, for the most part this existing work does not distinguish whether such encoding-related activity is a result of differences in perceptual input, differences in encoding strategies, or differences in how information connects to meaningful concepts or semantic knowledge. In particular, given that in the majority of memory tasks the same items tend to be remembered by each person (Isola et al., 2011; Bainbridge et al., 2013), in an experiment examining the neural correlates of subsequent memory performance, the stimuli that are subsequently remembered are likely to differ along a number of dimensions from those that are subsequently forgotten. Some work has specifically focused on the neural correlates of meaningful processing in subsequent memory. For example, when participants are instructed to make semantic (vs not) judgments of a stimulus at encoding, neural processing at encoding differs between these two conditions, and these differences predict subsequent memory performance (Sanquist et al., 1980; Hanslmayr et al., 2009; Dubé et al., 2013; Hanslmayr and Staudigl, 2014). However, despite examining the role of meaning and meaningful processing, this work involves an explicit manipulation of encoding strategy to require participants to focus on this aspect of the stimuli. Thus, the results are modulated by the performance of a different task by participants, rather than solely the particular processing of that stimulus that results in the most robust memory formation.

In the present study we test whether subsequent visual memory can be predicted by the extent to which a stimulus is initially processed in a meaningful way, independent of perceptual features, familiarity, or differences in encoding strategies. We use a set of visual stimuli, ambiguous “Mooney images” (Mooney, 1957), that are carefully matched such that different images tend to be remembered by different people, with no reliable predictability in which images are remembered. To examine whether subsequent memory performance can be predicted based on how the stimulus is processed initially, we record the brain’s electrophysiological activity during encoding of these stimuli and test whether single-trial EEG activity at encoding can predict whether a stimulus is later remembered or forgotten. In addition, we make use of a neural signature that is selective to face processing: the N170 component (George et al., 2005; Hadjikhani et al., 2009; Liu et al., 2014), which allows us to assess whether a stimulus was likely seen as a face or not at encoding. Overall, our data suggest that memory performance is significantly affected by how meaningful, in this case face-like, the stimulus was processed at encoding, even when the stimulus input and the encoding task is held constant across stimuli and individuals.

Materials and Methods

Participants. Data from 24 participants (of either sex) were included in the final data analysis; data from three additional participants had to be

excluded because <15 trials were available in at least one of the conditions because of artifacts in the EEG. All participants had normal or corrected-to-normal vision and were between 18 and 28 years of age. All studies were approved by the Institutional Review Board and all participants gave written informed consent before the experiment.

Experimental design. Participants were asked to remember a sequence of two-tone images (Mooney images; Mooney, 1957; Fig. 1). Stimuli consisted of 140 clear, unambiguous Mooney faces (stimuli designed so all participants would see them as a face), 140 ambiguous Mooney faces (stimuli designed so some participants would see them as a face and some would not see them as a face), and 140 non-face images. The unambiguous and ambiguous face images were chosen by using rankings collected on 20 independent naive individuals indicating whether they saw each stimulus as a face or not as a face (for details, see Schwiedrzik et al., 2018). The non-face images were created by segmenting continuous regions of black and white from the face images and scrambling these parts by placing them at random locations and random angles within a new image, thus preserving overall luminance and the shape of each of the regions while removing the holistic connection between regions that allows the images to be seen as faces. Note that although we refer to these as non-faces because they are scrambled and on average were not recognizable as faces, they could sometimes have some face features or be otherwise recognizable to participants as some other objects. Nevertheless, they are much less likely to be seen as meaningful than the unambiguous faces or ambiguous faces.

Half of the images of each set were randomly selected for each participant to be presented during the study blocks, and the remaining half of the images served as foils during the test block. During the study blocks, participants were shown a stream of images on a gray background. Each image was presented for 500 ms in the center of the screen, followed by a 3000 ms fixation cross. Participants were instructed to remember the images while keeping their eyes in the center of the screen. Participants were shown 210 images overall, divided into 10 study blocks (21 images per block, seven of each stimulus type). After each study block, participants performed a memory test by making an old/new judgment: Forty-two images were presented one at a time, and participants had to indicate for each test image whether they had seen it in the previous block (“old”) or not (“new”) by pressing one of two buttons on a keyboard. Half of the test images were old (i.e., they were presented during the previous encoding block), and the remaining half of the images were new (i.e., foils; participants had not seen them before). Study images from each stimulus set were selected at random for each participant and presented in random order. Participants were instructed to emphasize accuracy, not speed, in making their judgments.

Electrophysiological recordings. While participants performed the memory task, brain activity was recorded continuously from 32 Ag/AgCl electrodes arranged according to the 10–20 system, mounted in an elastic cap and amplified by an ActiChamp amplifier (BrainVision). The horizontal electro-oculogram was acquired using a bipolar pair of electrodes positioned at the external ocular canthi, and the vertical electro-oculogram was measured at electrode FP1, located above the left eye. All scalp electrodes were referenced to an electrode on the right mastoid online, and were digitized at a rate of 500 Hz. Signal processing was

performed with MATLAB (MathWorks) using the EEGLAB and ERPLAB toolboxes (Delorme and Makeig, 2004; Lopez-Calderon and Luck, 2014) and custom-written scripts. Continuous EEG data were filtered off-line with a bandpass of 0.01–112 Hz. Artifact rejection was performed for the interval –200 to +600 ms poststimulus onset and trials with horizontal eye movements, blinks, or excessive muscle movements were excluded from the analysis. Artifact-free data were re-referenced to the average of the left and right mastoids. For all further analysis, including decoding, data were epoched into trials and baselined to the 200 ms prestimulus interval, and only data from trials without artifacts were included in the analysis.

Statistical analysis

Behavioral analysis. To assess memory independent of response bias we calculated d' as our measure of sensitivity. We then used a repeated-measures analysis with image type as a within-subject factor, followed by planned t tests to assess the extent to which memory differs between our three conditions.

Decoding analysis. Our main electrophysiology analysis is based on a support vector machine (SVM) decoding approach (Cox and Savoy, 2003). In particular, we use EEG activity from each encoding trial to train a classifier to distinguish between the conditions and to predict subsequent memory. The input feature vector in each case consists 806 features: 31 electrodes \times 26 time points (a running average of 20 ms, sampled every 20 ms between 100 and 600 ms poststimulus). For situations where we train and test on the same conditions, we then apply a leave-one-out classification procedure with 500 iterations where a randomly chosen item from each category is left out. Performance is then averaged over iterations for each participant. For conditions where we train and test on orthogonal sets of data we simply train and test once on all of the data. In each case, we then perform one-sample t test versus chance performance (50%) to assess the statistical significance of the decoding accuracy. Our leave-one-out-hold-out set always held out one example of each of the two test categories (rather than just 1 item), such that 50% of tested items were of each kind regardless of the base rate in the data. Thus, chance is always 50%. To confirm this was effective in controlling for the base rate differences, we performed permutation analyses, shuffling the labels for the data for the training data, and find that chance is indeed exactly 50% for all of our decoding analyses.

We apply this SVM to training and testing based on: (1) whether an unambiguous face versus non-face was shown (decoding the stimulus itself); (2) whether an ambiguous face trial was subsequently remembered or forgotten; (3) whether training on the perceptual distinction of unambiguous face versus non-face transfers to testing on remembered versus forgotten ambiguous faces, as expected if this memory difference is driven by whether the items are perceived as faces; and (4) whether training on the distinction of remembered versus forgotten for (a) non-faces, and (b) unambiguous faces transfers to testing on remembered versus forgotten ambiguous faces, as would be expected if there were general memory strength signals, like attention.

ERP analysis. In addition to the classification analysis, we also conducted a planned analysis using event-related potentials (ERPs). In particular, we assessed the magnitude of the N170, a measure of face processing (Bentin et al., 1996; George et al., 1996; Eimer, 2000), while participants encoded the images. Our main question of interest was whether this neural signature of category-specific processing would predict which ambiguous-face images were remembered and which were forgotten. In particular, because the amplitude of the N170 is indicative of which images participants' see as a face (George et al., 2005; Hadjikhani et al., 2009; Liu et al., 2014), this component should provide a neural measure of category-specific processing for each image, which we hypothesized will predict which images will be memorable (Wiseman and Neisser, 1974; Bower et al., 1975). To assess this, ERPs were averaged separately for unambiguous face images, remembered-ambiguous face images, forgotten-ambiguous face images, and non-face images for each participant separately. ERPs were digitally low-pass filtered (–3 dB cut-off at 25 Hz) and the mean amplitude of the N170 component was measured between 140 and 180 ms at two right posterior electrodes PO8 and P8 (Eimer, 2000) with respect to a 200 ms prestimulus period. The

mean amplitudes of the filtered ERP waveforms were subjected to a repeated-measures analysis with image type as a within-subject factor. Planned pairwise comparisons were conducted to examine the N170 for unambiguous face images versus non-face images, and remembered-ambiguous face images versus forgotten-ambiguous face images. Summary data from all analyses are available at <https://osf.io/8b4gh/>.

Results

Behavioral results

We quantified memory performance primarily in terms of sensitivity (d'). As expected, participants remembered the unambiguous face images with the highest accuracy, followed by the ambiguous face images, and lastly the non-face images (Fig. 2A; $F_{(2,46)} = 47$, $p < 0.0001$, $\eta^2 = 0.67$). Pairwise comparisons revealed that all three conditions differed reliably from each other (unambiguous face vs non-face: $t_{(23)} = 8.19$, $p < 0.0001$, $\eta^2 = 0.75$; unambiguous face vs ambiguous face: $t_{(23)} = 3.45$, $p = 0.002$, $\eta^2 = 0.34$; ambiguous face vs non-face: $t_{(23)} = 6.28$, $p < 0.0001$, $\eta^2 = 0.63$). These effects were accounted for more by a difference in hit rates (74.9, 61.4, and 50.2%, respectively, for unambiguous-faces, ambiguous-faces and non-faces) from a difference in false alarm rates (24.3, 18.4, and 24.5%, respectively).

We found that participants were extremely inconsistent in which ambiguous face images they found easiest to distinguish from new images (Fig. 2B, C). In particular, the correlation between participants in which images were correctly categorized as old/new at test is extremely low ($r = 0.02$); <1% of the variance in which ambiguous faces were remembered was predicted by the images remembered by other participants (Fig. 2B). This is visualized in Figure 2C by showing each of the ambiguous Mooney face images as a column and indicating in green when a participant correctly categorized it as old/new at test. The lack of consistency is shown by the lack of vertical column structure (e.g., participants do not consistently get the same images correct or incorrect). Thus, on average, different participants remembered different ambiguous faces. This means that we can examine neural activity involved in successful memory encoding for ambiguous stimuli independent of the particular images.

In contrast to the current finding, there has been a significant amount of work showing that participants tend to be consistent in which images they report as old or as new (termed “memorability”; Bainbridge et al., 2013). This work is quite different from the current work in that it largely features a much more naturalistic and much larger context of images. In general, memory is deeply contextual (e.g., an image of an outdoor field may generally be not memorable and a man in a ball pit may generally be memorable, but in a context of 100 men in ball pits and only 1 field, exactly the opposite would be true). Thus, given our very narrow set of stimuli, all Mooney faces, and our shorter blocks (21 items, approximately half seen as faces and half not), we would expect much more variability between participants in our work than in this previous work on memorability. For example, imagine one ambiguous Mooney face that was an older male. For some participants, it would be seen as a face; for others not (introducing variance). For those who see it as a face, the context would also be highly variable (e.g., given the small number of other face stimuli in a block, of the 10 other “face” stimuli in a block, some participants might have all female faces; or all older male faces). Thus the uniqueness of the stimulus in a block and the things that differentiate a given item from its context are far more variable in our study than in the work on memorability that typically finds reliable patterns of which items are remembered across participants, likely explaining why we do not find the same

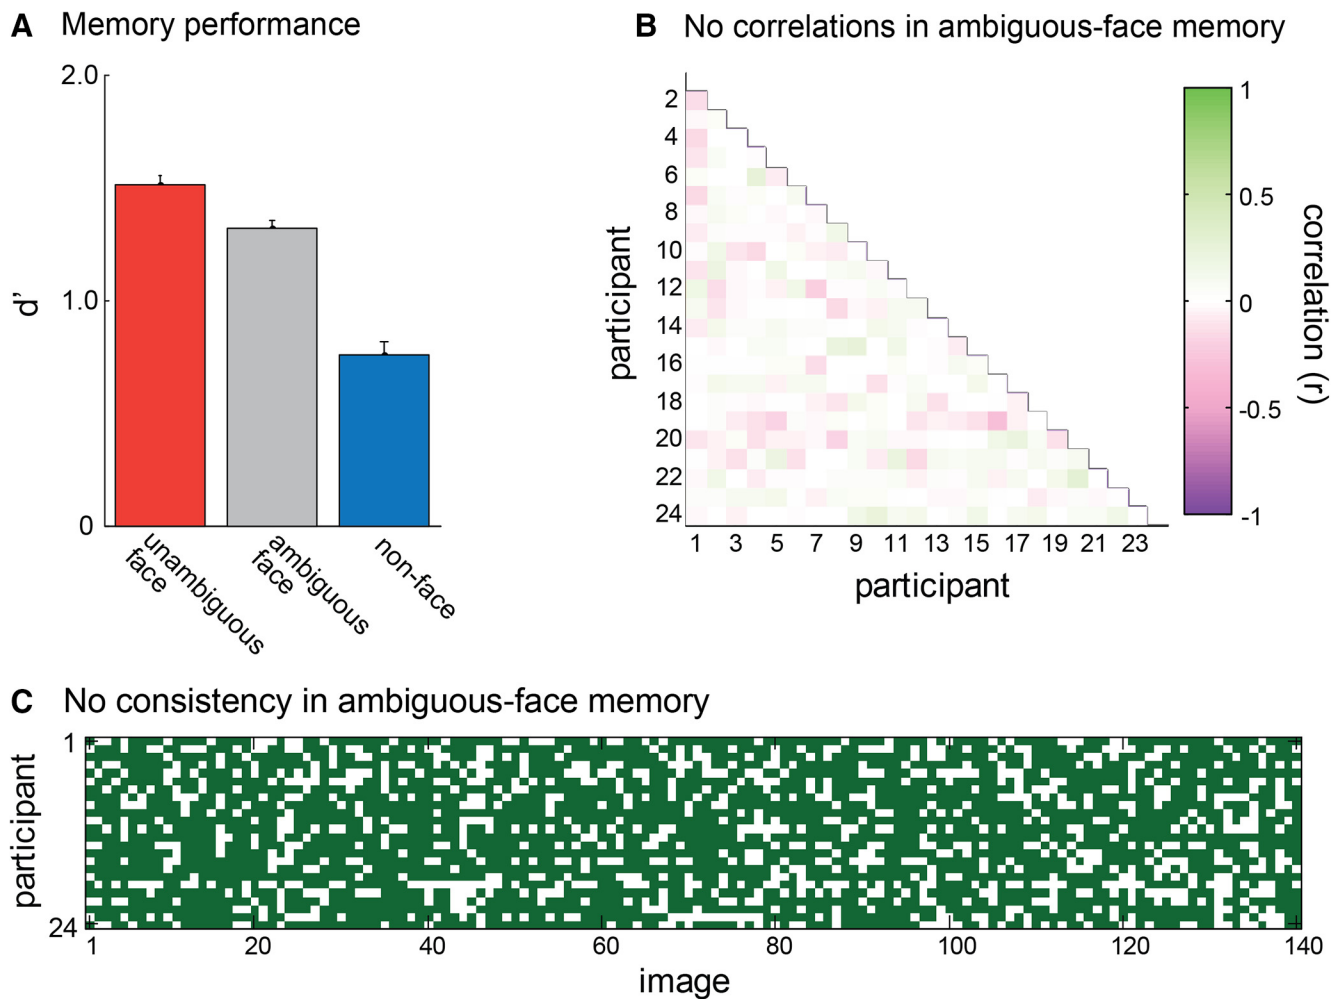


Figure 2. *A*, Memory performance for unambiguous faces, ambiguous faces and non-faces, in terms of d' (sensitivity). Participants were better able to remember unambiguous faces than non-faces, with ambiguous faces falling in between. Error bars represent ± 1 SEM. *B*, Correlations between participants in which ambiguous faces were correctly remembered. The lack of consistent correlation (mean $r = 0.02$) indicates that different participants remembered different ambiguous-faces. *C*, An alternative representation of this lack of consistency in which ambiguous-face images participants remembered, showing, for each of the 140 ambiguous-face images tested, which participants got the test trial correct (green) or incorrect (white). The lack of vertical columns indicates the lack of consistency.

reliability that they do. Importantly for our purposes, this lack of reliability in the current study means that we can examine neural activity involved in successful memory encoding for ambiguous stimuli independent of the particular images.

Decoding results

We first established a baseline for how well we could decode whether a stimulus currently being viewed is an unambiguous face or a non-face using single-trial EEG activity, effectively establishing an upper bound for decoding accuracy in our stimulus set. In particular, we trained an SVM to distinguish between unambiguous faces and non-faces based on EEG data averaged over 20 ms bins at all electrodes during the encoding interval (100–600 ms poststimulus), and then tested this classifier on a leave-one-out hold-out set (see Materials and Methods). On average across participants, we found decoding accuracy of 61.8% (SEM: 1.8%), well above-chance ($t_{(23)} = 6.63$, $p < 0.001$; Fig. 3A). Thus, despite the matched low-level features, we can decode whether a Mooney image that is currently being processed is a face or a non-face, suggesting that participants process face and non-face stimuli distinctly even when they are Mooney images (Kanwisher et al., 1998; Hsieh et al., 2010). Although this upper bound of

decoding accuracy is rather low, this is expected given noisy single-trial EEG data and well matched perceptual stimuli, in line with previous studies that have tried to decode subsequent memory or task performance in general using single-trial EEG activity (Leydecker et al., 2014; Noh et al., 2014; Höhne et al., 2016).

We can also use such decoding to ask whether our stimuli are well matched in simple factors like luminance. If the stimuli are well matched in terms of luminance and contrast, we should find that the classifier is incapable of decoding face versus non-face during the first 100 ms. We ran this analysis and found performance was near chance: 50.75% correct (SEM: 1.4%; $t_{(23)} = 0.55$, $p = 0.59$) with a Bayes factor providing evidence of 4.1:1 in favor of the null hypothesis (Rouder et al., 2009), suggesting that our matching process for face versus non-face was successful in equalizing the stimuli on low-level characteristics.

Our main question of interest was how participants encoded ambiguous face stimuli and whether ambiguous face stimuli that were subsequently remembered were processed distinctly from those that were subsequently forgotten. To assess this, we first examined whether subsequent memory could be decoded from the encoding-related activity to ambiguous faces, and then, using a variety of generalization tests, asked what dis-

tinctions in the way the stimuli were processed might lead to this decoding.

We found that we could decode subsequent memory status for the ambiguous face stimuli. That is, we can decode from the signal at the time of encoding whether an ambiguous face will be subsequently remembered or subsequently forgotten (54.5%, SEM: 1.5%). Although the overall performance level is modest, it is highly consistent across observers ($t_{(23)} = 2.91$, $p = 0.008$).

What drives this difference in subsequent memory for ambiguous faces? There are at least two non-mutually exclusive possibilities: first, there is the possibility that the ability to decode subsequent memory is a result of the ambiguous face stimuli sometimes being recognized as a face versus not (as shown in the behavioral results). Thus, the ambiguous face data may be made up of some encoding events that are similar to the unambiguous face trials (e.g., where participants saw a face) and others that are similar to the non-face trials (e.g., where participants did not see the face). If this were the case, then a classifier trained to distinguish between unambiguous faces and non-faces, without regard to memory performance should successfully generalize and without further training predict subsequent memory for ambiguous faces. That is, if the classifier believes the ambiguous face was “seen as a face”, this should predict that the item was subsequently remembered. Indeed, using this approach results in a decoding accuracy of 56.7% (SEM: 1.1%) for distinguishing subsequently remembered versus non-remembered ambiguous faces, significantly above-chance ($t_{(23)} = 6.06$, $p < 0.00001$). In fact, this generalization performance was not significantly different from attempting to decode memory directly (as above, $M = 54.5\%$ accuracy; $t_{(23)} = -1.30$, $p = 0.207$).

Another possibility that could help explain subsequent memory for ambiguous faces is a general subsequent memory effect. That is, independent of whether items are recognized as faces, some items may simply be better attended or further processed during encoding, resulting in a generic subsequent memory signal. To assess this possibility we trained a classifier to predict subsequent memory for non-face stimuli. We then asked whether this classifier generalizes to predict subsequent memory for ambiguous faces. We did not find evidence consistent with a generic subsequent memory signal in decoding the single-trial EEG data (generalization performance: 51.7%, SEM: 1.9%, $t_{(23)} = 0.90$, $p = 0.38$).

Last, we examined whether there was a more stimulus-specific subsequent memory signal, e.g., a signal of how strongly a stimulus was processed that was specific to face-like stimuli. To examine this, we next trained a classifier to predict subsequent memory for unambiguous face stimuli. We then asked whether this classifier generalizes to predict subsequent memory for ambiguous faces. We found evidence consistent with this possibility (generalization performance: 56.8%, SEM 1.9%, $t_{(23)} = 3.65$, $p = 0.001$).

Overall, our decoding results suggest that ambiguous faces that are processed as faces are subsequently remembered more often than those that elicit non-face-like activity at encoding. Thus, a classifier trying only to predict whether a stimulus is a face

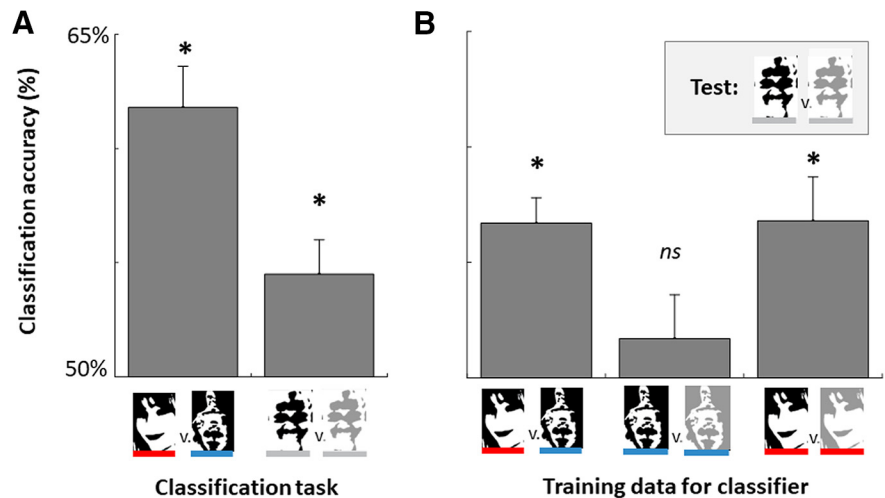


Figure 3. *A*, Decoding accuracy for a classifier applied to predicting at the time of encoding whether (left) people are seeing a face versus non-face and (right) whether they will subsequently remember an ambiguous face stimulus (i.e., decoding based on encoding activity ambiguous-remembered faces vs ambiguous-forgotten faces). *B*, Generalization tests for the classifier. To determine what predicts subsequent memory for ambiguous faces, we trained a classifier on other discriminations and asked whether it could predict subsequent memory for ambiguous faces. Left, A classifier trained to discriminate between seeing a face versus seeing a non-face predicted subsequent memory for ambiguous faces, such that images perceived as faces were more likely to be subsequently remembered. Middle, A classifier trained to discriminate subsequently remembered versus forgotten non-faces could not predict subsequent memory for ambiguous face stimuli. Right, A classifier trained to discriminate subsequently remembered versus forgotten unambiguous faces could predict subsequent memory for ambiguous face stimuli as well. * $p < 0.05$.

or non-face predicts subsequent memory for ambiguous faces even using only single-trial EEG data. In addition, we find that the same classifier that predicts the likelihood of an unambiguous face being remembered generalizes to predict the likelihood of an ambiguous face being remembered; but that the same is not true for a classifier trained to predict subsequent memory for non-faces. This is again consistent with the idea that subsequent memory for ambiguous faces is driven primarily by the strength of a face-specific response or a face-specific memory signal.

Importantly, this hypothesis is directly testable, because of the existence of a face-selective ERP component, the N170 (George et al., 2005; Hadjikhani et al., 2009; Liu et al., 2014). If the decoding of subsequent memory is largely driven by the strength of face-specific processing, we should be able to see the subsequent memory effect for ambiguous faces not only using single-trial decoding but also more selectively in the N170 component. This can help us localize the source of this decoding accuracy and pinpoint to what extent it is because of face-specific processing.

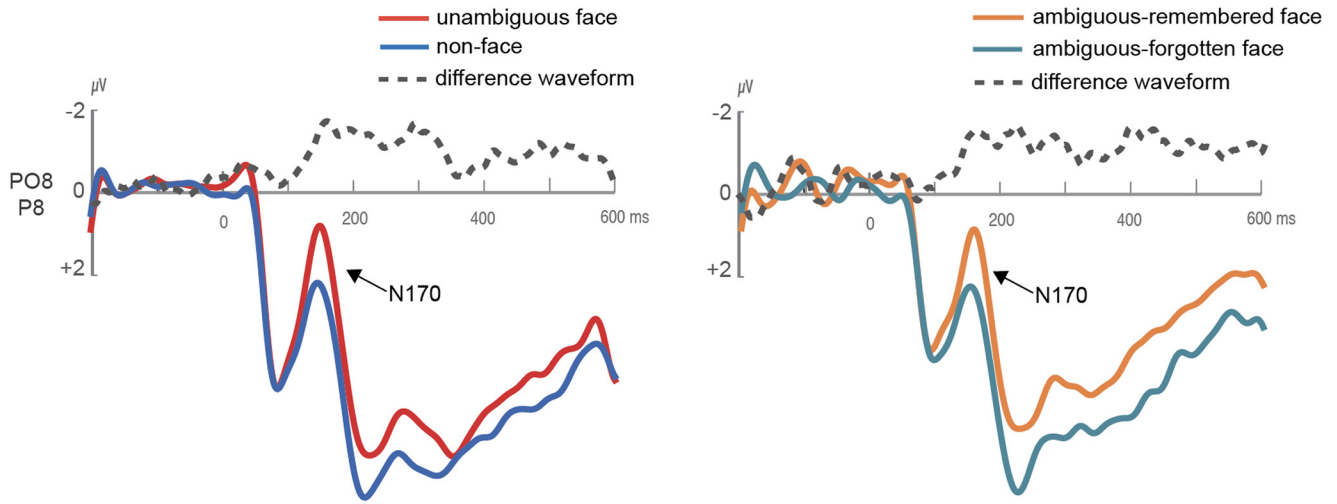
ERP results

Because the amplitude of the N170 is indicative of which images participants’ see as a face (George et al., 2005; Hadjikhani et al., 2009; Liu et al., 2014), this component provides a useful neural measure of category-specific processing for each image, allowing us to directly test the extent to which category-specific processing drives subsequent memory even with stimuli that are perceptually identical on average.

We find that the mean amplitude of the N170 elicited during the encoding period differed reliably for unambiguous faces, remembered ambiguous faces, forgotten ambiguous faces, and non-faces ($F_{(3,69)} = 4.64$, $p = 0.005$, $\eta^2 = 0.17$). As expected, the unambiguous face images elicited a larger N170 than the non-face images ($t_{(23)} = 2.54$, $p = 0.018$, $\eta^2 = 0.22$; Fig. 4A, left), verifying that the N170 is sensitive to faces even in our Mooney images.

Our main question of interest was whether the N170 amplitude differed for subsequently remembered and subsequently

A ERP waveforms during memory encoding period



B Mean N170 amplitude

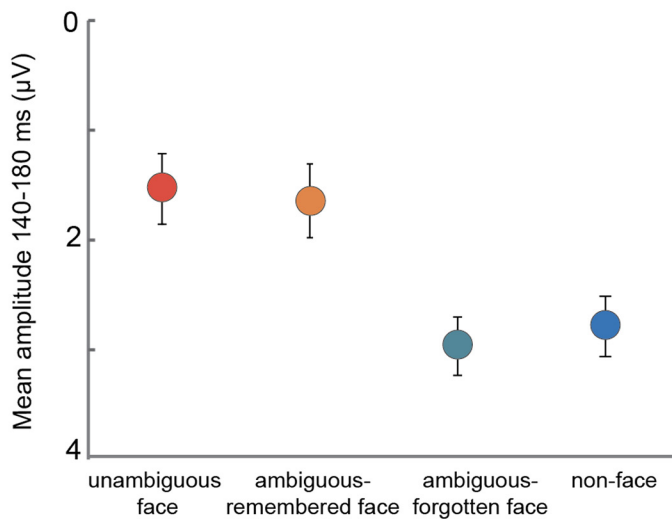


Figure 4. *A*, ERP waveforms during the memory encoding period and the differences between these waveforms. Left, The N170 is larger for unambiguous-faces relative to non-faces. Right, The N170 is also larger for ambiguous faces that were later remembered than for ambiguous faces that were later forgotten. *B*, Mean N170 amplitude for the 140–180 ms time window. Error bars represent ± 1 SEM.

forgotten ambiguous face images, even though across participants, different images were remembered or forgotten and thus the perceptual input was matched across the conditions. We found that the amplitude of the N170 did indeed predict memory performance for the ambiguous faces: For those images that participants later remembered correctly, the N170 was larger than the N170 for those images that participants later forgot ($t_{(23)} = 2.63, p = 0.015, \eta^2 = 0.23$; Fig. 4*A*, right). Furthermore, the N170 elicited by ambiguous-remembered faces was not different from that elicited by unambiguous faces ($p = 0.95$) but was larger than the N170 elicited by non-faces ($t_{(23)} = 2.32, p = 0.03, \eta^2 = 0.19$). Moreover, the N170 elicited by ambiguous-forgotten faces was not statistically different from that elicited by non-faces ($p = 0.67$) but was smaller than the N170 elicited by unambiguous faces ($t_{(23)} = 3.08, p = 0.005, \eta^2 = 0.29$).

Thus, a very strong predictor of whether an image was later remembered by a particular participant was whether that image elicited category-specific brain activity during the encoding period, as indexed by the amplitude of the N170.

We find a large effect of subsequent memory on the size of the N170 for ambiguous-face images, which can be either seen as a face or not seen as a face by particular participants. We interpret the increased N170 for subsequently remembered items as evidence that items that are seen as a face are more likely to be remembered. To ensure that the amplitude difference measured during the N170 time interval does not simply reflect a general effect of subsequent memory performance that is unrelated to seeing the stimulus as a face, we also examined the mean amplitudes for remembered and forgotten unambiguous faces and non-faces.

For the unambiguous stimuli, the N170 was larger for faces than non-faces, regardless of subsequent memory performance (main effect of stimulus type: $F_{(1,23)} = 5.83$, $p = 0.02$; $\eta^2 = 0.40$; no effect of memory accuracy: $p = 0.45$; $\eta^2 = 0.02$; no interaction: $p = 0.80$; $\eta^2 = 0.003$; Fig. 5), indicating that the N170 truly reflects whether participants see the images as a face and only predicts subsequent memory for the ambiguous-face stimuli. Bayes factors also preferred the null hypothesis in both cases (Rouder et al., 2009): The memory accuracy effect had a scaled JZS Bayes factor of 3.60:1 in favor of the null hypothesis; the interaction a Bayes factor of 4.5:1 in favor of the null hypothesis. Thus, we do not find evidence of general subsequent memory effects in the N170 data. However, this should not be taken as evidence that general subsequent memory effects do not exist; rather, it simply means that the early N170 difference is not a result of a general subsequent memory. Subsequent memory differences are often apparent at later time windows (>500 ms) that we cannot analyze in the current paradigm (Paller and Wagner, 2002). In particular, in the current study, because the images were presented sequentially one after the other, we instructed participants to blink between each image presentation (while not blinking the moment the image was presented); thus, the later time intervals after stimulus presentation have frequent eye movement artifacts, meaning the current study is not well suited to testing for general subsequent memory effect.

Relationship between the N170 and the decoding analyses

Do the decoding analyses pick up more than just the N170? Or is the N170 the only signal that plays a role in the successful classifier decoding we observe? To address this, we performed decoding analyses while excluding all times <200 ms, to ensure the only features in the classifier come from after the conclusion of the N170. We find very similar performance to the performance on the entire dataset. For example, training on the face/non-face distinction and testing on a hold-out set of face/non-face data reveals performance of 60.2% (SEM: 1.6%, $t_{(23)} = 6.44$, $p < 0.001$); similarly, training on face/non-face, and testing on ambiguous face memory reveals performance of 55.5% (SEM: 1.3%, $t_{(23)} = 4.17$, $p < 0.001$). These values are similar to the values including all of the data (61.8%, 56.7%). This suggests that there is significant information about face perception (and thus subsequent memory for ambiguous faces) outside the N170 component.

We can also ask whether doing feature selection to focus the classifier only on the N170 would result in even higher performance than using the full dataset. Training and testing a classifier on only the 140–180 ms window with only the P07/P08 electrodes, focused on using only information from the N170, gives considerably lower performance: rather than 61.8% (SEM: 1.8%) decoding of face/non-face in the full dataset, we find only 54.2% (SEM: 1.7%; $p = 0.02$) decoding in this analysis focused on the N170 alone. However, we believe this lower performance is hard to interpret, as there is much less data for the classifier to work

Mean N170 amplitude




Figure 5. Regardless of whether they are subsequently remembered, unambiguous face stimuli always evoke a strong N170, whereas non-faces do not. This suggests that N170's effect on subsequent memory is specific to the ambiguous face images, consistent with the role of meaning in visual memory. Error bars represent ± 1 SEM.

with. However, it is clear that this does not reflect a higher upper bound than the full dataset.

Thus, overall, we believe that the decoding analysis picks up more than just the N170, and that the upper bound of face/non-face performance using all of the data are not improved upon by using only the N170 time window and electrodes. This suggests that the decoding analysis and the N170 analysis are picking up on partially dissociable information.

Discussion

We examined the role of meaning in visual memory. In particular, we asked participants to remember the visual details of an image, which exact set of black and white shapes they saw, and examined whether memory performance depended on whether this information was understood by the participants and connected to a meaningful interpretation (i.e., seeing it as a face), or whether the exact same perceptual input would be equally well remembered regardless of how well understood it was. We first showed behaviorally that participants are not consistent in which particular ambiguous face stimuli they remember. Thus, any neural prediction of subsequent memory performance for ambiguous face stimuli must result purely from the role of meaningfulness or other more general encoding factors like attention, not from the perceptual properties of the stimuli or differences in encoding strategies set up by the experimenter.

Using single-trial EEG activity during stimulus encoding, we were able to predict whether a particular ambiguous face stimulus was later remembered or forgotten. Critically, using a variety of generalization tasks, we found that the decoding accuracy cannot be explained by general differences at encoding that are shared between stimuli conditions, but rather is explained by whether a particular image elicited face-like activity at encoding for that particular person. This is further confirmed by the additional analysis of the face-selective N170 component in the ERP waveform, which is unambiguously related to whether a stimulus was seen as a face or not.

Previous studies have shown that electrophysiological brain signals during stimulus encoding differ on average for items that are later remembered versus forgotten (Sanquist et al., 1980; Paller and Wagner, 2002). These differences are often apparent in amplitude or latency modulations of the ERP (Paller et al., 1987, 1988; Fukuda and Woodman, 2015) as well as distinct oscillations ranging from changes in the α , theta or gamma rhythms (Klimesch et al., 1996; Osipova et al., 2006; Hanslmayr et al., 2009; Hanslmayr and Staudigl, 2014). A few recent studies have also shown that single-trial EEG activity can predict subsequent memory by combining prestimulus and stimulus-locked activity during encoding, reaching similar decoding accuracies as in the present study (59.6%; Noh et al., 2014; Tan et al., 2015; Sun et al., 2016). However, in all of these studies, differences in encoding-related brain activity could be due to differences in the perceptual input (the same images tend to be remembered by the same people), differences in familiarity of the stimuli, distinct encoding strategies, or the attentional state of the observer at encoding (Klimesch, 2012; Dubé et al., 2013; Hanslmayr and Staudigl, 2014). In contrast, our results uniquely indicate that differential processing of the exact same stimuli, and in particular the extent to which these stimuli are processed in a category-specific and meaningful manner, predicts subsequent memory independent of other attentional or perceptual factors.

Our results fit with existing data showing that category-specific brain regions are involved in encoding information into long-term memory (Prince et al., 2009; Kuhl et al., 2012), but go beyond these data by carefully controlling for the perceptual properties of remembered and forgotten stimuli and by contrasting meaningful versus non-meaningful stimuli, as opposed to examining only memory for meaningful information. In particular, in previous data showing category-specific brain responses predict subsequent face memory, the perceptual properties of the stimuli are not matched across subsequent memory effects, and there are likely systematic differences between the faces that are remembered and those that are not. For example, same-race faces are both better remembered and provoke more fusiform face area activity during perception than other-race faces (Golby et al., 2001); and particular characteristics of faces predict which ones are better remembered in a way that is reliable across individuals (Bainbridge et al., 2013; Khosla et al., 2013). By contrast, our data show that perceiving an image as meaningful, as reflected in the category-specific brain response, rather than treating it as an unrecognizable set of shapes, is directly related to better subsequent memory, even for identical images. This result demonstrates the important role of meaning in visual memory.

Our data broadly suggest an important role for the meaning of a stimulus, rather than its perceptual properties, in explaining detailed visual memory performance (Wiseman and Neisser, 1974; Bower et al., 1975; Koutstaal et al., 2003; Konkle et al., 2010). Why does meaning play a crucial role in visual memory? It has long been argued that memory is semantically organized (e.g., in spreading activation models; Collins and Loftus, 1975), and that both elaborative encoding and building strong retrieval cues is critical to successful long-term memory. In particular, it is now thought that many items that are lost in long-term memory are likely not forgotten entirely; instead, retrieval fails because no distinctive retrieval cue allows access to the item (Wixted, 2004). A meaningful interpretation of an image may allow for a more elaborative initial encoding, creating more or stronger paths to access the memory (Bradshaw and Anderson, 1982). Thus, meaning may act as a “hook” to allow the retrieval of even visual details by creating specific retrieval cues (Konkle et al., 2010).

In terms of neural representation, meaning may play a role in so far as more neural populations can be recruited for meaningful stimuli than non-meaningful stimuli, resulting in more robust memory representations that are less limited by interference (Cohen, 1990). Because more relevant neural populations are available to support memory traces for faces and other meaningful stimuli than for nonsense blobs, these stimuli may have either more robust memory traces, or these traces may be more distinctive from each other and thus less likely to result in interference in memory. This is broadly consistent with the role of interference in long-term memory (Wixted, 2004). For example, cognitive studies have revealed that even with matched perceptual variation, categories that have a wider variety of conceptual information result in reduced interference between items in that category (Konkle et al., 2010), compatible with the idea that items from more conceptually broad categories may have more distinct memory traces.

Our data provide a neural analog and an important refinement of a classic result by Wiseman and Neisser (1974), who used a similar paradigm to the one we used in the current study. Wiseman and Neisser (1974) showed that when participants are shown faces for extended durations (5 s) and are explicitly asked to indicate whether an image is a face or not, they later remember the face images more accurately. In their study, which images were recognized as a face by different participants was not random, leaving open the possibility that the low-level visual features of the images were driving memory performance rather than solely whether they connected to meaningful interpretations for the participants. In addition, their long presentation times, required to perform the explicit face categorization task, means that subsequent memory effects could be caused by the fact that participants enjoy looking at faces more than unrecognizable shapes and so spend a larger portion of the 5 s encoding window focusing on stimuli they recognize as faces. In the current study we were able to pick up the brain response spontaneously elicited by briefly presented ambiguous-face images and to control for perceptual features in the remembered and non-remembered stimuli. Our results thus provide a much stronger test of the role of meaningfulness in subsequent memory.

Overall, the current data provide strong evidence for the involvement of meaning in visual memory, even in the absence of any perceptual differences between stimuli that evoke meaning and those that do not. Stimuli that evoke category-specific brain activity during encoding are better remembered than identical stimuli that do not evoke this brain activity. This provides a strong dissociation between perceptual factors and meaning, and suggests that even detailed visual memory is strongly supported by the meaning of a stimulus.

References

- Baddeley A, Eysenck MW, Anderson MC (2009) Memory. New York: Psychology.
- Bainbridge WA, Isola P, Oliva A (2013) The intrinsic memorability of face photographs. *J Exp Psychol Gen* 142:1323–1334. [CrossRef Medline](#)
- Bartlett FC (1932) Remembering: a study in experimental and social psychology. New York: Macmillan.
- Bentin S, Allison T, Puce A, Perez E, McCarthy G (1996) Electrophysiological studies of face perception in humans. *J Cogn Neurosci* 8:551–565. [CrossRef Medline](#)
- Bower GH, Karlin MB, Dueck A (1975) Comprehension and memory for pictures. *Mem Cognit* 3:216–220. [CrossRef Medline](#)
- Bradshaw GL, Anderson JR (1982) Elaborative encoding as an explanation of levels of processing. *J Verb Learn Verbal Behav* 21:165–174. [CrossRef](#)
- Brady TF, Konkle T, Alvarez GA, Oliva A (2008) Visual long-term memory has a massive storage capacity for object details. *Proc Natl Acad Sci* 105:14325–14329. [CrossRef](#)

- Brady TF, Konkle T, Gill J, Oliva A, Alvarez GA (2013) Visual long-term memory has the same limit on fidelity as visual working memory. *Psychological Science* 24:981–990. [CrossRef Medline](#)
- Cohen G (1990) Why is it difficult to put names to faces? *Br J Psychol* 81:287–297. [CrossRef](#)
- Collins A, Loftus E (1975) A spreading-activation theory of semantic processing. *Psychol Rev* 82:407. [CrossRef](#)
- Cox DD, Savoy RL (2003) Functional magnetic resonance imaging (fMRI) “brain reading”: detecting and classifying distributed patterns of fMRI activity in human visual cortex. *Neuroimage* 19:261–270. [CrossRef Medline](#)
- Daselaar SM, Prince SE, Cabeza R (2004) When less means more: deactivations during encoding that predict subsequent memory. *Neuroimage* 23:921–927. [CrossRef Medline](#)
- Delorme A, Makeig S (2004) EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *J Neurosci Methods* 134:9–21. [CrossRef Medline](#)
- Dubé C, Payne L, Sekuler R, Rotello CM (2013) Paying attention to attention in recognition memory: Insights from models and electrophysiology. *Psychological science* 24:2398–2408. [CrossRef Medline](#)
- Eimer M (2000) The face-specific N170 component reflects late stages in the structural encoding of faces. *Neuroreport* 11:2319–2324. [CrossRef Medline](#)
- Fukuda K, Woodman GF (2015) Predicting and improving recognition memory using multiple electrophysiological signals in real time. *Psychol Sci* 26:1026–1037. [CrossRef Medline](#)
- George N, Evans J, Fiori N, Davidoff J, Renault B (1996) Brain events related to normal and moderately scrambled faces. *Cogn Brain Res* 4:65–76. [CrossRef](#)
- George N, Jemel B, Fiori N, Chaby L, Renault B (2005) Electrophysiological correlates of facial decision: insights from upright and upside-down mooney-face perception. *Brain Res Cogn Brain Res* 24:663–673. [CrossRef Medline](#)
- Golby AJ, Gabrieli JD, Chiao JY, Eberhardt JL (2001) Differential responses in the fusiform region to same-race and other-race faces. *Nat Neurosci* 4:845–850. [CrossRef Medline](#)
- Hadjikhani N, Kveraga K, Naik P, Ahlfors SP (2009) Early (M170) activation of face-specific cortex by face-like objects. *Neuroreport* 20:403–407. [CrossRef Medline](#)
- Hanslmayr S, Staudigl T (2014) How brain oscillations form memories: a processing based perspective on oscillatory subsequent memory effects. *Neuroimage* 85:648–655. [CrossRef Medline](#)
- Hanslmayr S, Spitzer B, Bäuml KH (2009) Brain oscillations dissociate between semantic and nonsemantic encoding of episodic memories. *Cereb Cortex* 19:1631–1640. [CrossRef Medline](#)
- Höhne M, Jahanbekam A, Bauckhage C, Axmacher N, Fell J (2016) Prediction of successful memory encoding based on single-trial rhinal and hippocampal phase information. *Neuroimage* 139:127–135. [CrossRef Medline](#)
- Hollingworth A (2004) Constructing visual representations of natural scenes: the roles of short- and long-term visual memory. *J Exp Psychol Hum Percept Perform* 30:519–537. [CrossRef Medline](#)
- Hsieh PJ, Vul E, Kanwisher N (2010) Recognition alters the spatial pattern of fMRI activation in early retinotopic cortex. *J Neurophysiol* 103:1501–1507. [CrossRef Medline](#)
- Isola P, Xiao J, Torralba A, Oliva A (2011) What makes an image memorable? In: *Computer vision and pattern recognition (CVPR)*, IEEE Computer Society Conference on Computer Vision and Pattern Recognition, June 2011. Piscataway, NJ: IEEE Computer Society.
- Kanwisher N, Tong F, Nakayama K (1998) The effect of face inversion on the human fusiform face area. *Cognition* 68:B1–B11. [CrossRef Medline](#)
- Khosla A, Bainbridge WA, Torralba A, Oliva A (2013) Modifying the memorability of face photographs. *Proceedings of the IEEE International Conference on Computer Vision*, pp 3200–3207. Los Alamitos, CA: IEEE Computer Society.
- Klimesch W (2012) Alpha-band oscillations, attention, and controlled access to stored information. *Trends in cognitive sciences* 16:606–617. [CrossRef Medline](#)
- Klimesch W, Schimke H, Doppelmayr M, Ripper B, Schwaiger J, Pfurtscheller G (1996) Event-related desynchronization (ERD) and the Dm effect: Does alpha desynchronization during encoding predict later recall performance? *International Journal of Psychophysiology* 24:47–60. [CrossRef Medline](#)
- Konkle T, Brady TF, Alvarez GA, Oliva A (2010) Conceptual distinctiveness supports detailed visual long-term memory for real-world objects. *J Exp Psychol Gen* 139:558–578. [CrossRef Medline](#)
- Koutstaal W, Reddy C, Jackson EM, Prince S, Cendan DL, Schacter DL (2003) False recognition of abstract versus common objects in older and younger adults: testing the semantic categorization account. *J Exp Psychol Learn Mem Cogn* 29:499–510. [CrossRef Medline](#)
- Kuhl BA, Rissman J, Wagner AD (2012) Multi-voxel patterns of visual category representation during episodic encoding are predictive of subsequent memory. *Neuropsychologia* 50:458–469. [CrossRef Medline](#)
- Leydecker A, Biebmann F, Fazli S (2014) Single-trials ERPs predict correct answers to intelligence test questions. In: *2014 International Workshop Pattern Recognition in Neuroimaging*, pp 1–4. Piscataway, NJ: IEEE Computer Society.
- Liu J, Li J, Feng L, Li L, Tian J, Lee K (2014) Seeing Jesus in toast: neural and behavioral correlates of face pareidolia. *Cortex* 53:60–77. [CrossRef Medline](#)
- Lopez-Calderon J, Luck SJ (2014) ERPLAB: an open-source toolbox for the analysis of event-related potentials. *Front Hum Neurosci* 8:213. [CrossRef Medline](#)
- Mooney CM (1957) Age in the development of closure ability in children. *Can J Psychol* 11:219–226. [CrossRef Medline](#)
- Noh E, Herzmann G, Curran T, de Sa VR (2014) Using single-trial EEG to predict and analyze subsequent memory. *Neuroimage* 84:712–723. [CrossRef Medline](#)
- Osipova D, Takashima A, Oostenveld R, Fernández G, Maris E, Jensen O (2006) Theta and gamma oscillations predict encoding and retrieval of declarative memory. *Journal of Neuroscience* 26:7523–7531. [CrossRef Medline](#)
- Paller KA, Kutas M, Mayes AR (1987) Neural correlates of encoding in an incidental learning paradigm. *Electroencephalogr Clin Neurophysiol* 67:360–371. [CrossRef](#)
- Paller KA, McCarthy G, Wood CC (1988) ERPs predictive of subsequent recall and recognition performance. *Biol Psychol* 26:269–276. [CrossRef Medline](#)
- Paller KA, Wagner AD (2002) Observing the transformation of experience into memory. *Trends in Cognitive Sciences* 6:93–102. [CrossRef Medline](#)
- Prince SE, Dennis NA, Cabeza R (2009) Encoding and retrieving faces and places: distinguishing process- and stimulus-specific differences in brain activity. *Neuropsychologia* 47:2282–2289. [CrossRef Medline](#)
- Rouder JN, Speckman PL, Sun D, Morey RD, Iverson G (2009) Bayesian *t* tests for accepting and rejecting the null hypothesis. *Psychon Bull Rev* 16:225–237. [CrossRef Medline](#)
- Sanquist TF, Rohrbaugh JW, Syndulko K, Lindsley DB (1980) Electrocortical signs of levels of processing: perceptual analysis and recognition memory. *Psychophysiology* 17:568–576. [CrossRef Medline](#)
- Schwiedrzik CM, Melloni L, Schurger A (2018) Mooney face stimuli for visual perception research. *PLoS One* 13:e0200106. [CrossRef Medline](#)
- Sun X, Qian C, Chen Z, Wu Z, Luo B, Pan G (2016) Remembered or forgotten? An EEG-based computational prediction approach. *PLoS One* 11:e0167497. [CrossRef Medline](#)
- Tan ZHE, Smitha KG, Vinod AP (2015) Detection of familiar and unfamiliar images using EEG-based brain-computer interface. In: *2015 IEEE International Conference on Systems, Man, and Cybernetics*, pp 3152–3157. Los Alamitos, Calif: IEEE Computer Society.
- Wagner AD, Schacter DL, Rotte M, Koutstaal W, Maril A, Dale AM, Rosen BR, Buckner RL (1998) Building memories: remembering and forgetting of verbal experiences as predicted by brain activity. *Science* 281:1188–1191. [CrossRef Medline](#)
- Wiseman S, Neisser U (1974) Perceptual organization as a determinant of visual recognition memory. *Am J Psychol* 87:675–681. [CrossRef Medline](#)
- Wixted JT (2004) The psychology and neuroscience of forgetting. *Annu Rev Psychol* 55:235–269. [CrossRef Medline](#)