Behavioral/Systems/Cognitive

Implicit Perceptual Anticipation Triggered by Statistical Learning

Nicholas B. Turk-Browne,[1] **Brian J. Scholl,**[2] **Marcia K. Johnson,**[2] **and Marvin M. Chun**[2]

[1]Department of Psychology, Princeton University, Princeton, New Jersey 08540, and [2]Department of Psychology, Yale University, New Haven, Connecticut 06520

Our environments are highly regular in terms of when and where objects appear relative to each other. Statistical learning allows us to extract and represent these regularities, but how this knowledge is used by the brain during ongoing perception is unclear. We used rapid event-related fMRI to measure hemodynamic responses to individual visual images in a continuous stream that contained sequential contingencies. Sixteen human observers encountered these statistical regularities while performing an unrelated cognitive task, and were unaware of their existence. Nevertheless, the right anterior hippocampus showed greater hemodynamic responses to predictive stimuli, providing evidence for implicit anticipation as a consequence of unsupervised statistical learning. Hippocampal anticipation based on predictive stimuli correlated with subsequent processing of the predicted stimuli in occipital and parietal cortex, and anticipation in additional brain regions correlated with facilitated object recognition as reflected in behavioral priming. Additional analyses suggested that implicit perceptual anticipation does not contribute to explicit familiarity, but can result in predictive potentiation of category-selective ventral visual cortex. Overall, these findings show that future-oriented processing can arise incidentally during the perception of statistical regularities.

Introduction

While sensory input is complex and dynamic, it also contains regularities that have shaped sensory systems over phylogenetic time (Simoncelli and Olshausen, 2001). But regularities are also pervasive over more local timescales, in the sense that certain stimuli repeatedly precede, follow, or co-occur with other stimuli. By detecting and representing these regularities through statistical learning, we can parse complex sensory information into useful chunks, such as words, scenes, and events.

Statistical learning refers to an unconscious process by which regularities are automatically segmented from continuous environments, where the only cues for segmentation are statistics of co-occurrence between specific stimuli. This form of learning was first reported in studies of word learning: after passively listening to a brief syllable stream that contained unexpected repeated syllable subsequences ("words"), infants and adults could discriminate between words and nonwords based on transitional probabilities alone (Saffran et al., 1996; Aslin et al., 1998). Such learning also occurs for visual regularities consisting of repeated spatial and temporal configurations of shapes, locations, motions, and actions (Chun and Jiang, 1998; Fiser and Aslin, 2001, 2002; Olson and Chun, 2001). Statistical learning differs from other forms of learning because regularities are not presegmented into discrete trials (cf. paired associate learning), the underlying structure is stimulus specific (cf. artificial grammar learning), and the learning is automatic and incidental (cf. some varieties of category learning). These features allow statistical learning to operate over naturalistic sensory input without engaging deliberate/conscious effort.

Statistical learning is ubiquitous, and most research has focused on defining the scope of stimuli over which it operates. But what are the behavioral and neural consequences of statistical learning? The general belief has been that statistical learning produces memories of indivisible higher-order chunks (Orbán et al., 2008), such that later familiarity depends on complete input patterns (Fiser and Aslin, 2005; Turk-Browne et al., 2008). However, a long history of research suggests that organisms are not passive detectors, but rather use partial input to actively form expectations about the future (Tolman, 1932). Given that statistical learning can result in stimulus–stimulus associations (Turk-Browne and Scholl, 2009), the brain may use partial cues to implicitly anticipate upcoming perceptual events.

The primary brain system responsible for associative processing is the medial temporal lobe (Cohen and Eichenbaum, 1993). In particular, the hippocampus is involved in relational memory encoding (Mitchell et al., 2000; Davachi and Wagner, 2002; Chua et al., 2007) and several forms of implicit learning (Chun and Phelps, 1999; Schendan et al., 2003; Harrison et al., 2006), including statistical learning (Turk-Browne et al., 2009). The hippocampus may participate in the acquisition of statistical regularities, for example, by binding elements of events (Howard et al., 2005; Jensen and Lisman, 2005). Here we explore how the hippocampus and other regions may also participate in the expression of learning by reconstructing and anticipating the future based on learned associations (Marr, 1971; McClelland et al.,

1995; Norman and O'Reilly, 2003). Such anticipation would suggest that future-oriented processing can arise incidentally during the perception of statistical regularities.

Materials and Methods

Overview

We conducted an event-related fMRI study to explore implicit perceptual anticipation as a potential consequence of statistical learning, distinct from deliberate guessing or planning. fMRI is an ideal tool for studying anticipation, since we can monitor for anticipatory responses in the brain while observers perform an unrelated behavioral task. Previous fMRI studies of this type of statistical learning have examined responses to regularities, but could not examine anticipation because responses were collapsed across entire blocks or runs (McNealy et al., 2006; Turk-Browne et al., 2009). We reasoned that items at the beginning of predictive temporal regularities should engage anticipatory processes relative to items that fuel no prediction.

Participants viewed color photographs one at a time, and were required to make a categorical response to each one. Unbeknownst to them, the trial sequence in each run was constructed from four pairs of images, as well as four single images that were neither reliably predicted by the prior image nor predictive of the next image (Fig. 1). Within each pair, the Paired images were always presented consecutively and in the same order: the First image followed by the Second image. The pairs were randomly sequenced with Unpaired images, and the interval between individual images—both within and between pairs—was jittered and orthogonal to the pair structure. Thus, pairs existed only in terms of the greater transitional probabilities between Paired images. To assess anticipation, we compared hemodynamic responses for First images (that reliably predicted the next image) versus Unpaired images (that did not). Any differential responses necessarily reflect (un)certainty about future images, rather than surprise associated with the current image, since neither kind of image was itself predictable.

Participants

Sixteen naive observers (8 females; mean age: 24) participated in one fMRI session for monetary compensation. All were right handed with normal or corrected-to-normal vision.

Procedure

Stimuli and apparatus. Images were color photographs of male/female faces and indoor/outdoor scenes. They were displayed on a projection screen at the back of the scanner bore, viewed with a mirror attached to the head coil. Stimuli subtended ~15.9 × 15.9°. Each image was presented for 200 ms to discourage eye movements, and fixation was further aided by superimposing a dot on the center of all images and by cropping face stimuli such that the eyes were roughly centered in the image. Behavioral responses were collected with an MRI-compatible fiber-optic button box.

Trial sequence. Twelve novel images were used in each run, with three unique exemplars each of male faces, female faces, indoor scenes, and outdoor scenes. To introduce statistical regularities in the trial sequence, eight of these images were grouped into pairs: two pairs consisted of a particular face preceding a particular scene, and the other two consisted of a particular scene preceding a particular face. Based on pilot testing, we decided to use two pairs of each type rather than fully crossing category
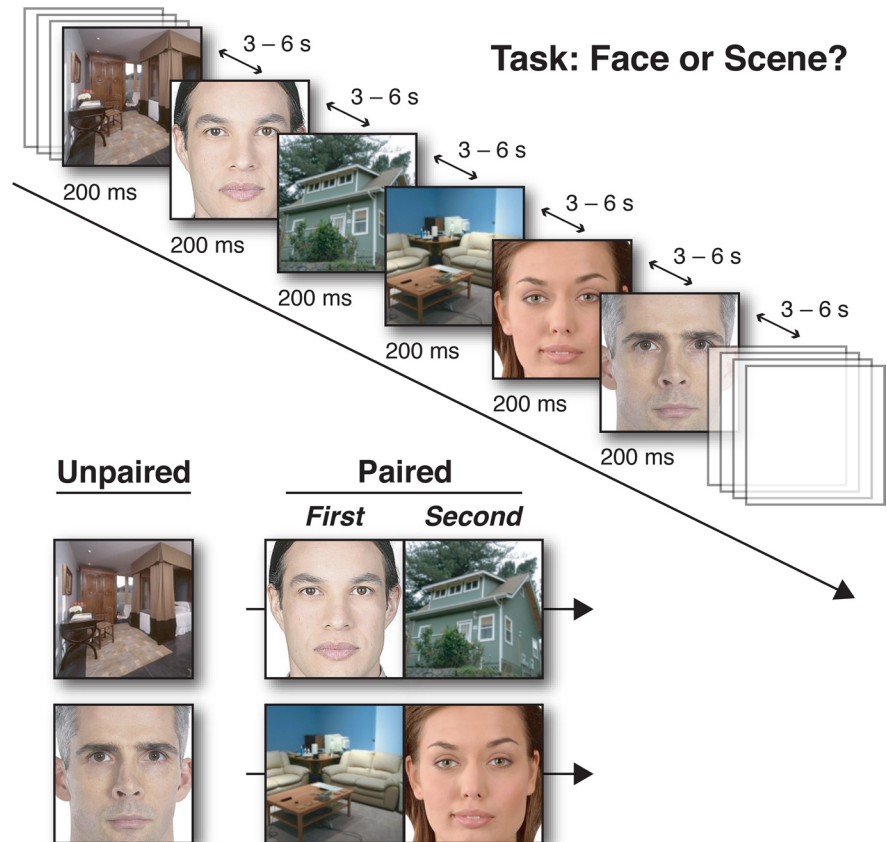
Figure 1. Experimental protocol. Portion of a trial sequence containing two example Pairs and two example Unpaired images. Paired images always appeared in the same ordered pair, but that pair was preceded and followed by other images at different times in the run. Participants responded "face" or "scene" to each image.

pairs to preserve statistical power (still only 12 trials of each condition per run); additionally, because of the unpaired images, the category of an upcoming stimulus could not be reliably predicted based on the category of the current stimulus, and results from the familiarity test suggest that participants were learning patterns of specific exemplars. The remaining four images (two faces and two scenes) were not paired. The sequence of images in each run was generated from six repetitions of each pair of images and each unpaired image (for a total of 72 trials/run) presented in a random order, with the sole constraint that the same pair or image could not be repeated back-to-back. The runs of four participants were longer and contained additional manipulations that we are pursuing separately. However, each of these runs began with the same full run of 72 trials that all other participants experienced, and thus were truncated during analysis such that the runs of all participants were structured identically.

It is worth explicitly noting two aspects of this experimental design. The first is that we tested only two levels of predictiveness in assessing anticipation: First images that deterministically predicted which exact image appeared next versus Unpaired images that weakly predicted which one of several possible images appeared next. This dichotomous treatment of transitional probabilities was used because this design is conventional across several previous behavioral investigations of statistical learning (e.g., Saffran et al., 1996; Fiser and Aslin, 2002; Kirkham et al., 2002; Turk-Browne et al., 2005; Brady and Oliva, 2008; Orbán et al., 2008). Since this was an initial rapid event-related fMRI study of this type of statistical learning, we designed our stimuli to connect directly with prior work. The second aspect of the design worth noting is that pairs were repeated only six times, because we assume that statistical learning can occur very quickly. Prior work demonstrates the remarkable speed of statistical learning (Turk-Browne et al., 2009), with evidence of learning emerging after only 2–3 repetitions of a regularity. Moreover, the small

number of repetitions allowed us to conduct several fMRI runs with pairs constructed from different images, both increasing our power and diminishing the likelihood of any item effects on learning.

Pairs were used to establish the image sequence, but only one image was presented at a time with images separated by jittered intervals of 3, 4.5, or 6 s. These intervals were sampled randomly from a distribution (50% 3 s, 30% 4.5 s, 20% 6 s), which helped psychologically in reducing the generic predictability of trial onsets, and methodologically in allowing the statistical separation of blood-oxygen level-dependent (BOLD) responses during analysis. It is worth noting that the pair structure could introduce correlations between the onsets of First and Second conditions. To this end, before each scan, optimized runs were generated by creating random jittered trial sequences, convolving these designs with a hemodynamic response function (HRF), and iterating until the pairwise correlations between any two regressors reached acceptable levels (all r values <0.3). Any residual collinearity could only hurt our statistical power by increasing the variability of parameter estimates, and even in such cases the difference between conditions (rather than their contribution relative to baseline) can be estimated efficiently.

Task. Participants were told that we were studying how the brain processes different types of images. On each trial, they responded as quickly and accurately as possible as to whether the image was a face or a scene by pressing one of two buttons using the index and middle fingers of their dominant right hand. Responses slower than 3 SDs above the run mean were excluded from the response time analysis. In addition, one run from one participant was excluded from analysis due a high proportion of missed responses and low overall accuracy (81%), indicating sleepiness; accuracy on this simple task in all other runs from this and all other participants was $>94\%$. In analyzing behavioral performance, we conducted two planned comparisons, First versus Unpaired (to assess the anticipation factor) and Unpaired versus Second (to assess the predictability factor). Before scanning, participants completed a short practice run that contained no pairs, and images from this run were not reused. After the anatomical scans, they completed five runs of the main task that lasted 316.5 s each.

Familiarity test. Following the last run of the main task, participants completed a familiarity test for the pairs presented in that run. Only these last run pairs were tested due to massive interference (and perhaps decay) for pairs from earlier runs, and time constraints in the scanning session. On every test trial, participants were presented with two images, and were required to judge whether the sequence of images was familiar by responding "old" or "new" with a button press. Half of trials contained a pair from the run that had been repeated six times, and the other half of trials contained a foil of two images that had never appeared sequentially in the last run, but that had individually been repeated six times as well. Critically, the foils were constructed by swapping the Second image between the two pairs of the same type: in other words, if Face1–Scene1 and Face2–Scene2 were pairs during the scanning run, then they were tested alongside the foils Face1–Scene2 and Face2–Scene1. All pairs and foils were tested an equal number of times to avoid any differential contribution to familiarity from the test items themselves. Since the categorical structure of the pairs was preserved for the foils and the novelty of individual images was equated, any ability to discriminate pairs from foils reflects statistical learning of the pair exemplar relations. We assessed discriminability by comparing hits (old responses to pairs) versus false alarms (old responses to foils) using A' (Grier, 1971; Aaronson and Watts, 1987). Each pair and foil was tested four times (32 trials total).

Localizer. Following the familiarity test, participants completed one run of a functional localizer. Blocks of 12 novel faces or scenes alternated every 24 s, with block order counterbalanced across participants. Participants judged whether faces were male or female and whether scenes occurred indoors or outdoors by pressing one of two buttons. Each image was presented for 200 ms with a 1500 ms stimulus onset asynchrony, resulting in 18 s blocks followed by 6 s of fixation. There were eight blocks of each type, and the run lasted 402 s.

Data acquisition

Neuroimaging data were collected on a 3T Siemens Trio scanner using an eight-channel head coil. Functional data were acquired with a T2*-weighted gradient-echo EPI sequence (TE = 25 ms; TR = 1500 ms; FA = 90°; matrix = 64×64). Covering the whole brain, 26 oblique axial slices aligned parallel to the anterior commissure/posterior commissure line (3.5×3.5 mm in-plane, 5 mm thickness) were acquired in an interleaved order. Each main task run contained 211 volumes, and the localizer run contained 268 volumes. Anatomical data consisted of two T1-weighted sequences: a coplanar FLASH sequence and a high-resolution 3D MPRage sequence.

Data analysis

Preprocessing. The first six volumes of each functional run were discarded to allow for T1 equilibration. Using Brain Voyager QX (Brain Innovation), data were then corrected for slice acquisition time, corrected for head motion, spatially smoothed (8 mm FWHM Gaussian kernel), detrended, and high-pass filtered with 128 s period cutoff. Functional runs were then registered to the coplanar anatomical scan, which was in turn registered to the high-resolution anatomical scan; data were normalized into Talairach space, and interpolated to 3 mm isotropic voxels.

Whole-brain analysis. To analyze our main task runs we used a summary statistic random effects approach. At the first (within-subjects) level, we estimated parameters encoding condition-specific activations. This involved specifying stimulus functions for each trial type, which were then convolved with a canonical hemodynamic response function to form regressors. Separate regressors were entered for faces and scenes, leading to a total of six regressors: First-face, First-scene, Second-face, Second-scene, Unpaired-face, and Unpaired-scene. For the whole-brain contrasts, estimates were collapsed across face and scene types. Although pairs could only be learned after one presentation, we included all trials to equate the number of item repetitions across conditions (this decision was conservative in that it could only hurt our chances of observing differences). See below for exploratory parametric analyses that modeled monotonic changes as a function of stimulus repetition and condition. Data were normalized with a percentage change transform by subtracting and dividing by the mean, and thus parameter estimates correspond to the percentage signal change in the BOLD response for each condition in every voxel for each participant.

The parameter estimates were then taken to a second (between-subjects) level for group inference. To assess regionally specific contrast effects, we performed one-sample t tests over linear combinations of the parameter estimates for different conditions across subjects. For the primary contrast of First > Unpaired, voxels were judged significant if their t value reflected $p < 0.001$ (two-tailed) and they were part of a cluster of least five contiguous significant voxels. This cluster size threshold was calculated based on Forman et al. (1995), but improved to take into account smoothness of 3D statistical maps (Goebel et al., 2006). In particular, the spatial smoothness of each group t test or correlation was computed separately, and we then performed Monte Carlo simulations by iteratively generating random maps, injecting the same spatial smoothness, thresholding at the predetermined statistical level, and identifying the number and size of clusters. To select a cluster threshold, we chose the cluster size for which the proportion of iterations containing at least one cluster of that size or larger was <0.05; this resulted in a corrected α rate of 0.05 at the cluster level. The simulations also provide an estimate of the corrected p values for each cluster size, which are reported in the main text for each cluster along with the center-of-mass Talairach coordinates.

We did not extract mean parameter estimates from clusters for further comparisons, since the manner in which the regions were defined was not independent of such comparisons (as all comparisons relied on the Unpaired baseline condition). We also did not perform contrasts between First and Second images because the results would be uninterpretable. In particular, we designed our experiment to be able to isolate anticipation induced by the First images. The Second images do not provide an appropriate baseline for testing anticipation because, while themselves not predictive, they are predictable. Thus, the comparison of First versus Second images would conflate predictiveness with predictability, and it would be impossible to interpret any findings. The appro-

priate baseline is provided by the Unpaired images, which were neither predictive nor predictable.

Correlation analyses. To assess relationships between indices of learning, we used correlation coefficients at the second (between-subjects) level. We conducted four analyses to explore the relationship between neural anticipation and subsequent facilitation (behavioral priming, neural priming, and behavioral familiarity). The first analysis was ROI based: we examined the relationship between anticipation scores (First–Unpaired parameter estimate difference) from the hippocampal region obtained in the group contrast described above and behavioral priming scores (Unpaired–Second RT difference) across individuals. The analysis assessed whether individual differences in hippocampal anticipation predicted individual differences in behavioral priming. The second analysis was conducted across the whole brain: we explored whether neural anticipation scores (First–Unpaired parameter estimate difference) in brain regions other than the hippocampus predicted an individual's behavioral priming score (Unpaired–Second RT difference). This brain-behavior correlation helped assess which anticipation signals were directly related to facilitated response times. The third correlation analysis involved a combined ROI and whole-brain approach: we explored whether the anticipation scores from the hippocampus—as used in the first analysis above—predicted an individual's neural priming score (Unpaired–Second parameter estimate difference) in any brain area. This brain–brain correlation helped assess which regions may be influenced by the hippocampus in terms of their subsequent processing of predicted items. Finally, a fourth analysis was also conducted across the whole brain: we examined whether behavioral familiarity scores (A' on the familiarity test) could be predicted from neural anticipation and/or neural priming scores obtained during the earlier face/scene task. This correlation examined how familiarity judgments about statistical regularities—the canonical measure of statistical learning—relate to the extent of an individual's anticipation based on First items and/or priming based on Second items. Because our sample size was small for detecting reliable correlations, we used a more liberal voxel threshold of $p < 0.005$ ($r > 0.66$) for all of the correlation analyses, but compensated with a larger minimum cluster size, which resulted in an actual α rate of 0.05 (with cluster correction computed independently for each correlation map due to differences in smoothness).

Parametric analyses. To examine monotonic changes in the brain as a result of statistical learning, we implemented a modified general linear model (GLM) that explicitly coded for an interaction between condition and time. For each of the six existing regressors in our primary model, we added a new regressor representing the parametric modulation of that regressor by weighting the predicted hemodynamic response on each trial by the number of times that the image had been presented before (including the current instance). To capture the possibility that changes could be linear or nonlinear, we modeled monotonic changes using one of two functions: (1) an increasing linear function in which the repetition numbers were used directly as the weights [1, 2, 3, 4, 5, 6], and (2) an increasing logarithmic function in which the natural logs of the repetition numbers were used as the weights [0, 0.6931, 1.0986, 1.3863, 1.6094, 1.7918]. These new regressors were included in two separate models (one for linear, one for logarithmic), both containing the existing unmodulated regressors that represented the main effect of each condition, with every trial equally weighted. Following standard practice (e.g., Büchel et al., 1998), the modulated regressors were orthogonalized with respect to the unmodulated regressors. Because the same linearly and logarithmically modulated regressors were fit for each subject and condition, we contrasted the resulting parameter estimates in the same manner as was done for the main analysis of anticipation.

ROI definition. We also examined BOLD responses in two a priori regions of interest (ROIs) identified in each participant from the functional localizer: the fusiform face area (FFA) (Kanwisher et al., 1997; McCarthy et al., 1997) and the parahippocampal place area (PPA) (Aguirre et al., 1998; Epstein and Kanwisher, 1998). A GLM including regressors for the Face and Scene blocks was used to model activation during the localizer. The contrast of Face versus Scene was used to define bilateral FFA and PPA. For each region and hemisphere, the voxel with the greatest t value in an anatomically restricted search was selected as the center

Face/Scene Task
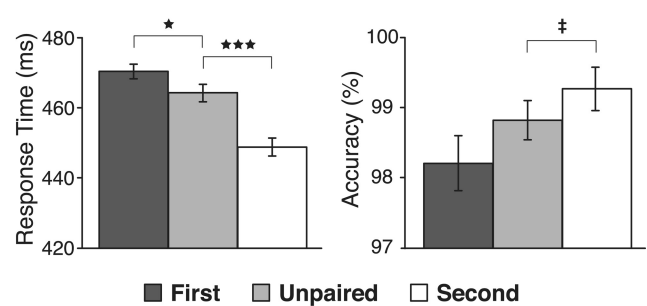
Figure 2. Task performance. Response time and accuracy of face/scene responses. Error bars reflect one within-subject SEM. ***$p < 0.001$, *$p < 0.05$, ‡$p < 0.1$.

of a 4 mm sphere ROI if it reached at least $p < 0.001$ uncorrected. Responses were collapsed across hemispheres for all subsequent analyses.

Results

Behavioral data

Online measures

Performance on the orthogonal categorization task was examined in terms of response time (RT) and accuracy (Fig. 2). If statistical learning of the predictive associations had taken place—and thus the Second image in a pair could be anticipated—then performance on Second trials should be facilitated. This pattern was observed in RT for Second versus Unpaired trials (449 ms vs 464 ms, respectively; $t_{(15)} = 4.73$, $p = 0.0003$), and could not be explained by speed–accuracy tradeoffs (99.27% vs 98.83%, respectively; $t_{(15)} = 1.77$, $p = 0.10$). In contrast with most statistical learning studies, these results provide robust online evidence of learning (see also Hunt and Aslin, 2001; Baker et al., 2004; Harrison et al., 2006; Abla et al., 2008; Abla and Okanoya, 2009; Turk-Browne et al., 2009). A second behavioral effect was entirely novel: RTs for First trials were slowed compared to Unpaired trials (470 ms vs 464 ms, respectively; $t_{(15)} = 2.61$, $p = 0.02$); again, this did not result from speed–accuracy tradeoffs (98.22% vs 98.83%, respectively; $t_{(15)} = 1.29$, $p > 0.2$). This slowing could be explained if First images engaged anticipatory processes that competed with the determination and report of the First image's category. This effect could only be observed because we used an Unpaired baseline; without such a baseline, differential responses to the First versus Last item in a pair or triplet (Abla et al., 2008; Abla and Okanoya, 2009) could reflect either costs or benefits (or both).

Offline measures

To mirror past studies, participants completed a familiarity test after the last run. On every test trial, participants were presented with two images sequentially, and judged whether they had previously appeared in that order by responding "old" or "new." Trials were equally likely to contain an actual pair or a foil consisting of two old images recombined into a novel order. Foils were constructed such that they could only be distinguished from pairs based on learning of image exemplar pairs. Participants reliably discriminated pairs from foils (mean $A' = 0.61$; $t_{(15)} = 2.28$, $p = 0.04$), providing additional evidence of statistical learning.

Debriefing responses

To help assess the implicitness of learning, participants were asked five questions outside the scanner: What do you think the experiment was about? Did you use any particular strategy? How do you think you did in the familiarity test? Have you done an experiment like this before? Did you notice any repeating pat-

First > Unpaired (Anticipation)

A   x = 30   R   y = -8

Anterior Hippocampus
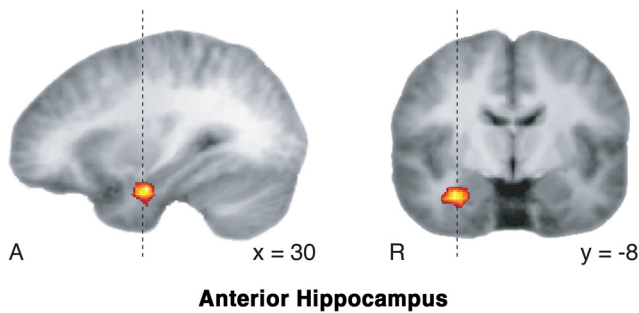
Figure 3. Implicit perceptual anticipation. Greater response to trials consisting of First versus Unpaired images in right anterior hippocampus ($p < 0.05$ corrected), displayed on participants' average high-resolution anatomical scan. No other clusters reached significance at this threshold.

terns during the face/scene task? Most relevantly, no participant reported being aware of the extensive pair structure in the main task. After being told about how the runs were constructed, 10 participants reported not being aware of any repeating patterns until at least the familiarity test (when they were asked to judge pairs); of the remaining group, the two participants who reported noticing the most pairs (6 pairs, 3 pairs) performed at or below chance on the familiarity test ($A' = 0.50$ and 0.44, respectively). Moreover, when asked whether they had used any particular strategies during the face/scene task, no participant reported attempting to predict the next stimulus. These responses suggest that participants did not engage in explicit anticipation—and indeed that the large majority did not even realize that this was a possibility.

fMRI data

Anticipation
Using a jittered rapid event-related design (where each image was an event), we compared BOLD responses for First images that reliably predicted the next image to Unpaired images that did not predict the next image. Importantly, First and Unpaired conditions were identical in all other respects: they contained equal numbers of faces and scenes, the images were individually repeated an equal number of times, and in neither case could the image itself be predicted from the preceding trial's image.

The right anterior hippocampus (center of mass in Talairach coordinates, 30, −8, −18; corrected significance of cluster, $p < 0.002$) responded more strongly to First than Unpaired trials (Fig. 3), and was the only region that responded differentially to these two conditions in either direction. This finding suggests that the hippocampus helps to mediate a form of implicit perceptual anticipation. We describe these effects as anticipation without assumptions about the underlying mechanism, although we do consider some possibilities in the Discussion. These findings provide converging evidence that the hippocampus is involved in prospective processing during online perception, separate from the explicit and deliberative tasks used in previous studies of prospection (Addis et al., 2007). It is unclear why this effect was specific to the right hippocampus, but note that other studies of implicit learning have found stronger, if not selective, effects in the right medial temporal lobe (Rose et al., 2002; Henke et al., 2003; Turk-Browne et al., 2009).

Implicit learning may depend on medial temporal lobe regions beyond the hippocampus proper (Manns and Squire, 2001). While no such regions were apparent at our a priori sta-

tistical threshold, we examined this possibility in an exploratory analysis with a more liberal threshold (Table 1). We observed a handful of additional regions with First > Unpaired and Unpaired > First. Of particular note, the hippocampal cluster described above grew larger and extended into the right perirhinal cortex, and a new cluster emerged in right inferior temporal cortex; both regions have been implicated in visual associative learning in primates (Miyashita, 1993; Erickson and Desimone, 1999). In addition, a region of medial frontal/orbitofrontal cortex emerged from this analysis; this region has been implicated in the rapid generation of predictions that can constrain posterior object recognition processes (Bar et al., 2006), but may also be involved in more general associative prediction (Bar, 2007) and reward prediction (Knutson and Cooper, 2005; Schultz, 2006).

In addition to the primary anticipation contrasts, we also conducted a whole-brain analysis of Unpaired > Second and Second > Unpaired contrasts. Only one region showed a significant effect at the $p < 0.05$ cluster-corrected threshold: left postcentral gyrus (−12, −38, 68; $p < 0.02$), which exhibited a stronger response to Unpaired versus Second items. This main effect, combined with the behavioral priming for Second items and several correlational and ROI results described below, suggests that anticipation has functional consequences for the processing of predicted items.

Anticipation correlations
In the previous analyses, we used *t* tests to find reliable activations across subjects. In what follows, we describe analyses of correlations between brain responses and various indices of learning across subjects. To explore the consequences of anticipation for behavior, we examined the relationship between neural anticipation based on predictive images and the subsequent response time facilitation when the predicted image appeared (behavioral priming). Specifically, across the whole brain, we correlated First–Unpaired parameter estimate differences (neural anticipation scores) with Unpaired–Second RT differences (behavioral priming scores) across participants. Several regions emerged from this analysis: right inferior intraparietal sulcus (26, −81, 35; $p < 0.001$ corrected), right precuneus (1, −59, 58; $p < 0.002$ corrected), right paracentral gyrus (−6, −40, 53; $p < 0.002$ corrected), left middle temporal gyrus (−58, −56, 2; $p < 0.02$ corrected), left middle occipital gyrus (−41, −83, 7; $p < 0.001$ corrected), and left cerebellum (−34, −75, −28; $p < 0.001$ corrected); two of these regions are shown in Figure 4A. In all cortical regions, greater neural anticipation was associated with greater behavioral priming (a positive correlation); a negative correlation was observed in the cerebellum.

The hippocampal region that showed a main effect of First > Unpaired did not correlate with behavioral priming across participants either in a voxelwise manner or in a focused ROI analysis ($r = 0.007$). The fact that the hippocampus engages in implicit anticipation but does not correlate with subsequent behavioral priming is consistent with the suggestion that it subserves generic revival of perceptual information based on associations, but that the fidelity with which the revived output is represented in posterior regions determines the extent of priming (Cabeza et al., 2008). Although we did not observe a direct link between anticipation in the hippocampus and subsequent behavioral priming, below we explore potential mediators of the effects of the hippocampus on behavior with the goal of constraining future models of these interactions.

Given the lack of a direct relationship to behavioral priming, we explored whether anticipation in the hippocampus may instead cor-

Table 1. Anticipation regions

Region	Hemisphere	BA	x	y	z	Extent (voxels)	Peak t
First > Unpaired							
Hippocampus/Medial temporal lobe	R		31	−8	−19	78	6.86
Inferior temporal cortex	R	20	44	−27	−14	12	4.14
Medial frontal gyrus	L	25	−4	11	−18	11	4.43
Unpaired > First							
Superior temporal gyrus	L	41/22	−53	−28	7	31	−4.54
Postcentral gyrus	R	3/4	12	−34	72	23	−4.50

Clusters resulting from exploratory analysis of First versus Unpaired ($p < 0.005$ voxel threshold; $p < 0.05$ cluster corrected). Coordinates are in Talairach space. L, Left; R, right.

relate with subsequent neural processing of predicted items in other brain regions. Specifically, we correlated First–Unpaired parameter estimate differences from the anterior hippocampal region discussed above (hippocampal anticipation scores) with Unpaired–Second parameter estimate differences (neural priming scores) in every voxel across participants. Two regions emerged from this correlation analysis (Fig. 4B): midline early visual cortex ($0, -86, -3; p < 0.002$ corrected), and left inferior parietal lobule ($-41, -69, 35; p < 0.05$ corrected). Note that since the former region was obtained at the group level, and since we did not conduct retinotopic mapping, there is no way to assess which precise visual areas exhibited this effect.

Interestingly, these regions had opposite relationships to the hippocampus: participants with greater hippocampal anticipation showed relatively reduced activation to Second images in parietal cortex and relatively greater activation to Second images in early visual cortex. These findings suggest that anticipation may reduce the need for top-down modulation of active representations from parietal cortex (Johnson et al., 2007), while enhancing visual extraction of information (Moores et al., 2003; J. J. Summerfield et al., 2006). The enhancement in early visual cortex was accompanied by category-specific effects, as described below. Overall, the fact that hippocampal anticipation correlated with both increased and decreased responses to predicted items across regions suggests that anticipation can have specific consequences for different processes, as opposed to causing a global change in readiness or arousal.

Familiarity correlations
The vast majority of studies of statistical learning involve a single measure of learning—offline tests of familiarity—which occur after the opportunity for statistical learning has passed (Fiser and Aslin, 2001, 2002; Turk-Browne et al., 2008). A handful of studies have used implicit measures based on response times (Chun and Jiang, 1998; Olson and Chun, 2001; Baker et al., 2004; Turk-Browne et al., 2005; Turk-Browne and Scholl, 2009) or neural responses (Abla et al., 2008; Abla and Okanoya, 2009; Turk-Browne et al., 2009). These studies have nevertheless observed patterns of results similar to when familiarity measures were used (Turk-Browne et al., 2005; Turk-Browne and Scholl, 2009), al-
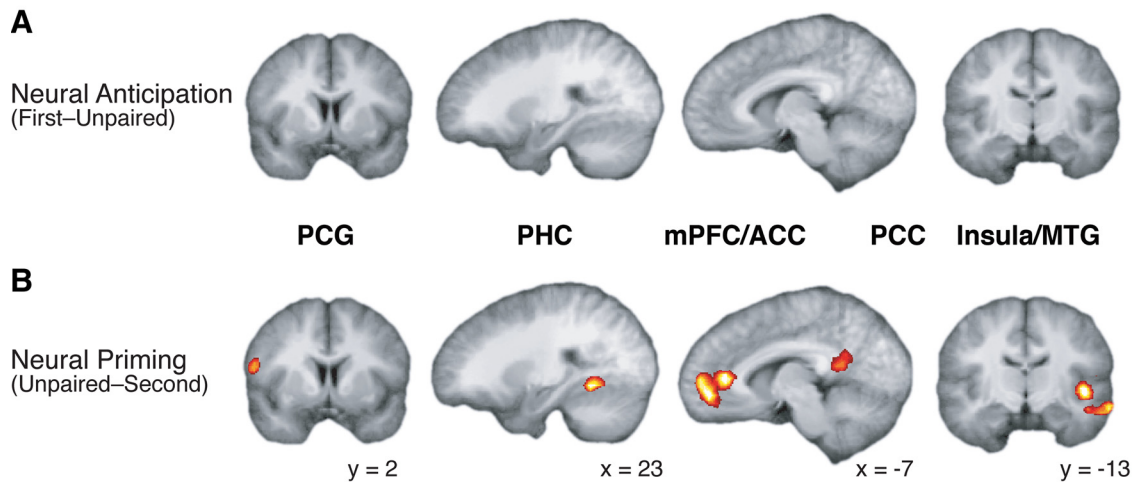
Anticipation Correlations
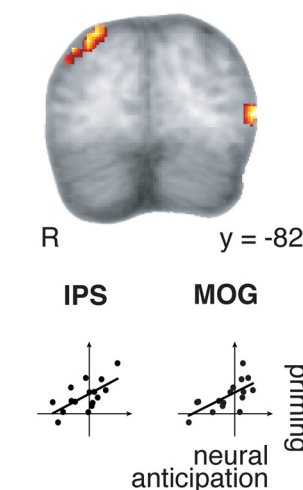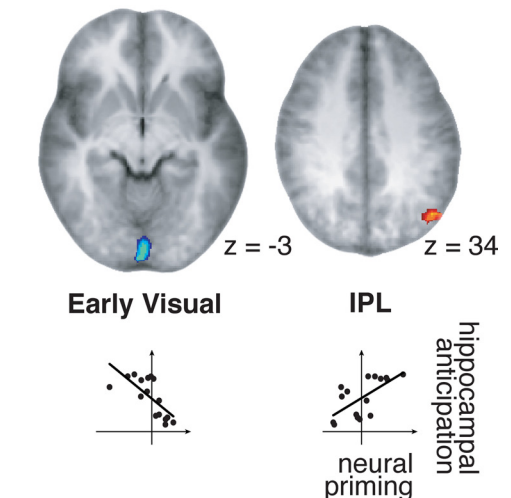
Figure 4. Consequences of anticipation. *A*, Correlations between neural anticipation (First–Unpaired parameter estimates) and subsequent behavioral priming (Unpaired–Second response times). Both right intraparietal sulcus (IPS) and left middle occipital gyrus (MOG) exhibited a positive relationship: greater BOLD responses to First trials correlated with faster behavioral responses to Second trials (see text for other regions). *B*, Correlations between hippocampal anticipation (First–Unpaired parameter estimates from hippocampus in Fig. 3) and neural priming (Unpaired–Second parameter estimates). In early visual cortex, greater hippocampal anticipation (expressed as greater BOLD response to First relative to Unpaired) correlated with reduced neural priming (expressed as greater BOLD response to Second relative to Unpaired). In inferior parietal lobule (IPL), greater hippocampal anticipation (expressed as greater BOLD response to First relative to Unpaired) correlated with greater neural priming (expressed as reduced BOLD response to Second relative to Unpaired). All correlation maps thresholded at $p < 0.05$ cluster corrected. The primary results are depicted in the brain slices, and scatter plots from the voxel with the weakest correlation in each cluster are shown for visualization purposes only, to give a sense of the distribution of participants' scores.

though the brain seems to provide an especially sensitive measure of statistical learning (Turk-Browne et al., 2009). However, the event-related design of the current study gave us a unique opportunity to distinguish between two potential consequences of statistical learning that have been confounded in previous studies: prospective anticipation based on the first item in a pair, and retrospective recognition of the pair when the second item appeared.

To assess which of these two effects may predict or underlie subsequent familiarity, we correlated participants' postscan familiarity A' scores with their neural anticipation score (First–Unpaired) and their neural priming score (Unpaired–Second) in every voxel. There was no significant relationship between anticipation and familiarity in any brain region. In contrast, as seen in Figure 5, several clusters showed robust correlations between neural priming and familiarity: right precentral gyrus ($47, -4, 19; p < 0.002$ corrected), right posterior parahippocampal cortex ($23, -51, -4; p < 0.02$ corrected), medial prefrontal cortex and anterior cingulate cortex ($-4, 45, 0; p < 0.002$ corrected), poste-

Familiarity Correlations

A

Neural Anticipation
(First–Unpaired)

PCG   PHC   mPFC/ACC   PCC   Insula/MTG

B

Neural Priming
(Unpaired–Second)

y = 2   x = 23   x = -7   y = -13

Figure 5. Familiarity Correlations. **A**, No regions showed a correlation between neural anticipation (First–Unpaired parameter estimates) and subsequent familiarity (A' values). **B**, Correlations between neural priming (Unpaired–Second parameter estimates) and subsequent familiarity (A' values). The right precentral gyrus (PCG), right parahippocampal cortex (PHC), medial prefrontal cortex (mPFC), anterior cingulate cortex (ACC), posterior cingulate cortex (PCC), insula, and middle temporal gyrus (MTG) exhibited a positive relationship, i.e., greater Unpaired–Second difference correlated with more familiarity. All correlation maps thresholded at $p < 0.05$ cluster corrected.

rior cingulate cortex (-12, -51, 18; $p < 0.002$ corrected), left anterior parahippocampal cortex (-28, -35, 1; $p < 0.05$ corrected), and left insula/middle temporal gyrus (-46, -16, 0; $p < 0.002$ corrected). There was a positive correlation in all regions, where greater familiarity was associated with more neural priming (reduced BOLD responses to Second relative to Unpaired). The lack of any correlations with anticipation and the abundance of correlations with neural priming suggest that familiarity was related to processing of the predicted Second items.

Given that neural priming was measured before the familiarity test, we cannot determine whether the regions that exhibited these correlations were involved in encoding that supported later familiarity, or whether they provided a signal of incidental retrieval success that might later have been co-opted to make familiarity judgments. Indeed, the specific set of regions that we obtained is consistent with both interpretations (Henson et al., 1999; Konishi et al., 2000), but note that these studies reported enhanced BOLD responses, and in our case, reduced processing of Second items (i.e., more priming) predicted greater familiarity. Our results may thus be more consistent with findings that deactivations (especially in midline regions) accompanying encoding and attenuated responses resulting from stimulus repetition can support better subsequent memory (Daselaar et al., 2004; Turk-Browne et al., 2006). Regardless, our results suggest that priming may support familiarity, but that anticipation per se may be a qualitatively distinct learning effect, unseen in previous studies of statistical learning where familiarity tests are the standard. In the face of our robust evidence of anticipation in behavior and the brain, familiarity tests may thus be ill-suited for studying the implicit and dynamic perceptual consequences of learning.

Parametric effects

We tested whether any brain regions showed a greater increasing trend for First versus Unpaired images. We used the Unpaired control condition for this contrast to account for any changes due to stimulus repetition alone (e.g., Reber et al., 2005). Voxels were judged significant if their t value reached $p < 0.001$, and they were part of a cluster of at least seven contiguous significant voxels

(this and all subsequent cluster thresholds were determined based on the smoothness of the statistical map to achieve a corrected α of 0.05). No brain regions emerged from this contrast for either the linear or logarithmic models, likely reflecting limited statistical power for observing changes over six trials. Indeed, as described in Materials and Methods, our experiment was designed based on prior work showing the surprising speed of statistical learning (Turk-Browne et al., 2009), and under the assumption that the analysis of First > Unpaired would be collapsed across the short timeframe.

For completeness, we examined the reverse of the test above—greater increasing trend for the Unpaired images versus the First images—which yielded no regions in the linear model, but two regions in the logarithmic model: subgenual anterior cingulate cortex (-4, 14, -7; $p < 0.001$ corrected), and left middle occipital gyrus (-35, -61, 6; $p < 0.001$ corrected). The pattern of responses in these regions is consistent with them reflecting uncertainty about the next stimulus, with reduced uncertainty over time in response to the First images. We also compared the modulated regressors for Unpaired images to those for Second images. In both the linear and logarithmic models, a region of left inferior frontal gyrus in BA 45 (-51, 36, 0; $p < 0.03$ corrected) showed a greater increasing trend for the Unpaired images versus the Second images. The pattern of responses in this region is consistent with it tracking prediction error, with increasingly accurate predictions over time of the Second images. The reverse contrast—greater increasing trend for the Second images versus the Unpaired images—revealed a region of the left cerebellum (-46, -48, -32; $p < 0.03$ corrected) in the linear model, and nothing in the logarithmic model.

Category-specific effects

To examine effects in ventral visual cortex, we first localized face- and scene-selective ROIs (FFA and PPA, respectively), and then ran the GLM for the main task in each participant. Importantly, we did not collapse across face and scene categories as in all other analyses. Results are shown in Figure 6 for the PPA. There were three results of particular interest: (1) First-faces, which predicted that a scene would appear next in the sequence, resulted in

a marginally greater PPA response versus Unpaired-faces ($t_{(15)} = 2.04$, $p = 0.06$); this activation may reflect implicit preparation for the upcoming stimulus, or preparation of resources relevant to its processing. (2) Second-faces, which could be predicted based on the previous item, resulted in an attenuated PPA response versus Unpaired-faces ($t_{(15)} = 2.74$, $p = 0.02$); this deactivation may reflect the fact that the PPA was suppressed in anticipation of a stimulus that was not expected to be in its preferred domain. (3) Second-scenes, which could be predicted based on the previous trial, resulted in a greater PPA response versus First-scenes ($t_{(15)} = 2.29$, $p = 0.04$); this activation may reflect potentiation of the PPA in anticipation of a stimulus from its preferred domain. The strongest test of this last claim is the contrast of Second- versus Unpaired-scenes, but this difference did not reach significance; nevertheless, the Second > First difference is suggestive.

In the FFA, differences between conditions for a given image category were not statistically reliable. This may be related to the fact that the FFA is a less pure probe of category-specific processing in that it responds to scenes as well as to faces. In contrast, the PPA generally shows no response to faces whatsoever, and thus is a purer probe of scene processing (Epstein and Kanwisher, 1998). This distinction was apparent in our data (collapsing over conditions): the FFA response to scenes was 32% of its response to faces, and the PPA response to faces was -8% of its response to scenes (i.e., a slight deactivation). Moreover, a similar lack of sensitivity to top-down effects for FFA versus PPA has been observed in other studies of the ventral stream (Johnson et al., 2007). Given the mnemonic role of parahippocampal cortex, however, it will be important for future work to test whether the PPA results generalize to other category-specific visual regions. However, it is worth noting that the scene-selective parahippocampal region (posterior parahippocampal gyrus and collateral sulcus) is typically spatially distinct from the more anterior parahippocampal region involved in contextual processing (Aminoff et al., 2007; cf. Epstein and Ward, 2010).

Discussion

We found evidence of implicit behavioral and neural anticipation (i.e., associative prediction) on the basis of regularities embedded surreptitiously in a continuous sequence of images. Participants were unaware of these contingencies, and yet statistical learning was manifested in behavioral priming for predictable images and behavioral costs for predictive images. We identified a potential neural mediator of this implicit anticipation, the anterior hippocampus, which showed enhanced responses on trials consisting of predictive images relative to trials consisting of images that were not predictive. This hippocampal anticipation correlated with neural differences in how the predicted image was processed in early visual cortex (greater activation for predicted items) and left inferior parietal lobule (less activation for predicted items).

Parahippocampal Place Area
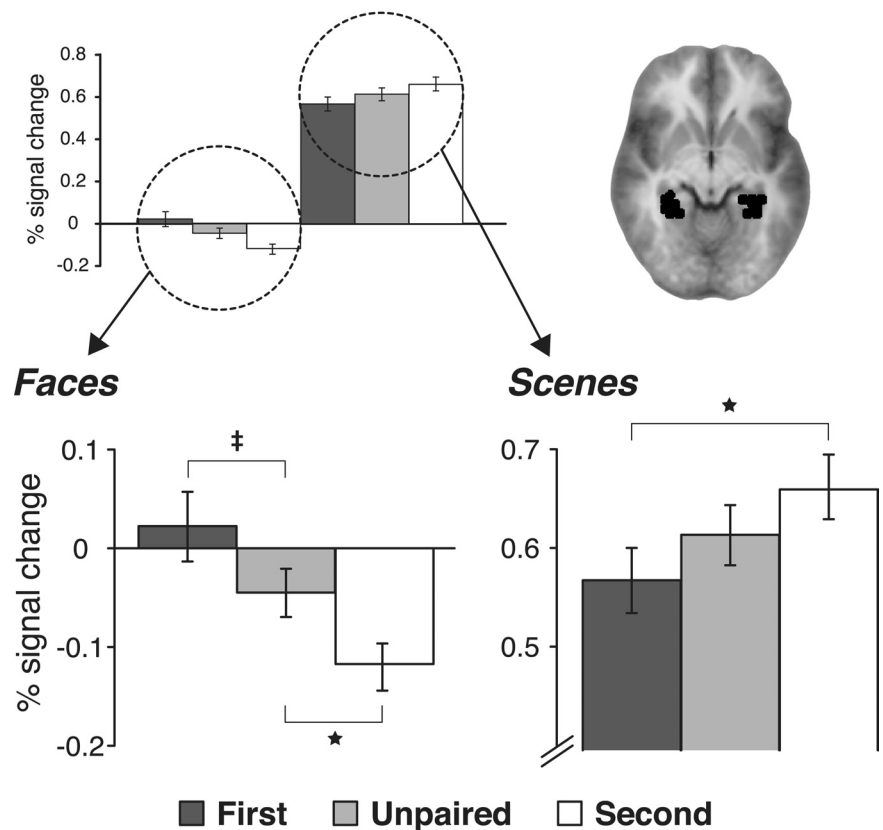
Figure 6. PPA Results. Parameter estimates for the main task extracted from bilateral PPA localized in each participant. Example ROIs are depicted as black spheres on inset brain. Error bars reflect one within-subject SEM. $*p < 0.05$, $^{\ddagger}p < 0.1$.

Anticipation in posterior brain regions, including right intraparietal sulcus and left middle occipital gyrus, correlated with behavioral priming. Additional correlation analyses suggest that anticipation and explicit familiarity may be dissociable. Finally, a region of interest analysis revealed predictive potentiation of a category-specific visual area in anticipation of stimuli from that category, along with suppression of the area when predictable stimuli came from a different category. These findings demonstrate an important functional consequence of this type of statistical learning for online behavior and neural processing—implicit perceptual anticipation.

The hippocampus and surrounding medial temporal lobe constitute the primary system involved in associative/relational memory (Cohen and Eichenbaum, 1993). For example, paired-associate learning is a form of explicit MTL-dependent learning. In such tasks, participants are instructed to study pairs of words or images, and learning is tested by providing one word or image and having participants recall (Shimamura and Squire, 1984) or recognize (Sakai and Miyashita, 1991) the associated item. While on the surface this protocol seems similar to ours in the sense that participants learn stimulus–stimulus pairs, there are at least three major differences. First, the two tasks involve different modes of learning. Paired associate learning is intentional because participants are told in advance to try to remember the pairs, whereas statistical learning in our study was incidental because participants were not oriented to the presence of pairs, they performed an unrelated cover task on individual images, and they ultimately did not report awareness of the pair structure. Second, although

the two tasks may result in similar kinds of associative knowledge, there are critical differences in how the pairs are learned. In paired associate learning, to-be-learned word/image pairs are studied together in discrete, clearly defined pair-units, whereas pairs in our study were embedded within a continuous trial sequence such that they could only be segmented and learned on the basis of statistics. Third, the two tasks differ with respect to how associates are revived. In paired associate learning, participants are prompted with one item and deliberately try to remember the paired item, whereas in our study learning was expressed without conscious effort or awareness during the orthogonal categorization task.

Despite these differences, our study may reveal a primitive and implicit mechanism that mediates learning about spatial and temporal contingencies in a variety of contexts, regardless of whether expression of learning is automatic or self-initiated and whether the revived content reaches awareness. Indeed, the consequences of associative learning in primate inferior/medial temporal neurons are similar whether associations are formed in explicit paired-associate tasks (Sakai and Miyashita, 1991), or as a result of exposure to task-irrelevant regularities, such as a fixed training order (Miyashita, 1988).

Thus far, we have focused on the role of the hippocampus in stimulus-specific learning and anticipation. However, the hippocampus may play a broader role in learning by detecting the presence of structure to be learned in the first place. For example, in the serial reaction time task, a sequence of visual shapes or locations is presented to subjects and they must execute the correct unique response to each visual stimulus (Nissen and Bullemer, 1987). If the sequence of stimuli is constructed using a table of transitional probabilities, several information-theoretic measures can be calculated for each stimulus. The left hippocampus has been observed to track one such measure, mutual information, which corresponds to a running representation of the conditional uncertainty between successive pairs of stimuli (Harrison et al., 2006); interestingly, this region did not reflect whether the uncertainty of the current stimulus was reduced by the prior stimulus (the reduction in surprise). In other words, activation in the left hippocampus reflected the extent to which preceding stimuli in the sequence were informative on average, regardless of whether the last stimulus was informative about the current stimulus (Harrison et al., 2006). Because the conditions in our study were distributed uniformly throughout each run, and every run had the same pseudo-randomized statistical structure, the difference we observed for First versus Unpaired cannot be explained by mutual information. Moreover, this difference also cannot reflect reduction in surprise, since both First and Unpaired images were unpredictable. Thus, the hippocampus may be involved both in anticipating future events on the basis of stimulus-specific regularities, and in tracking a general sense of whether such regularities exist in a given context.

The medial temporal lobe is not only involved in the encoding and retrieval of past events, but also in the deliberate imagination of future events (Addis et al., 2007; Hassabis et al., 2007). The current results reveal a new type of prospection that takes place during ongoing visual processing wherein the hippocampus can draw on incidentally learned associations to implicitly anticipate upcoming perceptual events. More broadly, implicit anticipation during ongoing behavior may be highly adaptive, helping to make optimal choices during decision making (Ferbinteanu and Shapiro, 2003; Johnson and Redish, 2007), prepare for and avoid aversive stimuli (Solomon et al., 1986; Cheng et al., 2008), and, as demonstrated here, recognize objects more quickly. These findings are consistent with cognitive models emphasizing the me-

morial consequences of sensory/perceptual processing as well as of more deliberate reflective processing (Johnson, 1983; Kolers and Roediger, 1984), and add to a growing body of work suggesting that future planning and simulation is a core function of cognition (Buckner and Carroll, 2007; Spreng et al., 2009).

We conclude by considering two potential benefits of implicit perceptual anticipation. First, while we have treated anticipatory responses as a stable consequence of statistical learning, anticipation may in turn help to refine learning when the resulting predictions are later confirmed or violated (Schultz and Dickinson, 2000). Indeed, temporal statistical learning is a particular case where sensitivity to prediction errors is the only way that regularities can be segmented (Perruchet and Pacton, 2006). Interestingly, the hippocampus also plays an important role in comparing expectations to outcomes (Kumaran and Maguire, 2007) and processing violations (Rose et al., 2005). Such violation responses in the hippocampus and in other regions that evaluate predictions (such as the putamen) may in turn influence the strength of learned associations by gating the connectivity between regions representing relevant stimuli and/or responses (den Ouden et al., 2010). In addition to questions about how changes in connectivity can influence learning, future research will also need to consider how learned associations and anticipation can prospectively alter connectivity and activation. For example, an auditory tone that predicts a visual image can result in early visual activation even when the image does not appear, with this unexpected event in turn increasing connectivity between auditory and visual cortex (den Ouden et al., 2009).

Second, expectations derived from learned regularities may facilitate perception by setting up templates (Hochberg, 1978). Such effects have been observed in tasks of object recognition, where the rapid extraction of low-spatial-frequency information by orbitofrontal cortex (Bar et al., 2006) and task sets represented in similar regions of ventromedial prefrontal cortex (C. Summerfield et al., 2006; Summerfield and Koechlin, 2008) constrain posterior object identification by modulating selective visual areas. These top-down influences on perception have been interpreted in a Bayesian framework that emphasizes the importance of encoding predictability when optimizing perceptual inference (Friston, 2005). For example, when stimulus repetitions are rare, and thus unpredictable, less repetition suppression is observed in category-selective ventral temporal cortex than when repetitions are frequent (Summerfield et al., 2008). In our study, we have addressed learning and anticipation in a conditional probabilistic context wherein individual images are equally predictable, but the transitions between specific images are not. These findings further demonstrate the sensitivity of the hippocampus to predictability in the environment (Strange et al., 2005; Harrison et al., 2006). In sum, while learning is a clear consequence of perception, our study highlights the reciprocal nature of this interaction, namely, that learning can result in implicit anticipation of the perceptual future.

References

Aaronson D, Watts B (1987) Extensions of Grier's computational formulas for A' and B'' to below-chance performance. Psychol Bull 102:439–442.

Abla D, Okanoya K (2009) Visual statistical learning of shape sequences: an ERP study. Neurosci Res 64:185–190.

Abla D, Katahira K, Okanoya K (2008) On-line assessment of statistical learning by event-related potentials. J Cogn Neurosci 20:952–964.

Addis DR, Wong AT, Schacter DL (2007) Remembering the past and imagining the future: common and distinct neural substrates during event construction and elaboration. Neuropsychologia 45:1363–1377.

Aguirre GK, Zarahn E, D'Esposito M (1998) An area within human ventral cortex sensitive to "building" stimuli: evidence and implications. Neuron 21:373–383.

Aminoff E, Gronau N, Bar M (2007) The parahippocampal cortex mediates spatial and nonspatial associations. Cereb Cortex 17:1493–1503.

Aslin RN, Saffran JR, Newport EL (1998) Computation of conditional probability statistics by 8-month-old infants. Psychol Sci 9:321–324.

Baker CI, Olson CR, Behrmann M (2004) Role of attention and perceptual grouping in visual statistical learning. Psychol Sci 15:460–466.

Bar M (2007) The proactive brain: using analogies and associations to generate predictions. Trends Cogn Sci 11:280–289.

Bar M, Kassam KS, Ghuman AS, Boshyan J, Schmid AM, Dale AM, Hämäläinen MS, Marinkovic K, Schacter DL, Rosen BR, Halgren E (2006) Top-down facilitation of visual recognition. Proc Natl Acad Sci U S A 103:449–454.

Brady TF, Oliva A (2008) Statistical learning using real-world scenes: extracting categorical regularities without conscious intent. Psychol Sci 19:678–685.

Büchel C, Holmes AP, Rees G, Friston KJ (1998) Characterizing stimulus-response functions using nonlinear regressors in parametric fMRI experiments. Neuroimage 8:140–148.

Buckner RL, Carroll DC (2007) Self-projection and the brain. Trends Cogn Sci 11:49–57.

Cabeza R, Ciaramelli E, Olson IR, Moscovitch M (2008) The parietal cortex and episodic memory: an attentional account. Nat Rev Neurosci 9:613–625.

Cheng DT, Disterhoft JF, Power JM, Ellis DA, Desmond JE (2008) Neural substrates underlying human delay and trace eyeblink conditioning. Proc Natl Acad Sci U S A 105:8108–8113.

Chua EF, Schacter DL, Rand-Giovannetti E, Sperling RA (2007) Evidence for a specific role of the anterior hippocampal region in successful associative encoding. Hippocampus 17:1071–1080.

Chun MM, Jiang Y (1998) Contextual cueing: implicit learning and memory of visual context guides spatial attention. Cognit Psychol 36:28–71.

Chun MM, Phelps EA (1999) Memory deficits for implicit contextual information in amnesic subjects with hippocampal damage. Nat Neurosci 2:844–847.

Cohen NJ, Eichenbaum H (1993) Memory, amnesia, and the hippocampal system. Cambridge, MA: MIT.

Daselaar SM, Prince SE, Cabeza R (2004) When less means more: deactivations during encoding that predict subsequent memory. Neuroimage 23:921–927.

Davachi L, Wagner AD (2002) Hippocampal contributions to episodic encoding: insights from relational and item-based learning. J Neurophysiol 88:982–990.

den Ouden HE, Friston KJ, Daw ND, McIntosh AR, Stephan KE (2009) A dual role for prediction error in associative learning. Cereb Cortex 19:1175–1185.

den Ouden HE, Daunizeau J, Roiser J, Friston KJ, Stephan KE (2010) Striatal prediction error modulates cortical coupling. J Neurosci 30:3210–3219.

Epstein R, Kanwisher N (1998) A cortical representation of the local visual environment. Nature 392:598–601.

Epstein RA, Ward EJ (2010) How reliable are visual context effects in the parahippocampal place area? Cereb Cortex 20:294–303.

Erickson CA, Desimone R (1999) Responses of macaque perirhinal neurons during and after visual stimulus association learning. J Neurosci 19:10404–10416.

Ferbinteanu J, Shapiro ML (2003) Prospective and retrospective memory coding in the hippocampus. Neuron 40:1227–1239.

Fiser J, Aslin RN (2001) Unsupervised statistical learning of higher-order spatial structures from visual scenes. Psychol Sci 12:499–504.

Fiser J, Aslin RN (2002) Statistical learning of higher-order temporal structure from visual shape sequences. J Exp Psychol Learn Mem Cogn 28:458–467.

Fiser J, Aslin RN (2005) Encoding multielement scenes: statistical learning of visual feature hierarchies. J Exp Psychol Gen 134:521–537.

Forman SD, Cohen JD, Fitzgerald M, Eddy WF, Mintun MA, Noll DC (1995) Improved assessment of significant activation in functional magnetic resonance imaging (fMRI): use of a cluster-size threshold. Magn Reson Med 33:636–647.

Friston K (2005) A theory of cortical responses. Philos Trans R Soc Lond B Biol Sci 360:815–836.

Goebel R, Esposito F, Formisano E (2006) Analysis of functional image analysis contest (FIAC) data with BrainVoyager QX: from single-subject to cortically aligned group general linear model analysis and self-organizing group independent component analysis. Hum Brain Mapp 27:392–401.

Grier JB (1971) Nonparametric indexes for sensitivity and bias: computing formulas. Psychol Bull 75:424–429.

Harrison LM, Duggins A, Friston KJ (2006) Encoding uncertainty in the hippocampus. Neural Netw 19:535–546.

Hassabis D, Kumaran D, Vann SD, Maguire EA (2007) Patients with hippocampal amnesia cannot imagine new experiences. Proc Natl Acad Sci U S A 104:1726–1731.

Henke K, Mondadori CRA, Treyer V, Nitsch RM, Buck A, Hock C (2003) Nonconscious formation and reactivation of semantic associations by way of the medial temporal lobe. Neuropsychologia 41:863–876.

Henson RN, Rugg MD, Shallice T, Josephs O, Dolan RJ (1999) Recollection and familiarity in recognition memory: an event-related functional magnetic resonance imaging study. J Neurosci 19:3962–3972.

Hochberg JE (1978) Perception. Englewood Cliffs, NJ: Prentice-Hall.

Howard MW, Fotedar MS, Datey AV, Hasselmo ME (2005) The temporal context model in spatial navigation and relational learning: toward a common explanation of medial temporal lobe function across domains. Psychol Rev 112:75–116.

Hunt RH, Aslin RN (2001) Statistical learning in a serial reaction time task: access to separable statistical cues by individual learners. J Exp Psychol Gen 130:658–680.

Jensen O, Lisman JE (2005) Hippocampal sequence-encoding driven by a cortical multi-item working memory buffer. Trends Neurosci 28:67–72.

Johnson A, Redish AD (2007) Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. J Neurosci 27:12176–12189.

Johnson MK (1983) A multiple-entry, modular memory system. In: The psychology of learning and motivation: advances in research and theory, pp 81–123. New York: Academic.

Johnson MR, Mitchell KJ, Raye CL, D'Esposito M, Johnson MK (2007) A brief thought can modulate activity in extrastriate visual areas: top-down effects of refreshing just-seen visual stimuli. Neuroimage 37:290–299.

Kanwisher N, McDermott J, Chun MM (1997) The fusiform face area: a module in human extrastriate cortex specialized for face perception. J Neurosci 17:4302–4311.

Kirkham NZ, Slemmer JA, Johnson SP (2002) Visual statistical learning in infancy: evidence for a domain general learning mechanism. Cognition 83:B35–B42.

Knutson B, Cooper JC (2005) Functional magnetic resonance imaging of reward prediction. Curr Opin Neurol 18:411–417.

Kolers PA, Roediger HL (1984) Procedures of mind. J Verb Learn Verb Behav 23:425–449.

Konishi S, Wheeler ME, Donaldson DI, Buckner RL (2000) Neural correlates of episodic retrieval success. Neuroimage 12:276–286.

Kumaran D, Maguire EA (2007) Which computational mechanisms operate in the hippocampus during novelty detection. Hippocampus 17:735–748.

Manns JR, Squire LR (2001) Perceptual learning, awareness, and the hippocampus. Hippocampus 11:776–782.

Marr D (1971) Simple memory: a theory for archicortex. Philos Trans R Soc Lond B Biol Sci 262:23–81.

McCarthy G, Puce A, Gore JC, Allison T (1997) Face-specific processing in the human fusiform gyrus. J Cogn Neurosci 9:605–610.

McClelland JL, McNaughton BL, O'Reilly RC (1995) Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. Psychol Rev 102:419–457.

McNealy K, Mazziotta JC, Dapretto M (2006) Cracking the language code: neural mechanisms underlying speech parsing. J Neurosci 26:7629–7639.

Mitchell KJ, Johnson MK, Raye CL, D'Esposito M (2000) fMRI evidence of age-related hippocampal dysfunction in feature binding in working memory. Brain Res Cogn Brain Res 10:197–206.

Miyashita Y (1988) Neuronal correlate of visual associative long-term memory in the primate temporal cortex. Nature 335:817–820.

Miyashita Y (1993) Inferior temporal cortex: where visual perception meets memory. Annu Rev Neurosci 16:245–263.

Moores E, Laiti L, Chelazzi L (2003) Associative knowledge controls deployment of visual selective attention. Nat Neurosci 6:182–189.

Nissen MJ, Bullemer P (1987) Attentional requirements of learning: evidence from performance measures. Cogn Psychol 19:1–32.

Norman KA, O'Reilly RC (2003) Modeling hippocampal and neocortical contributions to recognition memory: a complementary-learning-systems approach. Psychol Rev 110:611–646.

Olson IR, Chun MM (2001) Temporal contextual cuing of visual attention. J Exp Psychol Learn Mem Cogn 27:1299–1313.

Orbán G, Fiser J, Aslin RN, Lengyel M (2008) Bayesian learning of visual chunks by human observers. Proc Natl Acad Sci U S A 105:2745–2750.

Perruchet P, Pacton S (2006) Implicit learning and statistical learning: one phenomenon, two approaches. Trends Cogn Sci 10:233–238.

Reber PJ, Gitelman DR, Parrish TB, Mesulam MM (2005) Priming effects in the fusiform gyrus: changes in neural activity beyond the second presentation. Cereb Cortex 15:787–795.

Rose M, Haider H, Weiller C, Büchel C (2002) The role of medial temporal lobe structures in implicit learning: an event-related FMRI study. Neuron 36:1221–1231.

Rose M, Haider H, Büchel C (2005) Unconscious detection of implicit expectancies. J Cogn Neurosci 17:918–927.

Saffran JR, Aslin RN, Newport EL (1996) Statistical learning by 8-month-old infants. Science 274:1926–1928.

Sakai K, Miyashita Y (1991) Neural organization for the long-term memory of paired associates. Nature 354:152–155.

Schendan HE, Searl MM, Melrose RJ, Stern CE (2003) An FMRI study of the role of the medial temporal lobe in implicit and explicit sequence learning. Neuron 37:1013–1025.

Schultz W (2006) Behavioral theories and the neurophysiology of reward. Annu Rev Psychol 57:87–115.

Schultz W, Dickinson A (2000) Neuronal coding of prediction errors. Annu Rev Neurosci 23:473–500.

Shimamura AP, Squire LR (1984) Paired-associate learning and priming effects in amnesia: a neuropsychological study. J Exp Psychol Gen 113:556–570.

Simoncelli EP, Olshausen BA (2001) Natural image statistics and neural representation. Annu Rev Neurosci 24:1193–1216.

Solomon PR, Vander Schaaf ER, Thompson RF, Weisz DJ (1986) Hippocampus and trace conditioning of the rabbit's classically conditioned nictitating membrane response. Behav Neurosci 100:729–744.

Spreng RN, Mar RA, Kim ASN (2009) The common neural basis of autobiographical memory, prospection, navigation, theory of mind and the default mode: a quantitative meta-analysis. J Cogn Neurosci 21:489–510.

Strange BA, Duggins A, Penny W, Dolan RJ, Friston KJ (2005) Information theory, novelty and hippocampal responses: unpredicted or unpredictable? Neural Netw 18:225–230.

Summerfield C, Koechlin E (2008) A neural representation of prior information during perceptual inference. Neuron 59:336–347.

Summerfield C, Egner T, Greene M, Koechlin E, Mangels J, Hirsch J (2006) Predictive codes for forthcoming perception in the frontal cortex. Science 314:1311–1314.

Summerfield C, Trittschuh EH, Monti JM, Mesulam MM, Egner T (2008) Neural repetition suppression reflects fulfilled perceptual expectations. Nat Neurosci 11:1004–1006.

Summerfield JJ, Lepsien J, Gitelman DR, Mesulam MM, Nobre AC (2006) Orienting attention based on long-term memory experience. Neuron 49:905–916.

Tolman EC (1932) Purposive behavior in animals and men. New York: Appleton.

Turk-Browne NB, Scholl BJ (2009) Flexible visual statistical learning: transfer across space and time. J Exp Psychol Hum Percept Perform 35:195–202.

Turk-Browne NB, Jungé J, Scholl BJ (2005) The automaticity of visual statistical learning. J Exp Psychol Gen 134:552–564.

Turk-Browne NB, Yi DJ, Chun MM (2006) Linking implicit and explicit memory: common encoding factors and shared representations. Neuron 49:917–927.

Turk-Browne NB, Isola PJ, Scholl BJ, Treat TA (2008) Multidimensional visual statistical learning. J Exp Psychol Learn Mem Cogn 34:399–407.

Turk-Browne NB, Scholl BJ, Chun MM, Johnson MK (2009) Neural evidence of statistical learning: efficient detection of visual regularities without awareness. J Cogn Neurosci 21:1934–1945.