Supplemental Figure 1. Empirical QQ Plots of the Normalized Spike-Count Distributions at Retrieval Partitioned by Neuronal Excitability at Encoding with the Top Percentage of Data Removed

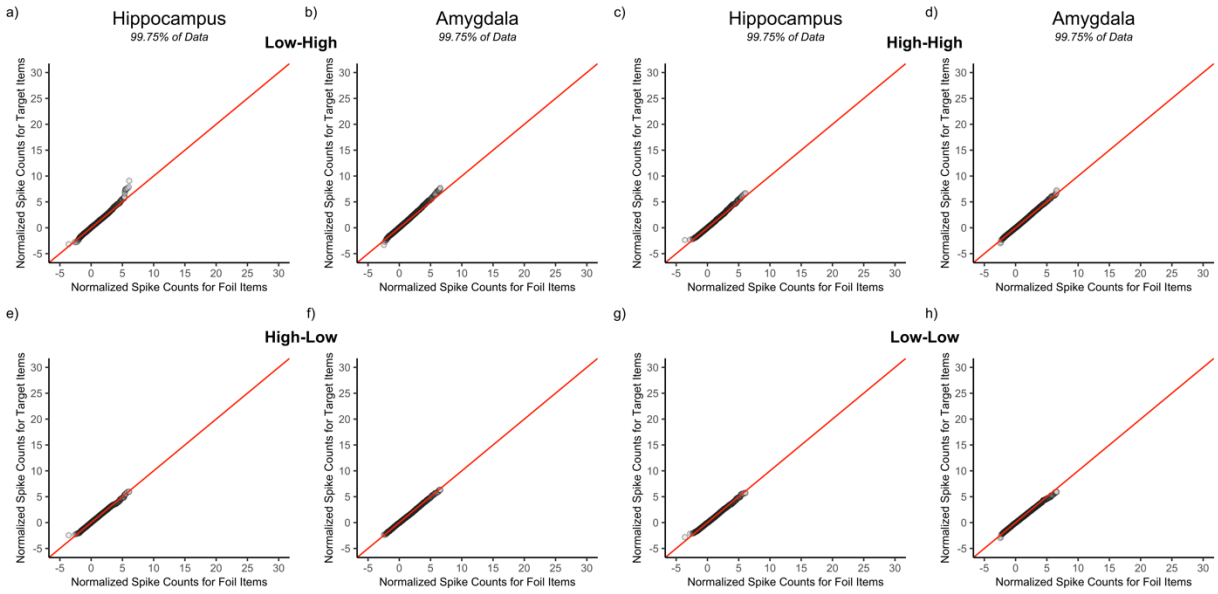

Note. QQ plots visualizing the shapes of the target-by-neuron and novel-by-neuron item distributions with the top 0.25% of the recordings from both distributions removed. Compare the QQ plots in this figure to the QQ plots of 100% of data plotted in Figure 4. For the hippocampus, most spike counts fell densely on the diagonal line with a sharp deflection upward towards the y-axis in Figure 4a, indicating the normalized spike count distribution for targets was more skewed than the distribution for novel items. In this figure, panel (a) demonstrates the deflection disappeared after removing the top 0.25% of data, indicating that relatively few hippocampal neurons fired strongly in response to targets mostly associated with excitability at encoding (Low-High) when compared to novel items. No differences between the target-by-neuron and novel-by-neuron distributions were observed for the amygdala (Figure 4b). None of the QQ plots indicated a difference in skewness between targets and novel item distributions for the subsets of targets in which spiking decreased (Panel e & f: High-Low), remained high (Panel c & d: High-High), or remained low (Panel g & h: Low-Low) for either the hippocampus or the amygdala, a similar pattern shown in Figure 4.

Supplemental Table 1. Statistical Results for Mean, Standard Deviation, and Kurtosis of the Target vs. Foil Distributions at Retrieval Partitioned by the Pattern of Target Firing at Encoding

		Mean			**Mean**		
	Interaction	Hippocampus			Amygdala		
	p	Targets	Novel	*p*	Targets	Novel	*p*
Low-High	0.597	0.13	0.08	0.001**	0.21	0.16	<0.0001***
High-High	0.077	0.11	0.08	0.037	0.16	0.16	0.874
High-Low	0.063	0.06	0.08	0.252	0.11	0.16	<0.0001***
Low-Low	0.104	0.04	0.08	0.013	0.10	0.16	<0.0001***
		Standard Deviation			**Standard Deviation**		
	Interaction	Hippocampus			Amygdala		
	p	Targets	Novel	*p*	Targets	Novel	*p*
Low-High	0.031	1.41	1.18	<0.0001***	4.0	1.23	0.0045**
High-High	0.476	1.21	1.18	0.541	1.26	1.23	0.400
High-Low	0.364	1.14	1.18	0.426	1.17	1.23	0.095
Low-Low	0.218	1.13	1.18	0.356	1.14	1.23	0.007
		Kurtosis			**Kurtosis**		
	Interaction	Hippocampus			Amygdala		
	p	Targets	Novel	*p*	Targets	Novel	*p*
Low-High	0.014	114.6	29.9	0.012*	28.7	34.9	0.736
High-High	0.505	17.7	29.9	0.694	28.9	34.9	0.703
High-Low	0.435	12.2	29.9	0.607	18.7	34.9	0.388
Low-Low	0.227	32.9	29.9	0.924	16.2	34.9	0.183

Note. No significant interactions were present for a given statistical moment between the target-vs.-foil normalized spike count distributions in one region (e.g., the hippocampus) compared to the corresponding difference the other region (e.g., the amygdala) according to bootstrap tests (k=10,000, $p < .05$, Bonferroni corrected). Within the hippocampus, a significant difference between the target vs. foil mean, standard deviation, and kurtosis was observed only for targets that were associated with excitability at encoding (Low-High; $p < .05$, Bonferroni corrected). Within the amygdala, a significant difference was present for the difference in means for the targets vs foils for those targets associated with excitability at encoding (Low-High), decreased excitability (High-Low), and those that remained low (Low-Low). Additionally, within the amygdala, a significant difference was present for the difference in standard deviation for the targets vs foils for only those targets associated with excitability at encoding (Low-High). The corresponding analysis for skewness is reported in Table 2 of the main text. Bonferroni corrected: * = $p < .05$, ** = $p < .01$, *** = $p < .001$

Supplemental Figure 2. Empirical QQ Plots for Remembered Targets vs. Foils and Forgotten Targets vs. Foils with the Top Percentage of Data Removed

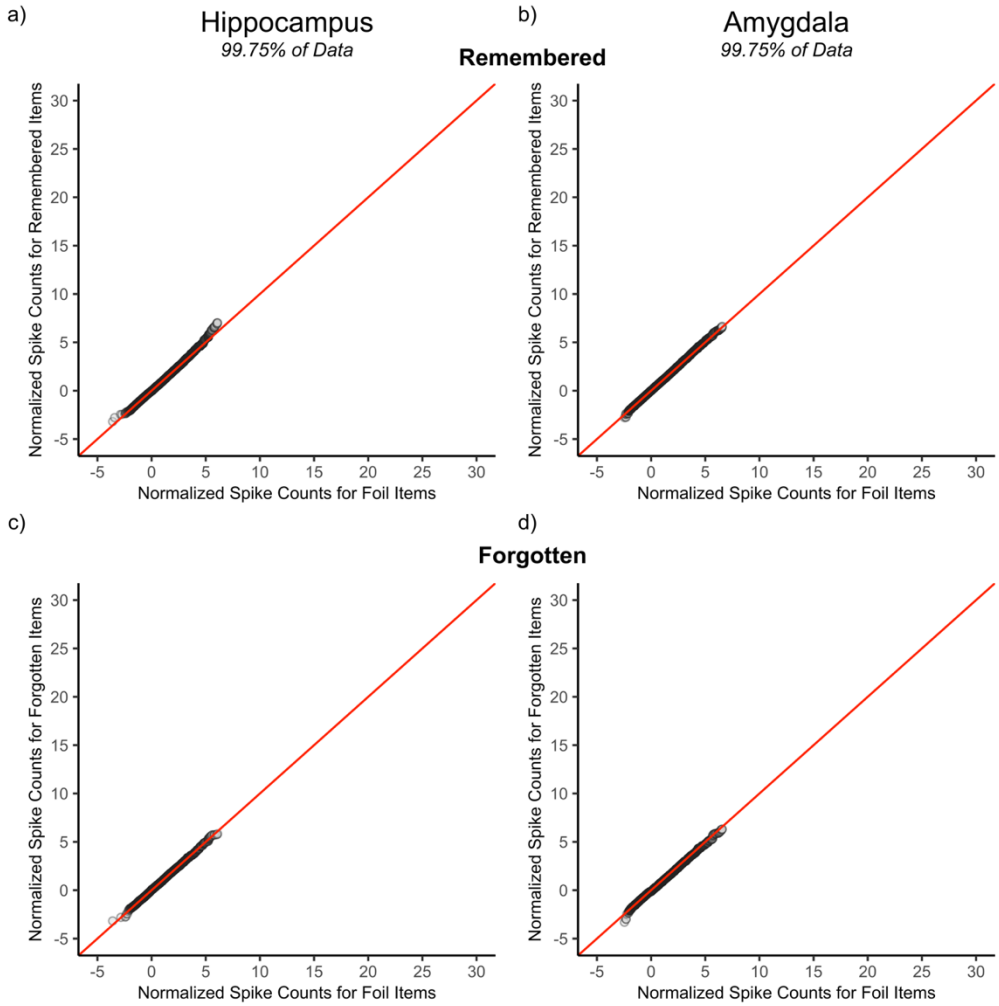

Note. QQ plots visualizing the shapes of the target-by-neuron and novel-by-neuron item distributions with the top 0.25% of the recordings from both distributions removed. Compare the QQ plots in this figure to the QQ plots of 100% of data plotted in Figure 5. For the hippocampus, most spike counts fell densely on the diagonal line with a sharp deflection upward towards the y-axis in Figure 5a, indicating the normalized spike count distribution for targets was more skewed than the distribution for novel items. In this figure, panel (a) demonstrates the deflection disappeared after removing the top 0.25% of data, indicating that relatively few hippocampal neurons fired strongly in response to targets when compared to novel items. No difference between the forgotten target-by-neuron and novel-by-neuron distributions was observed for the amygdala (Figure 5b). None of the QQ plots indicated a difference in skewness between targets and novel item distributions for forgotten items for either the hippocampus or the amygdala, a similar pattern shown in Figure 5c and 5d.

Supplemental Figure 3. Empirical QQ Plots for Remembered vs. Forgotten Targets with the Top Percentage of Data Removed

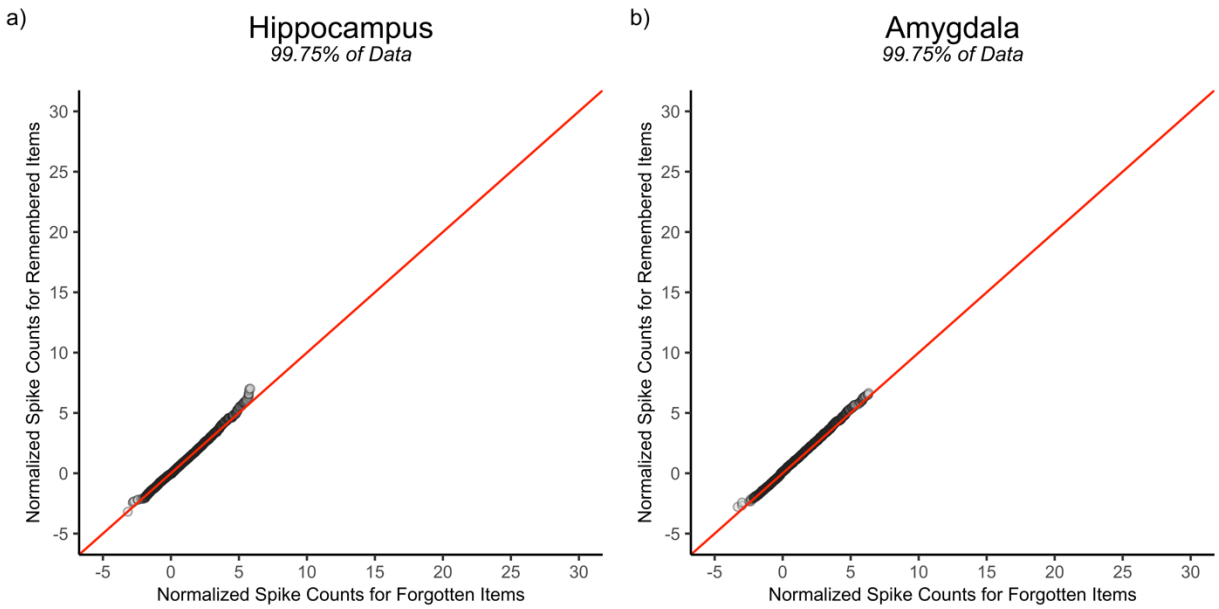

Note. QQ plots visualizing the shapes of the remembered target-by-neuron and forgotten target-by-neuron item distributions with the top 0.25% of the recordings from both distributions removed. Compare the QQ plots in this figure to the QQ plots of 100% of data plotted in Figure 6. For the hippocampus, most spike counts fell densely on the diagonal line with a sharp deflection upward towards the y-axis in Figure 6a, indicating the normalized spike count distribution for remembered targets was more skewed than the distribution for forgotten targets. In this figure, panel (A) demonstrates the deflection disappeared after removing the top 0.25% of data, indicating that relatively few hippocampal neurons fired strongly in response to targets that were remembered when compared to forgotten items. No difference between the remembered target-by-neuron and forgotten target-by-neuron distributions was observed for the amygdala (Figure 6b).

Supplemental Figure 4. Empirical QQ Plots for Targets Partitioned by Neuronal Excitability at Encoding as a Function of Subsequent Memory

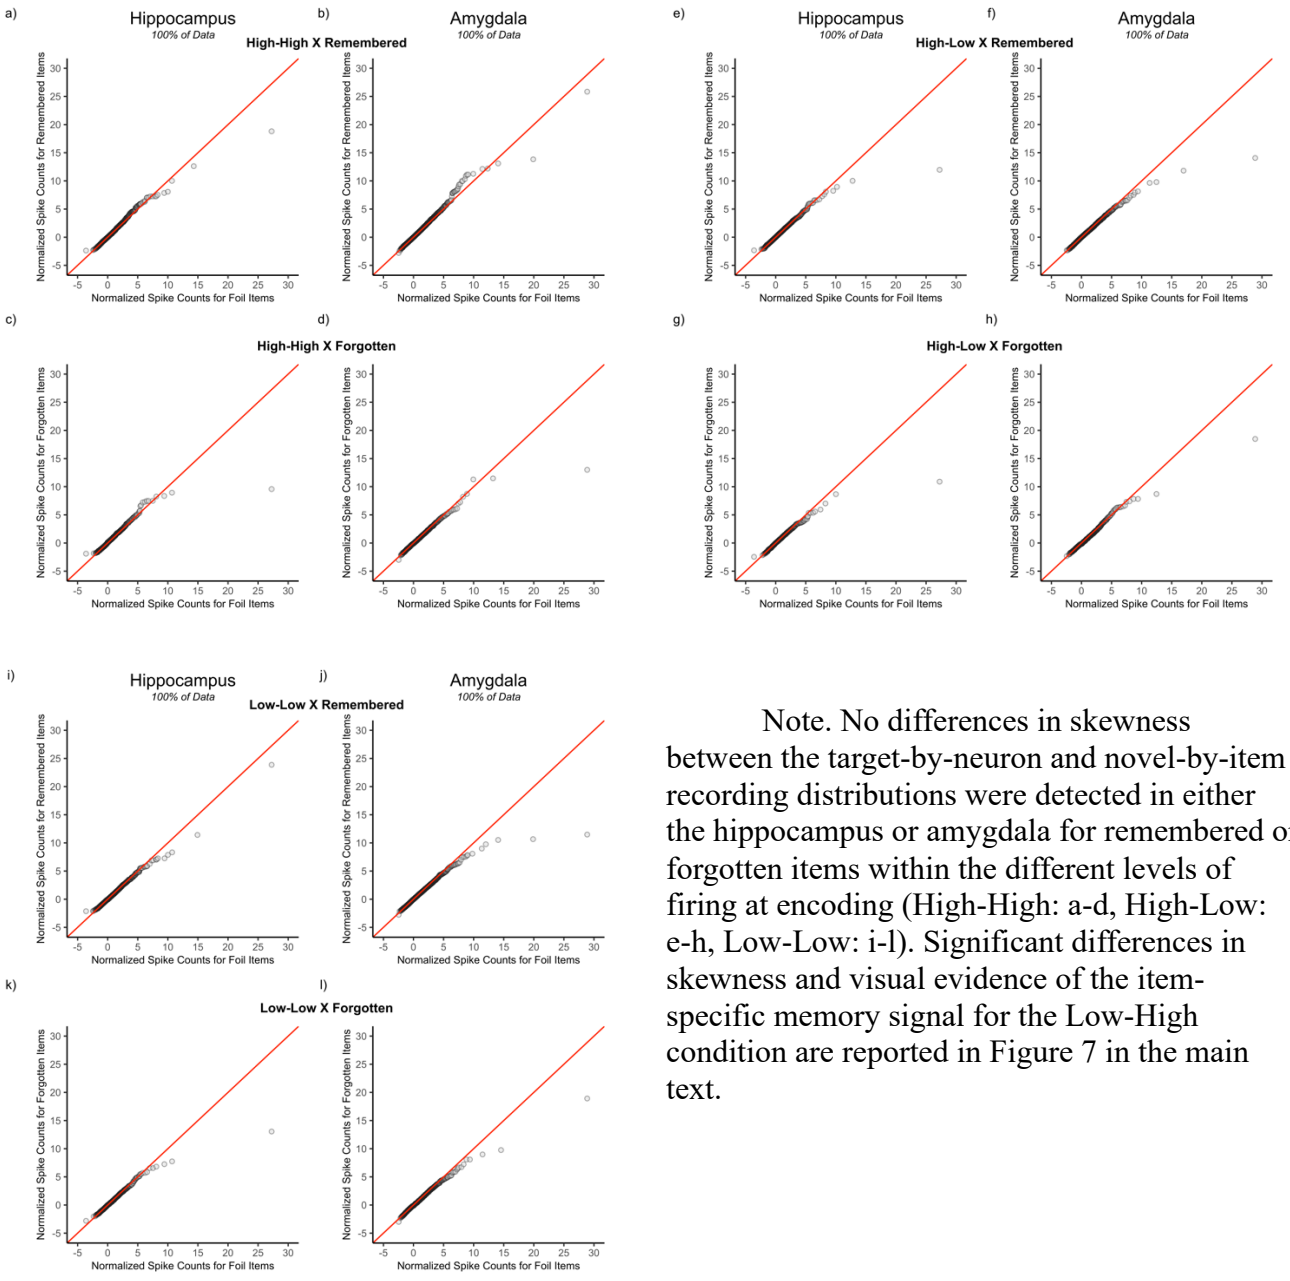

Note. No differences in skewness between the target-by-neuron and novel-by-item recording distributions were detected in either the hippocampus or amygdala for remembered or forgotten items within the different levels of firing at encoding (High-High: a-d, High-Low: e-h, Low-Low: i-l). Significant differences in skewness and visual evidence of the item-specific memory signal for the Low-High condition are reported in Figure 7 in the main text.

Supplemental Figure 5. Empirical QQ Plots for Targets Partitioned by Neuronal Excitability at Encoding as a Function of Subsequent Memory with the Top Percentage of Data Removed

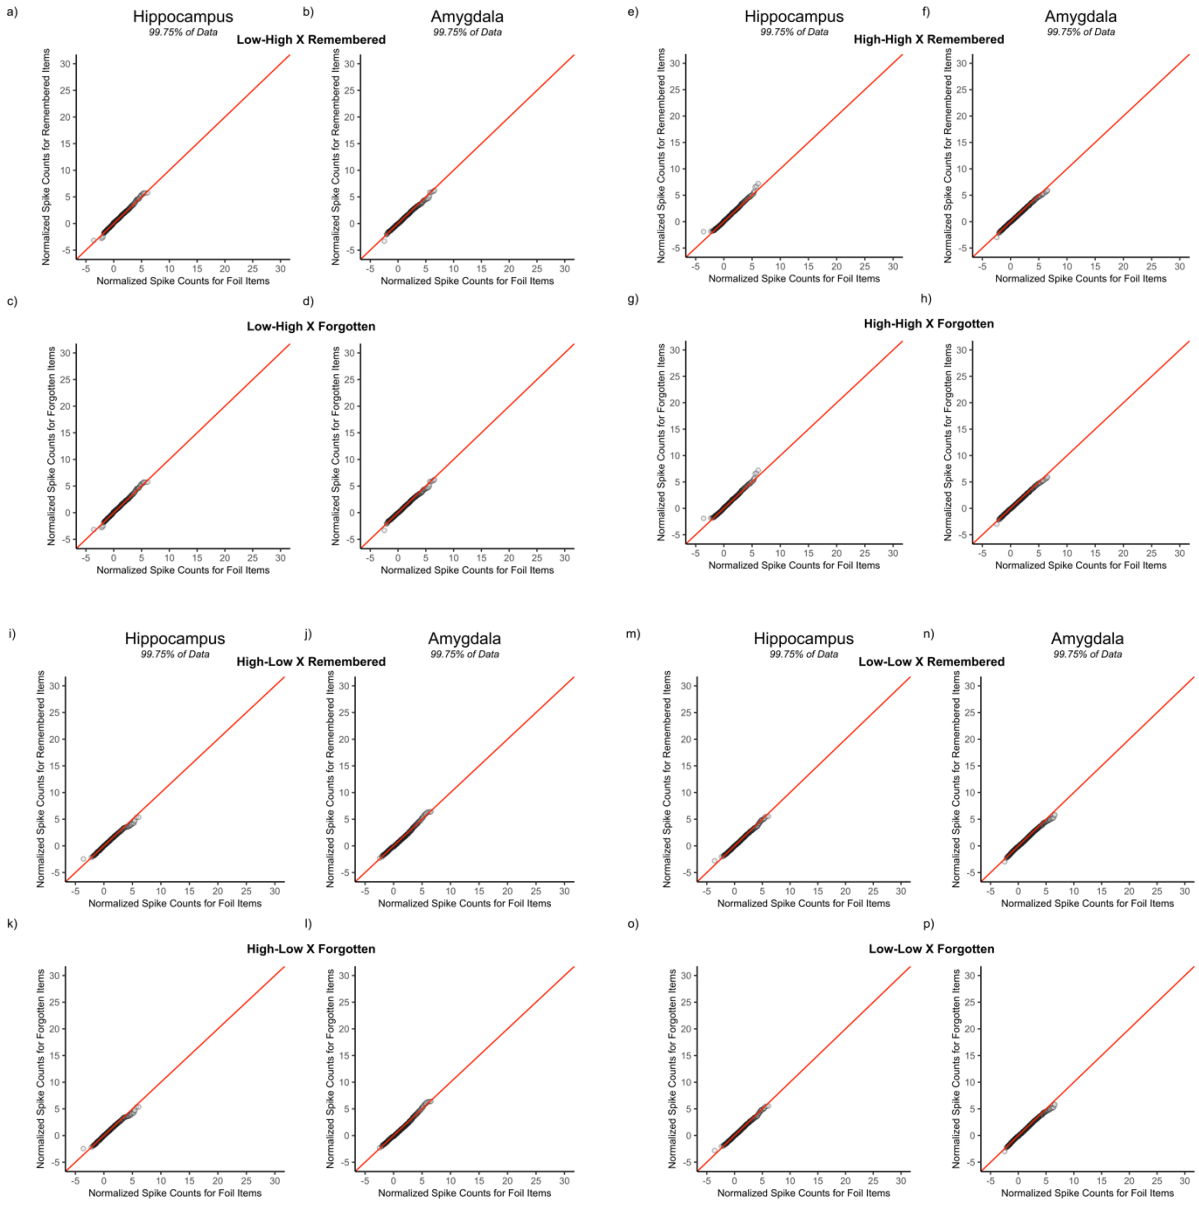

Note. In this figure, panel (a) demonstrates the deflection disappeared after removing the top 0.25% of data, indicating that relatively few hippocampal neurons fired strongly in response to targets that were remembered and were also associated with excitable neurons at encoding (compare to Figure 7a) when compared to novel items. No differences between the remembered target-by-neuron or forgotten target-by-neuron distributions and the novel-by-neuron distribution were observed for the hippocampus or the amygdala, which persisted when removing a small fraction from both distributions (b-p).

Generic Memory Signal

The generic recognition memory signal is (by definition) not item specific and instead consists of a difference in the average firing rate to old items compared to new items on the recognition test. This difference might be observed either at the level of single neurons (e.g., a neuron with a significantly higher firing rate to old items vs. new items) or at the population level (aggregated measurements of all recorded single neurons exhibiting a higher rate of responding to old items vs. new items). Bootstrapping analyses (see above) determined whether a higher firing rate for target or novel items was observed in the hippocampus or amygdala. For each single neuron, bootstrap trials (k=10,000 trials) compared the mean normalized firing rate for target and novel items. Each bootstrap trial 1) combined the target (N_{target}) and novel (N_{novel}) item normalized spike counts 2) randomly sampled with replacement from that combined set of measurements to generate a new set of N_{target} normalized "target" spike counts and a new set of N_{novel} normalized "novel" spike counts, then 3) calculated the difference in the means between the two sampled distributions, resulting in 10,000 mean different scores from the bootstrapped samples. The proportion of these trials in which the absolute value of the difference in means between the "target" and "novel" distributions was greater than or equal to the observed difference in the original target and novel distributions determined the P value. Significance was defined as $p < 0.05$ (two-tailed), unless otherwise specified.

As defined above, a difference in neural activity in response to previously studied items (targets) or novel items (novel) items is what we refer to as the "generic" episodic memory signal. The generic signal can be measured at the level of individual neurons or at the population level. Individual neurons that exhibit this property (i.e., firing rates that differ for targets vs. novel) are considered to be memory-selective. Memory-selective neurons that increase in spike count in response to targets compared to novel items are "repetition detectors", and those that increase in spike count in response to novel items compared to targets are "novelty detectors". The generic memory signal in single neurons has been observed in prior work when examining all responses irrespective of the behavioral response (Urgolites et al., 2022) or when excluding error trials (Rutishauser et al., 2010, 2015). We identified the generic memory signal in single neurons in the hippocampus and in the population firing of the amygdala.

Each recorded neuron across was tested for the generic memory signal was tested across all patients and session. The test compared spikes for targets vs. novel items during the poststimulus window of 200ms to 1-second after the stimulus presentation (see Methods; replication of Urgolites et al., 2022). Using this poststimulus window for all trials irrespective of behavioral response, the generic signal was not detected. More specifically, of the 736 hippocampal neurons, 32 (4.3%) were memory-selective, a proportion not significantly greater than the 36.8 expected by chance ($\alpha = 0.05$; $736 \times 0.05 = 36.8$). Among these, 18 neurons were repetition detectors, and 14 neurons were novelty detectors. Of the 1043 amygdala neurons, 50 (3.8%) were memory-selective, not significantly exceeding the 52.2 neurons expected by chance ($\alpha = 0.05$; $1043 \times 0.05 = 52.2$). Among these, 21 neurons were repetition detectors, and 29 neurons were novelty detectors.

Next, we aimed to replicate the identification of the generic memory signal previously reported with a subset of this dataset (Rutishauser et al., 2015). Following a similar approach, we extended our analysis to include spike counts within an extended poststimulus window of 200ms to 1.7-seconds after image onset, while also excluding error trials. Using this method, the generic

memory signal was detected in 58 of the 736 hippocampal neurons (7.9%), significantly exceeding the number of neurons expected at chance ($\alpha = 0.05$; $p = 0.0005$). Among these neurons, 34 were repetition detectors and 24 were novelty detectors. Of 1043 amygdala neurons, 50 (5.5%), were memory-selective, which was not significantly greater than chance. Of these, there were 23 repetition detectors and 27 novelty detectors. Therefore, we replicated the identification of the generic memory signal when examining correct responses during an extended post-stimulus window and found it was selective to the hippocampus.

Additionally, the difference in normalized response count for target and novel items in the hippocampus and amygdala was tested across all patients and sessions, for both poststimulus windows and trial types of the single neuron analyses of the generic memory signal. For the replication of Urgolites et al., 2022, mean firing in the amygdala was significantly greater for novel compared to target items (p=0.0374, mean target=0.14,mean novel=0.16,SD=0.005; Table 1). No significant difference in population firing was detected in the hippocampus between target and novel items (p=0.667, mean target=0.08,mean novel=0.08,SD=0.006; Table 1). For the replication of Rutishauser et al., 2015, mean firing in the amygdala was significantly greater for novel compared to target items (p=0.0027, mean hit = 0.15, mean correct rejection = 0.18). No significant difference in population firing was detected in the hippocampus between target and novel items (p=0.314, mean hit=0.07,mean correct rejection=0.06). We identified the generic memory signal in single neurons in the hippocampus and in the population firing of the amygdala

Evidence against the notion newly formed episodic memories are encoded via overlapping neural assemblies

The idea that episodic memories are represented by overlapping (not pattern-separated) neural assemblies was based on work involving “concept cells” (Quiroga, 2012). As noted by Quian Quiroga (2020), concept cells fire to a particular concept, such as a famous person (e.g., James Brolin) and not to other concepts. A concept cell is activated whether participants are looking at pictures of the person, reading the person’s name, or “…even when recalling or thinking about the person”(p. 1002). In short, any stimulus that triggers the thought of the famous person activates the concept cell. Concept cells reflect semantic memory, but their apparent role in episodic memory has been taken as evidence against pattern separation. For example, if the “preferred” (P) concept like James Brolin is experimentally associated with a different “non-preferred” (NP) stimulus like the Eiffel Tower in a single-trial episodic memory task (e.g., a picture of James Brolin standing next to the Eiffel Tower), some concept neurons immediately expand their tuning in such a way as to now also fire to the NP stimulus (Ison et al., 2015). Critically, immediately upon learning the association, the neuron continued responding to Brolin and now also fired to Brolin standing near the Eiffel Tower and to the Eiffel Tower without Brolin. These neurons were referred to as “pair-coding” neurons. The existence of these neurons was interpreted as being inconsistent with the notion that episodic memories are coded in distinct, pattern-separated assembles because the episodically learned associations were encoded by expanding the tuning of the neurons initially responding to a concept, not by recruiting new neurons to form non-overlapping (pattern-separated) representations.

However, Ison et al. (2015) pointed out another possible interpretation of their findings that did not require the assumption that the concept neuron episodically expanded its tuning at all. As they put it: “Two possible mechanisms can in principle account for the increased response to the NP stimuli after learning. On the one hand, neurons can rapidly change their tuning and start firing to the NP stimuli directly—that means, a neuron originally encoding the P stimulus starts encoding the NP stimulus after learning—in which case, the time courses of both P and NP signals are expected to be similar. On the other hand, the NP stimuli can act as a cue to evoke the representation of (and in turn the neuron’s firing to) the P stimuli” (p. 226). According to the second interpretation, the James Brolin concept neuron responded to non-preferred Eiffel Tower stimulus after associative learning only because it now triggered the thought of James Brolin. As noted above, a concept neuron is defined as a neuron that responds whenever any stimulus triggers a thought of its preferred concept. According to this interpretation, the fact that the James Brolin neuron now responds to the Eiffel Tower only means that it is still behaving as a concept neuron (not that it has expanded its turning based on an episodically learned association).

To distinguish between these two possible mechanisms, Ison et al. (2015) analyzed the response onset latencies of the pair-coding neurons to the P and NP stimuli. If the concept neuron expanded its tuning based on the newly learned episodic association, similar latency onsets for the P and NP stimuli should be observed. If the concept neuron did not expand its tuning but is instead activated by NP stimulus because it triggers the thought of the P stimulus, a longer latency onset should be observed for the NP stimulus compared to the P stimulus. Out of 21 pair-coding neurons, 13 showed no significant difference in latency between the P and NP stimuli (these were labeled Type 1 neurons), and 8 showed a significantly slower latency for NP stimuli compared to P stimuli (these were labeled Type II neurons), as if the NP stimuli merely triggered the thought of their corresponding P stimuli. The existence of Type I neurons was taken as evidence that some neurons did indeed expand their tuning as a result of the episodically learned association between the P and NP stimuli. This finding was interpreted as being inconsistent with pattern-separated episodic representations.

However, a null result for neurons labeled as “Type I” does not constitute strong evidence against a simpler explanation according to which the latencies for all 21 pair-coding neurons were longer for the NP stimuli than the P stimuli. Unless statistical power was very high, the so-called Type I neurons might simply reflect a failure to detect a true latency difference for these neurons as well. The idea that there are two distinct types of neurons based on the latency measures implies that the 21 P-NP latency difference-scores constitute a mixture distribution, one with a mean difference centered on 0 (Type I neurons) and another with a true mean difference greater than 0 (Type II neurons). However, no test for a mixture distribution was reported, so we performed such a test on the latency scores for the 21 pair-coding neurons estimated from Figure S4 of Ison et al. (2015). We first fit a single Gaussian distribution (free parameters = one mean and one standard deviation) to the 21 latency difference scores (P latency minus NP latency), and then fit a Gaussian mixture distribution (free parameters = two means, two standard deviations, and a mixing proportion) to the same data. The fits were performed using maximum likelihood estimation, and the quality of the fits was compared using both AIC and BIC to adjust for the difference in the number of free parameters. According to both goodness-of-fit measures, a single Gaussian distribution model provided a better fit than the two-Gaussian mixture model (AIC = 261.35 and BIC = 263.44 for the single-Gaussian model; AIC =

263.46 and BIC = 268.68 for the two-Gaussian mixture model). Thus, there is no evidence for two types of neurons based on the latency data.

In light of the above results, we next simply compared the average latency of the 21 pair-coding neurons in response to the P stimulus vs. the experimentally associated NP stimulus. This approach does not assume a categorical distinction between Type I and Type II neurons. Across all 21 pair-coding neurons, the average latency to the associated NP stimulus (322 ms) was significantly longer than the average latency to the P stimulus (259 ms), $P = .019$. This is the expected result if the experimentally learned NP stimulus triggers a thought of the P stimulus, which in turn, activates the concept neuron. According to these results, the activity of the concept neuron in response to the recently associated NP stimulus still reflects semantic memory, not episodic memory. That being the case, the data do not weigh against the notion that episodic memories are coded in non-overlapping neural assembles (assemblies that were not recorded in this experiment, which only recorded from semantic-memory concept neurons).