The Necessity of the Medial Temporal Lobe for Statistical Learning

Anna C. Schapiro[1], Emma Gregory[2], Barbara Landau[2],
Michael McCloskey[2], and Nicholas B. Turk-Browne[1]

Abstract

■ The sensory input that we experience is highly patterned, and we are experts at detecting these regularities. Although the extraction of such regularities, or statistical learning (SL), is typically viewed as a cortical process, recent studies have implicated the medial temporal lobe (MTL), including the hippocampus. These studies have employed fMRI, leaving open the possibility that the MTL is involved but not necessary for SL. Here, we examined this issue in a case study of LSJ, a patient with complete bilateral hippocampal loss and broader MTL damage. In Experiments 1 and 2, LSJ and matched control participants were passively exposed to a continuous sequence of shapes, syllables, scenes, or tones containing temporal regularities in the co-occurrence of items. In a subsequent test phase, the control groups exhibited reliable SL in all conditions, successfully discriminating regularities from recombinations of the same items into novel foil sequences. LSJ, however, exhibited no SL, failing to discriminate regularities from foils. Experiment 3 ruled out more general explanations for this failure, such as inattention during exposure or difficulty following test instructions, by showing that LSJ could discriminate which individual items had been exposed. These findings provide converging support for the importance of the MTL in extracting temporal regularities. ■

INTRODUCTION

The same items often reliably co-occur in time or space across repeated experiences. The mind can identify such statistical relationships by aggregating across these experiences and extracting probabilistic information—a process often referred to as statistical learning (SL). For example, 2 min of exposure to a novel language is sufficient for 8-month-old infants to acquire knowledge of the words embedded therein, based purely on the temporal statistics of the sequence of syllables (Saffran, Aslin, & Newport, 1996). This form of learning is also robust in adults and is observed for many modalities and types of stimuli (Baldwin, Andersson, Saffran, & Meyer, 2008; Conway & Christiansen, 2005; Fiser & Aslin, 2001, 2002a; Saffran, Johnson, Aslin, & Newport, 1999; Saffran et al., 1996).

The need to extract statistics over several exposures, along with the fact that SL is largely automatic and implicit (Kim, Seitz, Feenstra, & Shams, 2009; Turk-Browne, Junge, & Scholl, 2005), suggests that it might be accomplished in sensory or association cortex (Simoncelli & Olshausen, 2001; Gabrieli, Fleischman, Keane, Reminger, & Morrell, 1995; McClelland, McNaughton, & O'Reilly, 1995) or in the striatal memory system (Knowlton, Mangels, & Squire, 1996). In contrast, the medial temporal lobe (MTL) memory system, which specializes in declarative, rapid learning of individual episodes (Squire, Stark, & Clark, 2004), might not be expected to be involved.

There are several reasons, however, to suspect that the MTL, and hippocampus in particular, may be the most relevant neural circuitry. First, although the learning requires several exposures, it occurs on a timescale that is consistent with the rapid binding known to occur in the hippocampus (Cohen & Eichenbaum, 1993) and less consistent with the typically much slower cortical learning for novel information.

Second, fMRI studies have demonstrated the involvement of the hippocampus in SL. If two objects consistently appear in succession, the representations of these objects in the hippocampus and MTL cortex become more similar to each other (Schapiro, Kustner, & Turk-Browne, 2012). The hippocampus and MTL cortex are also responsive to various kinds of temporal statistical information (Bornstein & Daw, 2012; Durrant, Cairney, & Lewis, 2012; Turk-Browne, Scholl, Johnson, & Chun, 2010; Turk-Browne, Scholl, Chun, & Johnson, 2009; Harrison, Duggins, & Friston, 2006; Rose, Haider, & Buchel, 2005; Schendan, Searl, Melrose, & Stern, 2003).

The third reason to suspect that the MTL may be involved is that, despite surface appearances, SL may be closely related to other forms of learning that the MTL is thought to specialize in. The hippocampus, in particular, is crucial for the rapid encoding of relations between arbitrary elements of sensory information (Cohen & Eichenbaum, 1993). Learning and remembering regularities also requires binding stimuli that are temporally or spatially associated. In SL, this binding occurs over multiple experiences, and indeed, this is also a property of

[1]Princeton University, [2]Johns Hopkins University

hippocampal function, as shown in paradigms like the transitive inference task (Dusek & Eichenbaum, 1997). Unlike this task, however, the individual episodes being integrated in SL are not clearly defined—they are embedded in a continuous sequence of stimuli and must be segmented. Still, the relational and integrative aspects of MTL function are well suited to SL.

Although there are indications that the MTL is involved in SL, it remains unclear whether it is necessary. For instance, the fMRI activation seen during SL may be epiphenomenal or a byproduct of more essential learning elsewhere in the brain. To obtain stronger evidence for the necessity of the MTL for SL, we report a case study of LSJ, a patient with complete bilateral hippocampal loss and some additional damage to surrounding MTL cortex and left anterior temporal lobe. Across three experiments, LSJ was passively exposed to a continuous sequence of shapes, syllables, scenes, or tones containing temporal regularities. In all stimulus classes and for both relatively complex and simpler regularities, LSJ exhibited no evidence of SL. This finding contrasts with results from groups of age- and education-matched control participants, who showed reliable SL for all stimulus classes. LSJ was, however, able to remember the individual stimuli presented in the sequences, suggesting that her SL deficits did not result from general difficulties with task demands. These results provide converging support for the hypothesis that the MTL is necessary for the successful extraction of temporal regularities from the environment.

CASE REPORT

LSJ, a 62-year-old woman with a Bachelor of Fine Arts degree, was a highly successful illustrator and an accomplished amateur violist. In 2007 at age 57, she contracted herpes encephalitis, resulting in bilateral destruction of the hippocampus, bilateral damage to other MTL regions, and left anterior temporal lobe damage (Figure 1A).

Consistent with her extensive MTL damage and the neuropsychological profile of other cases of herpes encephalitis (e.g., Kapur et al., 1994), LSJ exhibits profound retrograde and anterograde memory impairment (see Gregory, McCloskey, & Landau, 2014, for a more detailed neuropsychological profile). The retrograde memory impairment encompasses both episodic (autobiographical) and semantic information. In the autobiographical domain, extensive testing has revealed no memory for specific life events; for example, LSJ is unable to report anything about her 10-year marriage. In testing of semantic memory, LSJ was severely impaired in identifying famous faces (2 of 60 correct) and famous landmarks (2 of 26; see Duchaine, Germine, & Nakayama, 2007; Duchaine, Nieminen-von Wendt, New, & Kulomaki, 2003, for control means). In testing with recall and recognition tasks developed to probe knowledge across a broad range of domains (e.g., commercial logos, familiar songs), including areas of premorbid expertise for LSJ (e.g., famous artworks, classical music), LSJ was uniformly and severely impaired (Gregory et al., 2014).

LSJ's profound anterograde amnesia was evident on the Wechsler Memory Scale II (Wechsler, 1997), where she scored below the first percentile on the General Memory Index. Anecdotally, LSJ consistently shows no recollection for tasks when questioned after completing them and expresses no recollection for experimenters across sessions, despite having met them many times.

In contrast to her impaired remote memory and inability to form new memories, LSJ scored in the normal range on tests of single-word reading and spelling (Wide Range Achievement Test III, Wilkinson, 1993: 58th and 55th
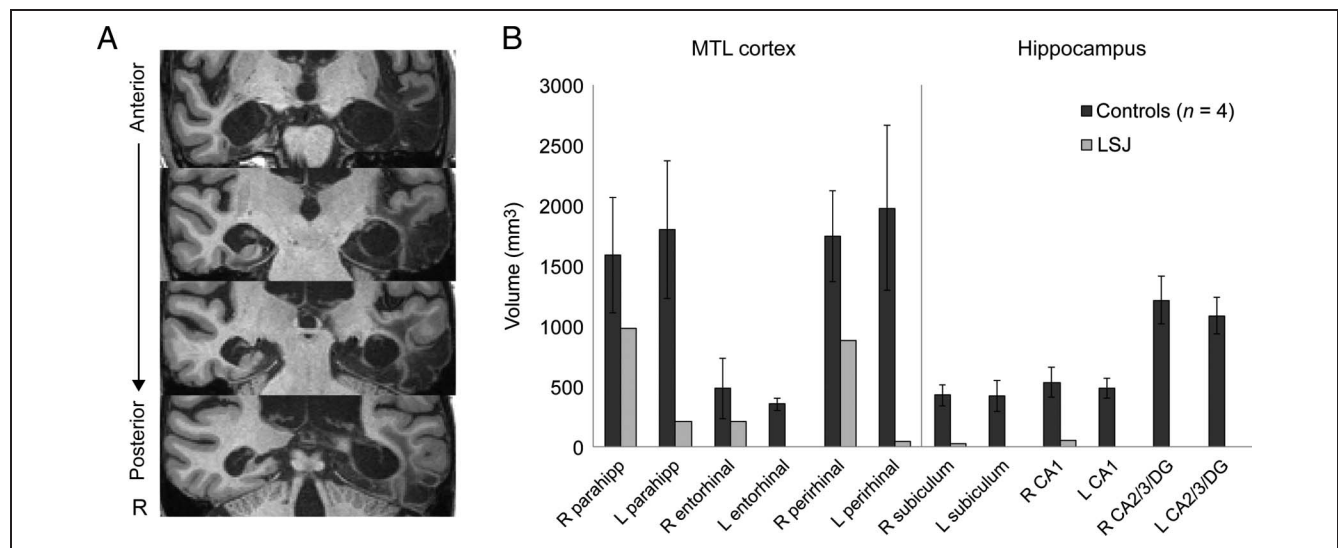
Figure 1. MTL anatomy. (A) Coronal slices from T1-weighted MP-RAGE sequence illustrating the lesions in LSJ. (B) Quantitative evaluation of MTL tissue volume in LSJ relative to controls. Segmented subregions of the MTL cortex and hippocampus include parahippocampal cortex (parahipp), entorhinal cortex, perirhinal cortex, subiculum, CA1, and CA2/CA3/dentate gyrus (CA2/3/DG). Error bars denote ± 1 SD. R = right; L = left.

percentile, respectively) and vocabulary (PPVT-R, Dunn & Dunn, 1981: 63rd percentile). Anecdotally, her language comprehension and production seem normal.

LSJ's visuospatial abilities also appear intact: Her performance was normal on the Block Design and Matrix Reasoning subtests of the WAIS-IV (Wechsler, 2008) and on all subtests of the Visual Object and Space Perception Battery (Warrington & James, 1991). She also scored in the normal range (five of six subtests) on the Montreal Battery of Evaluation of Amusia (Peretz, Champod, & Hyde, 2003). The only impairment was on a yes/no recognition memory subtest, which is uninformative, as she displayed a familiarity bias (answering "yes" to 29 of 30 items). Finally, LSJ continues to make art, including drawings and watercolor paintings, and she is able to sight read and play her viola.

MTL VOLUME ANALYSIS

To assess in detail the extent and location of LSJ's lesions, we acquired high-resolution anatomical MRI scans from LSJ and control participants.

Participants

Four controls were matched to LSJ on age (mean age = 62.0 years, range = 57–65 years) and education (college educated), had no history of neurological disease or disorder, and were native English speakers. Participants were compensated $20 per hour. In this MRI study and the behavioral experiments below, informed written consent was obtained from all participants (for LSJ, with power of attorney), and the protocol was approved by the Institutional Review Board for Human Subjects at Princeton University.

Data Acquisition

Data were acquired with a 3T Siemens Skyra scanner using a 16-channel head coil. We collected a high-resolution 3-D T1-weighted magnetization prepared rapid gradient echo sequence (0.9 × 0.9 mm in-plane, 0.9-mm thickness), shown for LSJ in Figure 1A. For manual segmentation of MTL ROIs, we collected a high-resolution T2-weighted turbo spin echo sequence (33 slices, 0.4 × 0.4 mm in-plane, 3-mm thickness), with oblique coronal slices aligned perpendicular to the anterior–posterior axis of the hippocampus.

Segmentation

For each participant, parahippocampal cortex, perirhinal cortex, entorhinal cortex, subiculum, CA1, and CA2/3/DG (CA2, CA3, and dentate gyrus cannot be separated at this resolution) were hand-segmented in each hemisphere on the T2-weighted anatomical images. The boundaries between regions were defined based on anatomical land-marks, following conventions used in previous studies (Schapiro et al., 2012; Carr, Rissman, & Wagner, 2010). Though it was difficult to identify certain anatomical landmarks for LSJ due to her damage, two coders performed independent segmentations as best as possible.

Volume Analysis

We calculated the volume of each MTL region based on the number of voxels falling within each hand-segmented ROI (Figure 1B). Volumes were averaged across control participants and across the segmentations performed by the two coders for LSJ. There is substantial remaining tissue in LSJ's right parahippocampal cortex, right entorhinal cortex, and right perirhinal cortex. Relative to controls, LSJ has very little, if any, remaining tissue in left parahippocampal cortex, left entorhinal cortex, and left perirhinal cortex and virtually no tissue in hippocampus bilaterally. In addition, regions that receive projections from the MTL may also be affected as a result of damaged connectivity.

EXPERIMENT 1

We began by testing LSJ on a canonical SL task (Fiser & Aslin, 2002a; Saffran et al., 1996). We repeated this experiment four times, with shapes, syllables, scenes, and tones, to test whether any deficits generalize across modalities and stimulus types.

Methods

Participants

LSJ and 28 healthy controls participated in exchange for monetary compensation ($12 per hour). One additional control participant was excluded because she reported that, during the test phase, she sometimes used the left and right side keys to indicate the first- and second-presented alternative, respectively, instead of the intended mapping to left and right screen/speaker. Note that the possibility for this sort of confusion did not arise for LSJ, because she responded by pointing to the side of the screen or the speaker where her preferred response alternative was presented (described below). Controls were matched to LSJ's age (mean age = 62.9 years, range = 57–67 years) and educational level. They had normal hearing and normal or corrected-to-normal vision, had no history of neurological disease or disorder, and were native English speakers.

The control participants were each randomly assigned to one version of the experiment—shape, syllable, scene, or tone (seven unique participants per version). LSJ participated in all four versions of the experiment. Note that testing a participant repeatedly is unusual for SL studies, which generally involve incidental learning and surprise tests. After participating in one SL task, normal participants

Figure 2. Visual (A) shape and (B) scene stimuli used in the experiments. Auditory syllables and tones were also used.
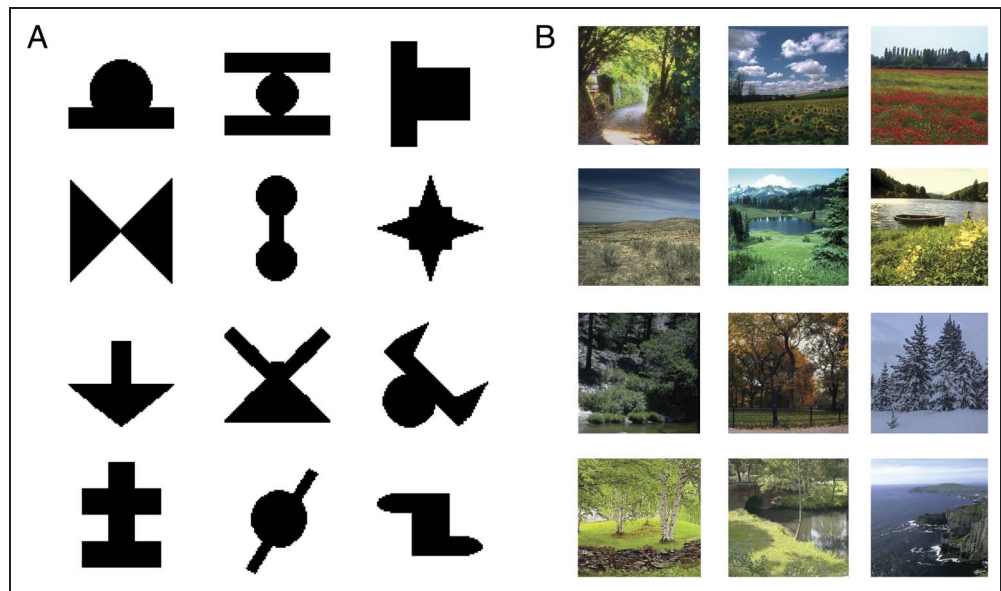
would know how the regularities were constructed (based on the structure of the alternatives in the test phase) and would expect to be tested on these regularities. However, we hypothesized that LSJ would have difficulty with SL, and so any benefit from repeated testing worked against our hypothesis. In any case, we saw no evidence that LSJ remembered participating in any of the SL tasks.

Apparatus and Stimuli

Visual and auditory stimuli were presented using Matlab with Psychophysics Toolbox (Brainard, 1997) on a MacBook laptop. Participants sat approximately 46 cm from the monitor without head restraint. Shape and scene stimuli subtended roughly 3.4° and 10.1° of visual angle, respectively. For tone and syllable versions, two portable speakers were placed on either side of the laptop, approximately 30 cm away from the edge of the laptop on each side.

Shape stimuli (Figure 2A) were 12 nonsense shapes used previously in SL studies (e.g., Fiser & Aslin, 2001), presented on a white background. Scene stimuli (Figure 2B) were 12 outdoor scene photographs, presented on a black background. Syllable stimuli were generated from the Mbrola (tcts.fpms.ac.be/synthesis/) speech synthesizer with a male voice and included: *bo*, *di*, *nu*, *ru*, *la*, *ze*, *me*, *ka*, *wu*, *va*, *pi*, and *fo*. Tone stimuli were pure tones at the following frequencies: 262, 278, 294, 311, 330, 349, 370, 392, 415, 440, 466, and 494 Hz. Stimulus duration was 0.5 sec, and the ISI was 0.5 sec for shapes, scenes, and tones. For syllables, a continuous sequence was synthesized with a syllable duration of 0.25 sec and no ISI.

Procedure

The experiment was conducted in two phases: exposure and test. In the exposure phase, participants were in-

structed to watch or listen to a sequence of shapes, syllables, scenes, or tones (with no mention of a memory test). Syllables and tones were played from both speakers. Shape, scene, and tone sequences were 288 items long and lasted 4.8 min, and syllables sequences were 1,152 items long and lasted 4.8 min. For each stimulus type, the 12 items (e.g., A, B, C, … L) were randomly assigned without replacement to the first, second, or third positions of four triplets (e.g., ABC, DEF, GHI, JKL). Items within a triplet always appeared in a fixed order consecutively in the sequence. The order of the triplets was random, except that no triplet could be repeated in immediate succession (e.g., ABCABC) and no pair of triplets could be repeated in immediate succession (e.g., ABCDEFABCDEF). These standard constraints were imposed to minimize the likelihood of explicit detection of the patterns. Control participants were given the same item-to-triplet assignments as LSJ in case certain items were more or less associable, but the order of triplets in the sequence was randomized for each participant.

In the test phase, participants completed 32 two-alternative forced-choice test trials, judging the relative familiarity of a triplet from the exposure phase compared with a foil. The experimenter explained how the task worked and how to make responses and then asked if the participant had any questions. The experimenter supervised the test phase to make sure the participant complied with instructions and to be available in case any questions arose. The experimenter encouraged participants to make their best guess when they were unsure, to facilitate expression of implicit recognition (Voss & Paller, 2010).

The four foils were composed of three unique items from three different triplets, such that the items could never have appeared consecutively during exposure (e.g., AEI, DHL, GKC, JBF). The four target triplets were paired

with each of the four foils once in the first half of test trials and again in the second half. Otherwise, the order of the test trials was random. Thus, target triplets and foils were presented equally often during test, and any relative familiarity for the target triplets must reflect SL from the exposure phase.

On each test trial of the visual versions of the experiment, the three items from the first alternative (counterbalanced to be the target triplet or foil) were presented on the left or right side of the screen, with the same timing as during exposure. After a 1-sec gap, the three items from the second alternative (target triplet or foil, whichever did not occur previously) were presented on the opposite side of the screen. The side of the screen on which the target triplet was presented was determined randomly on each trial. Control participants responded by pressing a key corresponding to the side on which the more familiar stimulus sequence had been presented. They then pressed a different key to initiate the next trial. LSJ indicated her responses by pointing to a side of the screen or speaker or by verbally saying "left" or "right." The experimenter pressed the keyboard button to initiate the next trial. We decided not to have LSJ use the keyboard for two reasons. First, we wanted to be sure to have the opportunity to remind her on every trial about the task instructions before the trial began. Second, this simplified the task by removing the need to remember response mappings—she simply pointed to or verbally indicated her choice. We wanted to do everything possible to facilitate expression of her knowledge. LSJ's responses were recorded by the experimenter, and the recorded responses were subsequently verified against a video recording of the test session.

All participants had the option of viewing the trial again before making their decision. Controls used a keypress to repeat the trial, and LSJ indicated verbally whether she would like to repeat the trial. We included this trial repeat option in the experiment to ensure that LSJ had sufficient opportunity to observe and consider the alternatives.

The procedure was identical for auditory versions of the experiment, except that sequences of syllables or tones were played from the left or right speaker instead of displayed on the left or right side of the screen. Additionally, arrows appeared on the screen for the duration of the alternative, pointing to the speaker that was playing the stimulus.

For controls, after the experiment was over, we asked (1) what they did while watching or listening in the exposure phase, (2) whether they used any particular strategy during the test (and if so, what strategy), and (3) whether they noticed any repeating patterns during the exposure phase (and if so, what the patterns were). For Question 1, the responses were grouped into four categories: simply watched/listened, looked for patterns in the sequence, counted the items, and assigned names to the items. For Question 2, the responses were again grouped into four categories: no particular strategy, tried to think back to exposure phase, used visual or auditory properties, and said the names of the items. For Question 3, we coded three levels of awareness: did not notice any patterns, noticed something about the patterns (could describe one pattern or knew that items came in groups but not which ones went together), and noticed a substantial amount about the patterns (could accurately describe two or more patterns).

For LSJ, the tasks were conducted in the following order: shapes, scenes, syllables, and tones. Shapes and

Figure 3. Test performance for control participants and LSJ in Experiment 1. For each of the shape, syllable, scene, and tone tasks, LSJ's score (black cross) and the individual control scores (gray dots) are plotted against chance (dashed line), along with the mean and standard error of the controls. To visualize LSJ's performance with respect to the entire sample of controls, the data in each task were converted to *z* scores and then collapsed across task (LSJ's four *z* scores were averaged). She fell 1.74 *SD*s below the mean on average.
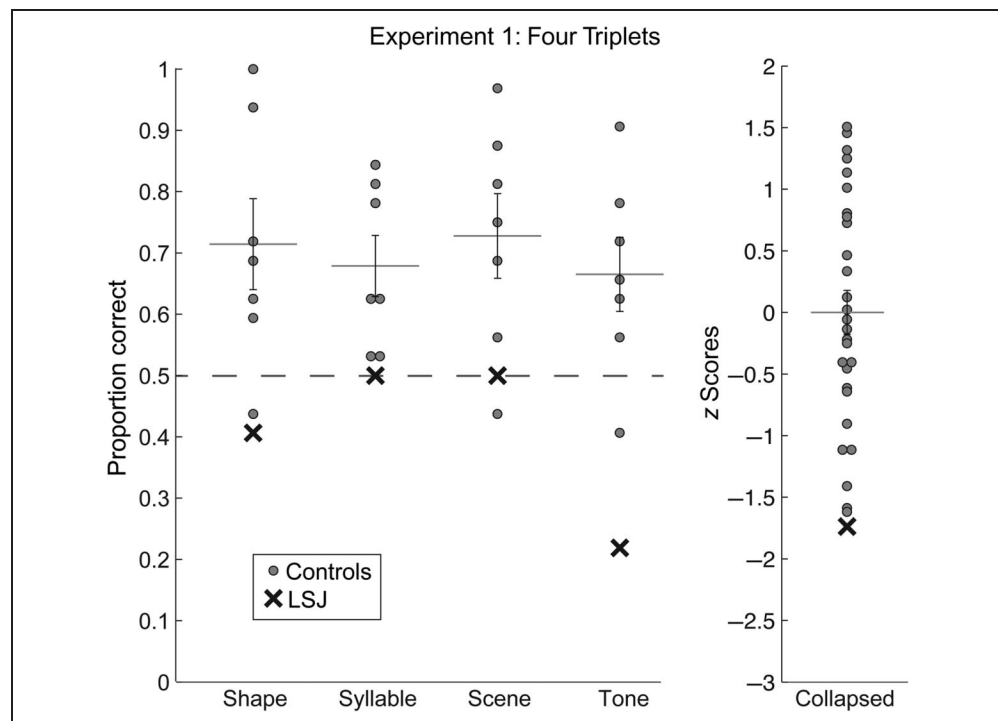
Table 1. Test Performance for LSJ and Controls

Task	Experiment 1 (32 Test Trials)	Experiment 2 (18 Test Trials)	Experiment 3 (18 Test Trials)
LSJ			
Shape	13	7	11
Syllable	16	10	14
Scene	16	10	17
Tone	7	11	9
Controls			
Shape	22.86 (2.37)	15.00 (0.63)	
Syllable	21.71 (1.60)	15.00 (1.67)	
Scene	23.29 (2.21)	14.40 (1.21)	
Tone	21.29 (1.94)	12.60 (1.08)	

For LSJ, the raw number of correct items is reported for each task in Experiments 1–3. For controls, the mean performance and ±1 *SEM* in parentheses are reported for each task and for the two experiments in which controls participated (Experiments 1 and 2).

scenes were administered on the first day, and syllables and tones were administered on a second day a month later. On each testing day, LSJ participated in several other experiments and had limited availability. We chose tasks that seemed maximally different from the other ones on that day to minimize interference.

Results and Discussion

Control groups exhibited reliable SL, accurately choosing the triplets as more familiar than the foils in the test phase of all versions (Figure 3, Table 1). Collapsing across version, controls were significantly above chance ($t[27] = 6.42, p < .001$). Examining reliability across the mean performance in each version, control accuracy was again significantly above chance ($t[3] = 13.37, p = .0009$). This was also true across participants within each version (shape: $t[6] = 2.89, p = .028$; syllable: $t[6] = 3.57, p = .012$; scene: $t[6] = 3.29, p = .017$; tone: $t[6] = 2.73, p = .034$). Indeed, there was no effect of version on performance ($F[3, 24] = 0.21, p = .888$). These control data, in addition to providing a comparison with LSJ, also provide information about older adults' SL abilities, which have been relatively unexplored (Campbell, Zimerman, Healey, Lee, & Hasher, 2012).

In contrast to the control groups, LSJ did not show evidence of learning the regularities in any version (Figure 3, Table 1). Across versions, her accuracy did not differ from chance ($t[3] = -1.41, p = .252$). She performed significantly worse than controls relative to their mean performance across version ($t[3] = -4.96, p = .016$), and when comparing her mean performance to the control distribution collapsed across versions ($t[30] = 3.41$,

$p = .002$). This provides an initial indication that LSJ has a deficit in SL.

The low performance in the tone version is surprising. During the tone test, LSJ indicated that she was choosing the sequences that sounded familiar from her past musical experience (although she seemed to understand that this was not what we were asking her to do). By chance, these familiar sequences were foils rather than triplets, leading to below-chance performance. We conclude that the results provide no indication that LSJ had learned the triplets in any version of the experiment.

LSJ did not show any bias to choose the alternative on the left versus right side of the screen or the alternative that appeared first versus second. Across tasks, the average proportion of trials on which she chose the left alternative was 0.563 (vs. chance: $t[3] = 1.633, p = .201$), and the average proportion of trials on which she chose the first alternative was 0.422 (vs. chance: $t[3] = 1.89, p = .155$). This suggests that she was not simply neglecting a side of the screen or showing a preference for a particular order or location. Across all versions, controls requested 7.86 trial repeats on average (*SD* across 28 participants = 9.29), and LSJ requested 21.5 trial repeats on average (*SD* across four tasks = 10.61).

To analyze the control participant debriefing responses, we collapsed across modalities and performed one-way ANOVAs to assess whether there were any differences in test performance across response types. For each of the three debriefing questions, we did not find any relationship between responses and performance level (ps > .175). This is of particular interest for the question regarding awareness of regularities. The participants who did not notice anything about the regularities (16 participants) performed significantly above chance (mean proportional correct = 0.660, $t[15] = 3.78, p = .002$) and no differently from the participants who did notice something about the patterns (12 participants; mean = 0.745, $t[26] = 1.39$, $p = .176$); no participants noticed a substantial amount about the patterns.

EXPERIMENT 2

Control performance was not especially high in Experiment 1. Although most controls performed well and LSJ performed poorly, we wanted to gain greater confidence that she was outside the normal range. We therefore repeated the same experiment using simpler regularities that we expected the controls to learn more robustly. If LSJ is fundamentally impaired in SL, her performance should remain at chance.

Methods

Participants

LSJ and 20 new controls participated in exchange for monetary compensation. Three additional control participants

were excluded because of their difficulty understanding response key instructions during the test phase: Two of these participants reported that they sometimes used the left and right side keys to indicate the first- and second-presented alternative, respectively, instead of the left- and right-side alternative; the third confused the trial repeat key with the right side key. Control participants were selected using the same criteria as in Experiment 1 (here, mean age = 62.1 years, range = 57–67 years). They were each randomly assigned to one version of the experiment—shape, syllable, scene, or tone (five unique participants per version). LSJ participated in all four versions of the experiment.

Stimuli

Stimuli were 6 of the 12 stimuli from each stimulus class in Experiment 1. The subset was chosen randomly for shapes, syllables, and scenes. For tones, the frequencies 262, 294, 330, 370, 415, and 466 Hz were chosen to span the initial frequency range from Experiment 1.

Procedure

The exposure phase was nearly identical to Experiment 1. The only differences were that six rather than 12 items were presented per stimulus class, the six items were divided into pairs instead of triplets, and there were no constraints on sequence randomization (repetitions were allowed). These modifications dramatically simplified the regularities, mirroring a design that has been used to examine temporal and spatial visual SL in infants (Fiser & Aslin, 2002b; Kirkham, Slemmer, & Johnson, 2002). The total number of items in each sequence and the duration of the sequence were matched to Experiment 1. Thus, in addition to using simpler regularities, each pattern was presented approximately 1.5 times more frequently than in Experiment 1. These changes were all intended to boost overall SL in controls.

The test phase was the same as Experiment 1, except that the test trials were generated from pairs. Now, three target pairs and three foils (where a foil consisted of two items from two different target pairs) were combined in all nine possible ways. Each combination was repeated once in the first and second halves of the test, for a total of 18 trials. The target pair and foil were again presented consecutively in a counterbalanced order, with the same timing as exposure, and with the left or right speaker or side of the screen.

For LSJ, the tasks were conducted in the following order: tones, shapes, syllables, and scenes. Tones and shapes were administered on the first day, and syllables and scenes were administered on a second day approximately a month later. The second test session for Experiment 1 and the first test session for Experiment 2 were approximately 4 months apart.

Results and Discussion

Control groups again exhibited reliable SL, accurately choosing the pairs as more familiar than the foils in the test phase of all versions (Figure 4, Table 1). Collapsing across version, controls were significantly above chance ($t[19] = 8.78, p < .001$) and significantly better than controls

Figure 4. Test performance for control participants and LSJ in Experiment 2. For each of the shape, syllable, scene, and tone tasks, LSJ's score (black cross) and the individual control scores (gray dots) are plotted against chance (dashed line), along with the mean and standard error of the controls. To visualize LSJ's performance with respect to the entire sample of controls, the data in each task were converted to z scores and then collapsed across task (LSJ's four z scores were averaged). She fell 2.32 *SD*s below the mean on average.
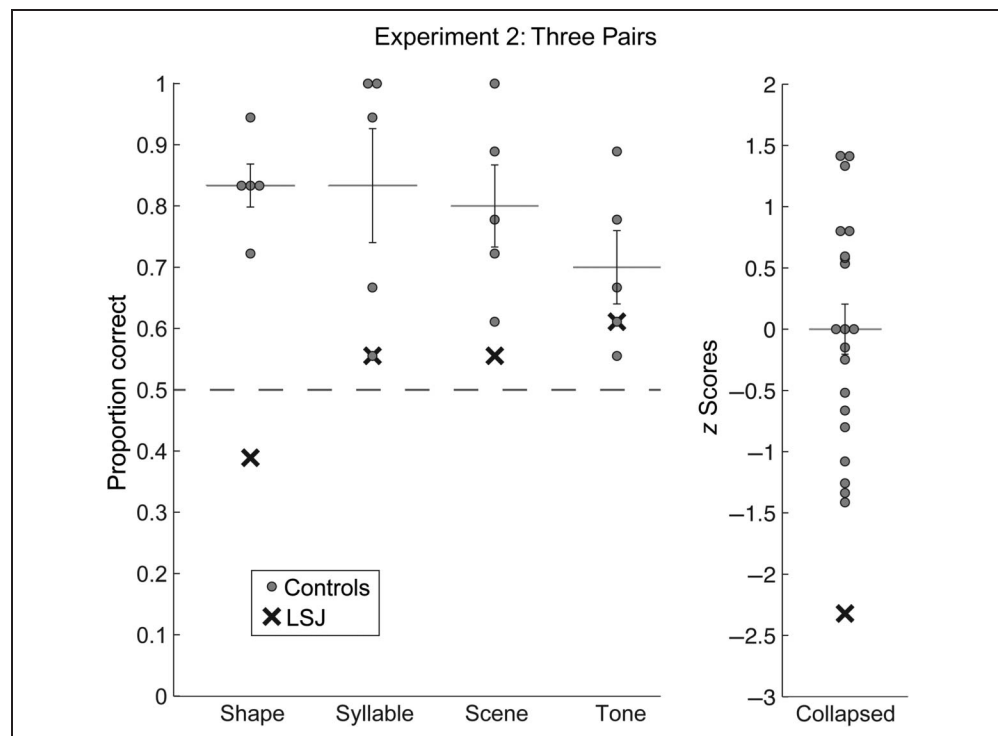
in Experiment 1 ($t[46] = 2.08$, $p = .043$). Examining reliability across the mean performance in each version, control accuracy was significantly above chance ($t[3] = 9.24$, $p = .003$) and significantly better than in Experiment 1 ($t[3] = 3.63$, $p = .036$). Test accuracy was also above chance across participants within each version (shape: $t[4] = 9.49$, $p = .0007$; syllable: $t[4] = 3.59$, $p = .023$; scene: $t[4] = 4.47$, $p = .011$; tone: $t[4] = 3.34$, $p = .029$). Indeed, there was no effect of version on performance ($F[3, 16] = 0.89$, $p = .469$).

LSJ was again unable to learn the regularities (Figure 4, Table 1). Across versions, her accuracy did not differ from chance ($t[3] = 0.577$, $p = .604$). She again performed significantly worse than controls relative to their mean performance across versions ($t[3] = 3.62$, $p = .036$), and when comparing her mean performance to the control distribution collapsed across versions ($t[22] = 3.38$, $p = .003$). There was no difference in LSJ's performance between Experiments 1 and 2 across versions ($t[3] = 1.32$, $p = .278$).

LSJ again did not show any bias to choose the alternative on the left versus right side of the screen or the alternative that appeared first versus second. Across tasks, the average proportion of trials on which she chose the left alternative was exactly 0.500, and the average proportion of trials on which she chose the first alternative was 0.528 (vs. chance: $t[3] = 0.322$, $p = .769$). Across all versions, controls requested 3.40 trial repeats on average (SD across 20 participants = 3.69), and LSJ requested 4.75 trial repeats on average (SD across four tasks = 2.22).

As in Experiment 1, to analyze the debriefing responses, we collapsed across modalities and performed one-way ANOVAs to assess whether there were any differences in test performance across response types. For each of the three debriefing questions, we did not find any relationship between responses and performance level (ps > .193). As in Experiment 1, the participants who did not notice anything about the patterns (seven participants) performed significantly above chance (mean = 0.762, $t[6] = 3.84$, $p = .009$) and no differently from the participants who noticed something (eight participants; mean = 0.771, $t[13] = 0.118$, $p = .908$) or noticed a substantial amount (five participants; mean = 0.867, $t[10] = 1.04$, $p = .325$).

These findings provide converging support for LSJ's deficit in SL by closely replicating Experiment 1. They are perhaps even more striking because the regularities were much simpler (fewer items, fewer groups of items, and more repetitions) and control participants performed better in this experiment.

EXPERIMENT 3

Across two experiments, LSJ was unable to discriminate regularities from foils at test. It remains possible, however, that this deficit reflects general difficulties with aspects of the task that are unrelated to SL per se. For example, LSJ may have been unable to pay attention during the exposure phase or could have had trouble following instructions or indicating her decisions in the test phase. To address these issues, we repeated the exposure phase from Experiment 2 but designed a different test phase that closely matched the task demands of the previous experiments but did not require knowledge of temporal statistics. Specifically, instead of choosing which of two sequences was more familiar, the task required choosing which of two individual items was more familiar based on exposure (where one item was from the exposure phase and the other was novel).

There is reason to believe that LSJ may be able to exhibit familiarity for items despite not being able to do so for regularities. In particular, rats with damage to the hippocampus cannot remember which of two odors came first in a sequence, but they can remember whether an odor was present in the sequence or not (Fortin, Agster, & Eichenbaum, 2002). Additionally, patients with selective hippocampal damage can have relatively intact item familiarity (Mayes, Holdstock, Isaac, Hunkin, & Roberts, 2002). Such familiarity may be supported by perirhinal cortex (Bowles et al., 2007; Eichenbaum, Yonelinas, & Ranganath, 2007; Brown & Aggleton, 2001) and parahippocampal cortex (Martin, McLean, O'Neil, & Kohler, 2013)—both of which are partially preserved in LSJ's right hemisphere. Moreover, visual cortex exhibits attenuated responses to repeated versus novel stimuli, which may also support item familiarity (Turk-Browne, Yi, & Chun, 2006; Henson, Shallice, Gorno-Tempini, & Dolan, 2002).

Nevertheless, our goal in this experiment was not to establish the neural basis of item familiarity nor to show that LSJ's item familiarity is intact relative to controls. Rather, we sought to demonstrate simply that LSJ can succeed at a task highly similar to those in Experiments 1 and 2. She may succeed for a variety of reasons, including reliance on different brain structures or reduced task difficulty. Regardless, above-chance performance would allow us to establish that LSJ can pay attention during exposure, follow instructions at test, and make reliable familiarity-based judgments.

To match previous experiments, we tested LSJ on item familiarity using all four stimulus sets. However, the tone test has a different character than the other three tests: To discriminate individual tones requires absolute or "perfect" pitch (sequences can benefit from relative pitch). Without such ability, her test performance cannot be easily interpreted. We therefore tested whether LSJ has absolute pitch by playing 24 computer-generated piano notes and asking her to provide the letter name corresponding to the note. She was no different than chance ($p = .646$, binomial test), correctly identifying only 3 of 24 notes. We therefore did not expect LSJ to be successful at the tone item test but collected the data for completeness.

Methods

Participants

LSJ participated in exchange for monetary compensation.

Stimuli and Procedure

The stimuli and exposure phases were identical to Experiment 2. The test phase was similar to Experiments 1 and 2, but two individual items were presented sequentially on each trial, instead of a triplet/pair and a foil. One of the items (target) had been presented during exposure, and the other item (foil) had not been presented during exposure. Foils were the six items from the Experiment 1 stimulus set that did not appear in the Experiment 3 exposure phase. All six targets and foils were tested three times each, with every target paired randomly with three different foils, for a total of 18 trials. The instructions were identical to the previous experiments, except that the judgment was now about an "image," "syllable," or "tone," instead of a "subsequence."

For LSJ, the tasks were conducted in the following order: shapes, scenes, syllables, and tones. They were administered all on one day. Two months elapsed between the second test session for Experiment 2 and this test session.

Results and Discussion

In contrast to her failure on the eight tests of regularities in Experiments 1 and 2, LSJ showed successful discrimination for items in the syllable and scene tests (Table 1). Her accuracy was significantly above chance in these versions by binomial test (syllable: $p = .031$, scene: $p = .0001$). Despite being numerically above chance in the shape version, this level of accuracy did not reach significance ($p = .481$). As anticipated, the tone test was unsuccessful ($p = 1.0$).

LSJ again did not show any bias to choose the alternative on the left versus right side of the screen or the alternative that appeared first versus second. Across tasks, the average proportion of trials on which she chose the left alternative was 0.514 (vs. chance: $t[3] = 0.190$, $p = .861$), and the average proportion of trials on which she chose the first alternative was 0.542 (vs. chance: $t[3] = 1.19$, $p = .319$). Across all versions, LSJ requested 4.75 trial repeats on average (SD across four tasks = 0.957).

Overall, her robust familiarity with the syllable and scene items, and an indication of some familiarity with the shapes, suggests that LSJ understood the structure of the test, was capable of indicating her decision, and had been paying attention during the exposure phases. These findings indicate that her lack of SL in Experiments 1 and 2 reflects an inability to extract and/or express familiarity for regularities rather than a more basic inability to follow the protocol.

GENERAL DISCUSSION

Taken together, our findings support the hypothesis that the MTL is necessary for rapid SL. Across three experiments, amnesic patient LSJ was passively exposed to a continuous sequence of shapes, syllables, scenes, or tones containing temporal regularities. Regularities in Experiment 1 consisted of four triplets, which LSJ was unable to discriminate from foils regardless of stimulus type. Although above chance, control performance was weak, suggesting that we may not have had sufficient sensitivity to detect learning in LSJ. We thus simplified the regularities to three pairs in Experiment 2, which boosted control performance but left LSJ at chance. This lack of learning did not seem to reflect general difficulties with the task, as LSJ successfully discriminated individual items in Experiment 3. Her item familiarity was near-perfect for scenes, providing a stark contrast between her knowledge of items and regularities for at least that stimulus class. This was also true to some extent for syllables, although not for shapes (nor for tones, which would have required perfect pitch). The collection of findings suggests that LSJ has a severe deficit in SL.

Under some accounts of hippocampal and MTL function, these findings are quite surprising. One reason is that the hippocampus is known to play an important role in conscious or declarative forms of memory (Squire et al., 2004). In contrast, SL is incidental and often implicit (Kim et al., 2009; Turk-Browne et al., 2005). Indeed, in the present experiments, roughly half of the participants did not report noticing any regularities and yet still exhibited robust SL. Another way in which these findings are surprising is that the hippocampus has the unique capability to encode individual experiences (McClelland et al., 1995), but SL requires the integration of multiple experiences. That is, a pair or triplet is impossible to segment from the continuous sequence of stimuli in which it is embedded after only one exposure—it exists only by way of the constituent items appearing together in various contexts.

There are other perspectives on hippocampal and MTL function that accommodate our findings more naturally. These accounts focus on the information processing capabilities of the hippocampus rather than awareness (Shohamy & Turk-Browne, 2013; Henke, 2010)—in particular, its ability to rapidly learn relational, configural, and contextual information (Ranganath, 2010; Rudy & Sutherland, 1995; Cohen & Eichenbaum, 1993). Several studies of patients with hippocampal damage support these perspectives. For example, with damage that includes the hippocampus, contextual cuing is impaired (Chun & Phelps, 1999) and eye movements no longer track associative memory or relational violations (Hannula, Ryan, Tranel, & Cohen, 2007; Ryan, Althoff, Whitlow, & Cohen, 2000). Our findings extend these prior studies by showing that the hippocampus is needed to extract regularities from a continuous sequence of stimuli, a key feature of learning in the real world.

Patient studies in artificial grammar learning (AGL)—a paradigm involving exposure to strings with certain letter transition probabilities and subsequent string "grammaticality" judgments—seem, on the surface, to contradict this narrative. Specifically, AGL can be intact in patients with MTL damage (Knowlton, Ramus, & Squire, 1992). Moreover, although AGL sometimes engages the MTL, it also engages the striatum (Lieberman, Chang, Chiao, Bookheimer, & Knowlton, 2004), suggesting that the striatum may be sufficient for learning when the MTL is damaged. However, the MTL may nevertheless contribute to AGL in normal individuals. These different brain systems seem to support different aspects of AGL, as revealed by their connection to different strategies for expressing knowledge of string grammaticality: Hippocampal activity is associated with comparing the surface similarity of groups of items at test to items during exposure (Opitz & Friederici, 2004), whereas the striatum is associated with grammatical rule adherence independent of similarity (Lieberman et al., 2004). Indeed, in a version of the AGL task in which both patients with MTL damage and controls always rely on a similarity-based strategy, patients do show a deficit in AGL (Channon et al., 2002). The SL paradigm, as used in the current experiments, also tests surface similarity between groups of items at test and exposure, which helps reconcile our findings with the AGL literature.

The serial reaction time task (SRTT), in which participants are cued to make button presses that have repeating temporal structure, represents another type of temporal learning that seems to be preserved after MTL damage (Reber & Squire, 1994). This task, like some forms of AGL, is thought to depend on the basal ganglia (Berns, Cohen, & Mintun, 1997; Rauch et al., 1997; Jackson, Jackson, Harrison, Henderson, & Kennard, 1995; Knopman & Nissen, 1991). Although the MTL does not appear to be required for the basic SRTT, it is nevertheless engaged during learning (Schendan et al., 2003), and versions of the task with higher-order temporal relationships are somewhat impaired after MTL damage (Curran, 1997). This suggests that temporal procedural learning, like more passive observational learning, may recruit the MTL for relatively complex relationships.

Our results also speak to a debate within the SL literature regarding whether SL is supported by modality-specific brain systems (Conway & Christiansen, 2005) or a domain-general system (Kirkham et al., 2002). LSJ has damage to a relatively focal area of her brain yet displayed deficits in two forms of visual and two forms of auditory SL. This finding strongly suggests that learning depends at least partially on a shared system (the MTL). Input to and output from the MTL is processed in modality-specific areas, however, raising the possibility that constraints there will modulate the expression of learning. In our experiments, we did not find any differences between modalities, although differences have been reported elsewhere (Conway & Christiansen, 2005, 2009). If exposed to regularities extensively (e.g., for several months), it is possible that LSJ would be able to learn them via modality-specific, local perceptual learning circuitry. However, the rapid SL of interest here seems to occur in a domain-general system.

It is important to note that LSJ's damage spanned hippocampus and MTL cortex bilaterally as well as the left anterior lateral temporal cortex. We thus cannot dissociate the role of MTL cortex, the hippocampus proper, and the anterior temporal cortex in producing the observed deficits. Previous work has implicated all of these regions in associative learning (Schapiro, Rogers, Cordova, Turk-Browne, & Botvinick, 2013; Schapiro et al., 2012; Li & DiCarlo, 2008; Miyashita, 1988), but there is evidence that MTL cortex may be necessary in order for the learning to occur in anterior temporal cortex (Higuchi & Miyashita, 1996). Future work, including with patients with more selective hippocampal damage, will be crucial to better understand the contributions of these different regions.

A related issue is that MTL damage likely has downstream consequences for other brain regions to which it projects. It is possible that these other regions are more important for SL than the MTL and that SL does not occur in LSJ because the functions of these regions are disrupted by a lack of MTL input. Even if true, the fact that MTL damage disrupted SL means that it is causally involved in the learning process. Moreover, prior whole-brain fMRI studies have reported selective engagement of the MTL, and hippocampus in particular, during SL (e.g., Turk-Browne et al., 2010; Harrison et al., 2006), suggesting that it does play a primary role.

Another limitation of the current study is that some individual control participants performed near chance. This raises the question of whether LSJ is truly impaired relative to controls: Perhaps some normal individuals are generally bad at SL, and LSJ simply falls into this category by coincidence. That is, the variation in control performance may reflect individual differences (Zimmerer, Cowell, & Varley, 2011; Visser, Raijmakers, & Pothos, 2009), as opposed to a learning process that is equally robust across individuals in the normal population but corrupted by various sources of noise. There are two ways to distinguish between these possibilities. First, the same control participants could be repeatedly tested to see if they consistently perform at the same level, consistent with stable individual differences. However, this is not possible in the case of SL, as repeated testing would make participants aware of the test phase and structure of the regularities, altering the nature of the process under investigation. Second, even with only one test, other measures could be identified that reliably predict SL performance, again consistent with stable individual differences. We explored whether debriefing responses could provide such a measure, but they did not relate to performance. Without such evidence, we have made the tentative assumption that control variation reflects samples drawn from a noisy unimodal distribution, and in this sense, LSJ falls outside the

distribution. Nevertheless, future studies with more patients and better individual difference measures are clearly needed.

Overall, the results of this study add to the literature demonstrating that the hippocampus can rapidly and incidentally form associations between temporally or spatially related objects in the environment. The unique contribution of our work is to show that this capability extends to a more naturalistic situation in which the associations of interest need to be extracted from a continuous sequence of information. In other words, the hippocampus can detect regularities even when there is no spatial or temporal segmentation—no trial structure—to define the associations. This study thus provides strong converging evidence with neuroimaging studies (Schapiro et al., 2012; Turk-Browne et al., 2009) that the MTL is important for SL.

REFERENCES

Baldwin, D., Andersson, A., Saffran, J., & Meyer, M. (2008). Segmenting dynamic human action via statistical structure. *Cognition, 106,* 1382–1407.

Berns, G. S., Cohen, J. D., & Mintun, M. A. (1997). Brain regions responsive to novelty in the absence of awareness. *Science, 276,* 1272–1275.

Bornstein, A. M., & Daw, N. D. (2012). Dissociating hippocampal and striatal contributions to sequential prediction learning. *European Journal of Neuroscience, 35,* 1011–1023.

Bowles, B., Crupi, C., Mirsattari, S. M., Pigott, S. E., Parrent, A. G., Pruessner, J. C., et al. (2007). Impaired familiarity with preserved recollection after anterior temporal-lobe resection that spares the hippocampus. *Proceedings of the National Academy of Sciences, U.S.A., 104,* 16382–16387.

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision, 10,* 433–436.

Brown, M. W., & Aggleton, J. P. (2001). Recognition memory: What are the roles of the perirhinal cortex and hippocampus? *Nature Reviews Neuroscience, 2,* 51–61.

Campbell, K. L., Zimerman, S., Healey, M. K., Lee, M. M., & Hasher, L. (2012). Age differences in visual statistical learning. *Psychology and Aging, 27,* 650–656.

Carr, V. A., Rissman, J., & Wagner, A. D. (2010). Imaging the human medial temporal lobe with high-resolution fMRI. *Neuron, 65,* 298–308.

Channon, S., Shanks, D., Johnstone, T., Vakili, K., Chin, J., & Sinclair, E. (2002). Is implicit learning spared in amnesia? Rule abstraction and item familiarity in artificial grammar learning. *Neuropsychologia, 40,* 2185–2197.

Chun, M. M., & Phelps, E. A. (1999). Memory deficits for implicit contextual information in amnesic subjects with hippocampal damage. *Nature Neuroscience, 2,* 844–847.

Cohen, N. J., & Eichenbaum, H. (1993). *Memory, amnesia, and the hippocampal system*. Cambridge, MA: MIT Press.

Conway, C. M., & Christiansen, M. H. (2005). Modality-constrained statistical learning of tactile, visual, and auditory sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31,* 24–39.

Conway, C. M., & Christiansen, M. H. (2009). Seeing and hearing in space and time: Effects of modality and presentation rate on implicit statistical learning. *European Journal of Cognitive Psychology, 21,* 561–580.

Curran, T. (1997). Higher-order associative learning in amnesia: Evidence from the serial reaction time task. *Journal of Cognitive Neuroscience, 9,* 522–533.

Duchaine, B., Germine, L., & Nakayama, K. (2007). Family resemblance: Ten family members with prosopagnosia and within-class object agnosia. *Cognitive Neuropsychology, 24,* 419–430.

Duchaine, B., Nieminen-von Wendt, T., New, J., & Kulomaki, T. (2003). Dissociations of visual recognition in a developmental agnosic: Evidence for separate developmental processes. *Neurocase, 9,* 380–389.

Dunn, L. M., & Dunn, L. M. (1981). *Peabody Picture Vocabulary Test-Revised*. Circle Pines, MN: American Guidance Service.

Durrant, S. J., Cairney, S. A., & Lewis, P. A. (2012). Overnight consolidation aids the transfer of statistical knowledge from the medial temporal lobe to the striatum. *Cerebral Cortex, 23,* 2467–2478.

Dusek, J. A., & Eichenbaum, H. (1997). The hippocampus and memory for orderly stimulus relations. *Proceedings of the National Academy of Sciences, U.S.A., 94,* 7109–7114.

Eichenbaum, H., Yonelinas, A. P., & Ranganath, C. (2007). The medial temporal lobe and recognition memory. *Annual Review of Neuroscience, 30,* 123–152.

Fiser, J., & Aslin, R. N. (2001). Unsupervised statistical learning of higher-order spatial structures from visual scenes. *Psychological Science, 12,* 499–504.

Fiser, J., & Aslin, R. N. (2002a). Statistical learning of higher-order temporal structure from visual shape sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 28,* 458–467.

Fiser, J., & Aslin, R. N. (2002b). Statistical learning of new visual feature combinations by infants. *Proceedings of the National Academy of Sciences, U.S.A., 99,* 15822–15826.

Fortin, N. J., Agster, K. L., & Eichenbaum, H. B. (2002). Critical role of the hippocampus in memory for sequences of events. *Nature Neuroscience, 5,* 458–462.

Gabrieli, J. D. E., Fleischman, D. A., Keane, M. M., Reminger, S. L., & Morrell, F. (1995). Double dissociation between memory-systems underlying explicit and implicit memory in the human brain. *Psychological Science, 6,* 76–82.

Gregory, E., McCloskey, M., & Landau, B. (2014). Profound loss of general world knowledge in retrograde amnesia: Evidence from an amnestic artist. *Frontiers in Human Neuroscience, 8,* 287.

Hannula, D. E., Ryan, J. D., Tranel, D., & Cohen, N. J. (2007). Rapid onset relational memory effects are evident in eye movement behavior, but not in hippocampal amnesia. *Journal of Cognitive Neuroscience, 19,* 1690–1705.

Harrison, L. M., Duggins, A., & Friston, K. J. (2006). Encoding uncertainty in the hippocampus. *Neural Networks, 19,* 535–546.

Henke, K. (2010). A model for memory systems based on processing modes rather than consciousness. *Nature Reviews Neuroscience, 11,* 523–532.

Henson, R. N., Shallice, T., Gorno-Tempini, M. L., & Dolan, R. J. (2002). Face repetition effects in implicit and explicit memory tests as measured by fMRI. *Cerebral Cortex, 12,* 178–186.

Higuchi, S., & Miyashita, Y. (1996). Formation of mnemonic neuronal responses to visual paired associates in inferotemporal cortex is impaired by perirhinal and entorhinal lesions. *Proceedings of the National Academy of Sciences, U.S.A., 93,* 739–743.

Jackson, G. M., Jackson, S. R., Harrison, J., Henderson, L., & Kennard, C. (1995). Serial reaction time learning and Parkinson's disease: Evidence for a procedural learning deficit. *Neuropsychologia, 33,* 577–593.

Kapur, N., Barker, S., Burrows, E. H., Ellison, D., Brice, J., Illis, L. S., et al. (1994). Herpes simplex encephalitis: Long term magnetic resonance imaging and neuropsychological profile. *Journal of Neurology, Neurosurgery and Psychiatry, 57,* 1334–1342.

Kim, R., Seitz, A., Feenstra, H., & Shams, L. (2009). Testing assumptions of statistical learning: Is it long-term and implicit? *Neuroscience Letters, 461,* 145–149.

Kirkham, N. Z., Slemmer, J. A., & Johnson, S. P. (2002). Visual statistical learning in infancy: Evidence for a domain general learning mechanism. *Cognition, 83,* B35–B42.

Knopman, D., & Nissen, M. J. (1991). Procedural learning is impaired in Huntington's disease: Evidence from the serial reaction time task. *Neuropsychologia, 29,* 245–254.

Knowlton, B. J., Mangels, J. A., & Squire, L. R. (1996). A neostriatal habit learning system in humans. *Science, 273,* 1399–1402.

Knowlton, B. J., Ramus, S. J., & Squire, L. R. (1992). Intact artificial grammar learning in amnesia—Dissociation of classification learning and explicit memory for specific instances. *Psychological Science, 3,* 172–179.

Li, N., & DiCarlo, J. J. (2008). Unsupervised natural experience rapidly alters invariant object representation in visual cortex. *Science, 321,* 1502–1507.

Lieberman, M. D., Chang, G. Y., Chiao, J., Bookheimer, S. Y., & Knowlton, B. J. (2004). An event-related fMRI study of artificial grammar learning in a balanced chunk strength design. *Journal of Cognitive Neuroscience, 16,* 427–438.

Martin, C. B., McLean, D. A., O'Neil, E. B., & Kohler, S. (2013). Distinct familiarity-based response patterns for faces and buildings in perirhinal and parahippocampal cortex. *Journal of Neuroscience, 33,* 10915–10923.

Mayes, A. R., Holdstock, J. S., Isaac, C. L., Hunkin, N. M., & Roberts, N. (2002). Relative sparing of item recognition memory in a patient with adult-onset damage limited to the hippocampus. *Hippocampus, 12,* 325–340.

McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review, 102,* 419–457.

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

Opitz, B., & Friederici, A. D. (2004). Brain correlates of language learning: The neuronal dissociation of rule-based versus similarity-based learning. *Journal of Neuroscience, 24,* 8436–8440.

Peretz, I., Champod, A. S., & Hyde, K. (2003). Varieties of musical disorders. The Montreal Battery of Evaluation of Amusia. *Annals of the New York Academy of Sciences, 999,* 58–75.

Ranganath, C. (2010). A unified framework for the functional organization of the medial temporal lobes and the phenomenology of episodic memory. *Hippocampus, 20,* 1263–1290.

Rauch, S. L., Whalen, P. J., Savage, C. R., Curran, T., Kendrick, A., Brown, H. D., et al. (1997). Striatal recruitment during an implicit sequence learning task as measured by functional magnetic resonance imaging. *Human Brain Mapping, 5,* 124–132.

Reber, P. J., & Squire, L. R. (1994). Parallel brain systems for learning with and without awareness. *Learning and Memory, 1,* 217–229.

Rose, M., Haider, H., & Buchel, C. (2005). Unconscious detection of implicit expectancies. *Journal of Cognitive Neuroscience, 17,* 918–927.

Rudy, J. W., & Sutherland, R. J. (1995). Configural association theory and the hippocampal formation: An appraisal and reconfiguration. *Hippocampus, 5,* 375–389.

Ryan, J. D., Althoff, R. R., Whitlow, S., & Cohen, N. J. (2000). Amnesia is a deficit in relational memory. *Psychological Science, 11,* 454–461.

Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science, 274,* 1926–1928.

Saffran, J. R., Johnson, E. K., Aslin, R. N., & Newport, E. L. (1999). Statistical learning of tone sequences by human infants and adults. *Cognition, 70,* 27–52.

Schapiro, A. C., Kustner, L. V., & Turk-Browne, N. B. (2012). Shaping of object representations in the human medial temporal lobe based on temporal regularities. *Current Biology, 22,* 1622–1627.

Schapiro, A. C., Rogers, T. T., Cordova, N. I., Turk-Browne, N. B., & Botvinick, M. M. (2013). Neural representations of events arise from temporal community structure. *Nature Neuroscience, 16,* 486–492.

Schendan, H. E., Searl, M. M., Melrose, R. J., & Stern, C. E. (2003). An fMRI study of the role of the medial temporal lobe in implicit and explicit sequence learning. *Neuron, 37,* 1013–1025.

Shohamy, D., & Turk-Browne, N. B. (2013). Mechanisms for widespread hippocampal involvement in cognition. *Journal of Experimental Psychology: General, 142,* 1159–1170.

Simoncelli, E. P., & Olshausen, B. A. (2001). Natural image statistics and neural representation. *Annual Review of Neuroscience, 24,* 1193–1216.

Squire, L. R., Stark, C. E., & Clark, R. E. (2004). The medial temporal lobe. *Annual Review of Neuroscience, 27,* 279–306.

Turk-Browne, N. B., Junge, J., & Scholl, B. J. (2005). The automaticity of visual statistical learning. *Journal of Experimental Psychology: General, 134,* 552–564.

Turk-Browne, N. B., Scholl, B. J., Chun, M. M., & Johnson, M. K. (2009). Neural evidence of statistical learning: Efficient detection of visual regularities without awareness. *Journal of Cognitive Neuroscience, 21,* 1934–1945.

Turk-Browne, N. B., Scholl, B. J., Johnson, M. K., & Chun, M. M. (2010). Implicit perceptual anticipation triggered by statistical learning. *Journal of Neuroscience, 30,* 11177–11187.

Turk-Browne, N. B., Yi, D. J., & Chun, M. M. (2006). Linking implicit and explicit memory: Common encoding factors and shared representations. *Neuron, 49,* 917–927.

Visser, I., Raijmakers, M. E., & Pothos, E. M. (2009). Individual strategies in artificial grammar learning. *American Journal of Psychology, 122,* 293–307.

Voss, J. L., & Paller, K. A. (2010). What makes recognition without awareness appear to be elusive? Strategic factors that influence the accuracy of guesses. *Learning and Memory, 17,* 460–468.

Warrington, E. K., & James, M. (1991). *The Visual Object and Space Perception Battery*. Bury St. Edmunds, UK: Thames Valley Test Company.

Wechsler, D. (1997). *Wechsler Memory Scale* (3rd ed.). San Antonio, TX: Psychological Corporation.

Wechsler, D. (2008). *Wechsler Adult Intelligence Scale* (4th ed.). San Antonio, TX: Pearson Assessment.

Wilkinson, G. S. (1993). *WRAT3 Wide Range Achievement Test: Administration manual*. Wilmington, DE: Wide Range, Inc.

Zimmerer, V. C., Cowell, P. E., & Varley, R. A. (2011). Individual behavior in learning of an artificial grammar. *Memory & Cognition, 39,* 491–501.