Sparse and distributed coding of episodic memory in neurons of the human hippocampus

John T. Wixted[a,1], Larry R. Squire[a,b,c,d,1], Yoonhee Jang[e], Megan H. Papesh[f], Stephen D. Goldinger[g], Joel R. Kuhn[a], Kris A. Smith[h], David M. Treiman[i], and Peter N. Steinmetz[h,i]

Departments of [a]Psychology, [b]Psychiatry, and [c]Neurosciences, University of California, San Diego, La Jolla, CA 92093; [d]Veterans Affairs Medical Center, San Diego, CA 92161; [e]Department of Psychology, University of Montana, Missoula, MT 59812; [f]Department of Psychology, Louisiana State University, Baton Rouge, LA 70803; [g]Department of Psychology, Arizona State University, Tempe, AZ 85287; and Departments of [h]Neurosurgery and [i]Neurology, Barrow Neurological Institute, Phoenix, AZ 85013

Contributed by Larry R. Squire, May 14, 2014 (sent for review March 20, 2014)

Neurocomputational models hold that sparse distributed coding is the most efficient way for hippocampal neurons to encode episodic memories rapidly. We investigated the representation of episodic memory in hippocampal neurons of nine epilepsy patients undergoing intracranial monitoring as they discriminated between recently studied words (targets) and new words (foils) on a recognition test. On average, single units and multiunits exhibited higher spike counts in response to targets relative to foils, and the size of this effect correlated with behavioral performance. Further analyses of the spike-count distributions revealed that (*i*) a small percentage of recorded neurons responded to any one target and (*ii*) a small percentage of targets elicited a strong response in any one neuron. These findings are consistent with the idea that in the human hippocampus episodic memory is supported by a sparse distributed neural code.

recognition memory | intracranial recording | amygdala

The hippocampus is known to play a fundamental role in declarative memory (1–4), but it is not known how mnemonic information is coded by the activity of individual hippocampal neurons. At least three different coding schemes have been considered: a localist coding scheme, a fully distributed coding scheme, and a sparse distributed coding scheme (5). In a localist coding scheme, an individual neuron (sometimes referred to as a "grandmother cell") codes only one memory, and each memory is coded by the activity of only one neuron. In a fully distributed coding scheme, each memory is coded instead by a pattern of activity across many hippocampal neurons. Falling between these two extremes is a sparse distributed coding scheme in which each memory is coded by the activity of a small proportion of hippocampal neurons, and each neuron contributes to the representation of only a few memories. Sparse distributed coding has long been hypothesized to be the most efficient way for hippocampal neurons to encode episodic memories (remembering events) in rapid succession without overwriting previously stored memories (6–8).

Most prior work concerned with the coding of declarative memory in the human hippocampus has focused on the neural representation of semantic memories (remembering facts), such as memory for famous people or landmarks (9, 10). The results of these studies suggest that long-established semantic memories may be represented by fewer than 1% of neurons in the hippocampus (11). However, neurocomputational theories are concerned with the representation of episodic memories. The purpose of our study was to test predictions of these neurocomputational theories about how episodic memories are represented by neurons of the hippocampus.

The representation of episodic memory in the hippocampus typically has been investigated using recognition procedures. In recognition, the task is to discriminate between familiar items presented earlier in the experimental session (targets) and novel items not previously presented (foils). An episodic memory signal is evident when neurons exhibit different levels of activity for targets (old items) vs. foils (new items). The first recognition studies with humans (12, 13) and monkeys (14–16) failed to detect evidence of episodic memory in neurons of the hippocampus, but more recent studies have identified hippocampal neurons that differentiate targets from foils (17–21). However, these studies did not investigate how the representation of individual targets is distributed across neurons of the hippocampus. Instead, the aim was to find cells that distinguish the class of targets from the class of foils.

We investigated the representation of individual targets in neurons of the human hippocampus. The participants were nine patients with pharmaco-resistant epilepsy requiring the implantation of intracranial wire electrodes for clinical evaluation and localization of seizure foci for possible surgical resection. Among them, the patients completed a total of 18 recognition memory tasks in which they first studied 32 words and then attempted to distinguish between the 32 targets that had appeared on the study list and 32 foils that had not. Each of the 64 items on the recognition test was presented only once, a format that differs from many other neurophysiology studies that present individual stimuli multiple times to identify neurons with reliable stimulus-specific firing properties. The multiple-presentation method is well-suited to the study of semantic memory (e.g., a neuron that is found to respond reliably to six presentations of the word "baby" likely is responding to its long-established semantic meaning) but is not well-suited to the study of episodic memory. When targets and foils are presented only once on a recognition test, the targets, but not the foils, are represented by an episodic

Significance

The ability to form episodic memories in rapid succession depends on the hippocampus, but how do hippocampal neurons represent such memories? Most neurocomputational models envision a sparse distributed coding scheme in which individual neurons each participate in the coding of a few memories, and each memory is coded by a small fraction of hippocampal neurons. We investigated this issue with epilepsy patients undergoing intracranial monitoring and found evidence consistent with sparse distributed coding in the hippocampus. These findings shed new light on the basic neural mechanisms that underlie the ability to remember events. A detailed characterization of those mechanisms is an essential part of the larger effort to understand memory loss associated with normal aging and dementia.

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

The authors declare no conflict of interest.

[1]To whom correspondence may be addressed. E-mail: jwixted@ucsd.edu or lsquire@ucsd.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1408365111/-/DCSupplemental.

memory formed earlier at the time of learning. Under these conditions, any difference in neural activity associated with targets and foils would indicate episodic memory. Note that, if the test items were presented again, the targets and foils no longer would be clearly differentiated because even the foils would be represented by a recently formed, context-specific episodic memory. Accordingly, instead of using multiple stimulus presentations during the recognition test, we examined the distribution of activity associated with once-presented targets vs. once-presented foils across all recorded neurons. The different coding schemes under consideration here make distinct predictions about the expected distributions of neural activity.

Results

Behavioral Data. Recognition decisions were made using an 8-point confidence scale (1 = Sure New … 8 = Sure Old). Confidence ratings of 5 through 8 were counted as hits for target words and as false alarms for foils. Hit rate, false alarm rate, percent correct, and discriminability (d') scores were computed for each patient. For patients who completed more than one recognition test, these measures were computed separately for each test and then averaged. All nine patients exhibited above-chance memory for every test (mean = 63% correct), although the performance for patient 8 was close to chance (Table S1).

Reaction time (RT) was defined as the interval between the onset of a test item and the mouse click indicating the confidence rating for that test item. All confidence ratings were made 1,500 ms or more after the presentation of the test item, that is, after the occurrence of the spikes counts that were analyzed here (200–1,000 ms after test stimulus presentation; see Fig. S1).

Neural Data. We recorded neural activity bilaterally from 220 units in the hippocampus (34 single units + 186 multiunits) and 300 units in the amygdala (68 single units + 232 multiunits) over the course of the 18 recognition memory tests. In the hippocampus, the mean firing rate of the 34 single units during the prestimulus baseline period (1,000–200 ms before stimulus presentation) was 1.7 Hz (SD = 2.76), and the mean firing rate of the 186 multiunits was 21.2 Hz (SD = 10.32). In the amygdala, the mean firing rate of the 68 single units during the prestimulus baseline period was 2.0 Hz (SD = 2.02), and the mean firing rate of the 232 multiunits was 20.2 Hz (SD = 11.77). Before analysis, poststimulus spike counts for each unit (200–1,000 ms after test stimulus presentation) were normalized based on the mean and SD of the unit's spike counts during the prestimulus baseline period across all 64 test items.

The normalized spike counts were subjected to three levels of analysis. First, for each patient, we measured the difference between test period activity for targets vs. foils averaged across all units recorded from the hippocampus and, separately, from the amygdala, and we compared that neural difference score with a standard behavioral measure of recognition memory performance (d'). Second, we analyzed the activity of individual units in the hippocampus and the amygdala to identify those for which the average activity level across the 32 targets differed significantly from the average activity level across the 32 foils. Third, at the most fine-grained level of analysis, we compared the distribution of all spike counts recorded for individual targets and compared it with the distribution of all spike counts recorded for individual foils. This analysis most directly addresses the question of how individual episodic memories are represented by neurons of the human hippocampus.

In the first (patient-level) analysis, for each patient (i), a single difference score (D_i') was computed consisting of the mean of all normalized spike counts to targets, μ_{Target_i}, minus the mean of all normalized spike counts to foils, μ_{Foil_i}. As an example, 17 hippocampal units (six single units and 11 multiunits) were recorded from patient 3 in one recognition test session. Because the recognition test consisted of 32 targets and 32 foils, 544 normalized scores were averaged to compute μ_{Target_3} (i.e., 17 units × 32 targets = 544 target scores), and 544 normalized scores were averaged to compute μ_{Foil_3} (i.e., 17 units × 32 foils = 544 foil scores). The individual neural difference score for patient 3 (D'_3) was equal to $\mu_{Target_3} - \mu_{Foil_3}$. D' scores for all nine patients were computed using a similar procedure. For patients who completed more than one recognition test, these measures were computed separately for each test and then averaged. Using the same procedure, a second D' score was computed for each patient based on recordings made from the amygdala. The results indicated the detection of an episodic memory signal in the hippocampus but not in the amygdala (Fig. 1*A*). Moreover, for recordings made from the hippocampus, the larger the difference between the average normalized spike counts between targets and foils for a particular patient, the better recognition memory performance was for that patient as measured by d' (Fig. 1*C*). When this same test was performed on recordings from the amygdala, no such relationship was observed (Fig. 1*D*). These findings are consistent with prior work showing that hippocampal lesions impair recognition memory performance, whereas amygdala lesions do not (1), but they differ from other single-unit studies that have detected episodic memory signals in the amygdala (19, 20).

In the second (unit-level) analysis, we identified units with spike counts that significantly differentiated targets from foils. For each unit j, where j = 1–220 in the hippocampus and j = 1–300 in the amygdala, a t test was performed comparing the mean normalized spike count across the 32 targets (μ_{Target_j}) with the mean normalized spike count across the 32 foils (μ_{Foil_j}). Under the null hypothesis of no difference, and assuming independence,

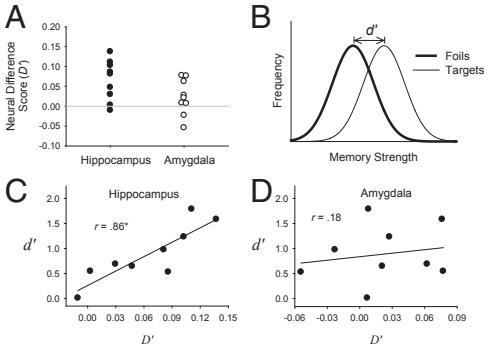

Fig. 1. Fully aggregated analyses. (*A*) For each patient i (where i = 1–9), a neural difference score (D'_i) was computed from recordings made from the hippocampus and, separately, from recordings made from the amygdala. The difference score represents the normalized spike counts in response to targets (averaged across all units and all 32 targets) minus the normalized spike counts in response to foils (averaged across all units and all 32 foils). For patients who participated in more than one recognition test, this value was computed separately for each session and the values were then averaged. The mean of the distribution of difference scores shown in *A* was significantly greater than 0 in the hippocampus [t (8) = 3.9, $P < 0.01$], but not in the amygdala [t (8) = 1.5, P = 0.17]. (*B*) An illustration of the standard signal-detection model of recognition memory in which the behavioral measure d' reflects the theoretical difference between the average memory strength of the targets minus the average memory strength of the foils. (*C*) A plot of the relationship between the neural D'_i scores from the hippocampus and the corresponding behavioral d'_i score (one pair of D'_i and d'_i scores for each of the nine patients). The correlation between these two measures (r = 0.86) was significant ($P < 0.01$). (*D*) A plot of the relationship between the neural D'_i scores from the amygdala and the corresponding behavioral d'_i scores. The correlation between these two measures (r = 0.18) was not significant (P = 0.64). The difference between these two correlation coefficients was marginally significant, z = 1.92, P = 0.055.

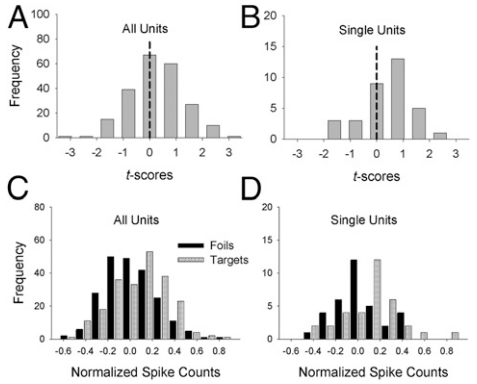

Fig. 2. Unit-level analyses. (*A* and *B*) For each unit *j*, where *j* = 1–220 for all units (*A*) and *j* = 1–34 for the single units (*B*), a *t* test was performed comparing the mean normalized spike count across the 32 targets (μ_{Target_j}) with the mean normalized spike count across the 32 foils (μ_{Foil_j}). The mean of the *t* distribution was significantly greater than 0 (dashed vertical line) both for all units [$M = 0.24$, $t(219) = 3.44$, $P < 0.001$] and for the single units [$M = 0.36$, $t(33) = 2.24$, $P < 0.05$]. (*C* and *D*) Distribution of target means (i.e., the distribution of μ_{Target_j} values) and foil means (i.e., the distribution of μ_{Foil_j} values) for all units (*C*) and for single units separately (*D*). For all units (*C*), the mean of the target means ($\bar{\mu}_{Target} = 0.08$) was significantly greater than the mean of the foil means ($\bar{\mu}_{Foil} = 0.01$), $t(219) = 3.69$, $P < 0.001$. Similarly, for the single units (*D*), the mean of the target means ($\bar{\mu}_{Target} = 0.11$) was significantly greater than the mean of the foil means ($\bar{\mu}_{Foil} = 0.003$), $t(33) = 2.28$, $P < 0.05$. Note that statistical tests for all units are reported for the sake of completeness, but they should be interpreted with caution because the measures may not be independent (i.e., the measures from two or more multiunits from the same patient may reflect partially overlapping neural activity).

5% of these tests would be significant on the basis of chance alone. In our data, only 12 of the 220 units (5.5%) in the hippocampus and only 8 of the 300 units (2.7%) in the amygdala significantly differentiated targets from foils, effects that occurred no more often than would be expected from chance. A similar result was obtained when the analysis was limited to the single units. Only 1 of the 34 single units recorded from the hippocampus (2.9%) significantly differentiated targets from foils, and none of the 68 single units recorded from the amygdala did so. Thus, analyzed in this fashion, the data offer no evidence of a neural episodic memory signal in the hippocampus or amygdala.

The apparent absence of an episodic memory signal at the level of individual units in the hippocampus parallels earlier results from similar episodic memory tests with epilepsy patients (12, 13) and with monkeys (16). However, an an episodic memory signal was evident in our data when the spike counts from the hippocampus were aggregated across units (Fig. 1 *A* and *C*). This apparent contradiction is resolved by analyzing the distribution of 220 *t* scores, instead of simply counting the number of units associated with a significant *t* score. A null effect would be characterized by a distribution of *t* scores with a mean of zero. However, when the full distribution of 220 *t* scores was analyzed, the mean was slightly but significantly greater than 0 (Fig. 2*A*). Similarly, the mean of the 34 *t* scores for the single units considered separately also was significantly greater than 0 (Fig. 2*B*). Although the distribution of *t* scores was elevated above 0 in the hippocampus, no such effects were evident in the amygdala (i.e., the mean of the *t* distribution did not differ significantly from 0).

The numerator of the *t*-score formula for a particular unit *j* is $\mu_{Target_j} - \mu_{Foil_j}$. Thus, the results shown in Fig. 2 *A* and *B* suggest that the entire distribution of μ_{Target_j} values was significantly elevated relative to the distribution of μ_{Foil_j} values, and indeed this was the case (Fig. 2 *C* and *D*). These results demonstrate that even though the number of units with significant *t* scores was no greater than would be expected on the basis of chance, most or all of the units were involved in the coding of episodic memory. In that sense, a distributed memory signal was identified, arguing against the notion of a strictly localist (i.e., "grandmother" cell) coding scheme. Under a localist scheme, the likelihood of detecting any memory signal would be extremely low. However, although a distributed signal was identified (Fig. 2 *C* and *D*), this analysis does not distinguish between a fully distributed coding scheme (Fig. 3*A*) and a sparse distributed coding scheme (Fig. 3*B*).

In the third and most detailed analysis, we investigated whether the significant increase in the μ_{Target_j} values in hippocampal neurons (Fig. 2 *C* and *D*) occurred because, for each unit, most or all of the 32 targets exhibited elevated spiking relative to the foils (consistent with a fully distributed coding scheme), or because, for each unit, only a few targets elicited elevated spiking relative to the foils (consistent with a sparse distributed coding scheme). To address this issue, we compared the full distribution of individual normalized spike counts for the targets with the full distribution of normalized spike counts for the foils. For the analysis involving all units, there were 7,040 individual target spike counts (32 targets × 220 units) and 7,040 individual foil spike counts (32 foils × 220 units). For a separate analysis involving only single units, there were 1,088 individual target spike counts (32 targets × 34 single units) and 1,088 individual foil spike counts (32 foils × 34 single units). If a modest memory signal were added to all the targets (a fully distributed coding

		A (Fully Distributed)								B (Sparse Distributed)							
Item Type	Test Items	--------Units--------								--------Units--------							
		1	2	3	4	5	6	7	8	1	2	3	4	5	6	7	8
Foils	*Earth*	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	Cup	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	Diamond	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	Paint	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	Game	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	Church	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	Weapon	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	Square	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Targets	*Clock*	0.1	0.2	0.2	0.1	0.2	0.2	0.2	0.1	0	0	0	0	**2**	0	0	0
	Chief	0.2	0.1	0.2	0.2	0.1	0.1	0.2	0.2	0	**2**	0	0	0	0	0	0
	Map	0.2	0.1	0.1	0.1	0.1	0.2	0.1	0.2	0	0	0	0	0	0	**2**	0
	Umbrella	0.2	0.2	0.2	0.2	0.1	0.1	0.2	0.2	**2**	0	0	0	0	0	0	0
	Leather	0.1	0.1	0.2	0.2	0.2	0.1	0.1	0.2	0	0	0	0	0	**2**	0	0
	Navy	0.2	0.2	0.2	0.1	0.1	0.1	0.2	0.1	0	0	**2**	0	0	0	0	0
	Milk	0.2	0.1	0.2	0.2	0.1	0.2	0.2	0.1	0	0	0	**2**	0	0	0	0
	Coffee	0.1	0.2	0.1	0.1	0.2	0.2	0.2	0.2	0	0	0	0	0	0	0	**2**

Fig. 3. Illustration of sparse distributed and fully distributed coding schemes. The figure presents a simplified illustration of what these coding schemes predict about normalized spike counts for eight units in response to the presentation of eight foils and eight targets. The normalized spike counts reflect hypothetical mean values (observed values would be distributed randomly about those means). Because the foils are not associated with a recent episodic memory, under both coding schemes the normalized spike counts for all hippocampal units would be drawn from a distribution with a mean of 0 (i.e., no response relative to baseline). Thus, both schemes predict a large unimodal foil distribution centered on 0. Both schemes also predict that the mean normalized spike count for the targets will be slightly greater than that of the foils. (*A*) For the fully distributed coding scheme, the targets should be associated with a large, right-shifted unimodal target distribution centered between 0.1 and 0.2. (*B*) For the sparse distributed coding scheme, each target would elicit a strong response in one unit (mean normalized spike count = 2) but would elicit no response in the remaining units. Thus, when the normalized spike counts for the targets are pooled together, a bimodal distribution would result, consisting of one large target distribution centered on 0 (coincident with the large unimodal foil distribution) and one much smaller target distribution centered on 2.

scheme; Fig. 3*A*), then, for both analyses, the entire target distribution would be shifted slightly rightward relative to the foil distribution. If, instead, a small proportion of the targets generated an especially strong response in a small proportion of neurons, with the large majority of targets eliciting no differential response relative to foils (a sparse distributed coding scheme; Fig. 3*B*), then a bimodal target distribution of normalized spike counts should be observed. That is, the large majority of target spike counts should coincide with the distribution of foil spike counts, but a small percentage of targets should be drawn from a distribution with a much higher mean.

The spike-count frequency distributions (Fig. 4 *A* and *B*) do not show clear visual evidence of a bimodal distribution for the targets. However, the bimodal target distribution predicted by the sparse distributed account would be difficult to detect visually in a frequency distribution because only a few targets would be expected to yield values associated with the upper distribution, and those few values would not necessarily be tightly organized in a visually apparent distribution. Some targets (and fewer foils) do fall in the far right tails of the frequency distributions. To determine whether the target distributions in Fig. 4 *A* and *B* are slightly right-shifted or instead are bimodal, we constructed empirical quantile–quantile (Q–Q) plots (22).

An empirical Q–Q plot simply displays one rank-ordered dataset (i.e., the sorted normalized spike counts for the targets) against another rank-ordered dataset (i.e., the sorted normalized spike counts for the foils). This graphical analysis method provides more accurate information about the relative shapes of two distributions than can be obtained from a visual inspection of the frequency distributions alone. A shifted distribution would yield a linear pattern of scores elevated above the diagonal line (illustrated with hypothetical data in Fig. 4*C*), whereas a bimodal distribution would yield a pattern of scores characterized by a sharp departure from the diagonal line (illustrated with hypothetical data in Fig. 4*D*). The empirical Q–Q plots (Fig. 4 *E* and *F*) show clear evidence of a bimodal distribution for the targets, as predicted by a sparse distributed coding scheme (Fig. 3*B*).

We next tested whether the apparent departure from a shifted function in the Q–Q plots was statistically significant. Specifically, we tested how often a departure that large would have occurred if the target distribution were simply right-shifted and not bimodal. To do so, we conducted a bootstrap analysis of the normalized spike count data. One bootstrap analysis was performed on the data from all units combined, and a second bootstrap analysis was performed only on the data from the single units. On each iteration of the bootstrap analysis, a foil distribution was constructed by randomly sampling (with replacement) n normalized foil spike count scores (n = 7,040 for the full analysis, and n = 1088 for the single unit analysis). Next, a target distribution was constructed by independently sampling a second set of n normalized foil spike count scores and then adding a constant α to each score. The value of α was set to the average difference between the target and foil distributions in the empirical data (α = 0.06 for the full analysis, and α = 0.10 for the single units). Except for random error introduced by the sampling-with-replacement process, the target distribution created in this manner was not bimodal but instead was right-shifted relative to the scores used to represent the foil distribution. A Q–Q plot then was constructed for the hypothetical foil and target values, and a statistic was formed by computing the sum of the squared differences between the paired target and foil values. Ten thousand bootstrap iterations were run, and the proportion of these iterations in which the sum-of-squared differences was larger than the sum-of-squared differences in the empirical data was calculated. The results showed that sum-of-squared differences as large as observed in the empirical data occurred with probability 0.004 in the bootstrap trials for all of the units and with probability 0.011 in the bootstrap trials for the single units. Thus, the apparent evidence for a bimodal target distribution in the Q–Q plots is unlikely to have occurred by chance. The results of these bootstrap statistical tests were nearly identical when α

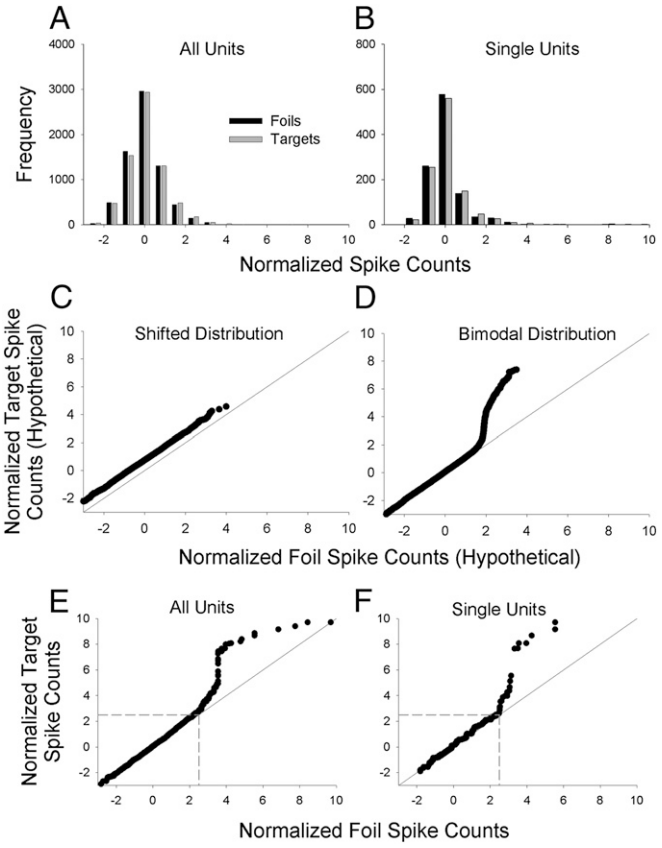

Fig. 4. Distributional analyses of normalized spike counts. (*A*) Frequency distribution of normalized spike counts for each of 7,040 normalized target spike counts (32 targets × 220 units; gray bars) and 7,040 normalized foil spike counts (32 foils × 220 units; black bars). The means of the target and foil distributions are necessarily the same as the means of the target and foil means shown in Fig 2*C* ($\bar{\mu}_{Target} = 0.08$ and $\bar{\mu}_{Foil} = 0.01$, respectively). (*B*) Single-unit frequency distribution of normalized target spike counts for each of 1,088 normalized spike counts (32 targets × 34 single units; gray bars) and 1,088 normalized foil spike counts (32 foils × 34 single units; black bars). Again, the means of the target and foil distributions are necessarily the same as the means of the target and foil means shown in Fig. 2*D* ($\bar{\mu}_{Target} = 0.11$ and $\bar{\mu}_{Foil} = 0.003$, respectively). (*C*) Hypothetical Q–Q plot illustrating the expected pattern of results for a shifted distribution. The plot is based on simulated data drawn from one Gaussian distribution with an arbitrary mean of 0.75 and SD of 1 (target distribution) vs. another Gaussian distribution with a mean of 0 and SD of 1 (foil distribution). (*D*) Hypothetical Q–Q plot illustrating a bimodal distribution. The plot is based on simulated data drawn from a bimodal mixture distribution with a mean of 0 and SD of 1 for 95% of the scores and a mean of 5 and a SD of 1 for 5% of the scores (target distribution) vs. a unimodal Gaussian distribution with a mean of 0 and SD of 1 (foil distribution). (*E* and *F*) Empirical Q–Q plots for the frequency distributions shown in *A* and *B*, respectively. The Q–Q plots suggest a bimodal distribution of target (but not foil) values. In the Q–Q plot for all units (*E*), there are 112 target values (1.6% of the total) that account for the upward trending portion of the curve that begins at ∼2.5 on the *x* and *y* axes (boundaries that are indicated by dashed gray lines). Although it is not obvious, the remaining 6,960 scores (with *x* and *y* values below 2.5) fall close to the diagonal line of equality. For the single-unit Q–Q plot (*F*), there are ∼30 target values (2.8% of the total) that account for the upward trending portion of the curve that begins at ∼2.5 on the *x* and *y* axes (indicated by dashed gray lines). The remaining 1,048 values with *x* and *y* values below 2.5 fall close to the diagonal line of equality.

was a random variable (instead of being a constant) drawn from exponential distributions with means of 0.06 and 0.10 for the full and single unit analyses, respectively.

For the Q–Q plot with all units combined (Fig. 4*E*), there are ∼112 target values (1.6% of the 7,040 target-by-unit spike counts) that account for the upward trending portion of the curve that begins at ∼2.5 on the *x* and *y* axes (indicated by the dashed gray lines). In the single-unit Q–Q plot (Fig. 4*F*), there are ∼30 target values (2.8% of the 1,088 target-by-unit spike counts) that account for the upward trending portion of the curve. For both plots, the remaining target values (>97%) appear to be coincident with the foil values. Such a small percentage of target-by-unit spike counts exhibiting elevated activity is consistent with the percentage of active units in hippocampal subfields in the rat, which have been found to range from 0.5% in dentate gyrus to 2.5% in CA1 and CA3 (23).

In Fig. 4*F*, the ∼30 elevated target-by-unit spike counts are spread across the 34 single units. Seven of the nine patients and 22 different single units are represented in the top 30 single-unit target responses. Two targets appear four times ("shallow" and "hand"), one target appears three times ("family"), two targets appear twice ("organ" and "sweat"), and 15 targets appear once (a total of 20 different words). Thus, generally speaking, each single unit was responsive to only a few targets, the pattern that is anticipated by the sparse distributed account (Fig. 3*B*). Note that the strong responses in some units to these target items cannot be attributed to random spiking activity because they occurred significantly more often in response to targets than to foils.

Despite their critical contribution to the significantly elevated single-unit *t* distribution (Fig. 2*B*), these 30 normalized spike counts were not associated with what usually would be regarded as a large increase in the absolute level of spiking activity relative to baseline. For example, for the 30 target-by-unit normalized spikes found to be elevated during the test period, the average number of raw spikes increased from 2.3 during the 800-ms prestimulus baseline period to 4.1 during the 800-ms poststimulus test period. This increase represents a difference of only 1.8 spikes, not even double the baseline count (see Fig. S2 for a representative raster plot).

Discussion

Several previous studies in humans and monkeys failed to identify any memory-related neurons in the hippocampus (e.g., 12, 15). Other studies identified only a few such neurons. For example, a continuous recognition study in the macaque (24) found that only 2.3% of hippocampal neurons (15 of 660) significantly differentiated repeated items from nonrepeated items. Similarly, we found that only 2.9% of the single units we recorded (1 of 34) significantly differentiated targets from foils. Findings such as these have been taken to mean that a small proportion of hippocampal neurons is involved in recognition, but our findings suggest otherwise. Even though only a few neurons yielded statistically significant differences in their firing rates in response to targets vs. foils, most or all of the single units we recorded appeared to be involved in recognition memory for at least some of the target items. Their involvement was suggested by the shift of the entire distribution of *t* scores to above 0 (Fig. 2 *A* and *B*). By itself, that finding does not distinguish between the three coding schemes under consideration here (localist, fully distributed and sparse distributed), although, as noted earlier, the mere fact that we detected neural evidence of a memory signal weighs against a strictly localist scheme. A further analysis of the individual target and foil spike-count distributions—an analysis that has not been performed in prior studies—showed that our data (Fig. 4 *E* and *F*) accord with predictions made by the sparse distributed coding scheme illustrated in Fig. 3*B*.

Although studies often do not find a greater-than-expected number of hippocampal neurons involved in episodic memory, some previous studies have identified hippocampal neurons that significantly differentiated targets from foils in numbers that were greater than would be expected on the basis of chance. This result has been observed in studies with humans (19, 20) and monkeys (21). However, these studies did not examine how spiking activity was distributed across test items, so it is not known if the neurons were responding to a general class of items (e.g., to all previously seen items) or to a relatively small subset of items that happened to generate strong enough responses to yield significant *t* tests. Conceivably, a bimodal distribution of spiking activity occurred in those studies as well (consistent with sparse distributed coding), in which case those findings would accord with the results reported here.

Another difference between the results of our study and those of previous studies is that we found no evidence of neurons responding to stimulus novelty (i.e., a stronger response to foils than to targets), but several prior studies have reported this effect (19–21). A possible explanation for this discrepancy is that the stimuli used in prior studies consisted of unfamiliar pictures, whereas the stimuli we used consisted of familiar words. Presumably, a novelty response is more likely to be detected when the foils are truly novel. This difference in stimulus materials (pictures vs. words) also may explain why prior studies have detected an episodic memory signal in the amygdala as well as in the hippocampus (19, 20), whereas we detected an episodic signal only in the hippocampus.

The pattern of results shown in Fig. 4 *E* and *F* indicates sparse distributed coding of episodic memory, and the pattern was the same whether the analysis was based largely on multiunits (Fig. 4*E*) or was limited to single units (Fig. 4*F*). In both cases, a bimodal distribution of spiking activity associated with recently encoded targets was observed, with the upper distribution consisting of a small percentage of recorded target activity. The most straightforward interpretation of why single units and multiunits exhibit the same pattern is that the target items are represented in the hippocampus by distributed clusters of localized neural activity. Under those conditions, a bimodal target distribution would be evident for single units and multiunits alike.

Previous work with humans has suggested that the representation of semantic memory in the hippocampus is relatively sparse (11). In addition, one study (10) found that episodic memory of a particular video clip (tested using recall) was preceded by the selective reactivation of the same neuron that had reliably responded to the presentation of that clip on an earlier test of semantic memory. This finding suggested that episodic memory (like semantic memory) might be represented by a small fraction of highly selective hippocampal neurons. However, episodic memory generally involves the retrieval of both specific episodic details (e.g., memory for context) as well as general semantic knowledge (25, 26). Thus, the neural activity measured when a clip was recalled could easily reflect the same semantic memory signal that was activated by the initial presentation of the clip. The goal of our study was to measure neural activity associated with a series of recently encoded memories that were unambiguously episodic in nature and that were formed rapidly on a single learning trial (and then tested only once).

Most neurocomputational models dating back to Marr (6) hold that episodic memory representations in the dentate gyrus/CA3 region of the hippocampus are supported by a sparse distributed neural code (Fig. 3*B*). Although our electrodes were not localized to particular hippocampal subfields, our findings nevertheless are consistent with this idea. Other evidence consistent with a sparse code in the hippocampus has been reported in studies using rats (e.g., 27, 28). However, these studies involved tasks in which memories were acquired over an extended period (allowing for the development of place fields), not tasks in which multiple memories were formed in rapid succession on a single

trial. The same is true of prior evidence for sparse coding of semantic memory in the human medial temporal lobe (11). Our findings suggest that, as has long been predicted, rapidly formed episodic memories are supported by a sparse distributed code in the human hippocampus.

Materials and Methods

Participants. The participants were nine 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 focus. The mean age of the patients was 39 y (range 19–50 y), five were female, eight were right-handed, and all had temporal lobe epilepsy. All patients provided informed consent to participate in the research using a protocol that was approved by the Institutional Review Board of St. Joseph's Hospital and Medical Center.

Materials. Stimuli for the experimental trials consisted of 192 words taken from the Medical Research Council Psycholinguistic database (29), three to seven letters in length, with a range in concreteness rating of 550–700. Half of the words (i.e., 96) were high-frequency words, and half were low-frequency words. The 192 words were randomly divided into three unique sets for each patient (64 words per set). Each set consisted of 32 targets (words that would appear on the study list and again on the recognition test) and 32 foils (words that would appear only on the recognition test) with equal representation of high- and low-frequency words. Therefore, participants could perform up to three recognition memory study/test cycles with different words.

Memory Task. Participants were told that they would be presented with a series of words and that, following the presentation of the list, their memory would be tested. During the study phase, a trial began with a fixation cross that appeared in the center of the computer screen for 750 ms, followed by the presentation of a word for 2 s. Half of the study words were presented at the top of the screen, and the other half were at the bottom, randomly selected on each trial for each participant.

During the test phase, 2 min later, 32 targets and 32 foils were presented in a randomly determined order. Test trials began with a 450-ms fixation cross, which appeared in the center of the screen, followed by a centrally presented test word. After 500 ms, a confidence rating scale appeared at the bottom of the screen, with boxes labeled from 1 (very sure new) to 8 (very sure old). To indicate their memory decision, participants clicked one of the boxes using a computer mouse. When responses were equal to or greater than 5 ("old" decision), participants then made a Remember-Know-Guess (R-K-G) judgment about their subjective memory experience. All R-K-G judgments were followed by "source" judgments, in which participants clicked one of two boxes (labeled "top" and "bottom") to indicate whether the word had been studied at the top or the bottom of the screen. The confidence ratings, R-K-G judgments, and source memory decisions are not directly relevant to the issue of sparse vs. distributed coding and are not analyzed here. All responses were self-paced.

Participants received practice trials to familiarize themselves with the task. A session (involving a single 32-item list and a 64-item recognition test) required ~20 min to complete. Four participants completed three recognition memory tests in separate sessions, one participant completed two tests in separate sessions, and the remaining four participants completed one test (Table S1). Thus, there were 18 recognition tests in all. Different sets of words were used for each test.

Microwire Recordings. Microwire implantation, recording, and spike-sorting details are described in *SI Materials and Methods*. The test period during which spike counts were recorded (200–1,000 ms after the presentation of the test item) was chosen because a previous study (9) found that selective responses of hippocampal neurons began ~300 ms after stimulus presentation and because the large majority of behavioral responses occur after 1 s. Fig. S3 shows waveforms for a single unit recorded from the hippocampus.

Because some patients received up to three recognition memory tests on three separate days, it is possible, although unlikely, that some units contributed to recordings made in more than one session (the microwires are fixed to the skull 3–4 cm from their tips). However, each recognition test involved an entirely different set of words for both targets and foils. In addition, our analyses were not based on selected units (30). Thus, our findings were not disproportionately influenced by the activity of selected units that might have been responsive to particular words.

ACKNOWLEDGMENTS. We thank the patients for participating in these experiments. We also thank Tom Albright for comments on an earlier version of this paper and Elaine Cabrales for technical assistance. This work was supported by the Medical Research Service of the Department of Veterans Affairs, National Institute of Mental Health Grant 24600, National Institute on Deafness and Other Communications Disorders Grant 1R21DC009781, the Barrow Neurological Foundation, and the University of California, San Diego Kavli Institute for Brain and Mind.

1. Squire LR, Zola-Morgan S (1991) The medial temporal lobe memory system. *Science* 253(5026):1380–1386.
2. Eichenbaum H, Cohen NJ (2001) *From Conditioning to Conscious Recollection: Memory Systems of the Brain* (Oxford Univ Press, New York).
3. Squire LR (1992) Memory and the hippocampus: A synthesis from findings with rats, monkeys, and humans. *Psychol Rev* 99(2):195–231.
4. Squire LR, Wixted JT (2011) The cognitive neuroscience of human memory since H.M. *Annu Rev Neurosci* 34:259–288.
5. Rolls ET, Treves A (2011) The neuronal encoding of information in the brain. *Prog Neurobiol* 95(3):448–490.
6. Marr D (1971) Simple memory: A theory for archicortex. *Philos Trans R Soc Lond B Biol Sci* 262(841):23–81.
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(3):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(4):611–646.
9. Quiroga RQ, Reddy L, Kreiman G, Koch C, Fried I (2005) Invariant visual representation by single neurons in the human brain. *Nature* 435(7045):1102–1107.
10. 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(5898):96–101.
11. Waydo S, Kraskov A, Quian Quiroga R, Fried I, Koch C (2006) Sparse representation in the human medial temporal lobe. *J Neurosci* 26(40):10232–10234.
12. Heit G, Smith ME, Halgren E (1988) Neural encoding of individual words and faces by the human hippocampus and amygdala. *Nature* 333(6175):773–775.
13. Heit G, Smith ME, Halgren E (1990) Neuronal activity in the human medial temporal lobe during recognition memory. *Brain* 113(Pt 4):1093–1112.
14. Brown MWF, Wilson FA, Riches IP (1987) Neuronal evidence that inferomedial temporal cortex is more important than hippocampus in certain processes underlying recognition memory. *Brain Res* 409(1):158–162.
15. Riches IP, Wilson FAW, 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(6):1763–1779.
16. Xiang JZ, Brown MW (1998) Differential neuronal encoding of novelty, familiarity and recency in regions of the anterior temporal lobe. *Neuropharmacology* 37(4-5):657–676.
17. Fried I, MacDonald KA, Wilson CL (1997) Single neuron activity in human hippocampus and amygdala during recognition of faces and objects. *Neuron* 18(5):753–765.
18. 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(10):1654–1662.
19. Rutishauser U, Mamelak AN, Schuman EM (2006) Single-trial learning of novel stimuli by individual neurons of the human hippocampus-amygdala complex. *Neuron* 49(6):805–813.
20. 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(1):329–334.
21. Jutras MJ, Buffalo EA (2010) Recognition memory signals in the macaque hippocampus. *Proc Natl Acad Sci USA* 107(1):401–406.
22. Chambers JM, Cleveland WS, Kleiner B, Tukey PA (1983) *Graphical Methods for Data Analysis* (Wadsworth, Pacific Grove, CA).
23. O'Reilly RC, McClelland JL (1994) Hippocampal conjunctive encoding, storage, and recall: Avoiding a trade-off. *Hippocampus* 4(6):661–682.
24. Rolls ET, Cahusac PMB, Feigenbaum JD, Miyashita Y (1993) Responses of single neurons in the hippocampus of the macaque related to recognition memory. *Exp Brain Res* 93(2):299–306.
25. Hemmer P, Steyvers M (2009) Integrating episodic memories and prior knowledge at multiple levels of abstraction. *Psychon Bull Rev* 16(1):80–87.
26. Nelson AB, Shiffrin RM (2013) The co-evolution of knowledge and event memory. *Psychol Rev* 120(2):356–394.
27. Jung MW, McNaughton BL (1993) Spatial selectivity of unit activity in the hippocampal granular layer. *Hippocampus* 3(2):165–182.
28. Leutgeb JK, Leutgeb S, Moser MB, Moser EI (2007) Pattern separation in the dentate gyrus and CA3 of the hippocampus. *Science* 315(5814):961–966.
29. Coltheart M (1981) The MRC psycholinguistic database. *Q J Exp Psychol A* 33:497–505.
30. Steinmetz PN, Thorp C (2013) Testing for effects of different stimuli on neuronal firing relative to background activity. *J Neural Eng* 10(5):056019.

Supporting Information

Wixted et al. 10.1073/pnas.1408365111

SI Materials and Methods

Microwire Implantation and Recordings. Electrode implantation was performed stereotactically (Medtronic StealthStation), and the position was confirmed by coaligning the postoperative CT or MRI (using the Statistical Parametric Mapping toolkit, www.fil.ion.ucl.ac.uk/spm/) with the preoperative structural MRI. This procedure localizes the tips of the microwires to within 2 mm (1). Bundles of nine platinum-iridium microwires 38 μm in diameter (California Fine Wire) were introduced through a lumen within the clinical intraparenchymal electrode during surgery. The implantation sites were chosen according to clinical criteria, limiting the potential recording sites. For the nine patients studied here, the sites included the hippocampus and amygdala, bilaterally. In the hippocampus, the wires usually were targeted to be in the midbody of the hippocampus, just behind the head of the hippocampus, opposite the apex of the cerebral peduncle. All patients received a postimplantation CT scan as a check to ensure there was no bleeding after the operation. These scans do not have sufficient resolution to resolve hippocampal subfields.

The extracellular potentials corresponding to single-unit activity and multiunit activity were recorded from the tips of the microwires. At each site, the potential difference between eight of the microwires was recorded relative to a ninth microwire in the same bundle using a headstage amplifier of custom design. This amplifier provided a 400× gain and was connected to signal-conditioning electronics and analog-to-digital converters (model DT9834; Data Translation) via a 1-m tether cable. Each signal channel was preconditioned with a high-pass filter (0.5-Hz corner) followed by a 10-kHz antialiasing filter and a computer-controlled 1–16× adjustable gain amplifier (custom-designed signal-conditioning board). The conditioned signal was digitized at 29,412 Hz with 16-bit resolution.

Data Analysis. Possible action potentials (APs) were detected by filtering twice (forward and backward, acausally) with a 24th-order digital IIR bandpass filter, 300–3,000 Hz, with a −100-dB stop band and −12-dB notches at 1, 2, and 3 kHz followed by a two-sided threshold detector (threshold 2.8 times each channel's SD) to identify AP times. The original signal then was high-pass filtered (100 Hz, single-pole Butterworth, applied causally) to capture the shape of the AP waveform in windows of 32 samples (1.1 ms) with the absolute peak value aligned at the ninth sample.

Because more than one neuron may be recorded near any given electrode, APs were grouped into several clusters of similar waveform shape. This clustering was performed using the open-source clustering program KlustaKwik (Klustakwik.sf.net), which is a modified implementation of the Classification Expectation-Maximization clustering algorithm (2). The first principal component of all waveform shapes recorded from a channel was the waveform feature used for sorting. After sorting, each cluster was graded as being noise, multiunit activity, or single-unit activity based on criteria including the waveform shape, size of the waveform relative to noise, evidence of a refractory interval, and lack of powerline interference, as previously described (3). Fig. S3 illustrates a typical cluster of single-unit activity after spike sorting.

SI Results

In the hippocampus, the mean normalized spike count in response to targets (μ_{Target}) averaged across the nine patients was 0.11 [which marginally exceeded the baseline value of 0, $t(8) = 2.18$, $P = 0.061$], and the mean normalized spike count in response to foils (μ_{Foil}) averaged across the nine patients was 0.04 (which was not significantly different from 0, $P = 0.43$). As noted in the main text, however, the average difference score, D', was significantly greater than 0 ($P < 0.01$). In the amygdala, the mean normalized spike count in response to targets (μ_{Target}) averaged across the nine patients was 0.13, and the mean normalized spike count in response to foils (μ_{Foil}) averaged across the nine patients was 0.11. Neither value was significantly different from 0, and the average difference score, D', also was not significantly different from 0 ($P > 0.16$).

Our analyses were based on all recorded clusters, not on a subset of clusters that were deemed to be task-relevant based on any indication of responsiveness to the study or test stimuli. However, because such subset analyses have become common practice, we also asked whether a higher percentage of significant clusters would be identified when the analysis was limited to only those clusters deemed to be responsive based on a significant change in average spike counts (relative to baseline) across the 32 items presented during the study phase. Using an alpha level of 0.10 to identify responsive units, we found that 15 of 205 clusters exhibited a significant change in firing relative to the prestimulus baseline during the initial presentation of the list items, a number that is not greater than would be expected from chance (expected = 0.10 × 205 = 20.5) (Although 220 clusters were analyzed during the recognition test, only 205 of those clusters yielded spike counts during study). Thus, we did not find evidence of neurons that were generally responsive to study items.

The quantile–quantile (Q–Q) plots for all clusters combined (Fig. 4*E*) and for the single units considered separately (Fig. 4*F*) show points associated with normalized spike counts of 10 or less. Five of the 1,088 values from single units (0.5%) exceeded 10 and were deemed to be too large to reflect true responses (and therefore were excluded from analysis). Some of the values less than 10 also may reflect measurement error. For example, the apparent return to the diagonal in Fig. 4*E* reflects the fact that, on rare occasions, extremely high values (e.g., 7 SDs or more above baseline firing) occurred for both targets and foils with approximately equal frequency. It is not clear how these extremely high scores should be interpreted, but they may simply reflect measurement error (which would occur for targets and foils with equal frequency).

The overall pattern of results is consistent with the bimodal target distribution predicted by the sparse distributed account, but the same Q–Q pattern could be generated by a continuous target distribution that is extremely skewed compared with the foil distribution. An extremely skewed distribution would reflect the strong response generated by a few items in the tail of the target distribution. This distribution, too, would correspond to sparse distributed coding. That is, no matter whether the data reflect a bimodal target distribution or an extremely skewed unimodal distribution, the results are consistent with the idea that a small percentage of targets generated a strong response in a small fraction of hippocampal neurons. Note that visual evidence for a bimodal target distribution was apparent when normalized spike counts were examined, but not when raw spike counts were examined. This is not surprising because large differences in the baseline firing rates of the recorded units swamp any evidence of the modestly elevated firing that occurred in response to a few targets for each unit.

1. Mehta AD, et al. (2005) Frameless stereotactic placement of depth electrodes in epilepsy surgery. *J Neurosurg* 102(6):1040–1045.
2. Celeux G, Govaert G (1995) Gaussian parsimonious clustering models. *Pattern Recognition* 28(5):781–793.
3. 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(1):016001, 10.1088/1741-2560/10/1/016001.

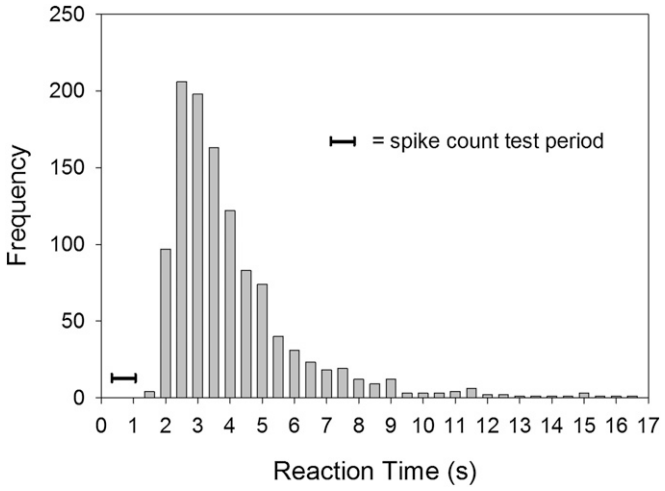

Fig. S1. Reaction time (RT) frequency distribution pooled over 18 recognition tests, where RT = the interval between the onset of a test item and the mouse click indicating the confidence rating for that test item. Spike counts were recorded from 200–1,000 ms after stimulus onset, before all overt responses.

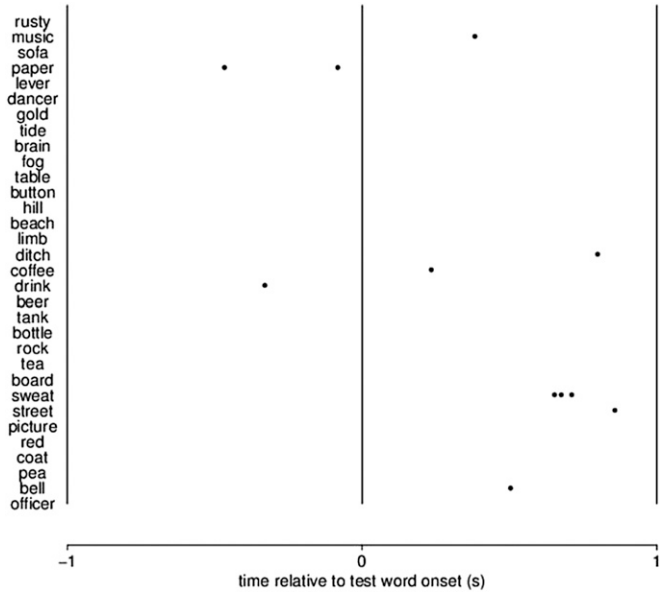

Fig. S2. Representative example of a raster plot of prestimulus and poststimulus activity for one of the single units that was responsive to one target. Although this unit yielded a strong normalized response to the target item "sweat" ($z = 7.67$), it yielded only a moderately strong response when measured in absolute terms. More specifically, three spikes occurred in response to that target during the 800-ms test period, reflecting a modest elevation in firing for a unit that had a baseline spike count mean and SD of 0.11 and 0.36, respectively. This result is fairly typical of the 30 targets that yielded notably elevated (off-diagonal) responses evident in the single-unit Q–Q plot (Fig. 4*F*). These results suggest that episodic memory may not be characterized by the kind of conspicuously elevated single-unit responding (in absolute terms) that is observed when single units exhibit an elevated response to repeatedly presented images of famous people and landmarks (e.g., ref. 1). It is important to emphasize that, using our single-presentation design, no single instance of bursting that coincides with the presentation of a test item (such as the response to "sweat" illustrated here) can be attributed confidently to the presentation of that item. The raster plot is intended only to illustrate the kind of bursting that occurs significantly more often in response to targets than to foils, not to suggest that the bursting that occurred in conjunction with the word "sweat" on this particular trial was necessarily triggered by the presentation of that word.

1. Quiroga RQ, Reddy L, Kreiman G, Koch C, Fried I (2005) Invariant visual representation by single neurons in the human brain. *Nature* 435(7045):1102–1107.

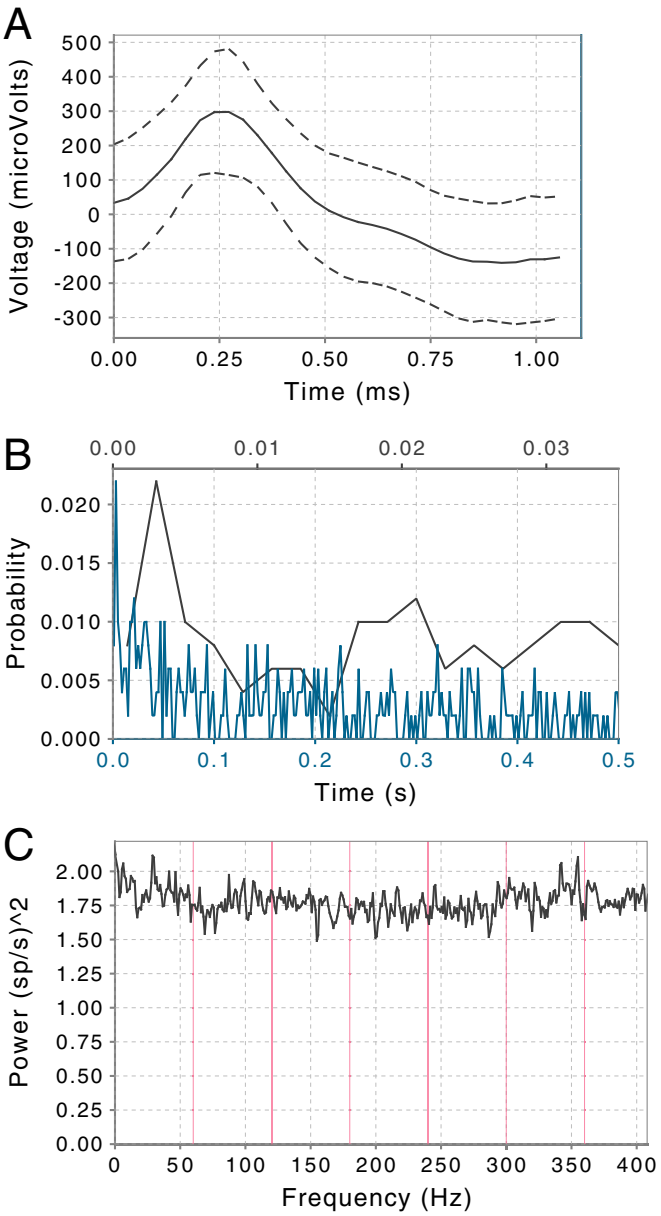

Fig. S3. Waveforms in a cluster identified as single-unit activity after sorting. This cluster was recorded from the left hippocampus. (*A*) The *y* axis shows the waveform shape on a scale of −300 to 500 μV. Dashed lines indicate ±1 SD at each sample point. (*B*) Distribution of interspike intervals on two time scales. The *y* axis shows the probability of interval. On the *x* axis the lower (blue) trace shows the duration of the interval for a range 0.00–0.5 s. The upper (black) trace shows the duration of interval for a range 0.00–0.035 s. Less than 2% of the interstimulus intervals are shorter than 3 ms. (*C*) Power spectral density of event times. The *y* axis shows power spectral density in events^2/Hz. The *x* axis shows the frequency in Hertz. The vertical magenta lines indicate primary and harmonics of the powerline frequency (60 Hz).

Table S1. Behavioral performance measures and number of clusters recorded from each patient

Patient	No. of sessions	Hit rate	False-alarm rate	% correct	d'	Multiunits	Single units
P1	3	0.55	0.17	0.69	1.23	29	13
P2	2	0.45	0.28	0.59	0.53	57	2
P3	1	0.69	0.44	0.63	0.65	11	6
P4	3	0.45	0.21	0.62	0.69	40	4
P5	3	0.58	0.06	0.76	1.79	20	5
P6	1	0.09	0.03	0.53	0.54	5	1
P7	3	0.54	0.08	0.73	1.58	9	2
P8	1	0.06	0.02	0.52	0.01	7	0
P9	1	0.19	0.03	0.58	0.98	8	1
Average	2.0	0.40	0.15	0.63	0.89	20.7	3.8

For patients who completed more than one session, the behavioral measures were computed separately for each session and then averaged across sessions.