



Intracranial Entrainment Reveals Statistical Learning across Levels of Abstraction

Brynn E. Sherman¹, Ayman Aljishi², Kathryn N. Graves², Imran H. Quraishi², Adithya Sivaraju², Eyiyeemisi C. Damisah², and Nicholas B. Turk-Browne²

Abstract

■ We encounter the same people, places, and objects in predictable sequences and configurations. Humans efficiently learn these regularities via statistical learning. Importantly, statistical learning creates knowledge not only of specific regularities but also of regularities that apply more generally across related experiences (i.e., across members of a category). Prior evidence for different levels of learning comes from post-exposure behavioral tests, leaving open the question of whether more abstract regularities are detected online during initial exposure. We address this question by measuring neural entrainment in intracranial recordings. Neurosurgical patients viewed a stream of photographs with regularities at one of two levels: In the exemplar-level structured condition, the same photographs appeared repeatedly in pairs. In the category-level structured condition, the photographs were trial-unique but their

categories were paired across repetitions. In a baseline random condition, the same photographs repeated but in a scrambled order. We measured entrainment at the frequency of individual photographs, which was expected in all conditions, but critically also at half that frequency—the rate at which to-be-learned pairs appeared in the two structured (but not random) conditions. Entrainment to both exemplar and category pairs emerged within minutes throughout visual cortex and in frontal and temporal regions. Many electrode contacts were sensitive to only one level of structure, but a significant number encoded both levels. These findings suggest that the brain spontaneously uncovers category-level regularities during statistical learning, providing insight into the brain's unsupervised mechanisms for building flexible and robust knowledge that generalizes across input variation and conceptual hierarchies. ■

INTRODUCTION

Everyday experience is highly structured, and humans can learn this structure via a process known as statistical learning (Sherman, Graves, & Turk-Browne, 2020). This knowledge in turn lets us generate predictions and behave more efficiently when we encounter familiar environments in the future. For example, after repeatedly shopping at your local grocery store, you know where your favorite brands are located, when the shelves are stocked, and whether you need to bag your own groceries, all of which makes shopping smoother than in an unfamiliar store. At the same time, beyond the specifics of your local grocery store, many features of your experience reflect general properties of grocery shopping that generalize to most or all other stores, including that aisles are organized by food type, that the check-out counters tend to be at the front of the store, and so forth, meaning that experienced shoppers can still efficiently navigate even in a new store.

Prior behavioral studies have shown that statistical learning supports this kind of generalization across idiosyncratic experiences or exemplars (Jung, Walther, & Finn, 2021; Jun & Chong, 2018; Luo & Zhao, 2018;

Emberson & Rubinstein, 2016; Otsuka, Nishiyama, Nakahara, & Kawaguchi, 2013; Brady & Oliva, 2008). A common design in such studies is to expose participants to a sequence of images with regularities at the category level (e.g., images of beaches always followed by images of canyons); this differs from standard studies of statistical learning in which the regularities exist at the level of particular exemplar images that repeat in pairs or triplets. Evidence for category-level statistical learning has been assessed offline in a behavioral test after sequence exposure. For example, participants might be asked to rate the familiarity of a category pair to which they were exposed (e.g., beach → canyon) versus a foil (e.g., beach → farm, where farm was a category in another pair). The categories in these test items are often represented by novel exemplars or category labels, such that they can only be discriminated if the participants learned categorical regularities that they can generalize to these novel stimuli.

These prior studies usefully demonstrated that statistical learning can occur across levels of abstraction, but the use of offline tests limits insight into the learning process itself. Specifically, it is unclear how and when participants detect these regularities at the category level, and critically whether this occurs during learning at all. An alternative explanation is that participants only learn the specific,

¹University of Pennsylvania, Philadelphia, ²Yale University, New Haven, CT

exemplar-level regularities to which they were exposed, and then at test generalize these regularities to novel exemplars or labels through analogy or inference. For example, if exposed to a pair of exemplars during learning (e.g., beach1 → canyon1), participants may exhibit familiarity or discrimination for new exemplars of the same categories at test (e.g., beach2 → canyon2) either (1) because they had already learned a general category relation online during exposure that is ready to be applied at test or (2) because no category-level learning occurred in advance and they instead retrieve specific learned pairs at test and infer that the right answer will preserve the same category relation. This theoretical distinction of whether inference occurs during encoding or retrieval has been examined in other forms of learning and memory (Zhou, Singh, Tandoc, & Schapiro, 2023; Preston & Eichenbaum, 2013). A prior study from our laboratory provided some tentative behavioral evidence that category-level statistical learning might occur online, during exposure (Sherman & Turk-Browne, 2020), which prompted us to conduct a targeted study to measure online learning of category-level regularities directly.

For this purpose, we adopted a technique known as neural entrainment (or frequency tagging) that has found recent success in tracking statistical learning of auditory and visual regularities (Henin et al., 2021; Batterink, 2020; Choi, Batterink, Black, Paller, & Werker, 2020; Batterink & Paller, 2017; Ding, Melloni, Zhang, Tian, & Poeppel, 2016). This EEG-based method capitalizes on the fact that brain oscillations in sensory regions can exhibit phase locking, or entrainment, at the frequency of onset of rhythmic stimuli (Bauer, Debener, & Nobre, 2020; Norcia, Appelbaum, Ales, Cottareau, & Rossion, 2015). The presence of such entrainment can be used to detect whether and where in the brain the stimuli are processed, including when multiple stimuli are presented at different frequencies (De Rosa, Ktori, Vidal, Bottini, & Crepaldi, 2022; Ding et al., 2016; Störmer & Alvarez, 2014; Nozaradan, Peretz, Missal, & Mouraux, 2011). Indeed, statistical learning studies have found neural entrainment not only at the frequency of individual stimuli but also to the frequency of learned groupings of multiple stimuli (Henin et al., 2021), despite no explicit segmentation cues indicating these groupings.

For example, in a visual stream in which certain scenes follow each other with high transition probability constituting pairs, neural entrainment is expected at the frequency of individual scenes, reflecting visual-evoked responses, but also at half of that frequency reflecting the rate of learned pairs. Critically, this provides a measure of learning because the pairs only exist in the minds of participants who extracted them based on statistical regularities across repetitions (again, there is no explicit timing, instruction, or other cue in the stimuli about the existence of pairs). Because entrainment to learned regularities is measured incidentally while participants are passively exposed to the stream, this method provides a continuous

online measure of statistical learning not readily available in behavior. Beyond being a sensitive online measure of learning, neural entrainment can also provide mechanistic insight into the learning process. For example, it can help elucidate the timecourse of statistical learning by quantifying how much exposure is required for entrainment to emerge. Moreover, with the higher spatial resolution and coverage of deep-brain structures provided by intracranial EEG (iEEG), it is also possible to localize statistical learning effects in the brain.

Prior studies of statistical learning with neural entrainment employed stimuli that were identical across repetitions, leaving open the question of whether category-level statistical learning occurs online. Thus, we combine, for the first time, the method of neural entrainment as an online measure with a task design optimized for evaluating statistical learning over category exemplars. This task builds on our recent iEEG study that conflated regularities at the exemplar and category levels (Sherman et al., 2022). Here, we evaluate these two levels of abstraction separately in distinct conditions (relative to a random baseline condition), allowing us to quantify neural entrainment online during exemplar-level and category-level statistical learning.

Across task runs, we manipulated the nature of regularities in a sequence of scene images (Figure 1): In category-level structured runs, each image appeared once such that regularities could exist only at the level of categories (e.g., Category A → Category B); this differed from exemplar-level structured runs with repeating images that contained regularities at the level of individual exemplars (e.g., Scene A → Scene B). Both of these structured runs with regularities were compared with a random run in which images repeated without any regularities in their temporal order. Patients were not informed about these different conditions or about the presence of regularities, and they learned them incidentally through passive exposure. By capitalizing on the spatial and temporal resolution of iEEG, we tracked statistical learning of exemplar and category regularities across the brain, providing insight into how, when, and where learning occurs across levels of abstraction. We focused primarily on neural entrainment in visual cortex, given prior work demonstrating strong entrainment in sensory regions (e.g., Sherman et al., 2022; Henin et al., 2021). However, in exploratory analyses, we also capitalized on the broad coverage afforded by iEEG, as several brain regions have been implicated in the learning and representation of statistical regularities (Henin et al., 2021; Karuza et al., 2013; Turk-Browne, Scholl, Chun, & Johnson, 2009).

METHODS

Participants

We tested eight patients (one woman; age range: 21–61 years; mean age = 37.8 years) who had been surgically




Figure 1. Task design and experimental conditions. Participants viewed a rapid stream of scene images (left), with varying levels of temporal structure (right). In the category-level structured condition (top), participants encountered a series of trial-unique scene images, drawn from six scene categories. Scene categories were temporally paired (three pairs of two categories), such that an image from one category (e.g., beach) was always followed by an image from another category (e.g., canyon). In the exemplar-level structured condition (middle), participants encountered six scene images that appeared in temporal pairs. In the random control condition (bottom), participants again encountered six (novel) scene images but now in a random temporal order without pairs.

implanted with intracranial electrodes for localization of seizure onset zone (see Table 1 for patient demographics and details on implant). This sample size was chosen a priori based on Sherman et al. (2022). Two patients were tested a second time (2 days later) because their first data

Table 1. Patient Demographics and Electrode Placement

<i>Patient Information</i>					
<i>ID</i>	<i>Age, years</i>	<i>Sex</i>	<i>nContacts</i>	<i>Implant Type</i>	<i>Hemisphere</i>
1	28	M	217	Combined	Primarily right
2	21	M	119	Combined	Left
3	33	M	191	Combined	Primarily right
4	34	M	116	Combined	Right
5	58	M	182	Combined	Primarily left
6	44	M	168	Combined	Left
7	61	F	155	Depth	Bilateral
8	23	M	162	Depth	Left

Implant type indicates whether the implanted electrodes were only depth electrodes (depth) or a combination of depth electrodes and grid/strip electrodes on the cortical surface (combined). Hemisphere indicates the cerebral hemisphere into which the electrodes were implanted (see also Figure 2).

set was found to be unusable: One of these patients experienced severe eye irritation during the first testing session, and there was a technical error with the triggers for the other patient. Electrode placement was determined solely by the clinical care team to localize seizure foci. Patients were recruited through the Yale Comprehensive Epilepsy Center and provided informed consent in a manner approved by the Yale University Human Subjects Committee. All data were collected at Yale New Haven Hospital.

iEEG Recordings

EEG data were recorded on a NATUS NeuroWorks EEG recording system. Data were collected at a sampling rate of 4096 Hz. Signals were referenced to an electrode chosen by the clinical team to minimize noise in the recording. To synchronize EEG signals with the experimental task, a custom-configured data acquisition system was used to convert signals from the research computer to 8-bit “triggers” that were inserted into a separate digital channel.

iEEG Preprocessing

iEEG preprocessing was carried out in FieldTrip (Oostenveld, Fries, Maris, & Schoffelen, 2011). A notch filter was applied to remove 60-Hz line noise. No rereferencing was applied.

Data were downsampled to 256 Hz and segmented into trials using the triggers.

Electrode Localization

Electrode contact locations were identified using post-operative CT and MRI scans. Reconstructions were completed in BioImage Suite (Papademetris et al., 2006) and were subsequently registered to the patient's pre-operative MRI scan, resulting in contact locations projected into the patient's pre-operative space. The resulting files were converted from the Bioimagesuite format (.MGRID) into native space coordinates using FieldTrip functions. The coordinates were then used to create a mask in FMRIB Software Library (FSL; Jenkinson, Beckmann, Behrens, Woolrich, & Smith, 2012), with the coordinates of each contact occupying one voxel in the mask (Figure 2).

Given prior evidence for entrainment in sensory regions, we were interested in measuring neural responses in visual regions. We constructed a broad visual cortex ROI, as in Sherman et al. (2022), on the Montreal Neurological Institute (MNI) T1 2-mm standard brain by combining the occipital lobe ROI from the MNI Structural Atlas and the following ROIs from the Harvard–Oxford Cortical Structural Atlas: inferior temporal gyrus (temporo-occipital part),

lateral occipital cortex (superior division), lateral occipital cortex (inferior division), intracalcarine cortex, cuneal cortex, parahippocampal gyrus (posterior division), lingual gyrus, temporal occipital fusiform cortex, occipital fusiform gyrus, supracalcarine cortex, and occipital pole. Each ROI was thresholded at 10% and then concatenated to create a single mask of visual cortex.

To localize contacts, we registered each patient's pre-operative anatomical scan to the MNI T1 2-mm standard brain template using linear registration (FSL FLIRT; Jenkinson, Bannister, Brady, & Smith, 2002; Jenkinson & Smith, 2001) with 12 degrees of freedom. We then used this registration matrix to transform each electrode mask into standard space. We overlaid the electrode masks onto the visual cortex ROI and onto the Harvard–Oxford cortical and subcortical structural atlases (maximum probability, 0 threshold). All but one of the patients had contacts in the visual cortex ROI, resulting in a final sample size of seven participants for analyses of visual cortex.

Stimuli

Task stimuli consisted of 720 unique scene images drawn from 18 distinct outdoor scene subcategories (amphitheater,




Figure 2. Electrode coverage for each patient. Each dot represents a single contact depicted on a standard glass brain. Contacts could be localized to the visual cortex ROI (purple shaded region) in seven of the eight patients, as indicated by darker black dots.

amusement park, beach, bridge, canyon, desert, farm, forest, garden, highway, lake, lighthouse, marsh, mountain, park, sports field, town square, and waterfall; 40 images per subcategory). Six subcategories were randomly assigned to each of the exemplar-level structured, category-level structured, and random conditions (see below); this randomization was performed independently for each participant. All images were collected from Google image searches and were cropped to a resolution of 600×800 pixels. Stimuli were presented using MATLAB (The Math Works) with the Psychophysics toolbox (Brainard, 1997; Pelli, 1997).

Procedure

Participants completed the experiment on a laptop while seated in their hospital bed. The task consisted of at least one run of each of the three experimental conditions. During each run, participants passively viewed a rapid stream of scene images and were asked to pay attention to each image. To enable entrainment-based neural analyses, the SOA was fixed at 500 msec; each scene was presented for 250 msec, followed by a 250-msec ISI, during which a fixation cross appeared in the center of the screen. Each run sequence was 240 trials in length (2 min of viewing time).

The category-level structured runs were our key runs of interest, in which we probed online learning of categorical regularities. Participants viewed a sequence of trial-unique scene images drawn from six scene categories. To increase categorical processing of the scenes, patients were told in advance that they would be viewing images of scene categories and were given the names of the six categories. Importantly, they were not informed about the specific category pairs nor even that the sequence contained pairs or a predictable order. Unbeknownst to them, the six categories were assigned to three statistical pairs, such that a scene from one category (Category A) was always followed by a scene from its paired category (Category B; Figure 1, top right). Critically, these pairs existed only at the category level because exemplars never repeated, requiring that patients generalize across exemplars to learn the regularities. No pair was allowed to repeat back-to-back in the sequence. In total, participants viewed 40 exemplars from each scene category once (40 repetitions of each category pair).

The exemplar-level structured runs served as a key comparison, enabling us to examine statistical learning of stimulus regularities without need for generalization across instances, as in prior studies (Henin et al., 2021; Batterink & Paller, 2017). In this run, participants viewed a sequence containing multiple repetitions of six scene images, one each from six categories that did not overlap with the other conditions. Unbeknownst to them, the scenes were assigned to three statistical pairs (e.g., Scene A \rightarrow Scene B; Figure 1, middle right). No pair was allowed to repeat back-to-back in the sequence. Each exemplar/pair was repeated 40 times throughout the sequence.

The random control runs served as our baseline condition, in which we did not expect any learning-related neural entrainment. As in the exemplar-level structured runs, participants viewed a sequence containing 40 repetitions of six scene images from six non-overlapping categories. In contrast to the two structured conditions, the scenes were presented in a random order without reliable pairs that could be learned (Figure 1, bottom right). No individual scene was allowed to repeat within two images in the sequence. In neither the exemplar-level structured nor random runs were participants given the names of individual scenes or scene categories in advance because these conditions contained only six exemplars that repeated multiple times.

Prior work has demonstrated that the order of statistical learning tasks can impact performance. Namely, learning is worse when one set of regularities is shown after another set or after randomness (Gebhart, Aslin, & Newport, 2009; Jungé, Scholl, & Chun, 2007). Thus, to maximize our chance of detecting category-level neural entrainment, should it exist, especially given unexpected complications and interruptions in working with hospitalized patients, we tested the category-level structured condition first. We attempted to complete two of these runs back-to-back with the same sequence. When two runs were obtained (6 of 8 patients), we included data from both runs in all analyses. However, we also performed control analyses with only the first run, to equate the amount of data across conditions. After the category-level structured run(s), we completed one run of the exemplar-level structured and random conditions next, counterbalancing order across participants. We decided on this semifixed condition order (category-level structured first) ahead of time, accepting that it could complicate comparison between conditions. However, note that each condition contains a positive control of neural entrainment to the individual image frequency, allowing us to assess data quality and ensure that conditions tested later in the session did not suffer from fatigue or inattention.

Neural Entrainment Analyses

We conducted a phase coherence analysis to identify which electrode contacts entrained to our task. We examined entrainment at two frequencies: (1) the image frequency (2 Hz, corresponding to the 500-msec SOA between images), which reflects entrainment to the frequency of visual stimulation and should be present in all runs, and (2) the pair frequency (1 Hz, corresponding to the 1000-msec interval between pair onsets), which reflects entrainment to the statistical pairs and should only be present in the structured runs (Henin et al., 2021).

Phase coherence provides a measure of event-related phase locking and thus is sensitive to whether neural activity is time-locked to a particular experiment event (i.e., image or pair onset). We measured phase coherence, rather than spectral power, because it can be more sensitive to low-frequency fluctuations (Kabdebon, Peña,

Buiatti, & Dehaene-Lambertz, 2015; Forget, Buiatti, & Dehaene, 2010), has been used in prior studies of statistical learning (Henin et al., 2021; Choi et al., 2020; Batterink & Paller, 2017), and has been validated via simulations (Benjamin, Dehaene-Lambertz, & Fló, 2021). Nevertheless, we replicated our primary results with spectral power (see Appendix Figure A1).

To compute phase coherence, we epoched our data into blocks; we then computed phase coherence across those blocks. For some runs of the task, there was a computer-based timing error such that the first trial's ISI period was shorter than expected. Because the phase coherence analysis depends on reliable timing across trials, we excluded the first two trials from all analyses. The raw signals from the remaining 238 trials were segmented into 17 blocks comprising 14 trials. For patients with two runs in the category-level structured condition, the raw signals were concatenated across runs, yielding 34 blocks.

We then converted the raw signals for each block into the frequency domain via fast Fourier transform and computed the phase coherence across blocks for each contact using the formula $R^2 = \left[\frac{1}{N} \sum^N (\cos \varphi) \right]^2 + \left[\frac{1}{N} \sum^N (\sin \varphi) \right]^2$, where N is the number of blocks and φ is the phase at a given frequency (Henin et al., 2021; Ding & Simon, 2013). Phase coherence was computed separately for each contact in the brain. We computed the peaks at the image and pair frequencies as the coherence at those frequencies relative to the coherence at the two neighboring frequencies (± 0.14 Hz).

To assess statistical reliability across participants, we used a nonparametric, random-effects bootstrap resampling approach (Efron & Tibshirani, 1986). We first pooled the data across contacts and computed the effect of interest (e.g., mean or correlation coefficient). For each of 10,000 iterations, we randomly resampled the same sample size of participants with replacement, pooled across the resampled participants' contacts, and recomputed the effect of interest to populate a sampling distribution of the effect. For example, to assess the reliability of coherence at the pair frequency across electrode contacts in visual cortex, we pooled all contacts from the seven participants who had contacts localized to the visual cortex ROI and computed the average coherence peak (coherence at the pair frequency, relative to the two neighbors) across pooled contacts. For each of 10,000 iterations, we then randomly resampled these seven participants with replacement, pooled the contacts from the seven resampled participants, and recomputed the average coherence peak. This sampling distribution was used to obtain 95% confidence intervals and perform null hypothesis testing. We calculated the p value as the proportion of iterations in which the resampled effect had the opposite sign as the true effect; we then multiplied these values by 2 to obtain a two-tailed p value and compared it with a significance threshold of .05. This tests the null hypothesis that the true effect is centered at zero and thus equally likely to be positive or negative by chance. A significant

effect indicates that it did not matter which patients were resampled on any given iteration, and thus that the patients were interchangeable and the effect reliable across the sample. Across-participants resampling was performed in R (Version 4.1.3), and the random number seed was set to 12345 before each resampling test.

To assess the reliability of a coherence peak within an individual electrode contact, we performed a randomization test. Using the approach described in Prichard and Theiler (1994), we created phase-shuffled surrogate data sets. We first transformed the raw signal for each block into the frequency domain via fast Fourier transform. To scramble the phase information, we randomly rotated the phase angle at each frequency. We then applied the inverse Fourier Transform to generate a phase-shuffled surrogate data set. This approach randomizes the phase information at each frequency while preserving the power spectrum and temporal autocorrelations of the original data. We performed this phase randomization for each block 1000 times and recomputed the phase coherence across blocks of surrogate data. By comparing the observed coherence for a given contact to a null distribution of coherence values (that would be expected by chance in that contact), this procedure allows us to characterize whether an individual contact exhibits reliable entrainment. We computed the proportion of iterations that the true peak (coherence at the frequency of interest minus the neighboring frequencies) was larger than the null distribution of peaks to calculate the p value. Given that we had a directional hypothesis (i.e., higher coherence than baseline), we did not multiply these p values by 2. Within-contact randomization testing was performed in MATLAB, and the random number seed was set to 12345 for each contact.

Phase Coherence Timecourse Analyses

To assess how neural entrainment to statistical pairs changed over the course of exposure, we performed two complementary analyses.

We first quantified the time to a reliable response (number of cumulative blocks at which we observed a reliable pair frequency response across electrodes; Sherman et al., 2022; Henin et al., 2021). For this analysis, we recomputed the coherence over an increasing number of blocks (e.g., first computing the coherence only between the first and second blocks, all the way up to all 17 blocks). For each cumulative block, we compared the coherence peak relative to a phase-shuffled surrogate data set (described above) to compute the within-contact reliability. This resulted in a timecourse of p values, allowing us to determine how many blocks of exposure were required for reliable entrainment. We performed this analysis at both the image and pair frequencies. We expected coherence at the image frequency to become reliable rapidly, as it reflects entrainment to sensory stimulation and does not require learning, providing a baseline for helping to interpret the timecourse of coherence at the (learned) pair frequency.

We computed p value timecourses separately for the category-level and exemplar-level structured conditions, focusing on visual contacts that showed reliable entrainment by the final block. That is, within each structured condition, we averaged the timecourses of all contacts that exhibited reliable entrainment to the pair frequency by block 17. To equate opportunity for learning across patients, we only considered the first run of the category-level structured condition for patients with two of these runs. To assess statistical reliability, we calculated the nonparametric p value for a given number of cumulative blocks as the proportion of iterations in which the resampled p value was greater than .05. We then multiplied these values by 2 to obtain a two-tailed p value. This resampling test was done in R (Version 4.1.3), with a random number seed of 12345.

The above approach allows us to measure how many blocks were needed to produce a reliable response. However, because this analysis calculates cumulative coherence from the first block, it is insensitive to fluctuations in coherence responses over the course of learning. Thus, as an alternative approach, we computed the coherence over sliding windows of nine blocks. We opted for sliding windows of nine blocks because the cumulative approach above indicated reliable entrainment at the image frequency—which, in principle, should be present throughout—after eight to nine blocks. We thus considered this the shortest window length over which we would be able to detect coherence. Nevertheless, comparable results were obtained with even shorter window lengths, suggesting that this analysis was mostly robust to this decision.

For each block, we computed the peak in coherence at both the image and pair frequencies as the difference between the coherence at the frequency of interest, relative to the two neighboring frequencies. This yielded a timecourse of coherence peaks for both the image and pair frequencies. As with the cumulative analysis above, we computed timecourses separately for the category-level and exemplar-level structured conditions, focusing on visual contacts that showed reliable entrainment in the respective condition by the final block. Furthermore, we limited analysis to each participant's first category-level structured run. To assess statistical reliability, we calculated the nonparametric p value for a given sliding window as the proportion of iterations in which the resampled coherence was less than 0. We then multiplied these values by 2 to obtain a two-tailed p value. This resampling test was done in R (Version 4.1.3), with a random number seed of 12345.

RESULTS

Evidence for Category-level Statistical Learning in Visual Cortex

To assess whether the brain represents visual regularities online during learning, we capitalized on the fast, periodic

nature of visual stimulation in our task and measured neural entrainment to the frequency of both individual images and statistical pairs (Figure 3A). Given prior work demonstrating neural entrainment in sensory regions (Sherman et al., 2022; Henin et al., 2021), we focused our analyses on visual cortex. Specifically, we computed coherence within each contact localized to the visual cortex ROI (116 contacts) and averaged the coherence across contacts and participants. To quantify the expected peaks in coherence at the image and pair frequencies, we computed the difference in coherence between the frequency of interest (pair or image frequency) and the two neighboring frequencies. We then performed a bootstrap resampling test to assess the reliability of these peaks across participants.

As a validation of our paradigm, we expected strong phase coherence at the frequency of image presentation in all three conditions. We further expected phase coherence at the frequency of pair presentation in the exemplar-level structured condition, replicating prior work demonstrating that the brain entrains to the frequency of statistical regularities (Henin et al., 2021; Batterink & Paller, 2017). Critically, if the brain generalizes across these stimuli to learn higher-level, categorical regularities, we would expect phase coherence at the pair frequency in the category-level structured condition.

As shown in Figure 3B, we found reliable peaks in coherence at the image frequency in all three conditions (exemplar-level structured: mean difference, relative to neighboring frequencies = 0.53; 95% CI [0.47, 0.57], $p < .001$; category-level structured: mean difference = 0.54; 95% CI [0.46, 0.60], $p < .001$; random: mean difference = 0.55; 95% CI [0.47, 0.62], $p < .001$). Critically, the peak in coherence at the pair frequency was reliable in both the exemplar-level structured condition (mean difference = 0.16, 95% CI [0.12, 0.21], $p < .001$) and the category-level structured condition (mean difference = 0.10, 95% CI [0.041, 0.20], $p < .001$), but not in the random condition (mean difference = -0.0027 , 95% CI [-0.012 , 0.0048], $p = .50$), providing online evidence for rapid statistical learning of both exemplar and category pairs.

To further understand these effects, we compared the peaks in coherence across conditions. We expected that there would be no condition differences in the peak at the image frequency, but that the peak at the pair frequency would be higher in exemplar- and category-level structured conditions, relative to random. Consistent with this hypothesis, there were no pairwise differences in the image frequency across conditions (exemplar-level structured vs. random: mean difference = -0.026 , 95% CI [-0.066 , 0.0094], $p = .15$; category-level structured vs. random: mean difference = -0.018 , 95% CI [-0.047 , 0.015], $p = .28$; exemplar- vs. category-level structured: mean difference = -0.0073 , 95% CI [-0.042 , 0.020], $p = .61$; Figure 3C, bottom). Importantly, the peak in coherence at the pair frequency was reliably higher for both structured conditions than for the random condition




Figure 3. Phase coherence analysis. (A) Schematic of analysis and hypothesized neural oscillations. We expected entrainment of visual contacts at the frequency of image presentation in all conditions. In the two structured conditions (exemplar-level and category-level), we also expected entrainment at the frequency of (learned) pairs. (B) These hypotheses were confirmed: We observed reliable peaks in coherence at the image frequency in all three conditions, but only at the pair frequency for the category-level and exemplar-level structured conditions. Error shading indicates bootstrapped 95% confidence intervals. (C) Coherence peaks at the pair frequency (top) and image frequency (bottom) for each participant across the three runs. Each circle/line represents one participant.

(exemplar-level structured vs. random: mean difference = 0.17, 95% CI [0.13, 0.21], $p < .001$; category-level structured vs. random: mean difference = 0.10, 95% CI [0.043, 0.21], $p < .001$; Figure 3C, top). Interestingly, the peak in coherence at the pair frequency was marginally higher in exemplar- versus category-level structured condition (mean difference = 0.062, 95% CI [-0.0085, 0.11], $p = .075$), suggesting that stimulus regularities may be represented more robustly than category-level regularities in visual cortex, at least after a fixed and small amount of exposure.

The above analyses were performed on data concatenated across the two runs of the category-level structured condition (for participants with two runs). To confirm that evidence for categorical learning was not dependent on including more data, we repeated the

analysis only considering the first category-level structured run. Indeed, we found a comparable peak in coherence at the pair frequency (mean difference = 0.096, 95% CI [0.042, 0.19], $p < .001$); the peak in coherence at the image frequency remained reliably high as well (mean difference = 0.56, 95% CI [0.49, 0.62], $p < .001$). Furthermore, the peak in coherence at the pair frequency remained reliably higher than the random condition (mean difference = 0.099, 95% CI [0.044, 0.20], $p < .001$) and marginally lower than that of the exemplar-level condition (mean difference = -0.067, 95% CI [-0.0028, 0.12], $p = .061$).

Together, these results demonstrate robust representation of statistical regularities in visual cortex, across levels of abstraction. After only 2 min of exposure, the visual cortex entrained not only to regularities at the exemplar level

(with the same exact image pairs repeating), but also to regularities that existed only at the category level (requiring generalization over exemplars to uncover the categorical structure). Critically, these data demonstrate that category-level regularities can be learned and represented online during learning, extending prior behavioral work that relied on delayed, offline test measures to infer that such learning occurred.

Co-Representation of Exemplar and Category Regularities

Above, we found evidence that visual cortex represents both exemplar- and category-level regularities. However, it is unclear whether these two effects are related. One possibility is that the more basic ability to extract regularities in sensory stimuli is a precursor for learning more complex regularities, in which case we might expect the same contacts to exhibit both effects and for the strength of these effects to be related. Another possibility is that exemplar-level and category-level statistical learning are fundamentally distinct processes that may be implemented in different neural populations, and thus may be represented in different contacts and/or in the same contacts but in an unrelated manner.

To address this question, we first asked whether the strength of neural entrainment was correlated between conditions. Across all electrode contacts in the visual cortex ROI, we computed the Pearson correlation coefficient between the coherence peaks at the pair frequency. We found a reliable correlation in the pair frequency peak for category- and exemplar-level structured conditions ($r = .33$, 95% CI [0.019, 0.58], $p = .033$; Figure 4A). In contrast, there was no reliable correlation between the random condition and either the category-level structured condition ($r = -.10$, 95% CI [-0.24, 0.077], $p = .22$) or

the exemplar-level structured condition ($r = -.011$, 95% CI [-0.14, 0.14], $p = .82$). The modest correlation between coherence for exemplar pairs and category pairs suggests a degree of shared representation of regularities across levels of abstraction. Importantly, given that we did not find such correlations with the random condition, we can be confident that this correlation was not driven by generic across-contact factors such as baseline coherence or data quality.

As a further control, we computed the pairwise correlations for the image frequency peaks. Unlike the pair frequency, we did not expect these correlations to differ between conditions. Indeed, we found high correlations across the board (category- and exemplar-level structured, Figure 4B: $r = .83$, 95% CI [0.70, 0.92], $p < .001$; category-level structured and random: $r = .88$, 95% CI [0.81, 0.93], $p < .001$; exemplar-level structured and random: $r = .87$, 95% CI [0.79, 0.93], $p < .001$).

To further address the relationship between exemplar and category regularities, we labeled individual contacts according to whether they exhibited a reliable coherence peak at the frequencies of interest in each condition. This labeling was the result of a randomization test performed for each individual contact, in which the coherence at the frequency of interest was compared with a null distribution of coherence computed from phase-shuffled surrogate data sets; such an approach allows us to define whether a given contact exhibits reliable entrainment at the frequencies of interest. Of the 116 total electrode contacts in visual cortex, 67 exhibited entrainment to the pair frequency in one or both structured conditions; 27 entrained to the pair frequency in the exemplar-level structured condition only, 12 in the category-level structured condition only, and 28 in both structured conditions. To assess whether this is more overlap than would be expected by chance, given the number of reliable




Figure 4. Correlations across contacts. (A) Correlation between the coherence peak at the pair frequency in the category-level structured condition and the coherence peak at the pair frequency in the exemplar-level structured condition. (B) Correlation between the coherence peak at the image frequency in the category-level structured condition and the coherence peak at the image frequency in the exemplar-level structured condition. Each circle represents an electrode contact. Color coding indicates whether a given contact exhibits reliable entrainment at the pair frequency (A) or image frequency (B) in the exemplar-level structured condition only, category-level structured condition only, or both (black indicates that the contact does not exhibit reliable entrainment in either condition). Error shading indicates bootstrapped 95% confidence intervals.

contacts in each condition, we independently shuffled the correspondence between contacts and significance labels across conditions and recomputed the overlap. We found that the observed overlap was indeed reliable (mean null overlap = 19 contacts, 95% CI [14, 24], $p < .001$), indicating that some parts of visual cortex may exhibit a dual representation of both exemplar and category regularities.

To understand whether these dual-coding contacts were responsible for the correlations observed above, we recomputed the correlations after removing these contacts. Indeed, this eliminated the correlation (88 non-overlapping contacts: $r = -.086$, 95% CI [-0.18, 0.15], $p = .31$). However, there was also no reliable correlation when restricting the analysis to only the dual-coding contacts (28 overlapping contacts: $r = .041$, 95% CI [-0.36, 0.33], $p = .41$). This suggests that the original correlation benefited from variance in coding properties across contacts and/or from the greater sensitivity provided by a larger sample size of contacts.

Examining the Timecourse of Learning in Visual Cortex

We have presented evidence that populations of electrode contacts in visual cortex entrain to both exemplar and category-level regularities online during statistical learning. However, it is possible that statistical learning of category regularities requires more exposure than learning of simpler, stimulus-driven exemplar regularities. To assess the evolution of entrainment over the course of learning, and whether this timecourse differs across conditions, we used two related analysis approaches. For both analyses, we only analyzed each patient's first category-level structured run (to equate the opportunity for learning both across the exemplar-level and category-level structured conditions) and limited analysis to electrode contacts that exhibited reliable entrainment (as measured by within-contact randomization tests) to the pair frequency by the final block.

First, we analyzed the time required for a reliable entrainment response. We recomputed coherence over an increasing number of blocks (e.g., first computing the coherence only between the first and second blocks, then between the first, second, and third blocks, all the way up to 17 blocks) to determine the block count at which contacts exhibited reliable entrainment. In other words, this analysis asks how much exposure was required for contacts that exhibited reliable entrainment in the final block to reach a statistically reliable response.

In the category-level structured condition (Figure 5A, left), we found reliable entrainment only when computing coherence across 16 or more blocks (16 blocks: mean $p = .023$, 95% CI [0.0033, 0.044], $p = .0056$; 17 blocks: mean $p = .0097$, 95% CI [0.0023, 0.018], $p < .001$). In the exemplar-level structured condition (Figure 5A, right), entrainment appeared marginally after 14 blocks (mean $p = .029$, 95% CI [0.011, 0.052], $p = .057$) and reliably

for 15 or more blocks ($ps < .001$). These data suggest that exemplar and category regularities were learned at a similar timescale, with slightly faster acquisition for exemplar regularities.

To establish a floor of how quickly we might theoretically expect to see a reliable entrainment effect, we performed this same analysis for the image frequency (again, only considering contacts that exhibited reliable entrainment to the image frequency in the final block). Because entrainment to the images was given by the sensory input and not from statistical learning, we did not expect meaningful differences between conditions. In the category-level structured condition (Figure 5B, left), there was reliable coherence at the image frequency by Block 9 (mean $p = .028$, 95% CI [0.0072, 0.049], $p = .037$; all subsequent block counts, $ps < .001$). The exemplar-level structured condition (Figure 5B, right) followed a similar pattern, with reliable entrainment by Block 8 (mean $p = .026$, 95% CI [0.009, 0.041], $p = .0018$; all subsequent block counts, $ps < .001$). Finally, we also computed the timecourse of the image frequency effect in the random condition and found a similar pattern, with reliable entrainment by Block 9 (mean $p = .033$; 95% CI [0.016, 0.048], $p = .024$; block 10: mean $p = .034$, 95% CI [0.016, 0.049], $p = .036$; all subsequent block counts, $ps < .001$).

The analysis above provides a relatively stringent test, measuring the timecourse of statistically reliable coherence peaks. Such an analysis may not be sensitive to fluctuations in learning over time, as it only captures the cumulative coherence response. Thus, we conducted another kind of analysis in which we computed coherence in sliding windows of nine blocks over the course of the exposure phase. In contrast to the cumulative analysis, the sliding-window analysis revealed reliable entrainment throughout exposure, for both the category-level and exemplar-level structured conditions (all $ps < .001$; Figure 5C). We observed no clear trajectory in learning in either condition, nor distinct peaks that were consistent across participants. However, there was considerable variability across participants, raising the possibility of individual differences in the speed of learning.

As with the "time to reliable response" analysis, we additionally considered the timecourse of the image frequency coherence peak using the sliding-window approach. Similar to the pair frequency results, we found reliable entrainment at the image frequency throughout exposure (all $ps < .001$; Figure 5D).

Category-level Statistical Learning across the Brain

We initially focused on how visual cortex represents visual regularities given our prior work (Sherman et al., 2022), but a wide range of brain regions have been implicated in statistical learning (Henin et al., 2021; Batterink, Paller, & Reber, 2019). To examine online learning more broadly, we measured neural entrainment to exemplar and category regularities in an exploratory brain-wide analysis.




Figure 5. Timecourse analyses. (A) Emergence of a reliable phase coherence peak at the pair frequency across blocks in the category-level structured (left) and exemplar-level structured (right) conditions. For each cumulative block count N , we computed the proportion of iterations that the coherence peak across n blocks was greater than the peak across n blocks of phase-shuffled data to obtain a p value; we then determined the first block at which the permuted p value across contacts was reliably less than .05 (dashed line). (B) Emergence of a significant response at the image frequency across blocks in the category-level structured (left) and exemplar-level structured (right) conditions. (C) Coherence at the pair frequency over sliding windows of nine blocks. x Axis ticks indicate the middle of a window (e.g., $x = 6$ indicates a sliding window over blocks 2–10). (D) Coherence at the image frequency over sliding windows of nine blocks. Error shading indicates bootstrapped 95% confidence intervals. Thick line indicates mean across participants. Thin lines indicate the mean for each participant.




Figure 6. Exploratory brain-wide analyses. (A) Histogram: distribution of how many contacts would be expected to entrain to the pair frequency in *both* the exemplar- and category-level structured conditions by chance; red line indicates the observed overlap, indicating that many more contacts coded for both exemplar and category regularities than would be expected by chance. Inset: Venn diagram illustrating the total number of contacts that entrained to the pair frequency in both conditions and their overlap. (B) Map of contacts (across all patients) that entrained to the pair frequency in one or both conditions on a standard glass brain.

Table 2. Gray-matter ROIs in the Harvard–Oxford Cortical and Subcortical Atlases that Contained at Least Five Contacts with Reliable Entrainment at the Pair Frequency in the Category- and/or Exemplar-level Structured Conditions

<i>Localization of Task-sensitive Contacts</i>					
<i>ROI</i>	<i>Total</i>	<i>Category</i>	<i>Exemplar</i>	<i>Overlap</i>	<i>Image</i>
Frontal pole	176	19	7	0	69
Insular cortex	59	5	14	0	28
Middle frontal gyrus	65	7	5	1	36
Precentral gyrus	52	10	5	0	24
Temporal pole	51	0	7	0	18
Middle temporal gyrus, post	45	3	6	0	8
Postcentral gyrus	47	9	6	0	25
Lateral occipital cortex, sup	36	2	10	4	26
Lingual gyrus	13	1	5	4	13
Occipital pole	21	5	5	9	21

We also included the total number of contacts in each ROI (total) and the number of contacts that entrained at the image frequency in the category- and/or exemplar-level structured conditions (image).

First, as in the analysis restricted to visual cortex, we identified which contacts represented exemplar and/or category regularities by testing for reliable phase coherence at the pair frequency relative to neighboring frequencies (by comparing to a null distribution of phase-shuffled surrogate data). Of 1310 contacts across all patients, we found reliable entrainment at the pair frequency in 175 contacts for the exemplar-level structured condition and in 177 contacts for the category-level structured condition; 41 of these contacts overlapped. This amount of overlap was reliably greater than expected by chance (Figure 6A; mean null overlap = 24 contacts, 95% CI [16, 32], $p < .001$). Because this brain-wide analysis included visual cortex, it is possible that the reliable overlap was driven by visual contacts, which we earlier showed exhibited reliable overlap. We therefore repeated the brain-wide analysis after excluding contacts in the visual cortex ROI. Of the remaining 1194 contacts across all patients, we found reliable entrainment at the pair frequency in 120 contacts for the exemplar-level structured condition and 137 contacts for the category-level structured condition; 13 of these contacts overlapped. However, this amount of overlap was not reliably greater than what would be expected by chance (mean null overlap = 14 contacts, 95% CI [8, 20], $p = .52$), suggesting that dual coding of exemplar- and category-level regularities in individual contacts was restricted to visual cortex.

We next sought to localize these structure-sensitive contacts throughout the brain (Figure 6B). We mapped the contacts onto the Harvard–Oxford cortical and subcortical atlases and quantified how many contacts exhibited effects within each gray-matter atlas ROI. Table 2 summarizes the results by listing atlas ROIs that contained at least five contacts that entrained at an uncorrected level to the pair

frequency in at least one of the structured conditions. Consistent with our planned visual ROI, many of these contacts were located in visual cortex (e.g., lateral occipital cortex, lingual gyrus, occipital pole). However, we also observed entrainment to learned regularities in frontal and anterior temporal regions, some showing a preference for regularities available directly in the exemplar stimuli (e.g., temporal pole) and others for regularities that required generalization across category exemplars (e.g., frontal pole and precentral gyrus). Importantly, claims about localization in the brain are limited by the fact that we did not have full coverage of all brain regions, given that electrode placement was determined clinically.

DISCUSSION

In the current study, we capitalized on the high spatial and temporal precision of iEEG to explore how the brain learns and represents statistical regularities across varying levels of abstraction. Specifically, we contrasted the learning of exemplar-level regularities (defined by the transition probabilities between individual images) with the learning of category-level regularities (defined by the transition probabilities between image categories). We found robust representation of both kinds of regularities in visual cortex and throughout the brain during statistical learning. These findings speak to several issues in the statistical learning literature and raise questions for future research.

Online Evidence for Category-level Statistical Learning

In measuring neural entrainment to the frequency of regularities, we employed a covert, online measure of

statistical learning. This builds on a body of work that measured category-level statistical learning with offline behavioral tests, such as asking participants to judge their familiarity with pairs of images or categories (Jung et al., 2021; Emberson & Rubinstein, 2016; Otsuka et al., 2013; Brady & Oliva, 2008). However, it is unclear whether above-chance performance on these tests reflects online learning of category relationships during the learning process itself or the formation of specific stimulus associations during learning that enabled successful inferences about test items from the same categories. It is also possible that these behavioral studies engendered both online learning and inferences at test, yet it remains unclear which effect (or both) drove test performance. Further complicating the interpretation of offline behavioral performance as evidence of online learning, online and offline measures of statistical learning are not always correlated (Kiai & Melloni, 2021). The current study sought to skirt these interpretational challenges by measuring neural entrainment as an online neural index of statistical learning. The observed entrainment to category pairs provides novel evidence for rapid statistical learning across levels of abstraction.

One limitation of our study is that it is unclear how the neural entrainment measure of statistical learning relates to more canonical behavioral measures. Given our short testing time with each patient, their limited energy and attention span, and the small number of patients, we optimized our task design and testing time for neural rather than behavioral measures of learning. Relatedly, to ensure the participants processed the scene categories, we instructed them in advance that they would be seeing images from six categories. Selective attention to individual items has been shown to facilitate statistical learning between items, even when these regularities themselves are not mentioned or instructed (Sherman & Turk-Browne, 2022; Turk-Browne, Jungé, & Scholl, 2005). Indeed, we did not indicate to participants that categories were paired nor tell them to attend to the temporal order of images (although we note that some prior studies of statistical learning have explicitly cued participants to potential temporal regularities; Plate, Schapiro, & Waller, 2022; Bogaerts, Richter, Landau, & Frost, 2020; Siegelman, Bogaerts, Kronenfeld, & Frost, 2018; Arciuli, Torkildsen, Stevens, & Simpson, 2014). Future studies could perhaps use scalp EEG in a well-powered normative sample to help link neural and behavioral measures of category-level statistical learning, as well as to understand the role of categorical attention for both behavioral and neural measures of statistical learning. Future studies could further consider how neural entrainment during learning relates to both online (e.g., response time) and offline (e.g., familiarity) behavioral measures; that said, it may be difficult to develop online behavioral measures during a task designed for neural entrainment, given the fast presentation rates that such tasks require. Prior studies have demonstrated that neural evidence of statistical learning can

appear earlier and even in the absence of behavioral evidence of learning (Turk-Browne et al., 2009); thus, it is possible that our current results reflect a rapid sensitivity of the brain to category regularities.

Although we observed reliable entrainment to category-level regularities, the effects were somewhat attenuated relative to the exemplar-level condition. For example, across participants, we observed a marginally higher peak in coherence for exemplar-level, relative to category-level regularities. Furthermore, our timecourse analyses provided some evidence that exemplar-level learning may emerge as reliable faster than category-level learning. One interpretation of these results is exemplar-level regularities are easier to learn (given the exact stimulus repetition). However, a recent study of statistical learning in the auditory modality demonstrated that entrainment to statistical regularities might partially reflect a methodological artifact resulting from stimulus repetition (Pinto, Prior, & Zion Golumbic, 2022). Although it is unclear whether such an artifact applies to visual stimuli (the specific effect pertained to acoustic properties), such an effect could not explain our category-level results, as the exemplars in this condition varied across category repetitions.

Additional limitations apply in how to interpret the timecourse results. Although these results provide evidence that learning occurs quite quickly (less than 2 min) in both structured conditions, it is unclear how this maps onto the underlying trajectory of learning. Our two analysis approaches provided complementary, although somewhat divergent, results. First, analyzing the time to a reliable entrainment response, we found evidence of statistical learning for exemplar regularities two blocks earlier than for category regularities. Does this small difference in the amount of required exposure mean that specific stimulus associations must be learned before more abstract associations? Or, perhaps the individual images were represented both as exemplars and categories during perception and associations were learned at both levels in parallel? In this case, learning of category regularities may be slower because of the added complexity in dealing with greater input variability (e.g., in the extent to which a given exemplar was a prototype of a category). The “time to reliable response” measure has been used before (Henin et al., 2021), but conflates time in the experiment with the number of cumulative blocks included. Thus, seemingly early learning (like of exemplar regularities) could instead reflect a larger effect size that requires fewer blocks to reach significance. We thus adopted a second method that defines a window size with a fixed number of blocks and examines the emergence and dynamics of learning by sliding the window over the timecourse. Using this sliding-window approach, we observed reliable entrainment in the group average across all windows in both conditions. However, there was considerable variability across participants, which may reflect individual differences in learning rates across participants or increased noise in considering coherence across a small number of

blocks (Benjamin et al., 2021). Furthermore, for both analyses, we only included contacts that exhibited reliable entrainment to the pair frequency in the final block, as including nonreliable contacts would increase the noise in any timecourse measurement. However, this choice may have obscured heterogeneous timecourses for different aspects or stages of learning across the brain. Future work could clarify the timecourse of exemplar and categorical statistical learning by including an online behavioral measure of learning, which would help validate the neural timecourse analyses.

Finally, different aspects of learned structure can be measured. For example, memory for the temporal order of items within a statistical unit (e.g., triplet) can be dissociated from memory for the item groupings (Forest, Finn, & Schlichting, 2022; Park, Rogers, & Vickery, 2018), and these distinct types of memory may be supported by different underlying neural representations (Henin et al., 2021; Davachi & DuBrow, 2015). Although providing evidence of learning overall, the current study, and the basic neural entrainment design it employed, is insensitive to these differing underlying representations. Future studies could employ other neural measures, such as pre- and postlearning templates (Schapiro, Kustner, & Turk-Browne, 2012), to assess changes in the representations of the individual paired items. Such measures could be used to test hypotheses about how these constituent items are represented at different levels of abstraction as a function of statistical learning.

Local and Distributed Representations of Visual Regularities across the Brain

We focused our main analyses on visual cortex, which we hypothesized would show neural entrainment to visual regularities between visual images (Sherman et al., 2022; Henin et al., 2021). However, we also performed an exploratory brain-wide analysis to uncover where category and exemplar regularities were represented throughout the brain. This analysis largely confirmed our a priori choice to focus on visual cortex, but also revealed a distributed representation of structure, with many frontal (e.g., frontal pole, insula, middle frontal gyrus, and precentral gyrus) and temporal (e.g., temporal pole, middle temporal gyrus) regions also exhibiting entrainment to visual regularities. These findings are largely consistent with prior fMRI studies demonstrating sensitivity to structure in these regions (Karuza et al., 2013, 2017; Turk-Browne, Scholl, Johnson, & Chun, 2010; Turk-Browne et al., 2009).

This analysis revealed relatively little evidence that entire brain regions specialize at a particular level of abstraction. Although some regions exhibited a bias toward one level (e.g., more contacts in the frontal pole entrained only to category regularities, and more contacts in the insula entrained only to exemplar regularities), very few regions solely represented one level. The only exception was the temporal pole, which only exhibited

entrainment to exemplar-level regularities. Similarly, most contacts did not show a general sensitivity to structure regardless of the level of abstraction. The small (but reliable) number of such contacts representing both category and exemplar regularities were restricted to visual cortex (e.g., occipital pole). Still, the majority of visual contacts entrained to one level of structure or the other, but not both. At the level of entire brain regions, some regions contained distinct contacts that entrained selectively to category and exemplar regularities, yet no contacts that entrained to both. This raises the possibility that there may be distinct neural populations and cognitive processes even within the same brain region for statistical learning at varying levels of abstraction.

An important limitation to these exploratory brain-wide analyses is that they only had access to partial coverage of the brain. Although we had relatively broad coverage of cortical regions for an iEEG study, the electrode locations were chosen entirely for clinical purposes and were thus not always comprehensive or standardized across patients. However, this is an expected limitation for any iEEG-based study. Furthermore, we had insufficient coverage of the hippocampus in this sample (only five contacts across all patients), a region that has been consistently implicated in rapid statistical learning (Graves et al., 2022; Henin et al., 2021; Sherman & Turk-Browne, 2020; Covington, Brown-Schmidt, & Duff, 2018; Schapiro et al., 2012; Turk-Browne et al., 2009). Future studies could recruit a more targeted sample of iEEG patients (e.g., with hippocampal depth electrodes) or use fMRI for high-resolution hippocampal coverage potentially across a larger sample of individuals.

Relatedly, it is important to note that our recordings were from macroelectrodes, meaning that a single intracranial contact is pooling over many neurons. Thus, even in cases in which we observed a single contact representing both exemplar and category regularities (e.g., in occipital pole), this does not necessarily mean that these types of regularities are being represented by precisely the same neural population. Rather, these two types of regularities may be coded by distinct neural populations that happen to be measured by a single contact. Future work employing microelectrodes might help to resolve the spatial precision needed to make strong claims about the underlying neural populations supporting these learning computations.

CONCLUSIONS

Together, our results provide evidence for rapid and robust online representation of categorical regularities during statistical learning. This occurred heavily within visual cortex, suggesting a remarkable capability for the brain to aggregate across noisy, idiosyncratic instances to extract stable properties of the environment that can generalize to new situations.

APPENDIX




Figure A1. Spectral power analysis. To supplement the phase coherence analyses presented in the main text, we replicated our primary analysis using spectral power instead of phase coherence. To compute spectral power, we first averaged the raw EEG signal (for a given electrode contact) across blocks of a given condition. We then extracted power for the averaged timeseries via the fast Fourier transform. To normalize across frequencies, we transformed power into decibels (dB). (A) We observed reliable, positive peaks in spectral power at the image frequency in all three conditions ($p < .001$), but at the pair frequency only in the category-level structured condition (mean peak, relative to neighbors = 5.08, 95% CI [2.36, 8.92], $p < .001$) and exemplar-level structured condition (mean peak = 6.68, 95% CI [5.34, 8.06], $p < .001$); power at the pair frequency relative to neighbors in the random condition was suppressed (mean peak = -0.78 , 95% CI [-1.42 , -0.22], $p = .0052$), although the average difference was an order of magnitude smaller than the positive peaks in the structured conditions. Error shading indicates bootstrapped 95% CIs. (B) Spectral power peaks (frequency of interest minus neighboring frequencies) at the pair frequency (top) and image frequency (bottom) for each participant. Relative to random, the peak in power at the pair frequency was reliably higher for both the category-level structured condition (mean difference = 5.86, 95% CI [2.77, 10.00], $p < .001$) and exemplar-level structured condition (mean difference = 7.45, 95% CI [6.14, 8.93], $p < .001$), whereas the structured peaks did not differ from one another (mean difference = 1.60, 95% CI [-1.73 , 4.29], $p = .33$). Peaks in power at the image frequency did not differ across conditions (all $p > .18$). Each circle/line represents one participant.

Acknowledgments

We are grateful to the patients who participated in the study and to the physicians, fellows, and support staff who cared for them at the Yale New Haven Hospital. We thank Kun Wu for providing the electrode reconstructions; Simon Henin for advice on statistical analyses; Erica Busch for advice on the brain wide analyses; and Isabella Huang, Alene Sherman, and Bryna Shuman for contributing images used in Figure 1.

Reprint requests should be sent to Brynn E. Sherman, Department of Psychology, University of Pennsylvania, 425 S. University Avenue, Philadelphia, PA 19104-6243, or via e-mail: brynns@sas.upenn.edu.

Data Availability Statement

The raw iEEG data and anatomical information about electrode placement have been deposited in Dryad (<https://doi.org/10.5061/dryad.q573n5tph>).

Author Contributions

Brynn E. Sherman: Conceptualization; Formal analysis; Investigation; Methodology; Visualization; Writing—Original draft; Writing—Review & editing. Ayman Aljishi: Conceptualization; Investigation; Writing—Review & editing. Kathryn N. Graves: Investigation; Writing—Review & editing. Imran H. Quraishi: Supervision; Writing—Review & editing. Adithya Sivaraju: Conceptualization; Supervision; Writing—Review & editing. Eyiymisi C. Damisah: Conceptualization; Project administration; Supervision; Writing—Review & editing. Nicholas B. Turk-Browne: Conceptualization; Methodology; Project administration; Supervision; Writing—Original draft; Writing—Review & editing.

Funding Information

This project was supported by a National Science Foundation Graduate Research Fellowship (<https://dx.doi.org/10>

.13039/100000001), to B. E. S., and the Canadian Institute for Advanced Research (<https://dx.doi.org/10.13039/100007631>) to N. B. T.-B.

Diversity in Citation Practices

Retrospective analysis of the citations in every article published in this journal from 2010 to 2021 reveals a persistent pattern of gender imbalance: Although the proportions of authorship teams (categorized by estimated gender identification of first author/last author) publishing in the *Journal of Cognitive Neuroscience* (*JoCN*) during this period were $M(\text{an})/M = .407$, $W(\text{oman})/M = .32$, $M/W = .115$, and $W/W = .159$, the comparable proportions for the articles that these authorship teams cited were $M/M = .549$, $W/M = .257$, $M/W = .109$, and $W/W = .085$ (Postle and Fulvio, *JoCN*, 34:1, pp. 1–3). Consequently, *JoCN* encourages all authors to consider gender balance explicitly when selecting which articles to cite and gives them the opportunity to report their article's gender citation balance. The authors of this article report its proportions of citations by gender category to be as follows: $M/M = .313$; $W/M = .354$; $M/W = .104$; $W/W = .229$.

REFERENCES

- Arciuli, J., Torkildsen, J. v. K., Stevens, D. J., & Simpson, I. C. (2014). Statistical learning under incidental versus intentional conditions. *Frontiers in Psychology*, 5, 747. <https://doi.org/10.3389/fpsyg.2014.00747>, PubMed: 25071692
- Batterink, L. (2020). Syllables in sync form a link: Neural phase-locking reflects word knowledge during language learning. *Journal of Cognitive Neuroscience*, 32, 1735–1748. https://doi.org/10.1162/jocn_a_01581, PubMed: 32427066
- Batterink, L. J., & Paller, K. A. (2017). Online neural monitoring of statistical learning. *Cortex*, 90, 31–45. <https://doi.org/10.1016/j.cortex.2017.02.004>, PubMed: 28324696
- Batterink, L. J., Paller, K. A., & Reber, P. J. (2019). Understanding the neural bases of implicit and statistical learning. *Topics in Cognitive Science*, 11, 482–503. <https://doi.org/10.1111/tops.12420>, PubMed: 30942536
- Bauer, A.-K. R., Debener, S., & Nobre, A. C. (2020). Synchronisation of neural oscillations and cross-modal influences. *Trends in Cognitive Sciences*, 24, 481–495. <https://doi.org/10.1016/j.tics.2020.03.003>, PubMed: 32317142
- Benjamin, L., Dehaene-Lambertz, G., & Fló, A. (2021). Remarks on the analysis of steady-state responses: Spurious artifacts introduced by overlapping epochs. *Cortex*, 142, 370–378. <https://doi.org/10.1016/j.cortex.2021.05.023>, PubMed: 34311971
- Bogaerts, L., Richter, C. G., Landau, A. N., & Frost, R. (2020). Beta-band activity is a signature of statistical learning. *Journal of Neuroscience*, 40, 7523–7530. <https://doi.org/10.1523/JNEUROSCI.0771-20.2020>, PubMed: 32826312
- Brady, T. F., & Oliva, A. (2008). Statistical learning using real-world scenes: Extracting categorical regularities without conscious intent. *Psychological Science*, 19, 678–685. <https://doi.org/10.1111/j.1467-9280.2008.02142.x>, PubMed: 18727783
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, 10, 433–436. <https://doi.org/10.1163/156856897X00357>, PubMed: 9176952
- Choi, D., Batterink, L. J., Black, A. K., Paller, K. A., & Werker, J. F. (2020). Preverbal infants discover statistical word patterns at similar rates as adults: Evidence from neural entrainment. *Psychological Science*, 31, 1161–1173. <https://doi.org/10.1177/0956797620933237>, PubMed: 32865487
- Covington, N. V., Brown-Schmidt, S., & Duff, M. C. (2018). The necessity of the hippocampus for statistical learning. *Journal of Cognitive Neuroscience*, 30, 680–697. https://doi.org/10.1162/jocn_a_01228, PubMed: 29308986
- Davachi, L., & DuBrow, S. (2015). How the hippocampus preserves order: The role of prediction and context. *Trends in Cognitive Sciences*, 19, 92–99. <https://doi.org/10.1016/j.tics.2014.12.004>, PubMed: 25600586
- De Rosa, M., Ktori, M., Vidal, Y., Bottini, R., & Crepaldi, D. (2022). Frequency-based neural discrimination in fast periodic visual stimulation. *Cortex*, 148, 193–203. <https://doi.org/10.1016/j.cortex.2022.01.005>, PubMed: 35180482
- Ding, N., Melloni, L., Zhang, H., Tian, X., & Poeppel, D. (2016). Cortical tracking of hierarchical linguistic structures in connected speech. *Nature Neuroscience*, 19, 158–164. <https://doi.org/10.1038/nn.4186>, PubMed: 26642090
- Ding, N., & Simon, J. Z. (2013). Power and phase properties of oscillatory neural responses in the presence of background activity. *Journal of Computational Neuroscience*, 34, 337–343. <https://doi.org/10.1007/s10827-012-0424-6>, PubMed: 23007172
- Efron, B., & Tibshirani, R. (1986). Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy. *Statistical Science*, 1, 54–75. <https://doi.org/10.1214/ss/1177013815>
- Emberson, L. L., & Rubinstein, D. Y. (2016). Statistical learning is constrained to less abstract patterns in complex sensory input (but not the least). *Cognition*, 153, 63–78. <https://doi.org/10.1016/j.cognition.2016.04.010>, PubMed: 27139779
- Forest, T. A., Finn, A. S., & Schlichting, M. L. (2022). General precedes specific in memory representations for structured experience. *Journal of Experimental Psychology: General*, 151, 837–851. <https://doi.org/10.1037/xge0001104>, PubMed: 34780215
- Forget, J., Buiatti, M., & Dehaene, S. (2010). Temporal integration in visual word recognition. *Journal of Cognitive Neuroscience*, 22, 1054–1068. <https://doi.org/10.1162/jocn.2009.21300>, PubMed: 19583465
- Gebhart, A. L., Aslin, R. N., & Newport, E. L. (2009). Changing structures in midstream: Learning along the statistical garden path. *Cognitive Science*, 33, 1087–1116. <https://doi.org/10.1111/j.1551-6709.2009.01041.x>, PubMed: 20574548
- Graves, K. N., Sherman, B. E., Huberdeau, D., Damisah, E., Quraishi, I. H., & Turk-Browne, N. B. (2022). Remembering the pattern: A longitudinal case study on statistical learning in spatial navigation and memory consolidation. *Neuropsychologia*, 174, 108341. <https://doi.org/10.1016/j.neuropsychologia.2022.108341>, PubMed: 35961387
- Henin, S., Turk-Browne, N. B., Friedman, D., Liu, A., Dugan, P., Flinker, A., et al. (2021). Learning hierarchical sequence representations across human cortex and hippocampus. *Science Advances*, 7, eabc4530. <https://doi.org/10.1126/sciadv.abc4530>, PubMed: 33608265
- Jenkinson, M., Bannister, P., Brady, M., & Smith, S. (2002). Improved optimization for the robust and accurate linear registration and motion correction of brain images. *Neuroimage*, 17, 825–841. <https://doi.org/10.1006/nimg.2002.1132>, PubMed: 12377157
- Jenkinson, M., Beckmann, C. F., Behrens, T. E., Woolrich, M. W., & Smith, S. M. (2012). FSL. *Neuroimage*, 62, 782–790. <https://doi.org/10.1016/j.neuroimage.2011.09.015>, PubMed: 21979382
- Jenkinson, M., & Smith, S. (2001). A global optimisation method for robust affine registration of brain images. *Medical Image Analysis*, 5, 143–156. [https://doi.org/10.1016/S1361-8415\(01\)00036-6](https://doi.org/10.1016/S1361-8415(01)00036-6), PubMed: 11516708

- Jun, J., & Chong, S. C. (2018). Visual statistical learning at basic and subordinate category levels in real-world images. *Attention, Perception, & Psychophysics*, *80*, 1946–1961. <https://doi.org/10.3758/s13414-018-1566-z>, PubMed: 30014318
- Jung, Y., Walther, D. B., & Finn, A. S. (2021). Children automatically abstract categorical regularities during statistical learning. *Developmental Science*, *24*, e13072. <https://doi.org/10.1111/desc.13072>, PubMed: 33295082
- Jungé, J. A., Scholl, B. J., & Chun, M. M. (2007). How is spatial context learning integrated over signal versus noise? A primacy effect in contextual cueing. *Visual Cognition*, *15*, 1–11. <https://doi.org/10.1080/13506280600859706>, PubMed: 18725966
- Kabdebon, C., Peña, M., Buiatti, M., & Dehaene-Lambertz, G. (2015). Electrophysiological evidence of statistical learning of long-distance dependencies in 8-month-old preterm and full-term infants. *Brain and Language*, *148*, 25–36. <https://doi.org/10.1016/j.bandl.2015.03.005>, PubMed: 25865749
- Karuz, E. A., Emberson, L. L., Roser, M. E., Cole, D., Aslin, R. N., & Fiser, J. (2017). Neural signatures of spatial statistical learning: Characterizing the extraction of structure from complex visual scenes. *Journal of Cognitive Neuroscience*, *29*, 1963–1976. https://doi.org/10.1162/jocn_a_01182, PubMed: 28850297
- Karuz, E. A., Newport, E. L., Aslin, R. N., Starling, S. J., Tivarus, M. E., & Bavelier, D. (2013). The neural correlates of statistical learning in a word segmentation task: An fMRI study. *Brain and Language*, *127*, 46–54. <https://doi.org/10.1016/j.bandl.2012.11.007>, PubMed: 23312790
- Kiai, A., & Melloni, L. (2021). What canonical online and offline measures of statistical learning can and cannot tell us. *bioRxiv*. <https://doi.org/10.1101/2021.04.19.440449>
- Luo, Y., & Zhao, J. (2018). Statistical learning creates novel object associations via transitive relations. *Psychological Science*, *29*, 1207–1220. <https://doi.org/10.1177/0956797618762400>, PubMed: 29787352
- Norcia, A. M., Appelbaum, L. G., Ales, J. M., Cottareau, B. R., & Rossion, B. (2015). The steady-state visual evoked potential in vision research: A review. *Journal of Vision*, *15*, 4. <https://doi.org/10.1167/15.6.4>, PubMed: 26024451
- Nozaradan, S., Peretz, I., Missal, M., & Mouraux, A. (2011). Tagging the neuronal entrainment to beat and meter. *Journal of Neuroscience*, *31*, 10234–10240. <https://doi.org/10.1523/JNEUROSCI.0411-11.2011>, PubMed: 21753000
- Oostenveld, R., Fries, P., Maris, E., & Schoffelen, J.-M. (2011). Fieldtrip: Open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. *Computational Intelligence and Neuroscience*, *2011*, 156869. <https://doi.org/10.1155/2011/156869>, PubMed: 21253357
- Otsuka, S., Nishiyama, M., Nakahara, F., & Kawaguchi, J. (2013). Visual statistical learning based on the perceptual and semantic information of objects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *39*, 196–207. <https://doi.org/10.1037/a0028645>, PubMed: 22686848
- Papademetris, X., Jackowski, M. P., Rajeevan, N., DiStasio, M., Okuda, H., Constable, R. T., et al. (2006). Bioimage suite: An integrated medical image analysis suite: An update. *Insight Journal*, *2006*, 209. <https://doi.org/10.54294/2g80r4>, PubMed: 25364771
- Park, S. H., Rogers, L. L., & Vickery, T. J. (2018). The roles of order, distance, and interstitial items in temporal visual statistical learning. *Attention, Perception, & Psychophysics*, *80*, 1409–1419. <https://doi.org/10.3758/s13414-018-1556-1>, PubMed: 29956264
- Pelli, D. G. (1997). The videotoolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*, 437–442. <https://doi.org/10.1163/156856897X00366>, PubMed: 9176953
- Pinto, D., Prior, A., & Zion Golumbic, E. (2022). Assessing the sensitivity of EEG-based frequency-tagging as a metric for statistical learning. *Neurobiology of Language*, *3*, 214–234. https://doi.org/10.1162/nol_a_00061
- Plate, R. C., Schapiro, A. C., & Waller, R. (2022). Emotional faces facilitate statistical learning. *Affective Science*, *3*, 662–672. <https://doi.org/10.1007/s42761-022-00130-9>, PubMed: 36385906
- Preston, A. R., & Eichenbaum, H. (2013). Interplay of hippocampus and prefrontal cortex in memory. *Current Biology*, *23*, R764–R773. <https://doi.org/10.1016/j.cub.2013.05.041>, PubMed: 24028960
- Prichard, D., & Theiler, J. (1994). Generating surrogate data for time series with several simultaneously measured variables. *Physical Review Letters*, *73*, 951–954. <https://doi.org/10.1103/PhysRevLett.73.951>, PubMed: 10057582
- Schapiro, A. C., Kustner, L. V., & Turk-Browne, N. B. (2012). Shaping of object representations in the human medial temporal lobe based on temporal regularities. *Current Biology*, *22*, 1622–1627. <https://doi.org/10.1016/j.cub.2012.06.056>, PubMed: 22885059
- Sherman, B. E., Graves, K. N., Huberdeau, D. M., Quraishi, I. H., Damisah, E. C., & Turk-Browne, N. B. (2022). Temporal dynamics of competition between statistical learning and episodic memory in intracranial recordings of human visual cortex. *Journal of Neuroscience*, *42*, 9053–9068. <https://doi.org/10.1523/JNEUROSCI.0708-22.2022>, PubMed: 36344264
- Sherman, B. E., Graves, K. N., & Turk-Browne, N. B. (2020). The prevalence and importance of statistical learning in human cognition and behavior. *Current Opinion in Behavioral Sciences*, *32*, 15–20. <https://doi.org/10.1016/j.cobeha.2020.01.015>, PubMed: 32258249
- Sherman, B. E., & Turk-Browne, N. B. (2020). Statistical prediction of the future impairs episodic encoding of the present. *Proceedings of the National Academy of Sciences, U.S.A.*, *117*, 22760–22770. <https://doi.org/10.1073/pnas.2013291117>, PubMed: 32859755
- Sherman, B. E., & Turk-Browne, N. B. (2022). Attention and memory. In M. J. Kahana & A. D. Wagner (Eds.), *The Oxford handbook of human memory*. Oxford University Press. <https://doi.org/10.31234/osf.io/xx6db>
- Siegelman, N., Bogaerts, L., Kronenfeld, O., & Frost, R. (2018). Redefining “learning” in statistical learning: What does an online measure reveal about the assimilation of visual regularities? *Cognitive Science*, *42*, 692–727. <https://doi.org/10.1111/cogs.12556>, PubMed: 28986971
- Störmer, V. S., & Alvarez, G. A. (2014). Feature-based attention elicits surround suppression in feature space. *Current Biology*, *24*, 1985–1988. <https://doi.org/10.1016/j.cub.2014.07.030>, PubMed: 25155510
- Turk-Browne, N. B., Jungé, J. A., & Scholl, B. J. (2005). The automaticity of visual statistical learning. *Journal of Experimental Psychology: General*, *134*, 552–564. <https://doi.org/10.1037/0096-3445.134.4.552>, PubMed: 16316291
- Turk-Browne, N. B., Scholl, B. J., Chun, M. M., & Johnson, M. K. (2009). Neural evidence of statistical learning: Efficient detection of visual regularities without awareness. *Journal of Cognitive Neuroscience*, *21*, 1934–1945. <https://doi.org/10.1162/jocn.2009.21131>, PubMed: 18823241
- Turk-Browne, N. B., Scholl, B. J., Johnson, M. K., & Chun, M. M. (2010). Implicit perceptual anticipation triggered by statistical learning. *Journal of Neuroscience*, *30*, 11177–11187. <https://doi.org/10.1523/JNEUROSCI.0858-10.2010>, PubMed: 20720125
- Zhou, Z., Singh, D., Tandoc, M. C., & Schapiro, A. C. (2023). Building integrated representations through interleaved learning. *Journal of Experimental Psychology: General*. <https://doi.org/10.1037/xge0001415>