

Published in final edited form as:

Curr Biol. 2012 September 11; 22(17): 1622–1627. doi:10.1016/j.cub.2012.06.056.

Shaping of object representations in the human medial temporal lobe based on temporal regularities

Anna C. Schapiro¹, Lauren V. Kustner¹, and Nicholas B. Turk-Browne¹

¹Department of Psychology and Princeton Neuroscience Institute, Princeton University, Princeton, NJ, 08540, USA

Summary

Regularities are gradually represented in cortex after extensive experience [1], and yet they can influence behavior after minimal exposure [2,3]. What kind of representations support such rapid statistical learning? The medial temporal lobe (MTL) can represent information from even a single experience [4], making it a good candidate system for assisting in initial learning about regularities. We combined anatomical segmentation of the MTL, high-resolution functional magnetic resonance imaging (fMRI), and multivariate pattern analysis (MVPA) to identify representations of objects in cortical and hippocampal areas of human MTL, assessing how these representations were shaped by exposure to regularities. Subjects viewed a continuous visual stream containing hidden temporal relationships—pairs of objects that reliably appeared nearby in time. We compared the pattern of blood oxygen level-dependent (BOLD) activity evoked by each object before and after this exposure, and found that perirhinal cortex (PRC), parahippocampal cortex (PHC), subiculum, CA1, and CA2/CA3/dentate gyrus (CA2/3/DG) encoded regularities by increasing the representational similarity of their constituent objects. Most regions exhibited bidirectional associative shaping, while CA2/3/DG represented regularities in a forward-looking predictive manner. These findings suggest that object representations in MTL come to mirror the temporal structure of the environment, supporting rapid and incidental statistical learning.

Experimental Procedures

The experimental procedures are summarized briefly throughout the Results and in Figure 1, and are presented in complete detail in Supplemental Experimental Procedures.

Results

We often encounter new environments whose structure remains stable over subsequent experiences. For example, we pass the same landmarks while commuting to work, observe the same interactions in a social group, and practice the same plays when picking up a sport. Learning such regularities can help us effectively perceive and act in familiar environments.

Neurophysiological studies suggest that MTL cortex is involved in representing regularities. For example, after viewing a sequence of fractal patterns repeatedly, neurons in macaque

© 2012 Elsevier Inc. All rights reserved.

Please address correspondence to: Anna Schapiro, Department of Psychology, Princeton University, Green Hall, Princeton, NJ 08540
schapiro@princeton.edu.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

PRC respond more similarly to fractals nearby in the sequence [5]. Similarly, PRC neurons develop selectivity for paired stimuli in paired-associate tasks [6]. This pair coding occurs even when learning is unsupervised and pairs are task-irrelevant [7]. While sometimes attributed to area TE and inferior temporal cortex (IT), PRC shows stronger and earlier effects [8] and is necessary for learning to occur in IT [9,10].

How the learning of regularities alters stimulus representations in hippocampus proper has been largely unexplored (cf. [11–13]). Human neuroimaging and patient studies support the possibility that the hippocampus is involved in such learning [3,14,15]. For example, the hippocampus responds more strongly during incidental exposure to structured vs. random sequences [3]. However, prior studies of the intact human brain have relied on univariate activation and could thus not assess representational changes: Changes in activation do not necessarily reflect changes to the representations of specific associated objects, and might instead reflect a scalar signal that can apply generically to any object (e.g., mutual information [16], novelty [17]).

The present work was inspired by the following questions: What are the consequences of exposure to regularities for representations in the human MTL? What role might the hippocampus play in such rapid forms of statistical learning? How do representations change from before to after learning? To answer these questions, we used high-resolution fMRI to measure how multivariate representations of objects in the MTL are shaped by temporal regularities.

Behavior

Subjects viewed a 40-min stream of colorful fractals presented one-at-a-time. During this time, they performed an orthogonal cover task of detecting grayscale patches that appeared infrequently on the fractals. Detection performance was excellent (mean $A' = .893$, $SD = .038$; $t(16) = 42.9$, $p < .001$). Unbeknownst to subjects, the fractal stream was generated from temporal pairs (Fig. 1A). For half of these pairs, the two objects always appeared successively ('strong' pairs). For the other half, the objects appeared successively one-third of the time ('weak' pairs). These pairs were used to generate the stream, but the fractals appeared continuously with no cues to the boundaries between pairs other than the temporal statistics. Before and after this exposure, fractals were presented in a random order to assess representational change (see below). Because the pair structure was disrupted by the final random run, we conducted a separate behavioral experiment without this run. Subjects could discriminate strong vs. weak pairs ($t(11) = 2.40$, $p = .035$; Fig. S1), suggesting that learning took place.

ROI Analyses

We hypothesized that the representations of paired fractals would become more similar after exposure to temporal structure. This cannot be tested during pair exposure because two fractals appearing close in time will have similar representations simply due to the temporal autocorrelation of BOLD signal. To avoid this confound, we presented fractals in the same random order before and after pair exposure, and estimated the distributed neural representation of each fractal with a GLM (Fig. 1B).

We focused on the MTL, although other brain systems likely also participate in and benefit from learning about regularities [3]. MTL ROIs were segmented manually using anatomical landmarks and aligned into a common space with the functional data. Because ROIs were segmented at a higher anatomical resolution, the labeling of functional voxels is inherently inexact. Therefore, these ROIs provide an index of only the relative contribution of

underlying subregions. Patterns of parameter estimates for each fractal were extracted from every ROI and correlated before and after learning.

We first examined bilateral MTL cortex (Fig. 2), including PHC, PRC, and entorhinal cortex (ERC). The pattern correlation for fractals from strong pairs increased from before to after learning over the entire MTL cortex ($t[16]=3.01$, $p=.008$), relative to the change in correlation for recombinations of the same fractals ('shuffled' pairs). This increase was reliable in PHC ($t[16]=2.53$, $p=.022$) and PRC ($t[16]=2.45$, $p=.026$), but not in ERC ($t[16]=1.21$, $p=.245$). Strong pairs also exhibited a greater increase than weak pairs over the entire MTL cortex ($t[16]=4.08$, $p<.001$), and in PHC ($t[16]=3.04$, $p=.008$) and PRC ($t[16]=3.96$, $p=.001$), but only marginally in ERC ($t[16]=2.08$, $p=.054$). Weak pairs exhibited a decreased correlation relative to shuffled pairs over the entire MTL cortex ($t[16]=-2.40$, $p=.029$), and in PHC ($t[16]=-2.27$, $p=.037$) and PRC ($t[16]=-2.14$, $p=.048$), but not in ERC ($t[16]=-1.43$, $p=.173$).

We next examined bilateral hippocampus, including subiculum, CA1, and CA2/3/DG ROIs. The pattern correlation for fractals from strong pairs increased from before to after learning over the entire hippocampus ($t[16]=3.20$, $p=.006$), relative to shuffled pairs. This increase was reliable in subiculum ($t[16]=2.99$, $p=.009$), CA1 ($t[16]=3.61$, $p=.002$), and CA2/3/DG ($t[16]=3.48$, $p=.003$). Strong pairs also exhibited a greater increase than weak pairs over the entire hippocampus ($t[16]=3.82$, $p=.002$), and in subiculum ($t[16]=2.83$, $p=.012$), CA1 ($t[16]=4.38$, $p<.001$), and CA2/3/DG ($t[16]=4.64$, $p<.001$). Weak pairs exhibited a decrease in pattern correlation relative to shuffled pairs over the entire hippocampus ($t[16]=-2.15$, $p=.048$), and in CA2/3/DG ($t[16]=-2.21$, $p=.042$), but not in subiculum ($t[16]=-1.39$, $p=.183$) or CA1 ($t[16]=-1.71$, $p=.107$).

Searchlight Analyses

Above, we calculated correlations over all voxels in each ROI. However, representational changes may have also occurred in local subregions. To examine such local changes, we swept a cubic searchlight through all MTL ROIs. These ROIs were registered across subjects with non-linear deformation tools to allow for voxelwise statistical tests (permutation-test $p<.001$). We focused on the main comparison of strong vs. weak pairs, and found greater increases for strong pairs in right and left PHC, left PRC, left subiculum, right CA1, and right CA2/3/DG (Fig. 3). An exploratory searchlight analysis outside of the MTL revealed similar effects in four other regions (Table S1).

Nature of Representational Shaping

Increased pattern similarity after learning could reflect different types of changes in the underlying representations. Seeing one member of a pair may activate the other because their representations became unitized or connected. In this case, after learning pair AB, the representation of A should be similar to how B was initially represented, and the representation of B should be similar to how A was initially represented ('association' hypothesis). Alternatively, because pairs appeared in a fixed order, the first member may activate the second more than the second activates the first. In this case, the representation of A should be similar to how B was initially represented, but the representation of B may not change as much ('prediction' hypothesis).

The association and prediction hypotheses are indistinguishable after learning, since the AB pattern correlation should be higher in both cases. However, correlating the representation of one member after learning with the other before learning is helpful (Fig. 4): If the A representation moved towards the initial B representation, then the A pattern after learning (A_{post}) will resemble the B pattern before learning (B_{pre}). If this correlation is greater than

that for B after learning (B_{post}) with A before learning (A_{pre}), then the prediction hypothesis would be supported. We therefore defined an asymmetry index, which reflects prediction if positive: $\text{corr}(A_{\text{post}}, B_{\text{pre}}) - \text{corr}(B_{\text{post}}, A_{\text{pre}})$.

We limited this analysis to ROIs and conditions that showed overall increases, since it only makes sense to ask *how* representations changed if there was an overall change to begin with. Over the entire MTL cortex and hippocampus ROIs, only right hippocampus showed marginal evidence for asymmetric changes ($t[16]=2.06$, $p=.056$; right MTL cortex: $t[16]=1.33$, $p=.201$; left MTL cortex and hippocampus: $t < 1$). Among subregions, only right CA2/3/DG exhibited asymmetric changes ($t[16]=2.55$, $p=.022$), with a trend in right subiculum ($t[16]=1.82$, $p=.087$), consistent with the encoding of forward-looking predictions. The index was non-significant in all other ROIs ($t < 1$), consistent with them encoding bidirectional associations.

Discussion

We found that object representations in the MTL rapidly changed based on incidental exposure to temporal regularities. That is, the multivoxel representations of strongly paired objects became more similar in PRC, PHC, subiculum, CA1, and CA2/3/DG, with this increase occurring asymmetrically in CA2/3/DG. The discovery that distributed patterns in the MTL are tuned by regularities sheds light on the role of the human MTL in statistical learning, beyond what can be inferred from univariate methods [3,16,17]. This ability to track multivoxel representations of *specific* items [18], in particular, holds great promise for studying the formation and evolution of representations in memory [19,20].

Representational Changes

The plasticity observed in PRC/PHC may reflect similar changes in the selectivity of neurons to what has been reported in non-human primates [5]. The way that representations of paired items might change ranges on a continuum from completely indistinguishable or 'unitized' [21] to largely discrete but weakly associated. Our results are consistent with either possibility. Indeed, PRC exhibits both types of changes, with subregion A35 supporting unitized representations and A36 supporting associative representations [22]. Relatedly, PHC has been implicated in associative retrieval, contextual associations, relational memory, and event order memory [23–27], all processes that could be supported by the representational changes we observed.

The use of fMRI allowed us to assess representational changes in the MTL more broadly. We were especially interested in the hippocampus because of its ability to rapidly encode information [4], its circuitry specialized for sequential processing [28,29], and hints of its involvement in similar forms of learning [3,14,15]. Increases in representational similarity were widespread in the hippocampus, analogous to multivariate changes in MTL cortex. This discovery challenges the standard view that learning of regularities is the domain of cortex [4]. The generality of these changes suggests that multiple hippocampal subfields may be involved in representing regularities. However, an important limitation of the current study is that our fMRI resolution and preprocessing, combined with the point-spread function of the BOLD response, generate ambiguity in attributing voxel-level responses to particular subfields. Future studies with higher functional resolution and field strength are needed to address this issue, but as a first step we replicated our pattern of results very closely using only voxels for which we had the highest confidence in subfield membership (Figs. S3,S4).

The apparent generality may also result from the interactivity of subfields. For instance, subiculum and CA1 are thought to compare predictions generated by CA3 with current input

transmitted from ERC [30,31]. How patterns in CA3 and ERC relate to patterns in subiculum/CA1—and how these interactions generate match-mismatch signals—are interesting questions for future research. That said, not all subfields exhibited the same pattern of results. For example, only CA2/3/DG (and possibly subiculum) represented regularities predictively: The first member of strong pairs reinstated the second member more than the second reinstated the first. Within the CA2/3/DG ROI, CA2, CA3, and DG cannot be distinguished, and so these findings cannot be attributed to any one subfield. Nevertheless, they are consistent with the role of CA3 in predicting future states, such as when place cells in rat CA3 fire predictively to locations ahead of the current position during navigation [29]. CA3 may represent and retrieve temporal information via recurrent dynamics [28], allowing it to reproduce previously experienced sequences [32]. This predictive function may be a special case of 'pattern completion', whereby partial cues reinstate information that was present during encoding [33–35].

More generally, increases in representational similarity were not universal. ERC did not show increased pattern similarity for strong pairs in either ROI or searchlight analyses. In addition, across all ROIs, increased representational similarity only occurred when fractal pairings were arbitrary (as in main analyses): Strong pairs containing fractals that were visually similar to begin with did not show increases (Fig. S2). These visually similar pairs were presented under the same protocol as arbitrary pairs, suggesting specificity in our results not only to strong regularities, but also to regularities among dissimilar-looking objects.

Hippocampus and Regularities

While the hippocampus is often studied in terms of its important role in the formation of conscious or declarative memories [36], a different perspective on hippocampal function focuses instead on characterizing the kinds of learning processes that it can support—including relational, configural, and contextual learning [37]. We interpret our findings in this latter framework, with hippocampal changes reflecting the learning of temporal relationships. These relationships seemed to be learned incidentally: they were orthogonal to the behavioral task, they were not cued by instructions or timing, and no subjects reported noticing the pairs when debriefed. As a result, we have no reason to believe that the observed hippocampal changes require strategic or intentional processes. Our findings are thus consistent with the view that conscious accessibility may not be a necessary consequence of hippocampally-mediated learning [38]. While these issues are controversial [39], learning occurs rapidly in all accounts of hippocampal function [4,34]. Thus, irrespective of awareness, the hippocampus is well-suited to forms of statistical learning that can occur after minimal exposure [2,3].

Open Questions

Although we focused on increased pattern similarity for strong pairs, the opposite result was obtained for weak pairs. This intriguing finding suggests a non-monotonic relationship between regularity strength and representational change: Strong pairs appeared most frequently during exposure (average joint probability of 0.067 each) and showed increased pattern similarity, shuffled pairs were rare (0.006) and showed no change in similarity, and weak pairs appeared occasionally (0.022) but showed *decreased* similarity. How did intermediate probabilities lead to differentiated representations? One possibility is that during initial exposure to weak pairs, the first member generated a feeble prediction of the second member that was violated most of the time, leading to an eventual prediction of non-occurrence. This process might cause a reduction in pattern similarity for these items from before to after learning.

The role of the hippocampus and CA2/3/DG in particular in generating predictions raises interesting questions about what purpose these predictions serve. Such predictions may be compared against sensory input in the hippocampus, to refine existing knowledge [30,31]. Predictions may also output to more posterior temporal cortex. The same interlaminar circuits in PRC that provide feedforward input during sensory processing reverse during retrieval [40], providing feedback to IT [8]. Indeed, objects that afford predictions activate both the hippocampus and visual areas specialized for processing the predicted information [15]. Future work could relate patterns in MTL to the reinstatement of stimulus-specific patterns in visual cortex, and to facilitated sensory processing of predicted objects.

Conclusions

Our study demonstrates that temporal regularities cause representational changes in the MTL. This statistical learning can occur incidentally and rapidly, allowing the MTL to efficiently model the structure of the environment. Such representations may underlie the tendency in naturalistic settings for objects to effortlessly reactivate memories of other objects experienced in similar contexts.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We thank Chris Kelly for assistance with analyses, and Mike Arcaro, Ben Hutchinson, Ken Norman, and Alexa Tompary for helpful discussions. This work was supported by NSF GRFP DGE-0646086 (ACS), the Luce Foundation (LVK), and NIH R01-EY021755 (NTB).

References

1. Gilbert CD, Sigman M, Crist RE. The neural basis of perceptual learning. *Neuron*. 2001; 31:681–697. [PubMed: 11567610]
2. Saffran JR, Newport EL, Aslin RN. Word segmentation: the role of distributional cues. *J Mem Lang*. 1996; 35:606–621.
3. Turk-Browne NB, Scholl BJ, Chun MM, Johnson MK. Neural evidence of statistical learning: efficient detection of visual regularities without awareness. *J Cogn Neurosci*. 2009; 21:1934–1945. [PubMed: 18823241]
4. McClelland JL, McNaughton BL, O'Reilly RC. Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychol Rev*. 1995; 102:419–457. [PubMed: 7624455]
5. Miyashita Y. Neuronal correlate of visual associative long-term memory in the primate temporal cortex. *Nature*. 1988; 335:817–820. [PubMed: 3185711]
6. Miyashita Y. Inferior temporal cortex: where visual perception meets memory. *Annu Rev Neurosci*. 1993; 16:245–263. [PubMed: 8460893]
7. Erickson CA, Desimone R. Responses of macaque perirhinal neurons during and after visual stimulus association learning. *J Neurosci*. 1999; 19:10404–10416. [PubMed: 10575038]
8. Naya Y, Yoshida M, Miyashita Y. Backward spreading of memory-retrieval signal in the primate temporal cortex. *Science*. 2001; 291:661–664. [PubMed: 11158679]
9. Higuchi S, Miyashita Y. Formation of mnemonic neuronal responses to visual paired associates in inferotemporal cortex is impaired by perirhinal and entorhinal lesions. *Proc Natl Acad Sci U S A*. 1996; 93:739–743. [PubMed: 8570626]
10. Buckley MJ, Gaffan D. Perirhinal cortex ablation impairs configural learning and paired-associate learning equally. *Neuropsychologia*. 1998; 36:535–546. [PubMed: 9705064]

11. Wirth S, Yanike M, Frank LM, Smith AC, Brown EN, Suzuki WA. Single neurons in the monkey hippocampus and learning of new associations. *Science*. 2003; 300:1578–1581. [PubMed: 12791995]
12. Naya Y, Suzuki WA. Integrating what and when across the primate medial temporal lobe. *Science*. 2011; 333:773–776. [PubMed: 21817056]
13. Paz R, Gelbard-Sagiv H, Mukamel R, Harel M, Malach R, Fried I. A neural substrate in the human hippocampus for linking successive events. *Proc Natl Acad Sci U S A*. 2010; 107:6046–6051. [PubMed: 20231430]
14. Chun MM, Phelps EA. Memory deficits for implicit contextual information in amnesic subjects with hippocampal damage. *Nat Neurosci*. 1999; 2:844–847. [PubMed: 10461225]
15. Turk-Browne NB, Scholl BJ, Johnson MK, Chun MM. Implicit perceptual anticipation triggered by statistical learning. *J Neurosci*. 2010; 30:11177–11187. [PubMed: 20720125]
16. Harrison LM, Duggins A, Friston KJ. Encoding uncertainty in the hippocampus. *Neural Netw*. 2006; 19:535–546. [PubMed: 16527453]
17. Kumaran D, Maguire EA. An unexpected sequence of events: mismatch detection in the human hippocampus. *PLoS Biol*. 2006; 4:e424. [PubMed: 17132050]
18. Chadwick MJ, Hassabis D, Weiskopf N, Maguire EA. Decoding individual episodic memory traces in the human hippocampus. *Curr Biol*. 2010; 20:544–547. [PubMed: 20226665]
19. Rissman J, Wagner AD. Distributed representations in memory: insights from functional brain imaging. *Annu Rev Psychol*. 2012; 63:101–128. [PubMed: 21943171]
20. Johnson JD, McDuff SG, Rugg MD, Norman KA. Recollection, familiarity, and cortical reinstatement: a multivoxel pattern analysis. *Neuron*. 2009; 63:697–708. [PubMed: 19755111]
21. Haskins AL, Yonelinas AP, Quamme JR, Ranganath C. Perirhinal cortex supports encoding and familiarity-based recognition of novel associations. *Neuron*. 2008; 59:554–560. [PubMed: 18760692]
22. Fujimichi R, Naya Y, Koyano KW, Takeda M, Takeuchi D, Miyashita Y. Unitized representation of paired objects in area 35 of the macaque perirhinal cortex. *Eur J Neurosci*. 2010; 32:659–667. [PubMed: 20718858]
23. Hannula DE, Ranganath C. The eyes have it: hippocampal activity predicts expression of memory in eye movements. *Neuron*. 2009; 63:592–599. [PubMed: 19755103]
24. Prince SE, Daselaar SM, Cabeza R. Neural correlates of relational memory: successful encoding and retrieval of semantic and perceptual associations. *J Neurosci*. 2005; 25:1203–1210. [PubMed: 15689557]
25. Tubridy S, Davachi L. Medial temporal lobe contributions to episodic sequence encoding. *Cereb Cortex*. 2011; 21:272–280. [PubMed: 20494967]
26. Aminoff E, Gronau N, Bar M. The parahippocampal cortex mediates spatial and nonspatial associations. *Cereb Cortex*. 2007; 17:1493–1503. [PubMed: 16990438]
27. Turk-Browne NB, Simon MG, Sederberg PB. Scene representations in parahippocampal cortex depend on temporal context. *J Neurosci*. 2012; 32:7202–7207. [PubMed: 22623664]
28. Levy WB. A sequence predicting CA3 is a flexible associator that learns and uses context to solve hippocampal-like tasks. *Hippocampus*. 1996; 6:579–590. [PubMed: 9034847]
29. Lisman J, Redish AD. Prediction, sequences and the hippocampus. *Philos Trans R Soc Lond B Biol Sci*. 2009; 364:1193–1201. [PubMed: 19528000]
30. Duncan K, Ketz N, Inati SJ, Davachi L. Evidence for area CA1 as a match/mismatch detector: a high-resolution fMRI study of the human hippocampus. *Hippocampus*. 2012; 22:389–398. [PubMed: 21484934]
31. Chen J, Olsen RK, Preston AR, Glover GH, Wagner AD. Associative retrieval processes in the human medial temporal lobe: hippocampal retrieval success and CA1 mismatch detection. *Learn Mem*. 2011; 18:523–528. [PubMed: 21775513]
32. Fortin NJ, Agster KL, Eichenbaum HB. Critical role of the hippocampus in memory for sequences of events. *Nat Neurosci*. 2002; 5:458–462. [PubMed: 11976705]
33. Leutgeb S, Leutgeb JK. Pattern separation, pattern completion, and new neuronal codes within a continuous CA3 map. *Learn Mem*. 2007; 14:745–757. [PubMed: 18007018]

34. Norman KA, O'Reilly RC. Modeling hippocampal and neocortical contributions to recognition memory: a complementary-learning-systems approach. *Psychol Rev.* 2003; 110:611–646. [PubMed: 1459236]
35. Hoang LT, Kesner RP. Dorsal hippocampus, CA3, and CA1 lesions disrupt temporal sequence completion. *Behav Neurosci.* 2008; 122:9–15. [PubMed: 18298244]
36. Squire LR, Stark CE, Clark RE. The medial temporal lobe. *Annu Rev Neurosci.* 2004; 27:279–306. [PubMed: 15217334]
37. Cohen, NJ.; Eichenbaum, H. *Memory, Amnesia, and the Hippocampal System.* Cambridge: MIT Press; 1993.
38. Henke K. A model for memory systems based on processing modes rather than consciousness. *Nat Rev Neurosci.* 2010; 11:523–532. [PubMed: 20531422]
39. Kumaran D, Wagner AD. It's in my eyes, but it doesn't look that way to me. *Neuron.* 2009; 63:561–563. [PubMed: 19755098]
40. Takeuchi D, Hirabayashi T, Tamura K, Miyashita Y. Reversal of interlaminar signal between sensory and memory processing in monkey temporal cortex. *Science.* 2011; 331:1443–1447. [PubMed: 21415353]

Highlights

- Voxel patterns provide a window into how MTL represents object regularities
- Incidental exposure to objects paired in time increases similarity of their patterns
- Such changes occur in both cortical areas and hippocampal subfields of human MTL
- Regularities are represented in a forward-looking predictive manner in CA2/3/DG

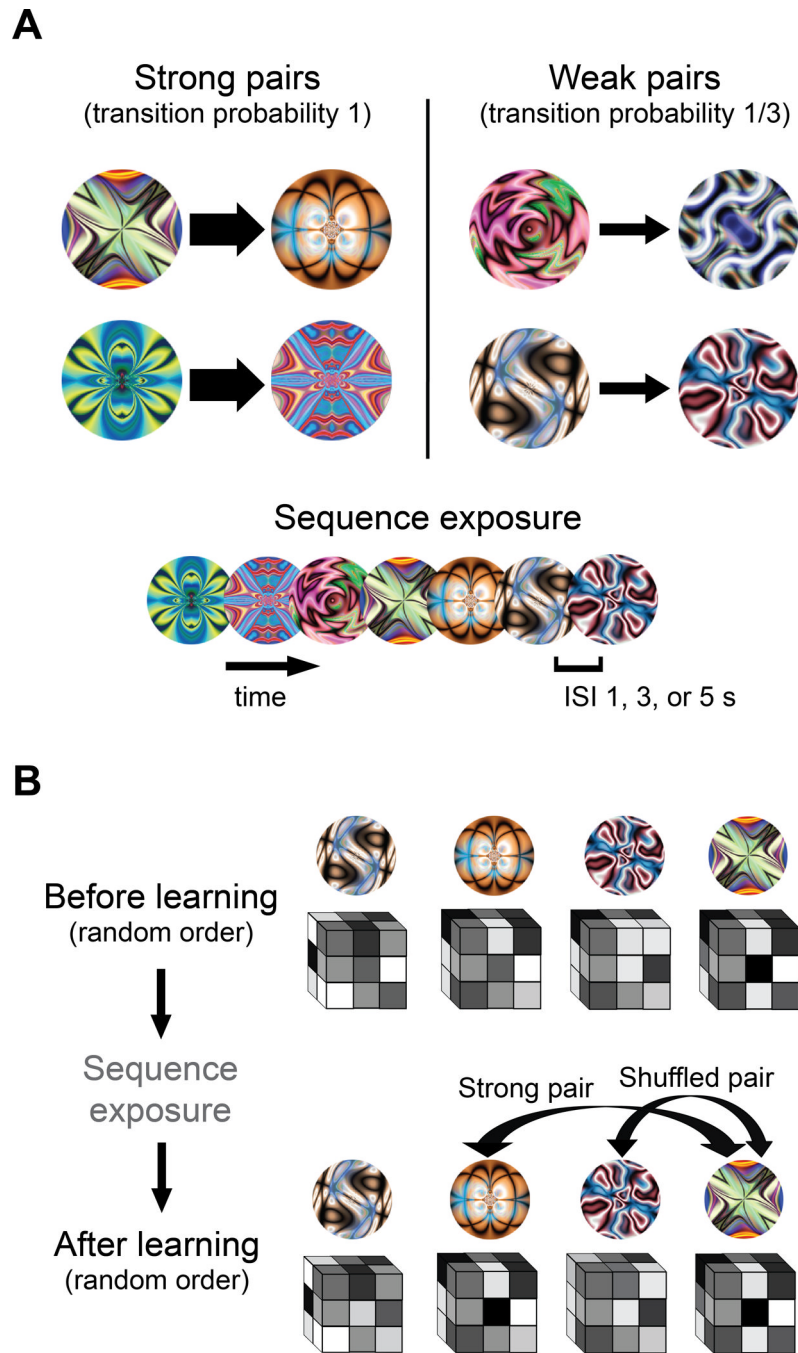
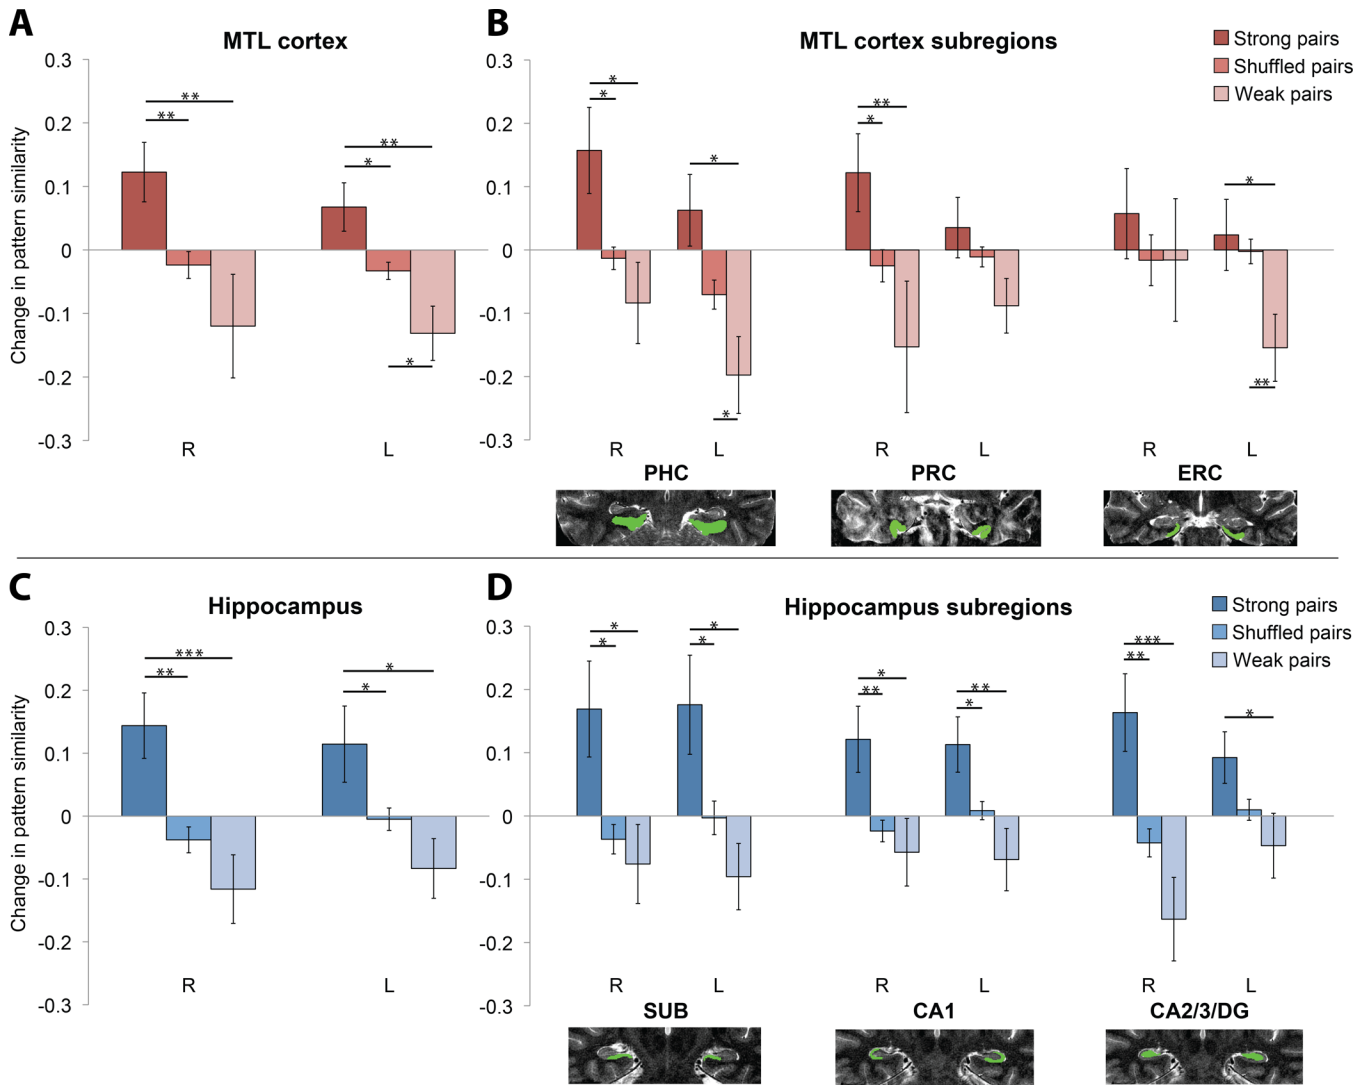


Figure 1.

(A) For each subject, fractals were randomly assigned to be the first or second member of 'strong' pairs or 'weak' pairs. During sequence exposure (middle runs), the order of fractals was generated from these pairs, with the constraint that no pair appeared twice in a row. For strong pairs, the first member was always followed by the second. For weak pairs, the first member was followed by the second on only one-third of trials. To equate the frequency of members from weak pairs, the second member was inserted into the trial sequence on its own on the remaining two-thirds of trials. Trials were presented continuously, with no grouping or segmentation cues to the pair structure other than the temporal regularities. (B)

Before and after sequence exposure (first and last run, respectively), the same fractals were presented in a random order. This allowed their representations to be measured while avoiding the concern that the temporal proximity of paired fractals would artificially increase their representational similarity. The same random order was used in the first and last run to equate for any spurious order effects or biases in modeling the BOLD response. For each fractal, the parameter estimates across voxels in each ROI were extracted and arranged into a vector. Pattern similarity was assessed by computing the Pearson correlation of vectors from different fractals. This produced three types of correlations: (1) between members of a strong pair, (2) between members of a weak pair, and (3) between members of different pairs ('shuffled' pairs). In all runs, fractals were presented for 1s, separated by a 1, 3 or 5s ISI. Subjects always performed an orthogonal cover task of detecting grayscale patches appearing randomly on 10% of otherwise colorful fractals. Subjects responded on every trial, indicating whether they saw a grayscale patch or not. See also Figure S1.

**Figure 2.**

Changes in pattern similarity (higher values indicate an increase) from before to after sequence exposure are shown for strong and weak pairs. The baseline (shuffled pairs) reflects the change in correlation for recombinations of fractal images into untrained pairs. Brain images show segmented ROIs on a T2 anatomical scan for a representative subject (R=right, L=left). See text for primary bilateral analyses. **(A)** MTL cortex: strong vs. shuffled pairs (R, $t[16]=3.27$, $p=.005$; L, $t[16]=2.27$, $p=.037$), strong vs. weak pairs (R, $t[16]=3.30$, $p=.004$; L, $t[16]=3.10$, $p=.007$), and weak vs. shuffled pairs (R, $t[16]=-1.26$, $p=.226$; L, $t[16]=-2.72$, $p=.015$). **(B)** MTL subregions: strong vs. shuffled pairs (R PHC, $t[16]=2.54$, $p=.022$; L PHC, $t[16]=1.95$, $p=.069$; R PRC, $t[16]=2.55$, $p=.021$; L PRC, $t[16]=1.13$, $p=.276$; R ERC, $t[16]=1.09$, $p=.291$; L ERC, $t[16]<1$), strong vs. weak pairs (R PHC, $t[16]=2.70$, $p=.016$; L PHC, $t[16]=2.83$, $p=.012$; R PRC, $t[16]=2.96$, $p=.009$; L PRC, $t[16]=1.88$, $p=.078$; R ERC, $t[16]<1$; L ERC, $t[16]=2.25$, $p=.039$), weak vs. shuffled pairs (R PHC, $t[16]=-1.05$, $p=.311$; L PHC, $t[16]=-2.62$, $p=.019$; R PRC, $t[16]=-1.28$, $p=.220$; L PRC, $t[16]=-1.79$, $p=.093$; R ERC, $t[16]<1$; L ERC, $t[16]=-3.31$, $p=.004$). **(C)** Hippocampus: strong vs. shuffled pairs (R, $t[16]=3.79$, $p=.002$; L, $t[16]=2.20$, $p=.042$), strong vs. weak pairs (R, $t[16]=4.27$, $p<.001$; L, $t[16]=2.51$, $p=.023$), weak vs. shuffled pairs (R, $t[16]=-1.46$, $p=.165$; L, $t[16]=-1.78$, $p=.095$). **(D)** Hippocampus subregions: strong vs. shuffled

pairs (R subiculum [SUB], $t[16]=2.77$, $p=.014$; L SUB, $t[16]=2.66$, $p=.017$; R CA1, $t[16]=3.03$, $p=.008$; L CA1, $t[16]=2.40$, $p=.029$; R CA2/3/DG, $t[16]=3.67$, $p=.002$; L CA2/3/DG, $t[16]=2.11$, $p=.051$), strong vs. weak pairs (R SUB, $t[16]=2.72$, $p=.015$; L SUB, $t[16]=2.58$, $p=.020$; R CA1, $t[16]=2.62$, $p=.019$; L CA1, $t[16]=3.27$, $p=.005$; R CA2/3/DG, $t[16]=4.26$, $p<.001$; L CA2/3/DG, $t[16]=2.47$, $p=.025$), weak vs. shuffled pairs (R SUB, $t<1$; L SUB, $t[16]=-1.86$, $p=.082$; R CA1, $t<1$; L CA1, $t[16]=-1.76$, $p=.098$; R CA2/3/DG, $t[16]=-2.03$, $p=.060$; L CA2/3/DG, $t[16]=-1.27$, $p=.221$). * $p<.05$; ** $p<.01$; *** $p<.001$. Error bars denote ± 1 SEM. See also Figure S2.




Figure 3.

ROIs for each subject were warped to a common template (green). Searchlights within right and left PHC, left PRC, left SUB, right CA1, and right CA2/3/DG showed a greater increase in pattern similarity for strong vs. weak pairs ($p < .001$ uncorrected). The lack of a left PRC increase in ROI analyses suggests that noise from other voxels may have swamped a local effect. The lack of searchlight effects in right PRC, right SUB, left CA1, and left CA2/3/DG despite increases in ROI analyses suggests that the underlying representations were distributed beyond the scope of searchlights, that we benefitted from the increased statistical power of ROI analyses and greater voxel sample sizes, and/or that specific locations of local changes were misaligned across subjects. For visualization, differences in pattern similarity for strong vs. weak pairs were assigned to the center voxel of each searchlight and resulting statistical maps were thresholded at $p < .01$. See also Figure S3 and Table S1.

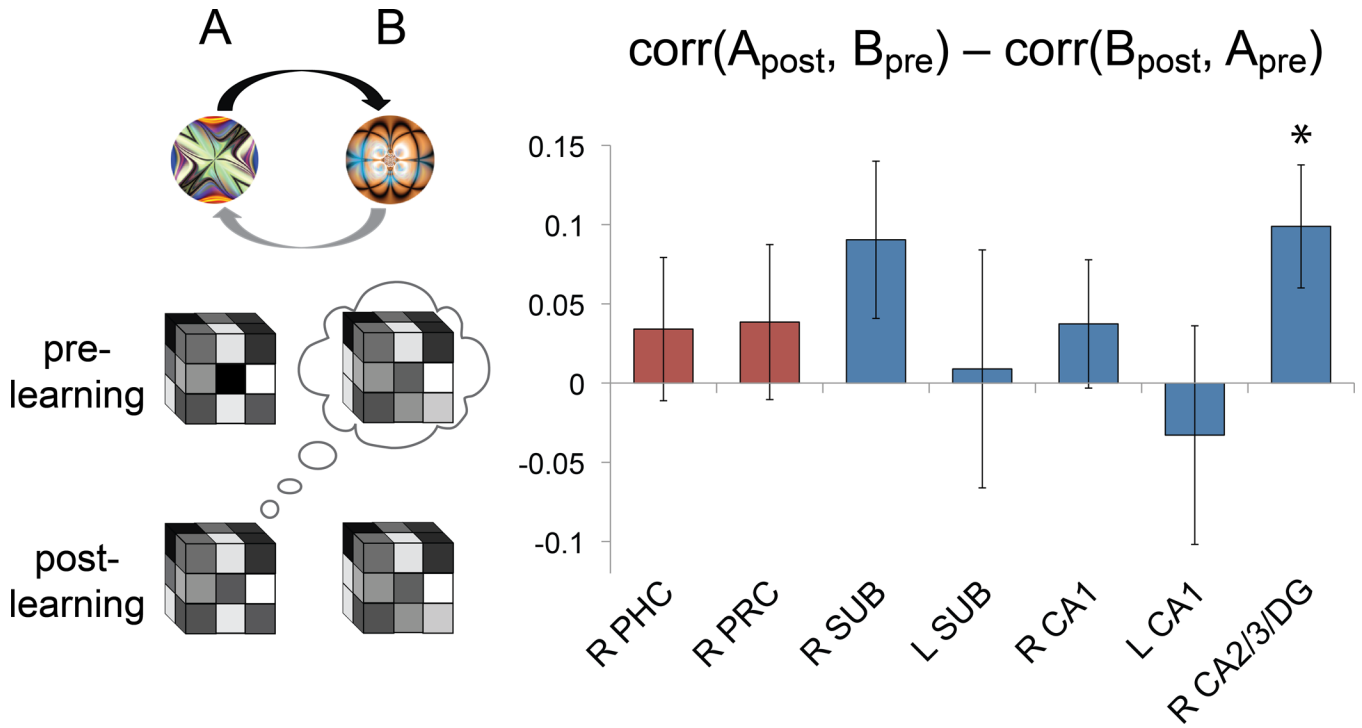


Figure 4. Asymmetry index. If the first member of a pair (fractal A) elicits the representation of the second member (fractal B) but not vice versa, then the correlation of A post-learning with B pre-learning should be greater than that of B post-learning with A pre-learning. Among the seven ROIs that showed increased pattern similarity for strong pairs, only right CA2/3/DG was reliably asymmetric. * $p < .05$. Error bars denote ± 1 SEM. See also Figure S4.