



Coding of episodic memory in the human hippocampus

John T. Wixted^{a,1}, Stephen D. Goldinger^b, Larry R. Squire^{a,c,d,e,1}, Joel R. Kuhn^a, Megan H. Pappesh^f, Kris A. Smith^g, David M. Treiman^h, and Peter N. Steinmetz^{g,h,2}

^aDepartment of Psychology, University of California, San Diego, La Jolla, CA 92093; ^bDepartment of Psychology, Arizona State University, Tempe, AZ 85287; ^cDepartment of Psychiatry, University of California, San Diego, La Jolla, CA 92093; ^dDepartment of Neurosciences, University of California, San Diego, La Jolla, CA 92093; ^eVeterans Affairs Medical Center, San Diego, CA 92161; ^fDepartment of Psychology, Louisiana State University, Baton Rouge, LA 70803; ^gDepartment of Neurosurgery, Barrow Neurological Institute, Phoenix, AZ 85013; and ^hDepartment of Neurology, Barrow Neurological Institute, Phoenix, AZ 85013

Contributed by Larry R. Squire, December 18, 2017 (sent for review September 19, 2017; reviewed by Elizabeth A. Buffalo and Earl K. Miller)

Neurocomputational models have long posited that episodic memories in the human hippocampus are represented by sparse, stimulus-specific neural codes. A concomitant proposal is that when sparse-distributed neural assemblies become active, they suppress the activity of competing neurons (neural sharpening). We investigated episodic memory coding in the hippocampus and amygdala by measuring single-neuron responses from 20 epilepsy patients (12 female) undergoing intracranial monitoring while they completed a continuous recognition memory task. In the left hippocampus, the distribution of single-neuron activity indicated that only a small fraction of neurons exhibited strong responding to a given repeated word and that each repeated word elicited strong responding in a different small fraction of neurons. This finding reflects sparse distributed coding. The remaining large fraction of neurons exhibited a concurrent reduction in firing rates relative to novel words. The observed pattern accords with longstanding predictions that have previously received scant support from single-cell recordings from human hippocampus.

hippocampus | episodic memory | single units | amygdala

Episodic memory affords the capacity to recollect past events that occurred at a particular time and place (1). In humans, episodic recollection allows for the reexperiencing of an event through a process of mental time travel (2). The ability to encode new episodic memories depends on the hippocampus, but it is not clear how episodic memories are coded by the activity of individual hippocampal neurons. We investigated the activity of isolated hippocampal neurons in epileptic patients undergoing intracranial monitoring while they encoded and retrieved episodic memories. Memory was tested using a continuous recognition procedure (3) in which words were presented in a continuous stream and were sometimes repeated. Throughout the task, patients were asked to classify each word as “new” upon its first presentation and as “old” if it was repeated. A correct old decision in response to a repeated word is an instance of successful episodic memory (i.e., memory for the prior occurrence of the word in the experimental context). Hippocampal lesions impair performance on continuous recognition tasks for words (4).

Neurocomputational models (5–8) have long posited that coding in the hippocampus is sparse and distributed. Thus, individual episodic memories are represented by the activity of small and typically nonoverlapping sets of neurons. Under such a coding scheme, activity associated with the retrieval of a specific episodic memory would be hard to detect because only a small proportion of hippocampal neurons would exhibit increased firing rates. Perhaps for this reason, single-unit studies of recognition memory in humans and nonhuman primates have often failed to detect any activity related to episodic memory in the hippocampus (9–12). Moreover, when activity related to episodic memory has been detected, the identified neurons responded nonspecifically, coding whether stimuli were novel or familiar (13–19), and leaving open the question of whether neurons can be found that sparsely code some recently studied items and not others.

A standard procedure for detecting stimulus-specific, single-unit activity involves repeatedly presenting a stimulus to determine if a neuron responds reliably only when that stimulus is presented. Notably, this approach has identified neurons that respond selectively to the presentation of a photo of a particular person or landmark (20, 21). In studies of episodic memory, one would expect to find not only neurons that code stable semantic knowledge about the material being learned, but also neurons that code aspects of the learning event itself. By definition, episodic memory involves retrieving an episode that occurred only once, i.e., a learning event such as remembering the earlier presentation of a word. Note that repeating a studied word not only prompts retrieval of its prior occurrence but also creates a new and distinct episodic memory, potentially coded by a different set of neurons. In that case, a neuron that responded the first time a word was repeated might not respond to its repetition. Moreover, if a given neuron did respond to every repetition of a word, the neuron might be responding to the word’s context-free semantic meaning, not to the word’s episodic occurrence in the experimental context. For these reasons, instead of searching for neurons that respond reliably to words repeated multiple times, we used an approach that is capable of detecting rare spiking events that theoretically signal episodic memory in response to words that were repeated only once.

Significance

Neurocomputational models hold that episodic memories are represented by sparse, stimulus-specific neural codes. In tests of episodic memory, single-unit recording studies of the human hippocampus have found neurons that operate as general novelty detectors or general familiarity detectors. Here, we investigated whether neurons can be found that sparsely code some recently studied items and not others. In the left hippocampus, but not the amygdala, we found that small fractions of neurons exhibited strong responses to specific repeated words. The remaining large fractions of neurons exhibited a concurrent reduction in firing rates relative to novel words. Both findings are consistent with predictions made by neurocomputational models of how episodic memory is coded in the hippocampus.

Author contributions: S.D.G., M.H.P., and P.N.S. designed research; S.D.G., M.H.P., K.A.S., D.M.T., and P.N.S. performed research; J.T.W., J.R.K., and P.N.S. analyzed data; and J.T.W., S.D.G., L.R.S., M.H.P., and P.N.S. wrote the paper.

Reviewers: E.A.B., University of Washington; and E.K.M., Massachusetts Institute of Technology.

The authors declare no conflict of interest.

Published under the [PNAS license](#).

¹To whom correspondence may be addressed. Email: jwixted@ucsd.edu or lsquire@ucsd.edu.

²Present address: NeurTex Brain Research Institute, Dallas, TX 75206-1910.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1716443115/-DCSupplemental.




Fig. 1. Behavioral forgetting function for the continuous recognition task. Discriminability (d') declined significantly as a function of the number of intervening items (lag) according to a repeated-measures ANOVA ($P < 0.001$). Each patient's d' score was first computed by averaging across recognition test sessions. Each point in the figure represents the average across all 20 patients. The smooth curve represents the least-squares fit of a power function, $d' = a \times \text{Lag}^{-b}$, where a and b are free parameters. Error bars represent SEs.

Results

Behavioral Performance. Invalid trials (13.6% of all trials) were excluded from analysis. These trials had either no responses, early responses, or multiple keys pressed. From the remaining valid trials, hit rates (proportions of correctly classified repeated words) and false alarm rates (proportions of novel words mistakenly classified as repeated) were computed for each session. These hit and false alarm rates were used to compute a standard discriminability measure (d'). Performance measures for the 37 sessions were computed separately for each patient (averaged across sessions) and were then averaged across patients. The average false alarm rate was 0.21. Unlike the false alarm rate, hit rates could be computed separately by lag and showed a monotonic decline as lag increased (hit rates = 0.88, 0.80, 0.76, 0.75, 0.71, and 0.70 for the six lags, respectively). The corresponding d' scores exhibited the typical power law of forgetting (Fig. 1).

Analysis of Single-Unit Activity in the Hippocampus and Amygdala.

Across all patients and all 37 sessions, we recorded 275 single units in the amygdala (161 left, 114 right) and 243 single units in the hippocampus (128 left, 115 right). The average background firing rates for these units were 2.20 and 1.60 spikes/s in the left and right hippocampus, respectively, and 1.30 and 1.04 spikes per second in the left and right amygdala, respectively.

In all four regions, some neurons exhibited spiking activity that significantly differed, on average, for repeated vs. novel items ("Significant units" in Table 1), but only in the left amygdala were significant units detected with a frequency (27 of 161) that exceeded chance expectations ($P < 0.0001$). This effect was largely attributable to increased firing rates to novel words. Of the 27 significant units in the left amygdala, 25 showed a novelty-detection pattern, whereas two showed the opposite pattern. Due to chance alone, under the null hypothesis, one would expect to find $\sim 0.05 \times 161 \sim 8$ significant units in the left amygdala, with equivalent counts of "novelty detectors" and "familiarity detectors." Thus, observing two familiarity detectors likely reflects chance alone, but this is unlikely to be the case for the 25 novelty detectors. Among the novelty detectors, the average normalized firing rate to novel items was 0.54 σ units above

baseline, whereas the average normalized firing rate to repeated items was only 0.14 σ units above baseline.

Analysis of Spike Count Distributions from the Hippocampus. In the hippocampus, units distinguishing repeated vs. novel items were not detected at a significant frequency. Yet, if a given neuron in the hippocampus strongly responds on only a handful of repeated trials (e.g., $<5\%$), as in a sparse distributed coding scheme, a significant difference in the overall average firing rate for novel vs. repeated items is unlikely to be detected. To detect such activity in the hippocampus, if it exists, one should instead examine the full distributions of normalized spike counts (pooled across single units recorded from all patients) from trials involving novel items and, separately, from trials involving repeated items. In the right hippocampus, no significant differences were observed in either the means (Fig. 2A) or the SDs (Fig. 2B) of the full distributions for repeated and novel items. In the left hippocampus, the means of these two distributions also did not differ significantly (Fig. 2A), but a reliable difference was observed in their SDs (Fig. 2B).

Two distributions that have similar means and different SDs can differ in more than one way (Fig. 3). To investigate the source of the SD difference between the distributions in the left hippocampus, we constructed empirical quantile-quantile (QQ) plots (22). An empirical QQ plot is a graphical method of analysis that essentially displays one rank-ordered dataset (i.e., the sorted normalized spike counts for the repeated items) against another independently rank-ordered dataset (i.e., the sorted normalized spike counts for the novel items). We recently used this approach in a study of episodic memory (23), but because only 34 single units were recorded, the analysis was based primarily on multiunits, and convincing evidence of sparse distributed coding at the level of single units was not demonstrated. The present analysis is based on a much larger sample of 243 single units, and no multiunits were included.

The QQ plot from left hippocampus (Fig. 4A) is consistent with a bimodal distribution for repeated items (as illustrated in Fig. 3B). The pattern is similar for trials in which patients made a correct response (hit or correct rejection) and trials in which they made an error (miss or false alarm, Fig. S1). As predicted by the sparse distributed coding account, the points fall mostly along the diagonal line and then exhibit a sharp upward deflection at the upper-right end of the plot. When broken down by lag, the pattern did not vary in any systematic way as lag increased (Fig. S2). In the right hippocampus, the QQ plot (Fig. 4B) shows no apparent departure from the diagonal line, which is consistent with the finding of similar means and SDs for the repeated- and novel-item distributions in the right hippocampus (Fig. 2, Right H). The data in Fig. 4A reflect an episodic memory signal in that the upward deflection at the upper-right end of the QQ plot indicates that some neurons responded strongly to a few repeated

Table 1. Recorded units in amygdala and hippocampus

Region	Side	Recorded units	Significant units	Fraction	adj P
Amygdala	L	161	27	0.17	<0.001
	R	114	4	0.04	0.827
Hippocampus	L	128	11	0.09	0.114
	R	115	7	0.06	0.469

The number of recorded units and number of significant units (i.e., units for which, using an unadjusted t test, mean spikes for repeated items differed significantly from mean spikes to novel items) from left (L) and right (R) amygdala and hippocampus. Fraction, significant units/recorded units. The P value (adj P) is the probability of observing that fraction by chance alone, after correcting for multiple testing using the Benjamini-Hochberg procedure. The same procedure was used to compute adjusted P values in the subsequent analyses.




Fig. 2. Mean and SD statistics associated with normalized spike counts (collapsed over patients and sessions) for repeated and novel items. Mean (A, left y axis) and SD (B, right y axis) of normalized spike counts associated with the full distributions (100% of the data) for repeated items ($n = 12,854$ spikes) and novel items ($n = 13,822$ spikes) in the left and right hippocampus (H) collapsed over lag. The normalized spike counts are expressed in SD units. In the left hippocampus, repeated words elicited a mean increase in firing that was 0.05 SD units above baseline (similar to novel words). However, the SD of the normalized spike counts was larger for repeated words than novel words (1.144 vs. 1.048). The P values represent the probability of obtaining the observed difference (for repeated vs. novel items) by chance, under the null hypothesis of no difference (adjusted for multiple comparisons). The SD effect tracked item status (repeated vs. novel), not the behavioral decision. More specifically, the SD scores for hits and misses (repeated items) were 1.141 and 1.156, respectively, and the corresponding values for correct rejections and false alarms (novel items) were 1.060 and 0.994, respectively.

words although the same neurons did not respond to those words when they were novel.

The significant difference in SDs for trials involving novel items vs. repeated items in the left hippocampus (Fig. 2B, Left H) is consistent with visual evidence of bimodality in the QQ plot from left hippocampus (Fig. 4A). If a small upper distribution for repeated items is responsible for both the increased SD and the visual signature of bimodality in the QQ plot (as illustrated in Fig. 3B), then removing a small percentage of scores from the upper tails of both the novel-item and repeated-item distributions should eliminate the difference in the SDs as well the visual evidence of bimodality in the QQ plot. In agreement with that prediction, when the highest 2.5% of spike counts were removed from both distributions, the QQ plot for the left hippocampus became essentially linear (Fig. 4C), and the difference in SD between novel and repeated item distributions was eliminated (Fig. 5B, Left H). Note that the pattern observed after removing the highest 2.5% of the scores remains evident when larger proportions of each distribution are removed (Fig. 6). These results indicate that, when 100% of the data are analyzed, the SD difference in the left hippocampus (Fig. 2B, Left H) arose because of strong neural responses that occurred on a small percentage of repeated-word trials (the same trials responsible for the nonlinear QQ plot in Fig. 4A). Our procedure was a verbal memory task, which likely explains why the effects were evident in only left hippocampus.

Removing the upper 2.5% of repeated- and novel-item scores from the aggregate distributions for the left hippocampus not only eliminated the SD difference between the repeated- and novel-item distributions, but also revealed another effect. Specifically, in the left hippocampus, the mean normalized spike count of the remaining 97.5% of repeated-item scores was now significantly reduced, relative to the mean of the remaining 97.5% of novel-item scores (Fig. 5A, Left H). This finding suggests that strongly activated neurons (representing an episodic memory trace) inhibited competing neurons, an effect that has been termed neural sharpening (8). Upon close inspection, this effect in the left hippocampus is visually apparent in the QQ plot (Fig. 4C), in that points on the left end of the plot consistently fall slightly below the diagonal line.

We next examined which patients and which repeated words contributed to the highest 2.5% of normalized spike counts. Of the 20 patients tested, single-unit activity was detected in the left hippocampus in 13 of them. Of those 13 patients, 11 yielded normalized spike counts in response to at least four unique words (mean = 22.3 words) that fell in the top 2.5% of normalized spike counts for repeated items. Thus, the increased SD associated with repeated items in the left hippocampus (Fig. 2B) was not caused by a single patient or a single repeated word but was instead a more general phenomenon.

Analysis of Spike Count Distributions from the Amygdala. As noted earlier, a general novelty signal in the left amygdala was strong enough to be detected at the level of individual single units (Table 1). How does that effect manifest itself in an analysis of the full novel-item vs. repeated-item distributions of single-unit recordings? In the left amygdala (but not in the right amygdala), the overall mean and SD of the full distributions were both significantly greater for novel compared with repeated items

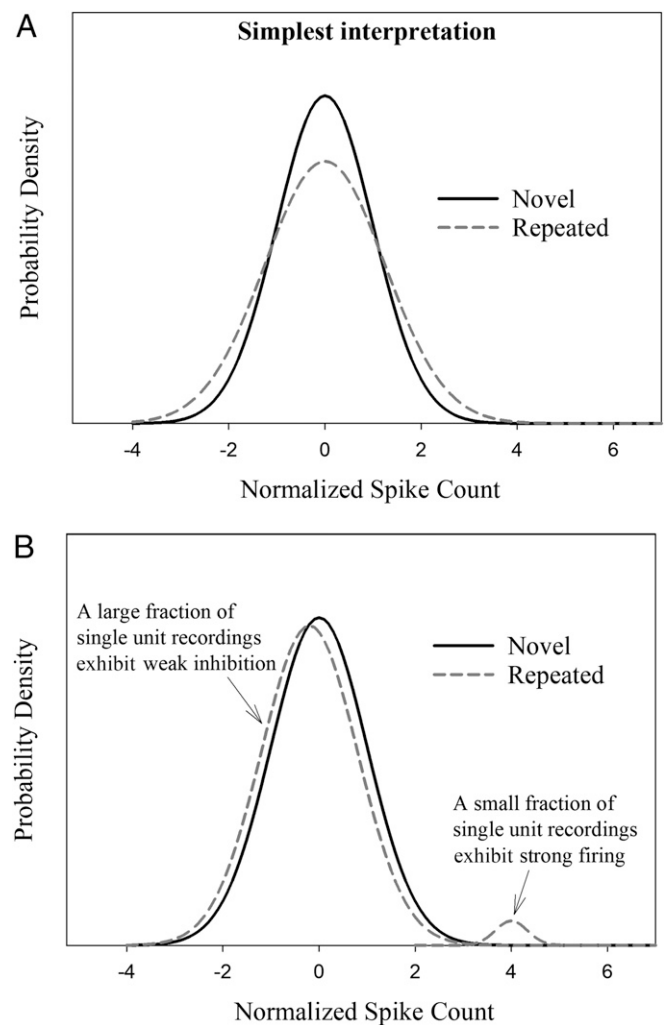


Fig. 3. Hypothetical novel- and repeated-item aggregate distributions of normalized spike counts with the same means but different SDs. (A) Distributions with the same shape but different SDs. (B) Distributions with different shapes and different SDs. As predicted by a sparse distributed coding account, a small percentage of recordings made to repeated items (~2.5%) would yield strong responses and the remainder (~97.5%) would yield weakly inhibited responses. The strong responses would increase the SD of the repeated-item distribution. The data conform to this pattern.

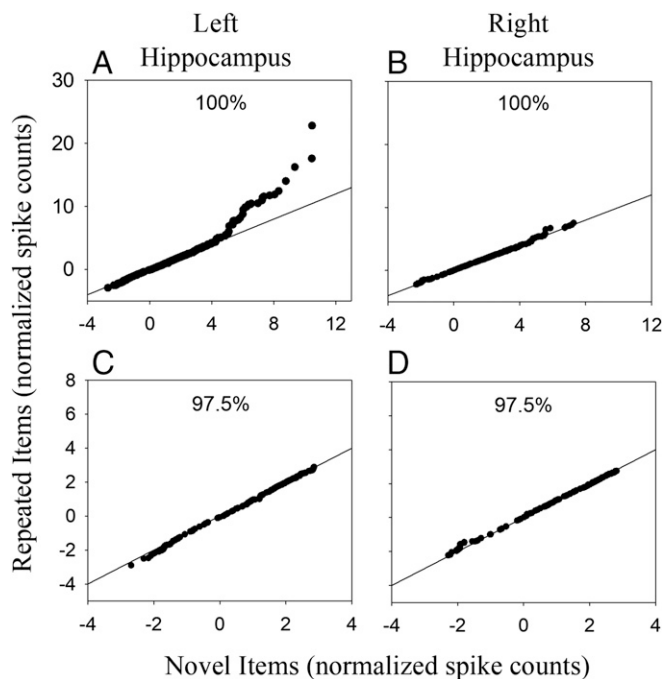


Fig. 4. QQ plots for the left and right hippocampus for 100% of the data (A and B, respectively) and after excluding 2.5% of the data with the highest spike counts from both the repeated-item and novel-item distributions (C and D, respectively). Each point on a QQ plot represents the normalized average spike count recorded on a single test trial. The plot displays those values aggregated across trials and patients. In the left hippocampus, the 100% plot displays 12,854 and 13,822 normalized spike counts for repeated and novel items, respectively. In the right hippocampus, the corresponding values are 11,089 and 11,955 normalized spike counts.

(Table S1). This overall novelty signal in the left amygdala is consistent with the fact that 25 single units were identified in the left amygdala that were general novelty detectors. The QQ plots for recordings made from left and right amygdala (Fig. S3 A and B, respectively) show no evidence of bimodality. When 2.5% of the highest scores were eliminated, the statistical pattern of results was unaffected, unlike in the left hippocampus. More specifically, whether 100% of the data are considered or 97.5% of the data are considered (Table S1), a significant difference is evident for both the mean and SD in the left amygdala. These differences remain significant even when 20% of the highest values are removed from the analysis. Thus, the pattern for both single-unit activity (Table 1), and the means and SDs of spike activity aggregated across single units (Table S1), differs for left amygdala compared with left hippocampus.

Discussion

In studies of semantic memory, single neurons in the medial temporal lobe have been identified that respond reliably to repeated presentations of a known place or landmark, such as the Eiffel Tower (20, 24, 25). Similarly, in studies of episodic memory (16–18), single neurons in the medial temporal lobe have been identified that respond reliably to repeated presentations of items drawn from a general stimulus class (e.g., novel items). However, repeated stimulus presentations cannot be used to identify stimulus-specific, episodic representations. A neuron that codes episodic memory for a previous, context-specific presentation of a particular stimulus should respond selectively to the first repetition of that stimulus in a recognition test (at which time retrieval of the original experience may occur), but it will not necessarily respond to any subsequent presentations

of that same stimulus. When repeated a second time, the stimulus may occasion retrieval of the first repetition (coded by different neurons), not the original experience. We therefore investigated stimulus-specific episodic coding in the human hippocampus, using an analysis performed on once-presented test items.

We found evidence for two complementary episodic memory signals in the human hippocampus, both of which have long been predicted by neurocomputational models (5–8). First, we identified a sparse-distributed memory signal, characterized by strong neural firing in response to repeated items (relative to novel items) for a small fraction of recordings (<2.5%). Typically, a neuron exhibited a strong response to only two or three repeated items but not to any of the other repeated items. Moreover, small fractions of neurons responded to different repeated items. Second, we identified a general suppression of firing rates in response to repeated items for the remaining large fraction of recordings (~97.5%). This neural-sharpening pattern was observed in the hippocampus (where sparse distributed coding of episodic memory is theorized to occur). A similar phenomenon, termed response suppression, has also been described in perirhinal/inferotemporal cortex in monkeys performing a recognition memory task (ref. 26, also see discussion in ref. 27). Suppression of firing rates was not observed in the amygdala. Instead, in the left amygdala, we identified individual neurons that function as general novelty detectors (16–18).

The complementary effects observed in the hippocampus may reflect Hebbian learning coupled with interneuron competition, now a cornerstone of neurocomputational models (8). Empirically, a pattern consisting of a small group of cells with high firing rates coupled with the global suppression of a large group of cells with much lower firing rates has been reported in area CA1 as rats formed memories of a novel maze (28). The pattern we observed in the human hippocampus may reflect similar effects with respect to the encoding and retrieval of episodic memories. That is, small neural assemblies, when active, spread inhibition across many other neurons. This interpretation accords with other findings showing that interneurons impose surprisingly widespread inhibition throughout cell layers (29, 30).

In our study, fewer than 2.5% of single units in the hippocampus were strongly activated when an item was repeated on the continuous recognition test. This 2.5% figure reflects a combination of lifetime sparseness (percentage of stimuli that a given neuron responds to) and population sparseness (percentage of neurons that respond to a given stimulus). These two measures of sparseness are typically assumed to be similar to

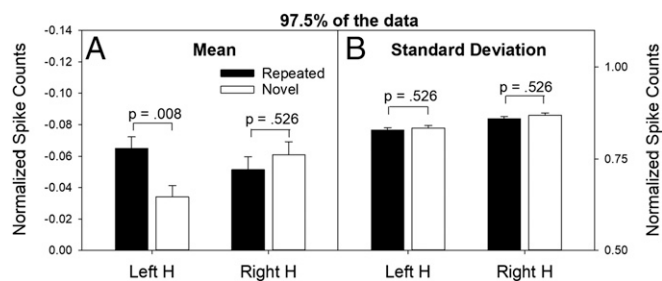


Fig. 5. Mean and SD statistics associated with 97.5% of normalized spike counts (collapsed over patients and sessions) for repeated and novel items. Mean (A, left y axis) and SD (B, right y axis) of normalized spike counts after excluding the highest 2.5% of the scores for each distribution (retaining 97.5% of the data) in the left and right hippocampus. Note in the left hippocampus that the mean firing rates were now significantly different for repeated vs. novel words, and the SDs were similar. Because the y axis in A covers a range of negative values, the mean for novel items is greater (i.e., closer to 0) than the mean for repeated items. The *P* values for these statistical tests were also adjusted for multiple comparisons.




Fig. 6. Statistics for the full distribution of scores represented as the difference in normalized firing rates for repeated vs. novel items in the left hippocampus. The figure shows difference scores for means and SDs as a function of the proportion of the scores from each distribution that was included in the analysis. The mean and SD difference scores for proportions of 1.00 and 0.975 correspond to the data for the left hippocampus shown in Figs. 2A and 5A, respectively. The asterisks indicate that the difference was significant at $P < 0.05$ (not corrected for multiple comparisons).

each other (31), and they have been found to be highly correlated in mouse V1 (32). Assuming the same is true of our data, we estimate that both lifetime sparseness and population sparseness in the human hippocampus are less than 2.5%.

Recent evidence suggests that the absolute number of neurons used to represent an experience is relatively stable between nonhuman primates and rats (33). Because nonhuman primates have a larger hippocampus than rats, the implication is that population sparsity (proportion of active cells) would be smaller in the nonhuman primate compared with the rat. The nonhuman primate population sparsity estimate in that study was ~4% in CA1, CA3, and DG (for rodents, the estimate was ~30%). Because humans have a larger hippocampus than nonhuman primates, our estimate of less than 2.5% in humans is consistent with the idea that the absolute number of neurons used to represent an episodic experience is evolutionarily preserved in humans, nonhuman primates, and rats.

Neurocomputational models predict sparse coding of episodic memory in the hippocampus but not in the amygdala. In accordance with that prediction, we found evidence of sparse coding only in the hippocampus. In the amygdala, a generalized novelty detection signal was observed. In contrast to the pattern we observed here, other studies of recognition memory in epilepsy patients reported evidence for both general novelty detectors and general familiarity detectors, in both the hippocampus and the amygdala (16–18, 34, 35). Similarly, another recent study of one-trial associative learning in epilepsy patients reported changes in stimulus-specific single-unit activity as a function of learning in both the hippocampus and the amygdala (36). The fact that similar memory signals were observed in both structures in these studies is somewhat surprising given that recognition memory is a hippocampus-dependent task (but not an amygdala-dependent task) and that neurocomputational models predict that memory-related neural activity associated with episodic memory will be detected in the hippocampus (but not in the amygdala). It is unclear why memory-related activity of single units in the hippocampus and amygdala are sometimes similar and sometimes different.

Although we previously found evidence for a sparse distributed memory code in the human hippocampus using an old/new

recognition procedure (23), we did not detect any evidence of either neural sharpening in the hippocampus or novelty detection in the amygdala, as in the present study. However, that study involved many fewer single units than we analyzed here, so there may have been insufficient power to detect an effect. Alternatively, the disparate pattern may reflect task differences. In continuous recognition memory, novel items carry greater significance, relative to study-test recognition, as the participant must simultaneously classify novel items as new and also encode them for later recognition. By this interpretation, the novelty signals we observed in the left amygdala may reflect the high task relevance of novel items on the continuous recognition task (36).

Methods

Participants. The participants were 21 patients with drug-resistant epilepsy requiring the implantation of depth electrodes (Ad-Tech Medical) for clinical evaluation and consideration of possible surgical resection of their seizure foci. The mean age of the patients was 40 (range 20–61 y), 12 were female, 20 were right-handed, and all had temporal lobe epilepsy. All patients provided informed consent to participate in the research, using a protocol approved by the Institutional Review Board of St. Joseph's Hospital and Medical Center. The final analysis included data from only 20 patients because the recognition memory performance of one patient was close to chance (see below).

Materials and Procedure. The patients were tested using a continuous recognition task with words as stimuli. The words were 120 one-syllable, 120 two-syllable, and 120 three-syllable words, all taken from the Medical Research Council (MRC) Psycholinguistic database (37). Each word was presented in either the Bradley or Impulse fonts (the font manipulation had no effect on any dependent measure, so we collapsed across fonts for all analyses). One set of stimulus materials consisted of 40 each of the one-, two-, and three-syllable sets in both fonts. Another 15 one-syllable words in each font were used as fillers and never repeated. There were three separate sets of stimulus materials that could be presented, and these were used for patients who volunteered for multiple sessions.

Each experimental session consisted of a sequence of 255 trials, including 15 filler trials. (Filler words were presented only once to make the overall probability of repetition equal 50%.) In each trial, a word was shown for 1.5 s, followed by a question mark. Up to 2 s was allowed for a key press, indicating either that the word was repeated (previously seen in this experimental session) or novel. Repeated words were presented after 0, 1, 3, 7, 15, or 31 intervening words. In total, we administered 45 recognition tests to 21 patients. Five patients took more than three tests and saw a stimulus set repeated one or two times, but repetition of the stimuli had no significant effect on performance (i.e., recognition accuracy was unaffected by having previously seen a particular stimulus set). Across patients, eight recognition tests resulted in poor recognition scores ($d' < 0.5$) and were excluded from further neural analysis, leaving 37 sessions to be analyzed from 20 patients.

Microwire Implantation. Electrode implantation was performed stereotactically (Medtronic StealthStation) using a preoperative structural MRI. This procedure localizes the tips of the microwires to within 2 mm (38). Bundles of nine 38- μ m-diameter platinum-iridium microwires (California Fine Wire) were introduced through a lumen within the clinical intraparenchymal electrode during surgery. The implantation sites were chosen according to clinical criteria, which limits the potential recording sites. For the 20 patients studied here, however, the sites included the hippocampus and amygdala, bilaterally. In the hippocampus, the wires were targeted to be in the mid-body of the hippocampus, just behind the head of the hippocampus, opposite the apex of the cerebral peduncle. In the amygdala, the wires were targeted to be in the center of that structure.

Filtering and Event Detection. Extracellular potentials were recorded from the tips of the microwires using techniques previously described (39) and digitized at 29,412 Hz with 16-bit resolution. Possible action potential events (APs) were detected using digital filtering and thresholding (39). Because more than one neuron may be recorded near any given electrode, APs were sorted into several clusters of similar waveform shape using the open-source clustering program KlustaKwik ([Klustakwik.sf.net](http://klustakwik.sf.net)). After sorting, each cluster was graded as being noise, multiunit activity (MUA), or single-unit activity (SUA) based on criteria such as the waveform shape (Fig. S4), size of the waveform relative to noise, evidence of a refractory interval, and lack of powerline interference, using the criteria described previously (39).

In our experience, this technique produces results comparable to prior reports in other laboratories (19) in terms of recorded waveform shapes, interspike intervals, and firing rates. While it is important to note that these and other reports of human single-unit recordings (40, 41) do not achieve the quality of unit separation achievable in animal recordings (42), they nonetheless represent neural activity at a much finer spatial and temporal scale than is achievable using other methods such as fMRI. Measurement at a fine spatial and temporal scale (not necessarily the measurement of single units per se) is necessary to test the predictions of neurocomputational models that assume a sparse distributed episodic memory coding scheme.

We recorded from a total of 1,546 clusters of events representing neural activity in the medial temporal lobe (amygdala and hippocampus), 518 of which satisfied the criteria for SUA and 1,028 of which were categorized as MUA. In this report, we focused on SUA (161 neurons in the left amygdala, 114 neurons in the right amygdala, 128 neurons in the left hippocampus, and 115 neurons in the right hippocampus). Poststimulus spike counts for each unit were recorded 200–1,000 ms after the onset of the test stimulus, and prestimulus (baseline) spike counts were recorded 200–1,000 ms before the onset of the test stimulus. The test period during which poststimulus spike counts were recorded was chosen because a previous study (20) found that selective responses of hippocampal neurons began ~300 ms after stimulus onset and because nearly all behavioral responses occurred after 1 s.

Data Analysis. For every recorded neuron, we computed normalized spike counts for each trial (i), where a “trial” refers to the presentation of a novel or repeated word. For each neuron (j), its baseline mean and SD of spike counts (μ_j and σ_j , respectively) were computed across all trials in that session. Normalized poststimulus spike counts for a given trial (N_{ij}) in which s_{ij} raw spike counts were recorded on trial i for neuron j is given by $N_{ij} = (s_{ij} - \mu_j)/\sigma_j$. Trials in which a behavioral response occurred during the 1.5-s stimulus presentation (and, therefore, before the signal to respond was presented) were denoted as “early” responses and were excluded from the analysis.

The data were analyzed separately for each of four brain regions (left hippocampus, right hippocampus, left amygdala, and right amygdala). First, we performed a conventional analysis on the normalized spike counts (using ANOVA) to identify individual neurons in the hippocampus and/or amygdala that were responsive to the general class of novel or repeated items, with word repetition status (novel vs. repeated) as the independent variable. Second, an aggregate analysis was performed on the full distributions of normalized single unit activity in a given region (collapsed over patients and sessions) for novel and repeated items. The question was whether the mean and/or SD of the repeated-item distribution differed significantly from the corresponding parameters of the novel-item distribution (e.g., in the left hippocampus). The statistical reliability of any difference in either the mean or the SD of the two distributions was tested using a bootstrap procedure. For each test (e.g., comparing the SDs for novel and repeated items in the left hippocampus), 10,000 bootstrap trials were performed in which (i) the data from all repeated words (n_{Repeated}) and all novel words (n_{Novel}) were combined, (ii) n_{Repeated} bootstrap “targets” and n_{Novel} bootstrap “foils” were randomly sampled with replacement from that combined dataset, and (iii) the difference between the statistic of interest (e.g., SD) of those two bootstrap samples was computed. The resulting P value was the proportion of bootstrap trials in which the absolute value of the difference was greater than the observed difference. A similar bootstrap analysis yielded the estimated SEs shown in Figs. 2, 5, and 6.

ACKNOWLEDGMENTS. We thank the patients who participated in this study as well as Kent Horne and Elaine Cabrales for technical assistance. This work was supported by a McKnight Foundation Memory and Cognitive Disorders Award (to J.T.W. and L.R.S.), Medical Research Service of the Department of Veterans Affairs Grant 5101CX000359 (to L.R.S.), National Institute of Child Health and Human Development Grant HD075800-04 (to M.H.P. and S.D.G.), and National Institute of Neurological and Communicative Disorders and Stroke Grant DC009781 (to P.N.S.).

1. Tulving E (1983) *Elements of Episodic Memory* (Clarendon, Oxford).
2. Tulving E (2002) Episodic memory: From mind to brain. *Annu Rev Psychol* 53:1–25.
3. Shepard RN, Teghtsoonian M (1961) Retention of information under conditions approaching a steady state. *J Exp Psychol* 62:302–309.
4. Stark CEL, Squire LR (2003) Hippocampal damage equally impairs memory for single items and memory for conjunctions. *Hippocampus* 13:281–292.
5. Marr D (1971) Simple memory: A theory for archicortex. *Philos Trans R Soc Lond B Biol Sci* 262:23–81.
6. Treves A, Rolls ET (1994) Computational analysis of the role of the hippocampus in memory. *Hippocampus* 4:374–391.
7. McClelland JL, McNaughton BL, O'Reilly RC (1995) Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychol Rev* 102:419–457.
8. Norman KA, O'Reilly RC (2003) Modeling hippocampal and neocortical contributions to recognition memory: A complementary-learning-systems approach. *Psychol Rev* 110:611–646.
9. Brown MW, Wilson FAW, Riches IP (1987) Neuronal evidence that inferomedial temporal cortex is more important than hippocampus in certain processes underlying recognition memory. *Brain Res* 409:158–162.
10. Heit G, Smith ME, Halgren E (1988) Neural encoding of individual words and faces by the human hippocampus and amygdala. *Nature* 333:773–775.
11. Heit G, Smith ME, Halgren E (1990) Neuronal activity in the human medial temporal lobe during recognition memory. *Brain* 113:1093–1112.
12. Riches IP, Wilson FA, Brown MW (1991) The effects of visual stimulation and memory on neurons of the hippocampal formation and the neighboring parahippocampal gyrus and inferior temporal cortex of the primate. *J Neurosci* 11:1763–1779.
13. Rolls ET, et al. (1989) Hippocampal neurons in the monkey with activity related to the place in which a stimulus is shown. *J Neurosci* 9:1835–1845.
14. Rolls ET, Cahusac PM, Feigenbaum JD, Miyashita Y (1993) Responses of single neurons in the hippocampus of the macaque related to recognition memory. *Exp Brain Res* 93:299–306.
15. Jutras MJ, Buffalo EA (2010) Recognition memory signals in the macaque hippocampus. *Proc Natl Acad Sci USA* 107:401–406.
16. Rutishauser U, Mamelak AN, Schuman EM (2006) Single-trial learning of novel stimuli by individual neurons of the human hippocampus-amygdala complex. *Neuron* 49:805–813.
17. Rutishauser U, Schuman EM, Mamelak AN (2008) Activity of human hippocampal and amygdala neurons during retrieval of declarative memories. *Proc Natl Acad Sci USA* 105:329–334.
18. Rutishauser U, et al. (2015) Representation of retrieval confidence by single neurons in the human medial temporal lobe. *Nat Neurosci* 18:1041–1050.
19. Viskontas IV, Knowlton BJ, Steinmetz PN, Fried I (2006) Differences in mnemonic processing by neurons in the human hippocampus and parahippocampal regions. *J Cogn Neurosci* 18:1654–1662.
20. Quiroga RQ, Reddy L, Kreiman G, Koch C, Fried I (2005) Invariant visual representation by single neurons in the human brain. *Nature* 435:1102–1107.
21. Quiroga RQ (2012) Concept cells: The building blocks of declarative memory functions. *Nat Rev Neurosci* 13:587–597.
22. Chambers JM, Cleveland WS, Kleiner B, Tukey P (1983) *Graphical Methods for Data Analysis* (Wadsworth Int Group, Pacific Grove, CA).
23. Wixted JT, et al. (2014) Sparse and distributed coding of episodic memory in neurons of the human hippocampus. *Proc Natl Acad Sci USA* 111:9621–9626.
24. Gelbard-Sagiv H, Mukamel R, Harel M, Malach R, Fried I (2008) Internally generated reactivation of single neurons in human hippocampus during free recall. *Science* 322:96–101.
25. Valdez AB, et al. (2015) Distributed representation of visual objects by single neurons in the human brain. *J Neurosci* 35:5180–5186.
26. Desimone R (1996) Neural mechanisms for visual memory and their role in attention. *Proc Natl Acad Sci USA* 93:13494–13499.
27. Grill-Spector K, Henson R, Martin A (2006) Repetition and the brain: Neural models of stimulus-specific effects. *Trends Cogn Sci* 10:14–23.
28. Karlsson MP, Frank LM (2008) Network dynamics underlying the formation of sparse, informative representations in the hippocampus. *J Neurosci* 28:14271–14281.
29. Berger TK, Silberberg G, Perin R, Markram H (2010) Brief bursts self-inhibit and correlate the pyramidal network. *PLoS Biol* 8:e1000473.
30. Kapfer C, Glickfeld LL, Atallah BV, Scanziani M (2007) Supralinear increase of recurrent inhibition during sparse activity in the somatosensory cortex. *Nat Neurosci* 10:743–753.
31. Kesner RP, Rolls ET (2015) A computational theory of hippocampal function, and tests of the theory: New developments. *Neurosci Biobehav Rev* 48:92–147.
32. Froudarakis E, et al. (2014) Population code in mouse V1 facilitates readout of natural scenes through increased sparseness. *Nat Neurosci* 17:851–857.
33. Thome A, et al. (2017) Evidence for an evolutionarily conserved memory coding scheme in the mammalian hippocampus. *J Neurosci* 37:2795–2801.
34. Fried I, MacDonald KA, Wilson CL (1997) Single neuron activity in human hippocampus and amygdala during recognition of faces and objects. *Neuron* 18:753–765.
35. Cameron KA, Yashar S, Wilson CL, Fried I (2001) Human hippocampal neurons predict how well word pairs will be remembered. *Neuron* 30:289–298.
36. Ison MJ, Quiroga R, Fried I (2015) Rapid encoding of new memories by individual neurons in the human brain. *Neuron* 87:220–230.
37. Coltheart M (1981) The MRC psycholinguistic database. *Q J Exp Psychol A* 33:497–505.
38. Mehta AD, et al. (2005) Frameless stereotactic placement of depth electrodes in epilepsy surgery. *J Neurosurg* 102:1040–1045.
39. Valdez AB, Hickman EN, Treiman DM, Smith KA, Steinmetz PN (2013) A statistical method for predicting seizure onset zones from human single-neuron recordings. *J Neural Eng* 10:016001.
40. Kreiman G, Koch C, Fried I (2000) Category-specific visual responses of single neurons in the human medial temporal lobe. *Nat Neurosci* 3:946–953.
41. Steinmetz PN (2009) Alternate task inhibits single-neuron category-selective responses in the human hippocampus while preserving selectivity in the amygdala. *J Cogn Neurosci* 21:347–358.
42. Hill DN, Mehta SB, Kleinfeld D (2011) Quality metrics to accompany spike sorting of extracellular signals. *J Neurosci* 31:8699–8705.

Supporting Information

Wixted et al. 10.1073/pnas.1716443115

SI Results

Left Hippocampus. Fig. S1 shows QQ plots from the left hippocampus broken down by correct and incorrect behavioral responses. The pattern is suggestive of bimodality regardless of the accuracy of the behavioral response. Fig. S2 shows QQ plots from the left hippocampus broken down by lag. The left column shows the results when 100% of recordings are included in the analysis, and the right column shows the results when the highest 2.5% of recordings for novel items and repeated items are excluded from the analysis. A pattern consistent with sparse distributed coding is evident across most of the lags with the exceptions of 8 and 16. In all cases, the QQ plot pattern is essentially linear when the highest 2.5% of the scores are eliminated from the analysis (as shown by the QQ plots in the right column).

Left Amygdala. As noted earlier, a general novelty signal in the left amygdala was strong enough to be convincingly detected at the level of individual single units (Table 1). How does that effect manifest itself in the full novel- and repeated-item distributions of single-unit recordings made from the amygdala? The mean and SD scores computed from the full distributions were both significantly greater for novel items compared with repeated items in the left amygdala but not in the right amygdala (Table S1). Again, the fact that these effects are observed in the left but not right amygdala may be attributable to the fact that a verbal memory task was used in our experiment. This overall novelty signal in the

left amygdala is consistent with the fact that 25 single units were identified in the left amygdala that were general novelty detectors (i.e., they responded significantly more, on average, to novel items than to repeated items). The fact that a subset of single units were novelty detectors would account for both the mean and SD differences in the full distributions in the left amygdala (Table S1).

The QQ plots for recordings made from the left and right amygdala (Fig. S3 *A* and *B*, respectively) both show slight evidence of novelty detection at the very high end of each plot. However, visual appearances notwithstanding, the number of points that fall visibly below the diagonal line toward the upper end is extremely small ($\sim 0.1\%$ of all points). These points did not have an appreciable effect on either the means or the SDs shown in Table S1, both because there were so few and because the degree to which they deviated from the diagonal line was relatively small. In fact, when 2.5% of the highest scores are eliminated (thereby eliminating all of the points diverging below the diagonal line and more), the statistical pattern of results does not change. In other words, whether 100% of the data are considered or 97.5% of the data are considered (Table S1), a significant difference is evident for novel and repeated items—for both the mean and SD—in the left amygdala but not in the right amygdala. In fact, those differences remain significant even when 20% of the highest values are removed from the analysis. Thus, the results suggest a general pattern of novelty detection.

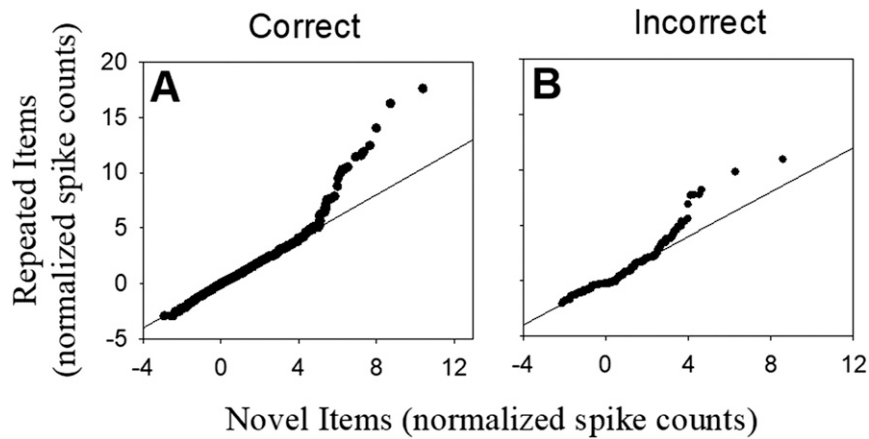


Fig. S1. QQ plots for left hippocampus for correct behavioral (A) and incorrect behavioral (B) decisions. Correct behavioral decisions consist of hits and correct rejections, and incorrect behavioral decisions consist of misses and false alarms. Each point on a QQ plot represents the normalized average spike count recorded on a single test trial. The plot displays those values aggregated across trials and patients. For correct decisions, the plot displays 10,695 and 11,163 normalized spike counts for hits and correct rejections, respectively. For incorrect decisions, the plot displays 2,159 and 2,659 normalized spike counts for hits and correct rejections, respectively.

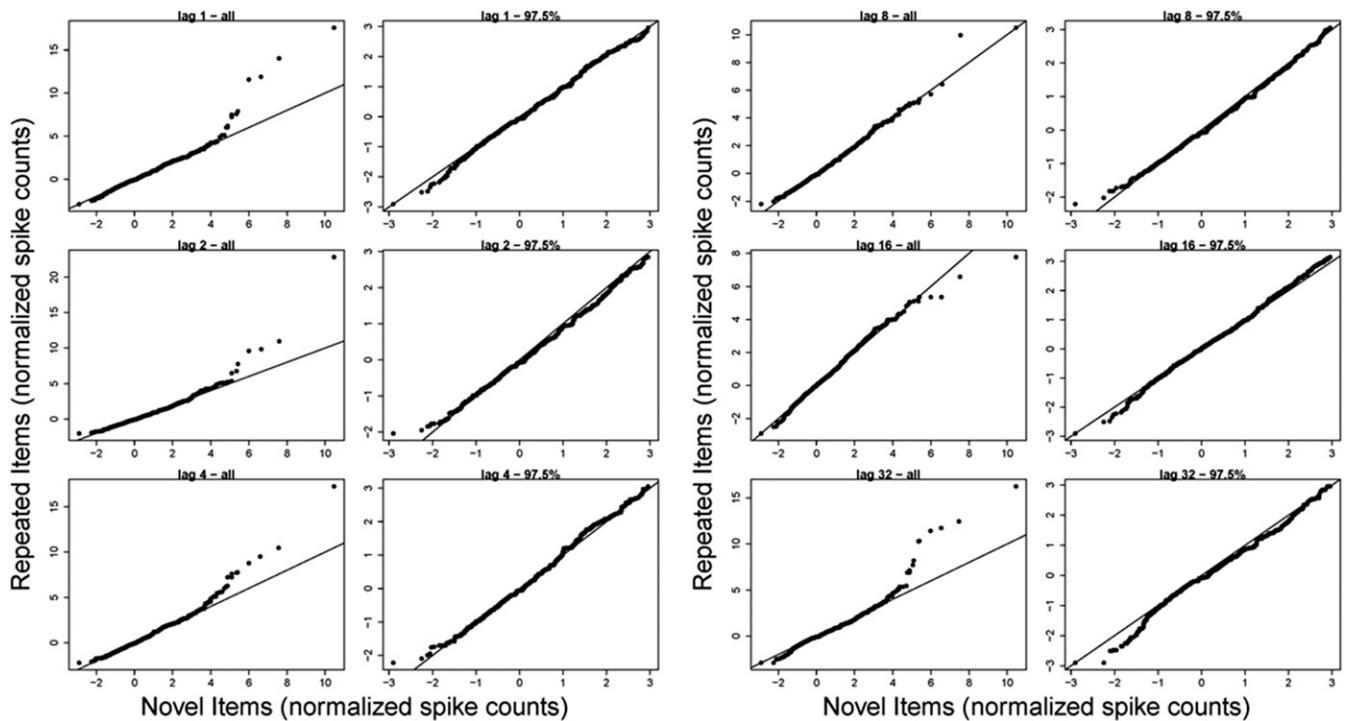


Fig. S2. QQ plots for left hippocampus for 100% of the data (*Left*) and after excluding 2.5% of the data with the highest spike counts from both the repeated-item and novel-item distributions (*Right*) separately for each lag.

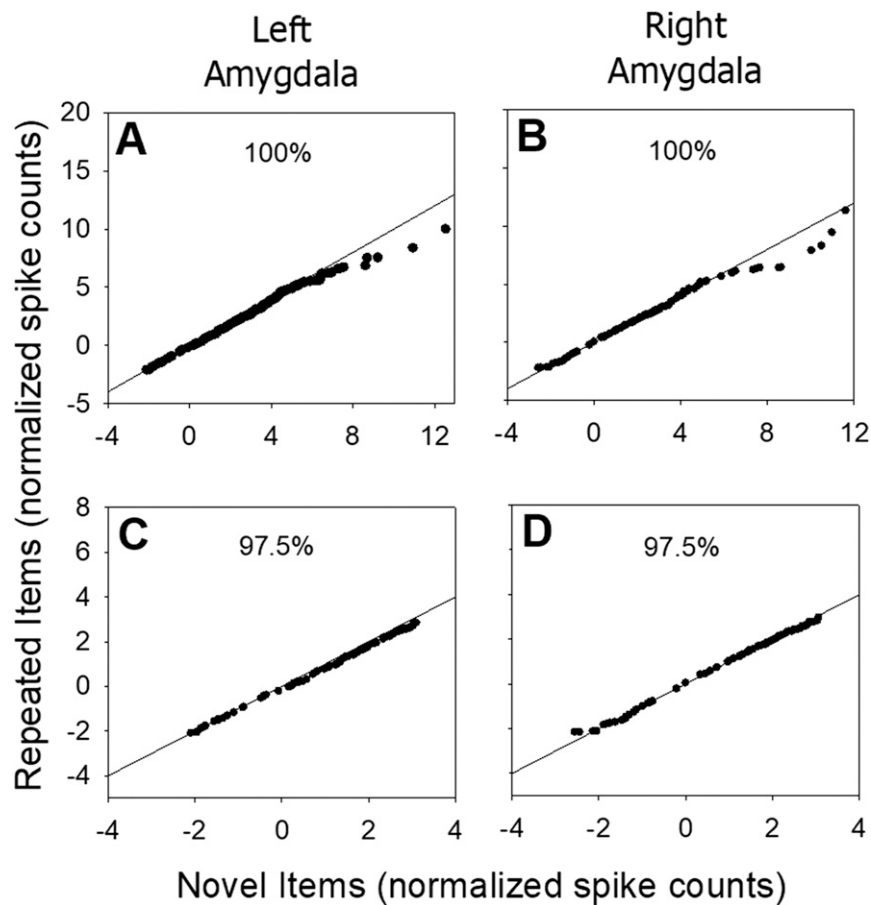


Fig. S3. QQ plots for the left and right amygdala for 100% of the data (*A* and *B*, respectively) and after excluding the upper 2.5% of the data (*C* and *D*, respectively). In the left amygdala, the 100% plot displays 15,121 and 16,191 normalized spike counts for repeated and novel items, respectively. In the right hippocampus, the corresponding values are 10,153 and 11,379 normalized spike counts.

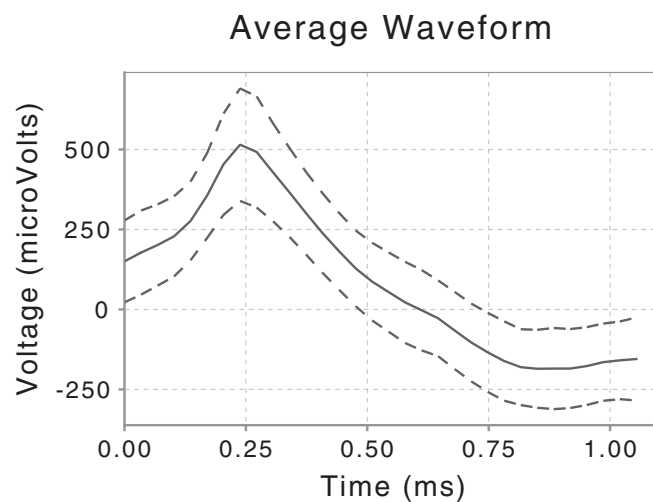


Fig. S4. Waveform shape of single-unit activity recorded in the left hippocampus. *x* axis: Time during waveform in milliseconds, with peak centered at 0.25 ms. *y* axis: Extracellular potential difference in microvolts; solid line shows the average value of all possible action potential events in the cluster, dotted lines show ± 1 SD.

Table S1. Distributional statistics (mean and SD) for recordings made from left and right amygdala

Data analyzed, %	Statistic	Left			Right		
		Repeated	Novel	adj <i>P</i>	Repeated	Novel	adj <i>P</i>
100	Mean	0.09	0.19	<0.001	0.01	0.00	0.686
	SD	1.08	1.15	<0.001	1.00	1.01	0.686
	<i>n</i>	15,121	16,191	—	10,153	11,379	—
97.5	Mean	−0.01	0.08	<0.001	−0.10	−0.11	0.609
	SD	0.88	0.95	<0.001	0.74	0.74	0.814
	<i>n</i>	14,742	15,786	—	9,899	11,094	—

Mean, SD, and number of normalized spike counts (*n*) associated with the distributions for repeated and novel items in the left and right amygdala (aggregated over all recorded units and trials) for 100% of the data (top 3 rows) and after excluding the highest 2.5% of the scores for each distribution, thereby retaining 97.5% of the data (bottom 3 rows).