Visual statistical learning is associated with changes in low-dimensional cortical architecture

Keanna Rowchan[1*], Daniel J. Gale[1], Qasem Nick[1,2], Jason P. Gallivan[1,2,3,‡], Jeffrey D. Wammes[1,2,‡]

[1]Department of Psychology, [2]Centre for Neuroscience Studies, and [3] Queen's University, Kingston, Ontario, Canada, K7L 3N6.

* Corresponding author
‡ The last two authors contributed equally to this work.

ABSTRACT

Our ability to automatically learn environmental regularities over time to make predictions about the world is central to human cognition and behavior. Despite its importance, the precise roles of cortical and subcortical regions in this statistical learning (SL) process remain debated, with limited understanding of how these different brain systems reorganize their activity during learning. To address this, we analyzed human functional MRI brain activity during a visual SL task, whereby individuals implicitly learned to associate pairs of images embedded within a larger sequence. By projecting individuals patterns of cortical and subcortical functional connectivity onto a low-dimensional manifold space, we found that SL was associated with changes along a single neural dimension describing connectivity across the visual-parietal and perirhinal cortex (PRC). During learning, we found that regions within the visual cortex expanded along this dimension, reflecting their increased segregation from other networks, whereas regions within the dorsal attention network (DAN) contracted, reflecting their integration with higher-order transmodal cortex. Notably, when SL was interrupted, we found that the PRC and entorhinal cortex, which did not initially show learning-related effects, now contracted along this same neural dimension, reflecting their increased integration with the default mode and DAN, and decreased covariance with visual cortex. While prior work has associated SL with either broader cortical, or medial temporal lobe activity changes, our findings suggest a more integrative view, whereby cortical regions reorganize during the formation of learned associations and the medial temporal lobe responds to their violation.

Keywords: statistical learning, manifold, connectivity, gradients, hippocampus, medial temporal lobe

SIGNIFICANCE STATEMENT

We investigated changes in whole-brain functional connectivity as participants implicitly learned temporal associations between visual stimuli. We found that areas of the visual cortex and dorsal attention network showed significant connectivity changes during learning, reflecting their segregation from, and integration with, other networks, respectively. Notably, when the learned associations were disrupted, medial temporal lobe regions that showed no evidence of initial learning, now increased their integration with higher-order cortex. Together, these results offer a unique, whole-brain perspective onto the neural basis of visual statistical learning, and suggest separable cortical and medial temporal lobe contributions for the encoding versus violation of learned associations.

INTRODUCTION

Each day, we are inundated with a constant influx of sensory information. Despite the apparent randomness of these inputs, we possess the remarkable ability to discern repeating patterns in our environment. This ability to implicitly extract environmental regularities over time is referred to as statistical learning (SL), and it enables us to make predictions about the world and modify our actions accordingly (1-3). One particularly well studied form of such learning is visual SL, which examines our ability to generate internal representations of, and extract probabilities for, visual stimuli that co-occur (e.g., in a temporal sequence or spatial array; 4-5).

At a neural level, it was traditionally assumed that SL engaged task-specific sensory cortical circuits, whereas episodic learning relied on the medial temporal lobe (MTL; primarily hippocampus, but also entorhinal cortex (ERC), perirhinal cortex (PRC) and parahippocampal cortex; 6-7). In recent years, however, these neural distinctions have eroded in light of new evidence. First, it has been shown that, in addition to recruiting sensory-specific cortices, SL also recruits higher-order association networks (8-10), suggesting that SL requires both task-specific processing (e.g., visual; 11-12, or auditory regions; 13-16) and task-general cognitive abilities that allow for the shifting of attention and changes in the allocation of resources (17-19). Second, contrary to traditional views that regularities were learned exclusively in cortical networks, recent studies reveal that the MTL and hippocampus exhibit representational change and differential responses during SL (11; 20-22), and that hippocampal volume predicts SL performance (23-24). Further challenging conventional views, studies in patients have shown that SL can be completely abolished following bilateral damage to the hippocampus and surrounding MTL areas (25), suggesting a critical role for these areas in SL.

While the above findings align with the MTL's established role in learning (26-28) and its sensitivity to temporal and spatial information (29-33), they suggest a need to update the classic cortical-MTL dichotomy. One model proposes a division of labor in the hippocampal system, in which the monosynaptic pathway (between the ERC and CA1) supports SL, and the trisynaptic pathway (between the ERC, CA1, CA3, and dentate gyrus, DG) supports episodic learning (34). Consistent with this model, direct stimulation of CA1 impairs SL (35), but focal DG lesions do not (36). Despite these findings being consistent with the hippocampus' sensitivity to

regularities, recent work suggests it may not be *necessary* for SL. Indeed, several patients with extensive bilateral hippocampal lesions show evidence of SL and at times, outperform controls, who often fail to demonstrate SL themselves (19; 37-40). These conflicting findings challenge the necessity of the hippocampus for SL and leave the relative contributions of the cortex and MTL unclear.

While many individual brain areas have been identified as contributors to SL (8-9; 13), it is widely recognized that such learning involves the coordination and integration of information processing across multiple brain regions (see 41-43). Indeed, while prior work points to the involvement of the MTL, as well as task-specific and task-general cortical networks in SL, these brain areas project widely across the cortex and are reciprocally interconnected with several distributed, whole-brain networks (44-45). Understanding these interconnections requires the study of functional connectivity between brain areas embedded within large-scale networks, as well as analytical approaches that are able to effectively characterize the changes in whole-brain functional architecture that occur during learning. Such an approach is not only important for characterizing whole-brain functional interactions, but also understanding at which specific stages of learning these interactions occur. This allows for critical questions to be addressed such as the relative extent to which changes in cortical or MTL connectivity occur during SL, the time windows over which such changes occur, and/or how these different systems might independently or cooperatively encode statistical regularities (19; 21; 37-38).

Recent advances in dimension reduction techniques have provided new tools for effectively studying large-scale patterns of brain activity among distributed networks (46-49). These methods allow high-dimensional patterns of functional connectivity to be characterized along a low-dimensional subspace or manifold, reflecting the main patterns of covariance across brain regions. This approach has not only provided fundamental insights into the principles that govern the collective activity of large-scale neural populations in many brain areas (50-51), but also into the intrinsic functional organization of the brain in both healthy and patient populations (52-54). More recently, this method has been extended to explore temporal changes in whole-brain functional organization as subjects perform various tasks (46; 55-56). Here, we adopted this approach to understand how cortical and MTL activity patterns are coordinated during various stages of visual SL. This allowed us to capture the evolving landscape of brain

activity during SL and understand how brain regions collectively organize their interactions over time to adapt to changing patterns in our environment.

RESULTS

The current study analyzed previously published functional magnetic resonance imaging (fMRI) data (22) from participants (N = 33) performing a classic visual SL task, in which they viewed a sequence of abstract images presented one at a time on screen (see Fig. 1A). The task was divided into eight blocks, where participants viewed a sequence of 80 images consisting of 5 repetitions of 16 unique images, resulting in a total of 640 image presentations. During six of these task blocks (blocks 2-7), image sequences contained an embedded set of regularities, where the 16 images were combined into 8 pairs, such that one image (A) was always followed by its pairmate (B; See Fig. 1A, 1B left). These SL blocks were bookended by two, unstructured task blocks (blocks 1 and 8), where images were presented in a completely random order (see Fig. 1B right). Importantly, participants were not explicitly aware of any embedded structure, and were only instructed to engage in a basic perceptual task that required the maintenance of attention, completely orthogonal to the pair structure (see Materials and Methods; 22).

Statistical learning results in neural changes at the regional level

Before examining any whole-brain changes in functional connectivity, we first sought to verify that subjects indeed learned temporal associations between stimuli during the SL blocks. Prior work, including our own, has shown that visual SL leads to changes in stimulus-specific patterns of neural activity in the hippocampus (20-22; 57). Specifically, depending on the visual similarity between paired images, learning can lead to either neural integration of representations, where two temporally coupled stimuli become more similar in their neural representations, or neural differentiation, whereby they become less similar. Recent theoretical (58-59), modeling (60-61) and empirical (22; 62-63) work indicates that this relationship follows a U-shaped function, as described by the nonmonotonic plasticity hypothesis (NMPH: 59-60). In our prior work based on these data (22), the strength of this U-shaped function in the DG provided the primary measure that learning had taken place. Consistent with this prior finding, when we performed a re-analysis on the current data including only the subset of subjects used for our connectivity analyses (see the Supplementary Extended Materials and Methods), we

observed a U-shaped learning effect in the CA1 (r=0.11, 95% CI=[-0.010, 0.232], p=0.0351) and DG (r=0.11, 95% CI=[-0.017, 0.248], p=0.0451), but not the CA2/3 (r=0.07, 95% CI=[-0.044, 0.180], p=0.115), or across the entire hippocampus when subfields were combined (r=0.07, 95% CI=[-0.032, 0.180], p=0.088; see Fig. 1C). These results replicate prior effects (22) and demonstrate that SL induced the hypothesized representational changes in hippocampus. Building on this foundational result, we next examined whether, and to what extent, functional interactions both within and outside of the MTL system are also altered during SL.

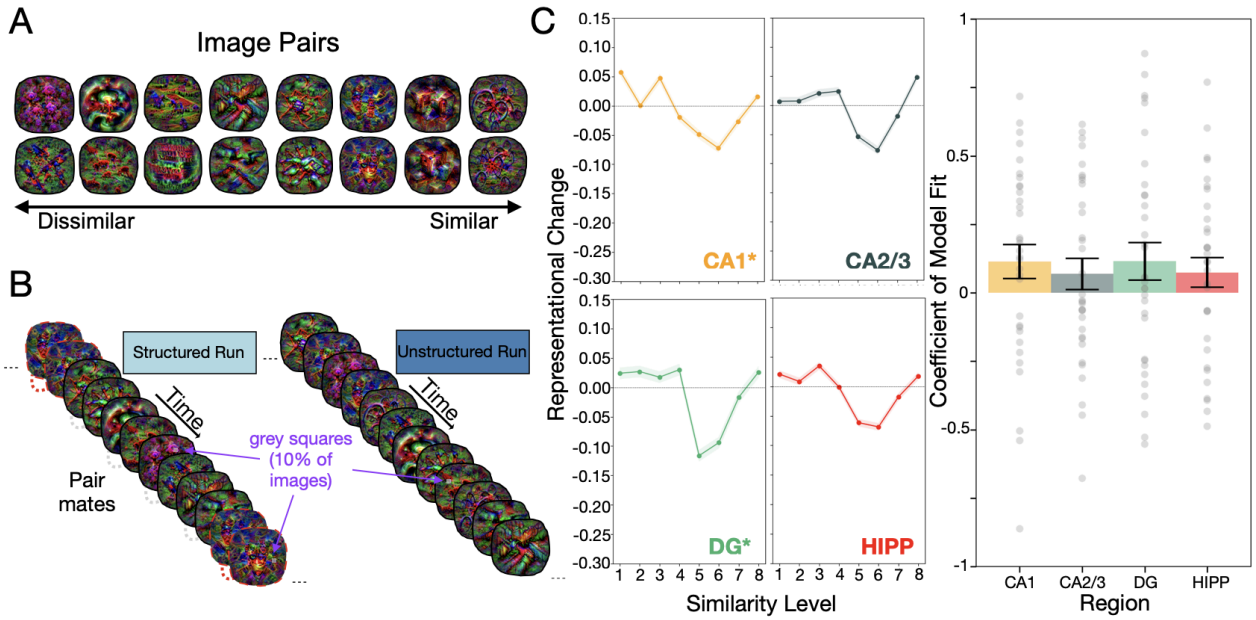

Figure 1. Overview of visual SL task and representational change predicted by the non-monotonic plasticity hypothesis. (A) Image pairs used in the visual SL task (22), ordered from most visually dissimilar, to most visually similar. (B) Visual SL task structure. Participants viewed images presented one at a time on screen. Structured runs contained an embedded pair structure, whereby an A image was always followed by its pairmate (B image; example of repeating pair mates denoted by red dashed line). By contrast, random runs presented images in a completely random order. Each run consisted of 80 image presentations (5 repetitions of each image or pair). Participants were tasked with pressing a button when a small faded gray square appeared in the image (10% of trials). (C) Left, Relationship between image similarity and representational change predicted by the nonmonotonic plasticity hypothesis (NMPH). For each region (CA1, CA2/3, DG and whole hippocampus; HIPP), we correlated parameter estimates derived from a general linear model for each A and B image in a pair, and then created comparison scores by subtracting Pre- from Post-learning correlation scores. This measure quantifies the degree of representational change that takes place between image pairs as participants underwent visual SL (i.e., from Early- to Late-learning), with positive scores indicating integration between two images in a pair, and negative scores indicating differentiation. The comparison scores were plotted as a function of similarity level (i.e.,

coactivation) for each of the eight image pairs (ordered from least to greatest visual similarity in paired items). Shaded areas depict the bootstrap resampled standard errors of the means (SEMs). Right, theory-constrained cubic models were fit to the representational change patterns in all but one held-out participant, then used to predict the held-out participants' scores. These individual-specific model fits are shown in the light gray scatter plots. Barplots display the group-level model fit, as measured by the correlation between actual and predicted values. At a group-level, representational change was reliably predicted by a theory-constrained U-shaped model in the CA1 and DG. Error bars depict SEM.

Investigating changes in functional brain organization during statistical learning

To investigate changes in whole-brain functional organization during visual SL, we selected one functional run representing each of four key epochs of the task. First, we defined Early- and Late-learning epochs based on the first and last structured blocks (runs 2 and 7) of the SL task, respectively. We next defined Pre- and Post-learning epochs based on the first and last unstructured blocks (runs 1 and 8), in which the images were presented in a completely random order. For each participant, we extracted the mean blood oxygenation level-dependent (BOLD) timeseries data for each cortical region pre-defined by the Schaefer-1000 parcellation (64), and MTL region using individual-specific parcellations generated using the Automated Segmentation of Hippocampal Subfields software package (ASHS; 65). We then computed covariance matrices (representing functional connectivity between each pairwise set of regions) for each of our task epochs (Pre-, Early-, Late-, and Post-learning). [see Materials and Methods, Fig. 2A; for full methodology, see Supplementary Extended Materials and Methods].

Prior work indicates that individual differences account for the majority of the variance in patterns of brain activity, and thus can easily obscure any task-related changes in functional connectivity (46, 66-67). We have recently developed a Riemannian manifold centering approach (67) to mitigate the effects of these individual differences, allowing task-related changes in brain structure to be more readily detected (see Methods and Materials; 68). To illustrate the beneficial impact of this centering method, we projected participants' covariance matrices before and after centering using uniform manifold approximation (UMAP; 69). UMAP effectively reduces the high-dimensional covariance matrix data onto a 2D plane, providing a visualization of the similarities/differences between the covariance matrices. Consistent with prior work, including our own (46; 66; 70), before centering we found covariance matrices mainly clustered according to subject-level identity (see Fig. 2B, left). However, after centering,

this subject-level clustering was completely abolished, allowing the possibility of teasing apart learning-related differences in functional brain structure over time (see Fig. 2B, right).

We used these centered matrices to estimate cortical-subcortical connectivity manifolds for each participant's Pre-, Early-, Late-, and Post-learning covariance matrices. Building from prior work (46; 70-72), we converted each matrix into an affinity matrix before applying PCA to the entire sample's affinity matrices. This reveals a set of principal components (PCs), which offer a low-dimensional representation of cortical and MTL structure (i.e., a cortical-MTL manifold). Finally, we aligned each manifold to a common Pre-learning template manifold, constructed using the mean of all participants' Pre-learning matrices (see Fig. 2A). This alignment served two important purposes: First, it provided a common target for manifold alignment, such that all participants and epochs could be directly compared within a common neural manifold space. Second, it allowed us to directly examine learning-specific deviations away from this template (Pre-learning) architecture, thus enhancing our sensitivity to detect subsequent learning-related changes in manifold structure (46).

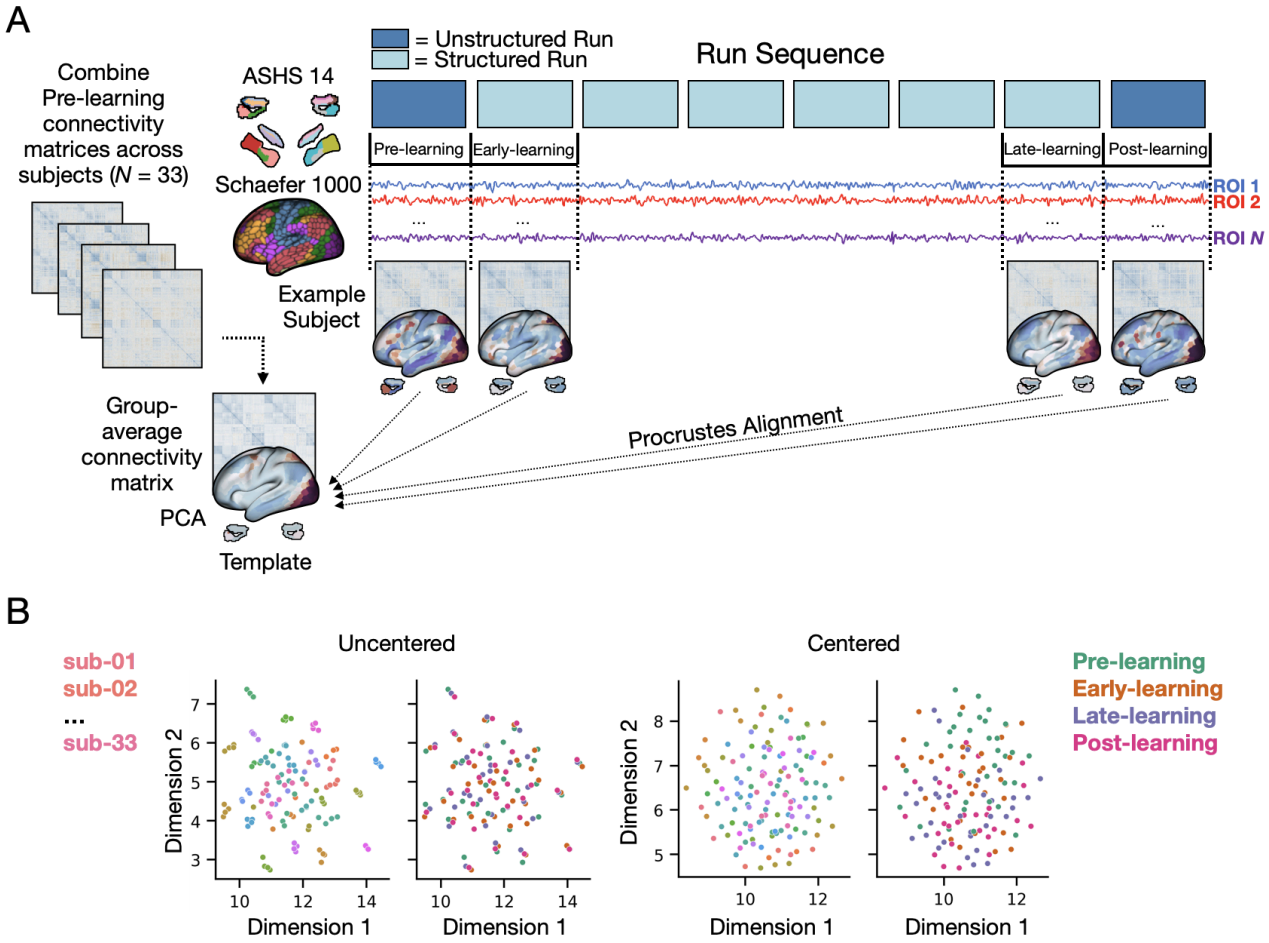

Figure 2. Overview of functional connectivity analysis. (A) Functional connectivity analysis approach. For each participant and task epoch (Pre-, Early-, Late-, and Post-learning), we computed functional connectivity matrices from parcel-wise timeseries data extracted from the Schaefer-1000 cortical parcellation (64) and participant-specific hippocampal parcellations derived from the Automated Segmentation of Hippocampal Subfields (ASHS; 65) software. We then estimated functional connectivity manifolds for each task epoch using PCA with centered and thresholded connectivity matrices (see Materials and Methods). All manifolds were aligned to a common template manifold created from the group-average Pre-learning connectivity matrix using Procrustes alignment. (B) Uniform Manifold Approximation (UMAP; 69) visualization of connectivity matrices, both before and after centering. Single data points represent individual matrices. For both uncentered and centered data, the left plots have the data points colored according to subject identity whereas the right plots have the data points colored according to task epoch. Note that prior to centering, matrices were mainly clustered according to subject identity whereas this cluster structure was abolished following centering.

Cortical and subcortical manifold structure prior to learning

The top three PCs of our Pre-learning template manifold describe the main dimensions of cortical and subcortical functional connectivity observed during the task (see Fig. 3A), and explain the majority of the variance in the data (60.0%; see Fig 3B). PC1 distinguishes the visual network, areas of the parietal dorsal attention network (DAN), and PRC (positive loadings in red) from other cortical regions (negative loadings in blue). PC2 separates transmodal, default mode network (DMN) regions (in red) from somatomotor regions (in blue). Lastly, PC3 distinguishes medial areas of the visual cortex (in red) and areas within the DAN (in blue). Together, PC1 and PC2 mainly differentiate visual and somatomotor regions from transmodal regions of the DMN (see Fig. 3C, left and middle), approximating the brain's intrinsic functional architecture, as observed in prior work using resting-state data (48-49). This tripartite architecture is thought to reflect a crucial feature of cortical brain organization, with the shift from unimodal to transmodal regions of the brain signifying a global processing hierarchy that ascends from lower-order sensory and motor systems, to higher-order association areas of the cortex, such as the DMN (73). By contrast, PC3 distinguishes association regions of the DAN from sensory-specific regions of visual cortex, previously implicated in visual SL (74; see Fig. 3C, right). For a complete visualization of this tripartite structure, we projected each brain region's loadings on all three PCs onto a three-dimensional manifold space (see Fig. 3D).

To better understand the typical functions associated with each individual dimension of connectivity that emerged during SL, we conducted a meta-analysis using the Neurosynth Decoder (75). This tool quantifies patterns of whole-brain activity against a comprehensive dataset of whole-brain maps from thousands of fMRI studies. The results of this analysis are in the form of correlation values, revealing the association between the whole-brain map for each PC, and various keywords which often reflect tasks or cognitive functions. For the purposes of visualization, we have transformed these correlation values into word clouds (see Fig. 3E), where the color and size of a keyword represents the direction and magnitude of the correlation. We used these word cloud visualizations to assign meaningful functional labels to each PC. We have termed PC1 the visual-objects axis, emphasizing the role of its constituent regions in object recognition and visual processing; PC2 the autobiographical-episodic axis,

given the role of its constituent regions in memory and time-related cognition; and PC3 the working-tasks axis, given the role of its constituent regions in multiple tasks.

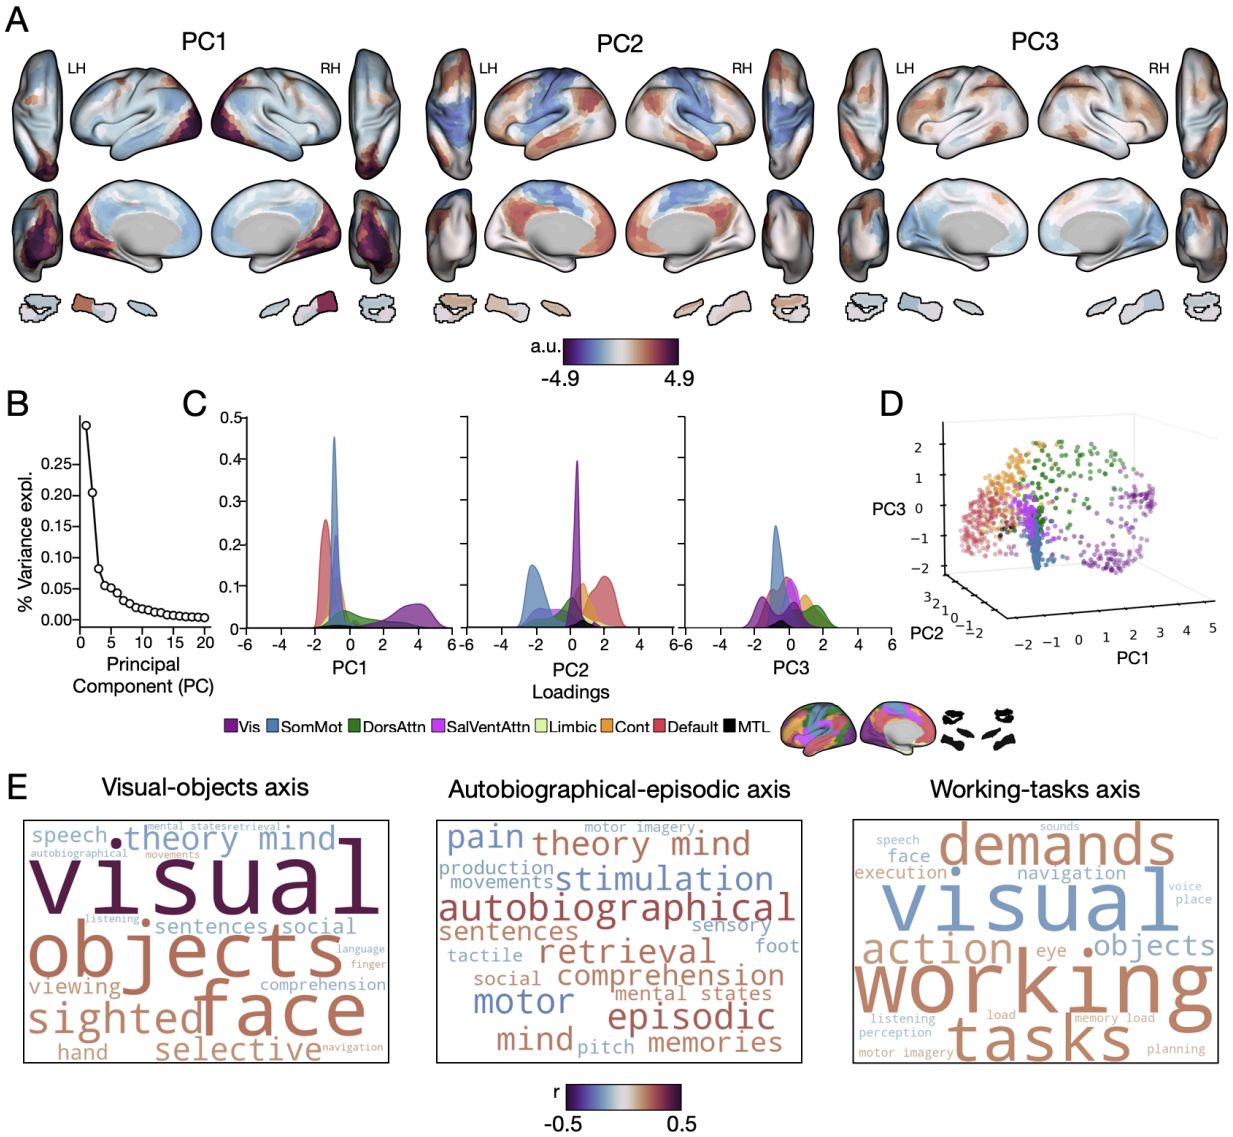

Figure 3. Functional organization of the Pre-learning epoch. (A) Cortical-Subcortical loadings for each of the top three principal components (PCs), projected in surface space. (B) Percent variance explained for the first 20 PCs. (C) Histograms displaying the distributions of each network's loading along each PC, colored according to its functional network assignment, denoted on the bottom (76). (D) Functional organization of the Pre-learning epoch. Each brain region is depicted as a point in a three-dimensional manifold space, with its location determined by its loading on the first three PCs. (E) Meta-analysis of PC loadings submitted to the Neurosynth Decoder tool (75). Word clouds show the top and bottom ten functions likely associated with each PC. The text size and colour within each word cloud describes their

correlation value (bigger words have higher correlation values and smaller words have smaller values) and their polarity (words in red are more positive and words in blue are more negative).

While together, the top three PCs recreate the tripartite structure seen in prior work (46-49; see Fig. 3D), the order of variance explained by these orthogonal dimensions offers insights into the specific neural demands of our task. Notably, prior resting-state studies using the same dimension reduction method (PCA) have identified our visual-object connectivity axis (PC1) as explaining the third-highest proportion of variance, while PC2 and PC3 explained the first and second highest variances, respectively (see 46-48). Considering the stimuli in our task were designed to specifically engage higher-order object-selective visual regions (22), it is intuitive that PC1 explains the greatest amount of variance in our data. Given this, we next examined whether visual SL selectively modulated connectivity motifs along this visual-object axis.

Changes along the visual-object connectivity axis during statistical learning

Figure 4A shows the average loading of each brain region along the autobiographical-episodic, and visual-object axes as a function of task epoch (Early-, Late- and Post-learning). Notably, we found that Late-learning appeared to be associated with a compression along the visual-object axis, which then expanded along its top edge once the learned associations were interrupted during Post-learning. To directly test whether this pattern of compression and expansion along the visual-objects axis was significant across task epochs, we calculated, for each participant, the min-max spread along this axis for each of our four task epochs (Pre-, Early-, Late-, and Post-learning) and performed a repeated measures ANOVA (rmANOVA) on these values. This analysis revealed a significant effect of epoch along the visual-objects axis ($F(3, 96)=5.10$, $p=0.008$), indicating changes (contraction and expansion) in the distribution of regional embeddings along this axis during SL (see 53 for a similar approach). Notably, when we performed this same analysis for the autobiographical-episodic (PC2) and working-task (PC3) axes we observed no significant changes in their expression across epochs (autobiographical-episodic axis: $F(3, 96)=0.61$, $p=0.61$; working-task axis: $F(3, 96)=2.47$, $p=0.10$).

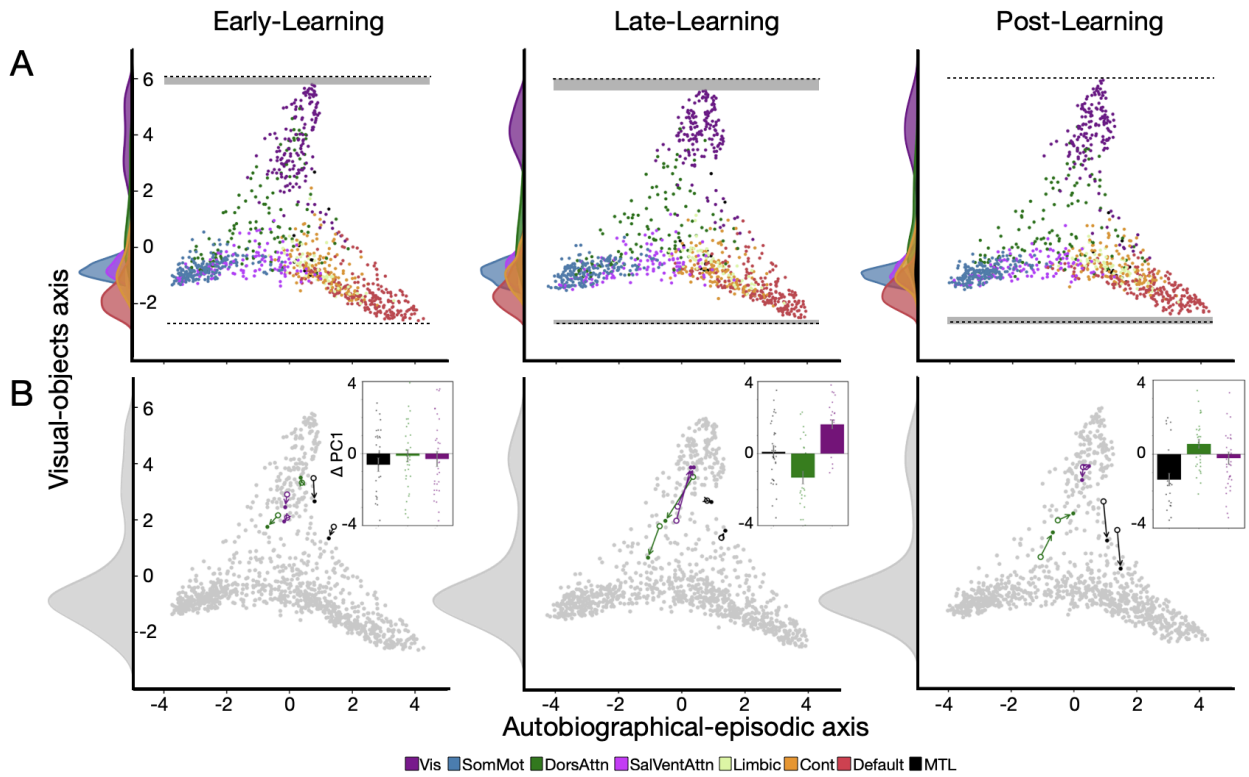

Figure 4. Changes in group-average manifold loadings along the visual-objects axis. (A) Changes in functional organization during Early-, Late- and Post-learning epochs. Each brain region is shown as a point in a two-dimensional manifold space along the visual-objects (PC1) and autobiographical-episodic (PC2) axis, with its color determined by its functional network assignment, denoted at the bottom (76). The horizontal dashed line and bars across epochs (in grey) allows for a direct visual comparison of the change in the expression of the visual-objects axis across task epochs. Histograms along the Y-axis show point densities for each epoch and functional network. (B) Temporal trajectories represent left and right hemisphere seed regions. The starting position of each seed region from the preceding epoch is marked by an open white circle, with arrows indicating the displacement of each region in the current epoch (colored-in endpoint; 76). Gray histograms along the Y-axis display the mean density of points across all regions. Inset bar plots within each task epoch display the average change (Δ) in loading for each pair of seed regions from the preceding to current epoch. Negative and positive values reflect decreases and increases along the visual-objects axis, respectively. Single data points depict single subjects. Error bars depict standard error of the mean (SEM).

To further visualize these changes in regional embeddings along the visual-object axis, in Fig. 4B we highlight three pairs of brain regions (a left and right hemisphere pair) for three distinct brain networks that load positively onto this axis (visual network regions are in purple, DAN in

green, and PRC in black). We highlight these sets of regions to show that, during SL, there appear to be large, yet differentiable shifts in the embedding of certain regions across epochs. For example, the visual network and DAN regions appear to exhibit their largest shifts along the visual-object axis from Early- to Late-learning (see Fig. 4B middle), albeit in different directions. By contrast, the PRC appears to exhibit the largest shifts along this axis from Late- to Post-learning (see Fig. 4B right). To directly test the statistical significance of these and any other regional effects, for each brain region we performed a rmANOVA on their loading scores onto the visual-object axis across the four task epochs. Following FDR-correction across brain regions for multiple comparisons ($q < 0.05$; 77), this analysis showed that 129 regions exhibited a significant main effect of task epoch (see Fig. 5A). Further inspection revealed that these significant regions were mainly distributed within two major cortical networks: the visual network (63 regions) and the DAN (37 regions); however, we also found much smaller clusters within the DMN (9 regions), cingulo-opercular network (8 regions), somatomotor network (6 regions), limbic network (2 regions), and salience and ventral attention network (1 region). Notably, in addition to these cortical regions, our rmANOVAs also revealed a handful of significant areas within the MTL, including left and right PRC, and the right ERC (see Fig. 5A). To visualize the trajectory changes for each of these significant regions during SL, we have plotted the change in embedding of each region along the visual-object axis in Fig. 5B.

To further characterize the above main effects, we performed follow-up FDR-corrected ($q < 0.05$; 77) paired-samples t-tests between epochs, both at the level of each significant brain region (determined via the rmANOVA), and aggregated across all regions (regardless of significance) within a given functional brain network (76). To test whether any regions significantly changed in their loading along the visual-object connectivity axis at the onset of learning (when participants were first exposed to a structured run after the first unstructured run), we performed an Early-learning > Pre-learning contrast. This contrast mainly revealed a pattern of contraction along the visual-objects axis in regions within the visual cortex and DAN (outlined blue regions; see Fig. 5C top), which could also be observed qualitatively when aggregated at the level of functional brain networks (i.e., when we averaged regional loadings according to their network assignments, see Fig. 5C bottom). Notably, however, the network-level loading changes in loadings between the Pre- and Early-learning epochs were not significant after we applied an FDR ($q <0.05$; 77) correction (see Fig. 5C). (Note that, for spider

plots, we used the 17-network cortical mapping in addition to mapping the hippocampus and broader MTL in order to capitalize on the improved spatial precision compared to the 7-network mapping; 64-65, 76).

Next, to examine whether brain regions changed in their embeddings along the visual-object connectivity axis from the onset of learning to the end of learning, we performed a Late-learning > Early-learning contrast. This contrast revealed that several regions within the visual cortex, which had originally exhibited contraction from Pre- to Early-learning, now exhibited significant expansion along the visual-object axis (red areas in Fig. 5D top). In addition to these changes, we also observed a further contraction of several areas of the DAN, which had initially begun modestly contracting from Pre- to Early-learning (compare Fig. 5D top with 5C top). However, when extrapolated to the whole-brain functional network level, we found that after FDR ($q < 0.05$) correction, it was only the visual network that showed significant changes along the visual-object axis from Early to Late learning (see Fig. 5D, bottom).

Although these above neural changes from Early- to Late-learning are presumed to reflect an intrinsic process related to SL, it is possible that these effects are merely attributable to the passage of time. That is, rather **than** being learning-related in nature, these neural changes may simply reflect a general habituation effect and/or decrease in attention as the task progresses. Behavior from the perceptual cover task indicates that this is unlikely to be the case, as participants actually displayed faster reaction times and greater accuracy as time went on in the task, improving from both Early- to Late-Learning, and from Pre- to Post-learning (both $ps < 0.05$). Nevertheless, to test whether this alternative explanation can account for our results at the neural level, we examined whether the observed pattern of changes (i.e., expansion of visual cortex regions) simply continued into the final, Post-learning epoch, when the learned sequence of structured images was interrupted. Notably, when we contrasted the Post-learning > Late-learning epochs, we found that regions within the MTL — specifically areas of the PRC and ERC — as well as the retrosplenial cortex exhibited significant contractions along the visual-object axis (see Fig. 5E, top). Consistent with these observations, at the whole-brian level we observed significant network changes in both the MTL and DMN-C (see Fig. 5E, bottom). Notably, the visual network which exhibited large changes from Early- to Late-learning changed only minimally during the Post-learning period — and in the opposite direction — indicating

that the observed learning effects were not simply due to the passage of time. Together, these results suggest that, once SL was interrupted (during Post-learning), retrosplenial cortex and regions within the MTL, which did not exhibit significant changes during the learning period, modified their connectivity along the visual-object axis.

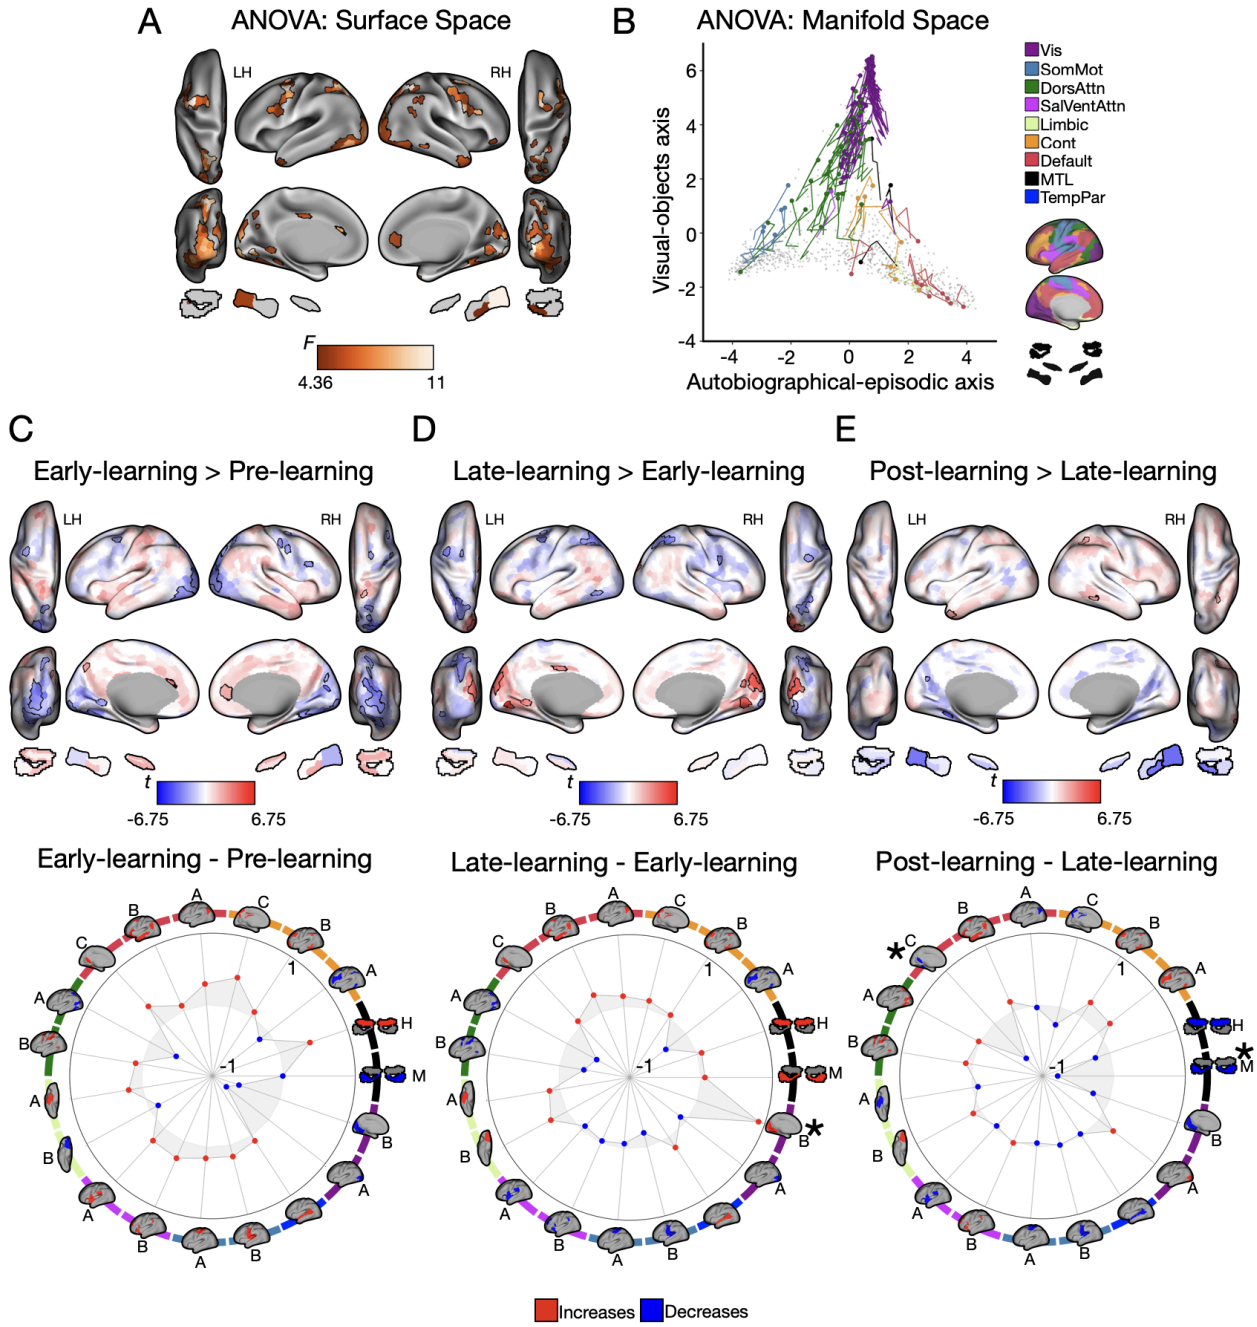

Figure 5. Visual SL is associated with changes along the visual-object connectivity axis. (A) Regions displaying significant changes in embedding across task epochs. (B) Temporal

trajectories of the significant regions from A. Each region is depicted as a single point in the two-dimensional axis space, with significant regions colored according to their functional network assignments on the right (76). Colored circles denote each significant regions embedding during the Pre-learning epoch, and each trace attached to the circle represents that region's displacement along the two axes throughout the subsequent learning epochs (Early-, Late- and Post-learning). (C-E) Top, Unthresholded maps of pairwise comparisons between task epochs. Outlined regions denote significant pairwise comparisons. Positive (red) values indicate increases in loadings (expansion along the axis) and negative (blue) values indicate decreases (contraction along the axis). Bottom, Polar plots depicting mean functional network loading differences for pairwise comparisons. Positive (red) and negative (blue) dots reflect average increases and decreases along the axis. The color behind each brain indicates its functional network assignment (legend in B), with letters depicting its constituent subnetwork assignment (76). Asterisks indicate FDR-corrected significance in network changes.

Changes in regional connectivity that underlie changes along the visual-objects axis

While significant changes in the embedding of regions along the visual-objects axis indicate a change to a brain region's overall connectivity profile, this does not itself reveal *how* these regions change their connectivity with other brain areas. To unpack this, we performed seed connectivity analyses using representative regions from left and right hemispheres of the visual network, DAN, and PRC. Together, these regions capture the main patterns of effects observed with respect to the expansions and contractions along the visual-objects axis from Early- to Late- and Late- to Post-learning (see Fig. 6 for seed analyses of the right hemisphere regions; see Fig. S1 for the left hemisphere analogues). Note that the seed regions used in our analyses were the same regions highlighted in Fig. 3B, enabling a direct link from understanding changes in whole-brain functional connectivity to the displacement of regions along the visual-object axis.

For each seed region, we characterized both the learning effect and the effect associated with the violation of learned associations by contrasting the associated whole-brain connectivity maps for Early- to Late-learning and Late- to Post-learning with region-wise paired-samples t-tests. In Fig. 6 we display (1) the unthresholded region-wise contrast maps for each seed region to allow for a complete visualization of a given region's connectivity changes (see Fig. 6 left), and (2) corresponding polar plots, depicting the network-level changes in connectivity for each region (see Fig. 6 right). Together, these plots provide a comprehensive description of the collective changes in connectivity that contribute to the changes in regional embedding along the visual-objects axis.

For the right visual seed region (see Fig. 6A), which exhibited significant expansion along the visual-object axis from Early- to Late-learning (see Fig. 5D), we found that its connectivity during SL was primarily characterized by increased covariance with other areas of the visual network, as well as with a few cortical regions belonging to the DMN and limbic networks. By contrast, from Late- to Post-learning, we observed an overall reversal in this pattern of connectivity, with the visual seed region now exhibiting decreased covariance with those initial areas as well as the MTL, and increased covariance with the Salience/Ventral Attention (SalVentAttn) network. We should note, however, that these reversals in connectivity did not result in a statistically significant contraction of this seed region along the visual-object axis during Post-learning; nevertheless, this general region of cortex did tend to exhibit contraction during this Post-learning period (see negative changes in blue in Fig. 5E). Together, these results suggest that during SL, the visual network's expansion along the visual-objects axis is mainly driven by increased within-network connectivity; however, once SL is interrupted, these patterns of connectivity appear to reverse, with visual cortex exhibiting increased connectivity with brain networks linked to exogenous attention (SalVentAttn network).

For the right DAN (see Fig. 6B), which in contrast to the visual seed region, exhibited significant *contraction* along the visual-object axis from Early- to Late-learning (see Fig. 5D), we found that its connectivity during this learning period was most prominently characterized by a decrease in connectivity with other DAN areas and the visual cortex, and an increase in connectivity with the transmodal regions belonging to other brain networks, such as the DMN B, Control B and SalVentAttn B subnetworks. By contrast, for Late- to Post-learning, where this same region now exhibited significant *expansion* along the visual-object axis, we observed a stark reversal in this overall pattern of connectivity. Specifically, during this Post-learning period we found that this DAN region now exhibited increased connectivity with areas belonging to the same network, as well as with regions of the visual network. Taken together, these results suggest that the *contraction* of the DAN during learning is driven by its decreased within-network connectivity and increased between-network connectivity with higher-order brain networks, whereas *expansions* during Post-learning is being driven by its increased within-network connectivity and reinstatement of connectivity with the visual cortex.

Finally, for the right PRC (see Fig. 6C), we found that this region exhibited significant contraction along the visual-object axis from Late- to Post-learning. Our seed connectivity maps revealed that, during Post-learning, PRC decreased its connectivity with visual cortex, while most prominently increasing connectivity with large swaths of parietal and prefrontal cortex (this effect did not appear to be unique to particular brain networks, see Fig. 6C). The fact that the left and right PRC only significantly changed their embeddings along the visual-object axis once SL was interrupted strongly suggests that the PRC is responsive to prediction errors, and communicates these errors with the parietal and prefrontal cortex.

Collectively, these seed-based analyses allow us to characterize, at the level of whole-brain functional connectivity, the selective changes in areas along the visual-objects axis both during SL and once the learned associations are violated (during Post-learning). In general, these findings lend support for the view that expansions along the visual-object axis can be interpreted as reflecting, at least in part, increases in *within-network* connectivity while contractions reflect increases in *between-network* connectivity.

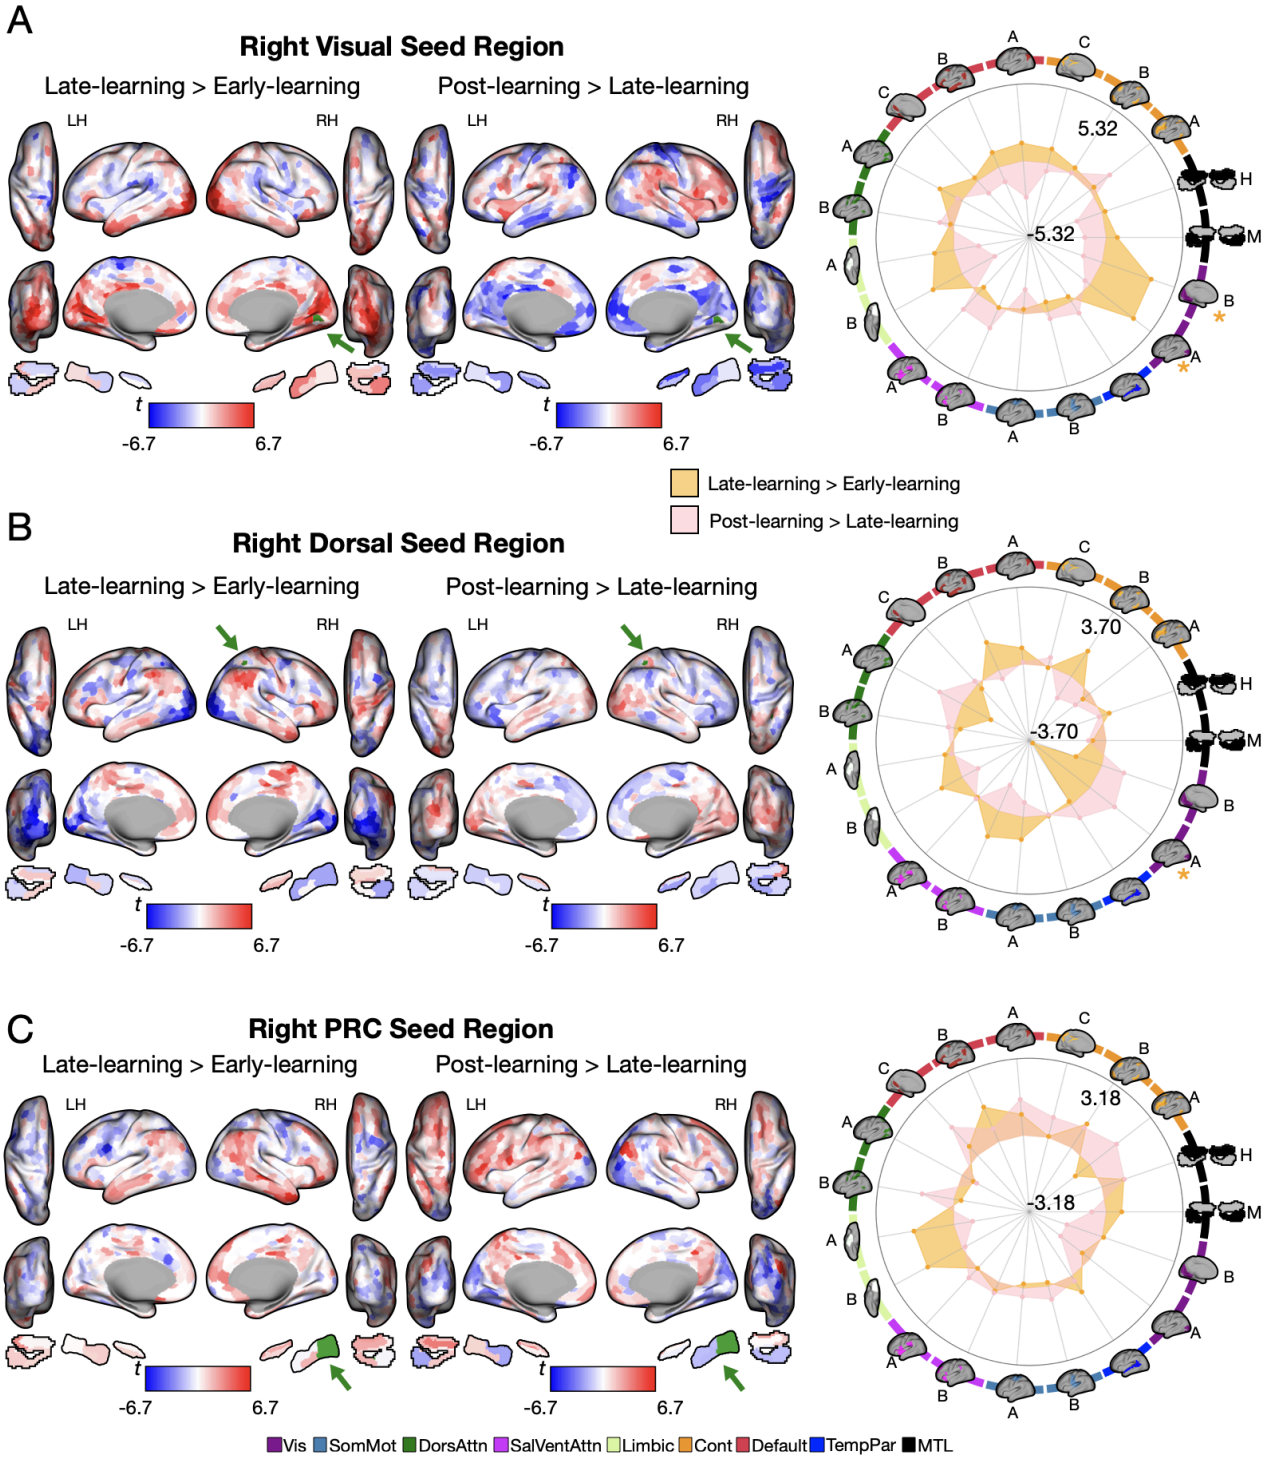

Figure 6. Patterns of connectivity underlying changes along the visual-objects axis throughout visual SL. (A) Connectivity changes of representative seed regions in right visual, DAN, and PRC areas (shown in green, as indicated by green arrow) from Late-learning > Early-learning and Post-learning > Late-learning. Positive (red) values show increases and negative (blue) values show decreases in loading across along the visual-objects axis. (B) Polar

plots show seed-based changes in connectivity between epochs at the network level (according to the Yeo 17-networks parcellation (76), as well as the additional MTL regions derived from the Automated Segmentation of Hippocampal Subfields (ASHS; 65) split into MTL (M) and hippocampus (H)). The color behind each brain indicates its functional network assignment (legend in B), with letters depicting its constituent subnetwork assignment (76). Asterisks indicate FDR-corrected significance in network changes between epochs.

The role of the hippocampus in statistical learning

In the previous sections, we characterized the regional changes along the visual-objects axis across the visual, DAN, and PRC during visual SL, as well as the underlying shifts in functional connectivity that may be driving these changes. While our findings in PRC suggest a role for this structure in processing the violation of predictions, there were no statistically significant changes in the embedding of hippocampal regions across our task; a conspicuous absence given prior literature implicating hippocampus in SL (24; 78), and that there was evidence for local representational change in our earlier analyses (Fig. 1C; 22). One possible explanation for this discrepancy may stem from our functional connectivity analysis including 998 cortical but only eight hippocampal regions. Since dimension reduction approaches like PCA are designed to explain the largest amount of variance in a relatively small number of orthogonal components, it is not surprising that our strongest loadings (PCs 1-3) were dominated by cortex, and not hippocampus (see Fig. 3A). It may also be the case that the changes that occur in the hippocampus are largely internal, and not observable using this high-level analysis that focuses on macro-scale changes in connectivity across whole-brain networks. Nevertheless, as prior work has identified a clear role for the hippocampus in visual SL (20-22), we next sought to more directly examine changes in the connectivity of hippocampal regions with cortex, as well as with each other. To this end, we derived aggregated seed-based hippocampal connectivity maps for subfields in left and right CA1, CA2/3, DG and subiculum for each task epoch, and then created contrast maps to compare Early- to Late-learning, and Late- to Post-learning, as we did with other seed regions. For an overview of how our hippocampal-seed analyses were performed, see the Supplementary Extended Materials and Methods.

Our hippocampal-based seed connectivity analysis (see Fig. 7A; for the left hippocampus, see Fig S2) revealed that from Early- to Late-learning, the right hippocampus showed modest qualitative increases in connectivity with all other networks, but specifically with other areas of the MTL, and limbic-A subnetwork (see Figure. 7B). However, once SL was interrupted during

Post-learning, these patterns tended to reverse, with the hippocampus most notably decreasing its connectivity with the visual cortex. In addition, the hippocampus also increased its connectivity with the contralateral hemisphere and with the retrosplenial subdivision of the DMN (i.e., DMN-C subnetwork). While this reversal in connectivity aligns with changes within the PRC and the overall decrease along the visual-objects axis (Fig. 5E), it's notable that no significant changes were observed in the hippocampus. This, in conjunction with the striking patterns seen in other medial temporal lobe (MTL) regions, suggests that the PRC and EC may play a more crucial role than previously thought in mediating information exchange between the cortex and hippocampus during learning and when the learned associations are violated.

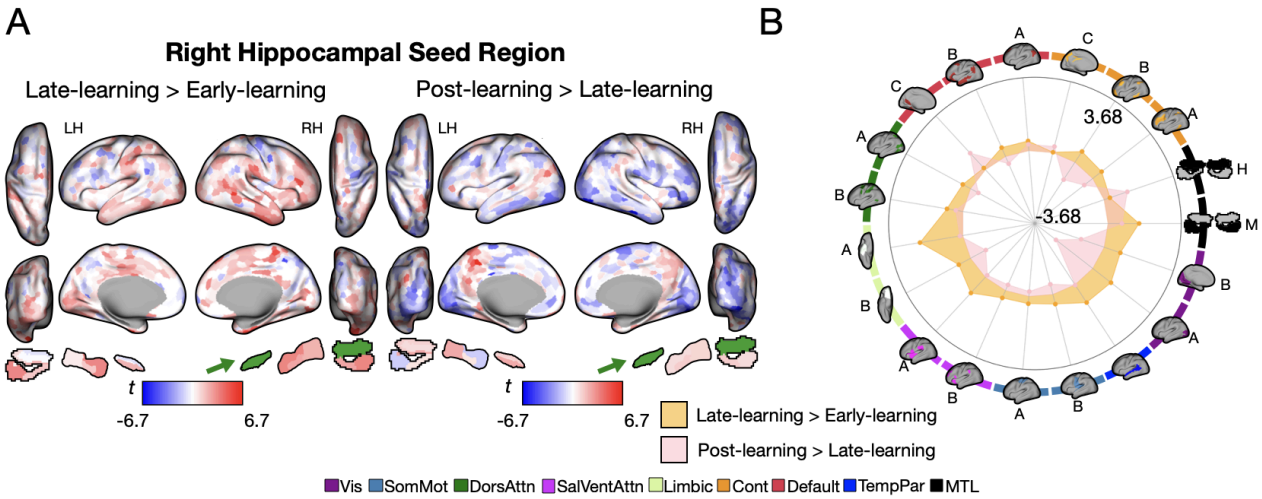

Figure 7. Hippocampal-Based Connectivity Analysis. (A) Connectivity changes for the right hippocampus (in green) from Late-learning > Early-learning and Post-learning > Late-learning. Positive (red) and negative (blue) values show increase and decreases in connectivity, respectively. (B) Polar plots show seed-based changes in connectivity between epochs at the network level (according to the Yeo 17-networks parcellation (76), as well as the additional MTL regions derived from the Automated Segmentation of Hippocampal Subfields (ASHS; 65) split into MTL (M) and hippocampus (H)). The color behind each brain indicates its functional network assignment, with letters depicting its constituent subnetwork assignment (76).

DISCUSSION

The current study aimed to characterize changes in the landscape of whole-brain functional activity that underlie visual SL. We used recent advances in dimension reduction techniques to build a low-dimensional manifold which characterized the main axes of whole-brain connectivity. Among these main axes, we found that only a single connectivity dimension — the visual-object axis — demonstrated a significant change in its expression across different stages

of the task, and thus we focused our analyses along this particular axis when exploring regional changes in functional organization. We found that during SL, regions within the higher-order visual cortex expanded along this axis, reflecting both their increased functional segregation from other networks, and increased within-network integration. Meanwhile, areas of the DAN contracted along this axis, reflecting their increased functional integration with other networks and decreased covariance with the visual cortex. Consistent with these effects being learning-related, we found that regions within the MTL, which did not show strong learning effects, now contracted along the axis when the learned associations were violated (in Post-learning). Finally, we found that while the hippocampus, consistent with prior work (20,22) showed significant local changes in representational organization following SL, there were no significant changes in the embedding of this region (or any of its subfields) along the visual-objects axis. Instead, we found only subtle, non-significant changes in connectivity with cortical networks during learning, which reversed once SL was interrupted. Together, our findings provide a unique window into the whole-brain changes that underlie visual SL, as well as the violation of learned associations. Our results suggest that, at the macro-scale level, SL is mainly associated with changes in the connectivity of cortical regions, whereas the MTL is associated with responding to violations in the learned associations, highlighting a more prominent role for extra-hippocampal MTL regions like PRC and ERC in responding to prediction errors following learning.

From Early- to Late-learning, we found that regions within the visual network significantly expanded along the visual-objects axis. This expansion was primarily characterized by increased within-network connectivity (i.e., modularity), likely reflecting a mechanism whereby the modularity of task-specific networks is increased to insulate encoded associations from additional change (79-82) and thus enhance the robustness of learning. In contrast, the DAN's contraction along the visual-objects axis during SL was driven by decreased connectivity with other DAN regions and the visual cortex, and increased connectivity with areas of transmodal cortex (Control and SalVentAttn networks). These findings are consistent with prior work showing reduced connectivity between visual and higher-order association regions as SL takes place (9), suggesting that as regularities are encoded, cross-talk between sensory areas involved in task demands, and areas involved in attentional control become less important for learning.

A novel aspect of our task was the inclusion of SL blocks book-ended by unstructured runs, allowing us to uniquely investigate changes in whole-brain activity when learned visual associations were disrupted (from Late- to Post-learning). Consistent with the idea that the network segregation effects were related to visual SL rather than merely the passage of time, we observed large reversals in several of the whole-brain seed connectivity patterns. Regions that had increased or decreased in covariance with the visual network and DAN during learning now showed opposite patterns after SL was interrupted. Specifically, the visual cortex, which had become more modular during learning, exhibited reduced within-network connectivity and increased integration with higher-order attention networks. Similarly, while the DAN became less modular during learning, it now showed increased within-network connectivity and greater connectivity with the visual network. These marked reversals in connectivity suggest that the brain's reorganization during SL is partially undone when learned patterns are disrupted, possibly reflecting a response to a prediction error signal (see below).

One particularly noteworthy finding was that the hippocampus, despite showing theory-consistent representational change as a result of SL, was not identified as exhibiting significant changes in its connectivity during or after learning. However, extra-hippocampal MTL regions that did not initially show learning effects — the PRC and ERC — exhibited significant contractions along the visual-objects axis during Post-learning. These contractions reflected an increased connectivity with areas of the DAN, DMN, and higher-order control network, as well as decreased covariance with the visual cortex. Current frameworks highlight the PRC as a specialized region positioned at the apex of the ventral visual stream, supporting object recognition (20; 83-87), and the ERC as a bridge between the ventral visual pathway and the hippocampus, via (in part) PRC (88-90). The PRC and ERC's contraction likely reflects a response to prediction error, whereby these regions recouple with higher-order cortical networks to reorient attention to incoming sensory information. Prediction error signals evoked in the hippocampus (91-93), propagate out to cortical regions implicated in perception and attention (94-95) to update existing schemas (96). When predictions are violated the outcome is often renewed communication between sensory and association areas (97), including between the hippocampus and its sensory inputs (98-99). Our findings not only extend this theoretical view of SL, but also provide a novel whole-brain perspective of how such prediction errors may be distributed across the cortex. While prior work has focused on the

role of the PRC and ERC in more explicit forms of learning (100-102), our results suggest a crucial role for these regions in detecting violations in implicit learning, perhaps precipitating the reversals observed in our visual and DAN seed regions. Similarly, while prior work has largely implicated the hippocampus in responding to prediction error, our findings indicate that other MTL regions likely play a larger role in controlling the flow of information than previously assumed.

The differential role of cortical and MTL networks in our main analysis (see Fig. 5) provides a more nuanced understanding of whether SL operates primarily through cortical or MTL mechanisms. Specifically, our results highlight that interactions within and between these regions dynamically update during SL, leading to different network configurations depending on the phase of the task. Both cortical and MTL circuits appear to learn in parallel: the hippocampus undergoes local representational change to encode paired stimuli (20-22; 25), while sensory cortical circuits become more modular, disconnecting from higher-order cognitive networks. However, when learned regularities are disrupted, the MTL detects this mismatch and reintegrates with these cognitive networks, while decreasing its connectivity with the primary visual cortex. Importantly, our current findings indicate that these above changes in the landscape of whole-brain activity occur largely independent of the hippocampus, suggesting a more prominent role for other MTL regions, like PRC and EC. This highlights the complexity of SL mechanisms, where cortical and MTL regions may contribute in distinct yet complementary ways. It is also worth considering that the level at which the hippocampus encodes learned associations may not be captured by our current methods, which focus on larger-scale neural network changes.

While the current study offers several unique insights into the whole-brain changes in functional connectivity associated with visual SL, there remain several important avenues for future research. First, while our findings point to distinct patterns of connectivity change associated with learning and the later violation of that learning, future work will be needed to directly track the behavioral correlates of these neural changes on a more moment-by-moment basis. Second, while the form of learning covered in the current study is implicit in nature, it remains underexplored how the neural changes observed here may relate to other forms of learning (e.g., auditory, spatial SL, error-based learning), as well as more general distinctions

made between implicit versus explicit forms of learning in the literature (103). For example, in the sensorimotor control literature, implicit learning is thought to be primarily cerebellar-dependent, and yet it is unknown whether statistical learning also relies on cerebellum, which contains similar functional networks to cortex (104). Future work will be needed to test these ideas. Third, given its established role in SL and prediction (105-106), it was surprising to us that the hippocampus did not significantly change its embedding along the visual-objects axis during any phases of the task. Nevertheless, when we conducted a seed connectivity analysis focused on the hippocampus directly, we observed a muted pattern of effects, comparable to what was observed in the PRC. In future work, it may be beneficial to construct a hippocampus-centric manifold to increase the sensitivity in detecting changes in its patterns of covariance; alternatively, our findings and whole-brain approach may highlight the relatively greater importance of other regions (outside the hippocampus) in guiding SL and the encoding of violations of learned associations. In summary, our results provide a whole-brain perspective for understanding connectivity changes during SL, and suggest that while cortical and MTL regions are involved in SL, this may occur via different mechanisms and during different time periods.

MATERIALS AND METHODS

We re-analyzed data (22) of 33 healthy individuals (21 female) between the ages of 18-35 years old, completing a visual SL task. Eight of the original dataset's 41 participants were excluded either because they completed fewer task runs than the majority of the sample (n = 7), or because there were issues with the quality of hippocampal or cortical segmentations (n = 1).

Participants completed eight experimental runs, where in each, 16 abstract images were presented one at a time on screen, each repeated five times, for a total of 80 image presentations. In all runs, the participants' unrelated cover task was to identify rare (10% of trials) instances of a small, faded grey square being overlaid on the images, by pressing a button. In the first experimental run (Pre-learning), the images were presented in a completely random, unstructured order. Following this run, participants completed six consecutive experimental runs (80 image presentations each, 480 total), in which unbeknownst to them, the 16 images were covertly combined into 8 pairs. This pair structure was embedded into the image sequence, such that a given image (A) was always followed by its pair-mate (B). The first

and last of these structured runs were used as Early- and Late-learning epochs respectively, for the present analysis. Finally, after SL, participants performed an additional Post-learning run, where they again viewed the 80 images in a completely random order (similar to Pre-learning).

For each participant and epoch, we obtained region-wise timeseries data by taking the mean BOLD activity for each of the 998 cortical regions pre-defined by the Schaefer-1000 parcellation (two regions were removed due to small parcel size; 64) and the 14 MTL regions, derived in an individual-specific way using ASHS software (65). Across the superset of these 1012 regions, we created centered functional connectivity matrices for each epoch, transforming them into cosine similarity affinity matrices after applying row-wise proportional thresholding. Next, we used PCA on affinity matrices to construct connectivity manifolds (46-48) and aligned all PCs to a group-average template Pre-learning manifold using Procrustes transformation.

Guided by prior work (53) we uncovered task-related variations within our top three PCs, by quantifying each subject's variability in PC loadings (min-max spread) across our task epochs. For each PC, we performed a rmANOVA on subjects' difference scores, and corrected for multiple comparisons using FDR correction ($q < 0.05$; 77). This analysis revealed a significant effect of epoch along a single neural dimension, the visual-objects axis. To evaluate learning-related changes in whole-brain functional connectivity over time, we computed each brain region's loadings along the visual-objects axis and assessed region-wise differences in loadings across epochs using rmANOVAs and follow-up paired t-tests. To explore alterations in whole-brain connectivity underlying changes along this axis, we performed seed connectivity contrasts on representative seed regions (belonging to the visual network, DAN, and PRC) that demonstrated representative effects both during and after SL. To examine learning-related changes relative to the hippocampus, we also performed similar seed connectivity contrasts for hippocampal regions to examine changes both during and after SL. For a comprehensive overview of the methods, please refer to the Supplementary Extended Materials and Methods.

Code availability. Imaging data were preprocessed using fmriPrep, which is open source and freely available. Our analysis is available at https://github.com/thelamplab/visualSLmanifolds.

REFERENCES

1. N. B. Turk-Browne, "Statistical Learning and Its Consequences" in The Influence of Attention, Learning, and Motivation on Visual Search, Nebraska Symposium on Motivation., M. D. Dodd, J. H. Flowers, Eds. (Springer New York, 2012), , pp. 117–146.
2. J. R. Saffran, "What Is Statistical Learning, and What Statistical Learning Is Not" in Neoconstructivism, 1st Ed., (Oxford University PressNew York, 2009), , pp. 180–194.
3. R. N. Aslin, Statistical learning: a powerful mechanism that operates by mere exposure. WIRES Cognitive Science 8, e1373, (2017).
4. N. B. Turk-Browne, J. A. Jungé, B. J. Scholl, The Automaticity of Visual Statistical Learning. Journal of Experimental Psychology: General 134, 552–564, (2005).
5. N. B. Turk-Browne, P. J. Isola, B. J. Scholl, T. A. Treat, Multidimensional visual statistical learning. Journal of Experimental Psychology: Learning, Memory, and Cognition 34, 399–407, (2008).
6. J. L. McClelland, B. L. McNaughton, R. C. O'Reilly, Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. Psychological Review 102, 419–457 (1995).
7. R. C. O'Reilly, R. Bhattacharyya, M. D. Howard, N. Ketz, Complementary Learning Systems. Cognitive Science 38, 1229–1248 (2014).
8. S. A. Huettel, P. B. Mack, G. McCarthy, Perceiving patterns in random series: dynamic processing of sequence in prefrontal cortex. Nat Neurosci 5, 485–490 (2002).
9. E. A. Karuza, et al., Neural Signatures of Spatial Statistical Learning: Characterizing the Extraction of Structure from Complex Visual Scenes. Journal of Cognitive Neuroscience 29, 1963–1976 (2017).
10. V. M. Karlaftis, et al., Multimodal imaging of brain connectivity reveals predictors of individual decision strategy in statistical learning. Nat Hum Behav 3, 297–307 (2019).
11. N. B. Turk-Browne, B. J. Scholl, M. M. Chun, M. K. Johnson, Neural Evidence of Statistical Learning: Efficient Detection of Visual Regularities Without Awareness. Journal of Cognitive Neuroscience 21, 1934–1945 (2009).
12. N. B. Turk-Browne, B. J. Scholl, M. K. Johnson, M. M. Chun, Implicit Perceptual Anticipation Triggered by Statistical Learning. J. Neurosci. 30, 11177–11187 (2010).
13. E. A. Karuza, et al., The neural correlates of statistical learning in a word segmentation task: An fMRI study. Brain and Language 127, 46–54 (2013).
14. K. McNealy, J. C. Mazziotta, M. Dapretto, Cracking the Language Code: Neural Mechanisms Underlying Speech Parsing. J. Neurosci. 26, 7629–7639 (2006).
15. T. Cunillera, et al., Time course and functional neuroanatomy of speech segmentation in adults. NeuroImage 48, 541–553 (2009).
16. T. Overath, et al., An Information Theoretic Characterisation of Auditory Encoding. PLoS Biol 5, e288 (2007).
17. J. R. Saffran, E. D. Thiessen, "Domain-General Learning Capacities" in Blackwell Handbook of Language Development, E. Hoff, M. Shatz, Eds. (Blackwell Publishing Ltd, 2007), pp. 68–86.
18. R. Frost, B. C. Armstrong, N. Siegelman, M. H. Christiansen, Domain generality versus modality specificity: the paradox of statistical learning. Trends in Cognitive Sciences 19, 117–125 (2015).

19. R. Frost, B. C. Armstrong, M. H. Christiansen, Statistical learning research: A critical review and possible new directions. Psychological Bulletin 145, 1128–1153 (2019).
20. A. C. Schapiro, L. V. Kustner, N. B. Turk-Browne, Shaping of Object Representations in the Human Medial Temporal Lobe Based on Temporal Regularities. Current Biology 22, 1622–1627 (2012).
21. A. C. Schapiro, N. B. Turk-Browne, K. A. Norman, M. M. Botvinick, Statistical learning of temporal community structure in the hippocampus: STATISTICAL LEARNING OF TEMPORAL COMMUNITY STRUCTURE. Hippocampus 26, 3–8 (2016).
22. J. Wammes, K. A. Norman, N. Turk-Browne, Increasing stimulus similarity drives nonmonotonic representational change in hippocampus. eLife 11, e68344 (2022).
23. M. L. Schlichting, K. F. Guarino, A. C. Schapiro, N. B. Turk-Browne, A. R. Preston, Hippocampal Structure Predicts Statistical Learning and Associative Inference Abilities during Development. Journal of Cognitive Neuroscience 29, 37–51 (2017).
24. A. S. Finn, M. Kharitonova, N. Holtby, M. A. Sheridan, Prefrontal and Hippocampal Structure Predict Statistical Learning Ability in Early Childhood. Journal of Cognitive Neuroscience 31, 126–137 (2019).
25. A. C. Schapiro, E. Gregory, B. Landau, M. McCloskey, N. B. Turk-Browne, The Necessity of the Medial Temporal Lobe for Statistical Learning. Journal of Cognitive Neuroscience 26, 1736–1747 (2014).
26. M. E. Mundy, P. E. Downing, D. M. Dwyer, R. C. Honey, K. S. Graham, A Critical Role for the Hippocampus and Perirhinal Cortex in Perceptual Learning of Scenes and Faces: Complementary Findings from Amnesia and fMRI. Journal of Neuroscience 33, 10490–10502 (2013).
27. L. Davachi, A. D. Wagner, Hippocampal Contributions to Episodic Encoding: Insights From Relational and Item-Based Learning. Journal of Neurophysiology 88, 982–990 (2002).
28. A. R. Preston, J. D. E. Gabrieli, Dissociation between Explicit Memory and Configural Memory in the Human Medial Temporal Lobe. Cerebral Cortex 18, 2192–2207 (2008).
29. H. Eichenbaum, Time (and space) in the hippocampus. Current Opinion in Behavioral Sciences 17, 65–70 (2017).
30. T. Geyer, F. Baumgartner, H. J. Müller, S. Pollmann, Medial temporal lobe-dependent repetition suppression and enhancement due to implicit vs. explicit processing of individual repeated search displays. Front. Hum. Neurosci. 6 (2012).
31. M. W. Howard, H. Eichenbaum, Time and space in the hippocampus. Brain Research 1621, 345–354 (2015).
32. N. B. Turk-Browne, The hippocampus as a visual area organized by space and time: A spatiotemporal similarity hypothesis. Vision Research 165, 123–130 (2019).
33. C. E. Westerberg, B. B. Miller, P. J. Reber, N. J. Cohen, K. A. Paller, Neural correlates of contextual cueing are modulated by explicit learning. Neuropsychologia 49, 3439–3447 (2011).
34. A. C. Schapiro, N. B. Turk-Browne, M. M. Botvinick, K. A. Norman, Complementary learning systems within the hippocampus: a neural network modelling approach to reconciling episodic memory with statistical learning. Phil. Trans. R. Soc. B 372, 20160049 (2017).
35. A. Aljishi, et al., Statistical learning in epilepsy: Behavioral, anatomical, and causal mechanisms in the human brain. [Preprint] (2023). Available at: http://biorxiv.org/lookup/doi/10.1101/2023.04.25.538321

36. H. S. Wang, et al., Dentate Gyrus Integrity Is Necessary for Behavioral Pattern Separation But Not Statistical Learning. Journal of Cognitive Neuroscience 35, 900–917 (2023).
37. L. J. Batterink, K. A. Paller, P. J. Reber, Understanding the Neural Bases of Implicit and Statistical Learning. Topics in Cognitive Science 11, 482–503 (2019).
38. H. Marlatte, Z. Belchev, M. Fraser, A. Gilboa, The effect of hippocampal subfield damage on rapid temporal integration through statistical learning and associative inference. Neuropsychologia 193, 108755 (2024).
39. D. E. Warren, T. C. Roembke, N. V. Covington, B. McMurray, M. C. Duff, Cross-Situational Statistical Learning of New Words Despite Bilateral Hippocampal Damage and Severe Amnesia. Front. Hum. Neurosci. 13, 448 (2020).
40. K. Schevenels, et al., The role of the hippocampus in statistical learning and language recovery in persons with post stroke aphasia. NeuroImage: Clinical 36, 103243 (2022).
41. J.(Park, K. Janacsek, D. Nemeth, H.-A. Jeon, Reduced functional connectivity supports statistical learning of temporally distributed regularities. NeuroImage 260, 119459 (2022).
42. E. Bullmore, O. Sporns, Complex brain networks: graph theoretical analysis of structural and functional systems. Nat Rev Neurosci 10, 186–198 (2009).
43. M. P. Van Den Heuvel, R. C. W. Mandl, R. S. Kahn, H. E. Hulshoff Pol, Functionally linked resting-state networks reflect the underlying structural connectivity architecture of the human brain. Human Brain Mapping 30, 3127–3141 (2009).
44. A. J. Barnett, et al., "Organization of cortico-hippocampal networks in the human brain" (Neuroscience, 2020).
45. E. Froudarakis, et al., The Visual Cortex in Context. Annu. Rev. Vis. Sci. 5, 317–339 (2019).
46. D. J. Gale, et al., Distinct patterns of cortical manifold expansion and contraction underlie human sensorimotor adaptation. Proc. Natl. Acad. Sci. U.S.A. 119, e2209960119 (2022).
47. C. Paquola, et al., Microstructural and functional gradients are increasingly dissociated in transmodal cortices. PLoS Biol 17, e3000284 (2019).
48. D. S. Margulies, et al., Situating the default-mode network along a principal gradient of macroscale cortical organization. Proc. Natl. Acad. Sci. U.S.A. 113, 12574–12579 (2016).
49. J. M. Huntenburg, P.-L. Bazin, D. S. Margulies, Large-Scale Gradients in Human Cortical Organization. Trends in Cognitive Sciences 22, 21–31 (2018).
50. S. Vyas, M. D. Golub, D. Sussillo, K. V. Shenoy, Computation Through Neural Population Dynamics. Annu. Rev. Neurosci. 43, 249–275 (2020).
51. K. V. Shenoy, M. Sahani, M. M. Churchland, Cortical Control of Arm Movements: A Dynamical Systems Perspective. Annu. Rev. Neurosci. 36, 337–359 (2013).
52. R. A. I. Bethlehem, et al., Dispersion of functional gradients across the adult lifespan. NeuroImage 222, 117299 (2020).
53. S.-J. Hong, et al., Atypical functional connectome hierarchy in autism. Nat Commun 10, 1022 (2019).
54. M. Xia, et al., Connectome gradient dysfunction in major depression and its association with gene expression profiles and treatment outcomes. Mol Psychiatry 27, 1384–1393 (2022).
55. A. Samara, J. Eilbott, D. S. Margulies, T. Xu, T. Vanderwal, Cortical gradients during naturalistic processing are hierarchical and modality-specific. NeuroImage 271, 120023 (2023).
56. H. Song, W. M. Shim, M. D. Rosenberg, Large-scale neural dynamics in a shared low-dimensional state space reflect cognitive and attentional dynamics. eLife 12, e85487 (2023).

57. A. C. Schapiro, T. T. Rogers, N. I. Cordova, N. B. Turk-Browne, M. M. Botvinick, Neural representations of events arise from temporal community structure. Nat Neurosci 16, 486–492 (2013).
58. D. Frank, M. A. Montemurro, D. Montaldi, Pattern Separation Underpins Expectation-Modulated Memory. J. Neurosci. 40, 3455–3464 (2020).
59. V. J. H. Ritvo, N. B. Turk-Browne, K. A. Norman, Nonmonotonic Plasticity: How Memory Retrieval Drives Learning. Trends in Cognitive Sciences 23, 726–742 (2019).
60. V. J. H. Ritvo, A. Nguyen, N. B. Turk-Browne, K. A. Norman, Differentiation and Integration of Competing Memories: A Neural Network Model. [Preprint] (2023). Available at: http://biorxiv.org/lookup/doi/10.1101/2023.04.02.535239
61. B. Midler, J. McClelland, Retrieval induced forgetting in a non-monotonic hippocampal model. [Preprint] (2023). Available at: http://biorxiv.org/lookup/doi/10.1101/2023.08.25.544249
62. J. A. Lewis-Peacock, K. A. Norman, Competition between items in working memory leads to forgetting. *Nat Commun* **5**, 5768 (2014).
63. G. J. Detre, A. Natarajan, S. J. Gershman, K. A. Norman, Moderate levels of activation lead to forgetting in the think/no-think paradigm. Neuropsychologia 51, 2371–2388 (2013).
64. A. Schaefer, et al., Local-Global Parcellation of the Human Cerebral Cortex from Intrinsic Functional Connectivity MRI. Cerebral Cortex 28, 3095–3114 (2018).
65. L. E. M. Wisse, et al., Automated Hippocampal Subfield Segmentation at 7T MRI. AJNR Am J Neuroradiol 37, 1050–1057 (2016).
66. C. Areshenkoff, et al., Neural excursions from manifold structure explain patterns of learning during human sensorimotor adaptation. eLife 11, e74591 (2022).
67. C. N. Areshenkoff, et al., Muting, not fragmentation, of functional brain networks under general anesthesia. NeuroImage 231, 117830 (2021).
68. Q. Zhao, D. Kwon, K. M. Pohl, "A Riemannian Framework for Longitudinal Analysis of Resting-State Functional Connectivity" in Medical Image Computing and Computer Assisted Intervention – MICCAI 2018, Lecture Notes in Computer Science., A. F. Frangi, J. A. Schnabel, C. Davatzikos, C. Alberola-López, G. Fichtinger, Eds. (Springer International Publishing, 2018), pp. 145–153.
69. L. McInnes, J. Healy, J. Melville, UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. (2018). https://doi.org/10.48550/ARXIV.1802.03426
70. Q. Nick, et al., Reconfigurations of cortical manifold structure during reward-based motor learning. [Preprint] (2023). Available at: http://biorxiv.org/lookup/doi/10.1101/2023.07.05.547880
71. R. Vos De Wael, et al., BrainSpace: a toolbox for the analysis of macroscale gradients in neuroimaging and connectomics datasets. Commun Biol 3, 103 (2020).
72. S.-J. Hong, et al., Toward a connectivity gradient-based framework for reproducible biomarker discovery. NeuroImage 223, 117322 (2020).
73. J. Smallwood, *et al.*, The default mode network in cognition: a topographical perspective. *Nat Rev Neurosci* **22**, 503–513 (2021).
74. B. Sherman, N. Turk-Browne, Attention and Memory. [Preprint] (2022). Available at: https://osf.io/xs6db [Accessed 5 May 2024].
75. T. Yarkoni, R. A. Poldrack, T. E. Nichols, D. C. Van Essen, T. D. Wager, Large-scale automated synthesis of human functional neuroimaging data. Nat Methods 8, 665–670 (2011).

76. B. T. Thomas Yeo, et al., The organization of the human cerebral cortex estimated by intrinsic functional connectivity. Journal of Neurophysiology 106, 1125–1165 (2011).
77. Y. Benjamini, D. Yekutieli, False Discovery Rate–Adjusted Multiple Confidence Intervals for Selected Parameters. Journal of the American Statistical Association 100, 71–81 (2005).
78. W. Tang, M. H. Christiansen, Z. Qi, Learning in chunks: A model of hippocampal representations for processing temporal regularities in statistical learning. [Preprint] (2022). Available at: http://biorxiv.org/lookup/doi/10.1101/2022.04.04.487009
79. M. Kirschner, J. Gerhart, Evolvability. Proc. Natl. Acad. Sci. U.S.A. 95, 8420–8427 (1998).
80. M.-A. Félix, A. Wagner, Robustness and evolution: concepts, insights and challenges from a developmental model system. Heredity 100, 132–140 (2008).
81. N. Kashtan, U. Alon, Spontaneous evolution of modularity and network motifs. Proc. Natl. Acad. Sci. U.S.A. 102, 13773–13778 (2005).
82. G. P. Wagner, L. Altenberg, PERSPECTIVE: COMPLEX ADAPTATIONS AND THE EVOLUTION OF EVOLVABILITY. Evolution 50, 967–976 (1996).
83. S. L. Kivisaari, L. K. Tyler, A. U. Monsch, K. I. Taylor, Medial perirhinal cortex disambiguates confusable objects. Brain 135, 3757–3769 (2012).
84. E. A. Murray, B. J. Richmond, Role of perirhinal cortex in object perception, memory, and associations. Current Opinion in Neurobiology 11, 188–193 (2001).
85. H. Chen, W. Zhou, J. Yang, Dissociation of the Perirhinal Cortex and Hippocampus During Discriminative Learning of Similar Objects. J. Neurosci. 39, 6190–6201 (2019).
86. E. B. O'Neil, A. D. Cate, S. Kohler, Perirhinal Cortex Contributes to Accuracy in Recognition Memory and Perceptual Discriminations. Journal of Neuroscience 29, 8329–8334 (2009).
87. M. J. Buckley, The Role of the Perirhinal Cortex and Hippocampus in Learning, Memory, and Perception. The Quarterly Journal of Experimental Psychology Section B 58, 246–268 (2005).
88. W. A. Suzuki, Y. Naya, The Perirhinal Cortex. Annu. Rev. Neurosci. 37, 39–53 (2014).
89. H. Eichenbaum, A. P. Yonelinas, C. Ranganath, The Medial Temporal Lobe and Recognition Memory. Annu. Rev. Neurosci. 30, 123–152 (2007).
90. M. J. Buckley, The Role of the Perirhinal Cortex and Hippocampus in Learning, Memory, and Perception. The Quarterly Journal of Experimental Psychology Section B 58, 246–268 (2005).
91. N. M. Long, H. Lee, B. A. Kuhl, Hippocampal Mismatch Signals Are Modulated by the Strength of Neural Predictions and Their Similarity to Outcomes. J. Neurosci. 36, 12677–12687 (2016).
92. L. Davachi, S. DuBrow, How the hippocampus preserves order: the role of prediction and context. Trends in Cognitive Sciences 19, 92–99 (2015).
93. J. Chen, P. A. Cook, A. D. Wagner, Prediction strength modulates responses in human area CA1 to sequence violations. Journal of Neurophysiology 114, 1227–1238 (2015).
94. H. C. Barron, R. Auksztulewicz, K. Friston, Prediction and memory: A predictive coding account. Progress in Neurobiology 192, 101821 (2020).
95. N. C. Hindy, F. Y. Ng, N. B. Turk-Browne, Linking pattern completion in the hippocampus to predictive coding in visual cortex. Nat Neurosci 19, 665–667 (2016).
96. S.-H. Wang, R. G. M. Morris, Hippocampal-Neocortical Interactions in Memory Formation, Consolidation, and Reconsolidation. Annu. Rev. Psychol. 61, 49–79 (2010).
97. L. H. Arnal, B. Morillon, C. A. Kell, A.-L. Giraud, Dual Neural Routing of Visual Facilitation in Speech Processing. J. Neurosci. 29, 13445–13453 (2009).

98. J. E. Lisman, A. A. Grace, The Hippocampal-VTA Loop: Controlling the Entry of Information into Long-Term Memory. Neuron 46, 703–713 (2005).
99. O. Bein, N. A. Plotkin, L. Davachi, Mnemonic prediction errors promote detailed memories. Learn. Mem. 28, 422–434 (2021).
100. W. Jeong, C. K. Chung, J. S. Kim, Episodic memory in aspects of large-scale brain networks. Front. Hum. Neurosci. 9 (2015).
101. B. C. Dickerson, H. Eichenbaum, The Episodic Memory System: Neurocircuitry and Disorders. Neuropsychopharmacol 35, 86–104 (2010).
102. K. Danieli, A. Guyon, I. Bethus, Episodic Memory formation: A review of complex Hippocampus input pathways. Progress in Neuro-Psychopharmacology and Biological Psychiatry 126, 110757 (2023).
103. D. Kahneman, Thinking, fast and slow (2011).
104. E. Y. Choi, B. T. T. Yeo, R. L. Buckner, The organization of the human striatum estimated by intrinsic functional connectivity. Journal of Neurophysiology 108, 2242–2263 (2012).
105. D. Shohamy, A. D. Wagner, Integrating Memories in the Human Brain: Hippocampal-Midbrain Encoding of Overlapping Events. Neuron 60, 378–389 (2008).
106. M. T. R. Van Kesteren, D. J. Ruiter, G. Fernández, R. N. Henson, How schema and novelty augment memory formation. Trends in Neurosciences 35, 211–219 (2012).