Title: Hippocampal neurons code individual episodic memories in humans

Authors: L. D. Kolibius[1,4*], F. Roux[1,4], G. Parish[1,4], M. Ter Wal[4], M. Van Der Plas[1,4], R. Chelvarajah[2], V. Sawlani[2], D. T. Rollings[2], J. Lang[3], S. Gollwitzer[3], K. Walther[3], R. Hopfengärtner[3], G. Kreiselmeyer[3], H. Hamer[3], B. P. Staresina[5,6], M. Wimber[1,4], H. Bowman[4,7], S. Hanslmayr[1,4*]

Affiliations:

[1] Centre for Cognitive Neuroimaging, School of Neuroscience and Psychology, University of Glasgow; Glasgow, United Kingdom.

[2]Neuroscience Department, Queen Elizabeth Hospital Birmingham, Complex Epilepsy and Surgery Service; Birmingham, United Kingdom.

[3]Department of Neurology, University Hospital Erlangen, Epilepsy Centre; Erlangen, Germany.

[4] Centre for Human Brain Health, School of Psychology, University of Birmingham; Birmingham, United Kingdom.

[5]Department of Experimental Psychology, University of Oxford; Oxford, United Kingdom.

[6]Oxford Centre for Human Brain Activity, Wellcome Centre for Integrative Neuroimaging, Department of Psychiatry, University of Oxford, Oxford, United Kingdom

[7]Centre for Cognitive Neuroscience and Cognitive Systems and the School of Computing, University of Kent; Canterbury, United Kingdom.

*Corresponding author.
Email: l.kolibius.1@research.gla.ac.uk (Luca D. Kolibius)
Email: simon.hanslmayr@glasgow.ac.uk (Simon Hanslmayr)

Abstract:

The hippocampus is an essential hub for episodic memory processing. However, how human hippocampal single neurons code multi-element associations remains unknown. Some argue that each hippocampal neuron codes for an invariant element within an episode. Instead, others have proposed that hippocampal neurons bind together all elements present in a discrete episodic memory. Here, we provide evidence for the latter. We show that individual neurons, which we term Episode Specific Neurons (ESNs), code discrete episodic memories using either a rate code or a temporal firing code. We find evidence for these neurons exclusively in the hippocampus. Importantly, these ESNs do not reflect the coding of a particular element in the episode (i.e.,

concept or time). Instead, they code for the conjunction of the different elements that make up the episode.

One-Sentence Summary:

Individual neurons in the hippocampus code for discrete episodic memories.

Introduction:

Episodic memory refers to our ability to reinstate the what, where and when of past experiences (Tulving, 2002). This ability is thought to depend on the reinstatement of neural activity that was present at memory encoding (Pacheco Estefan et al., 2019). It is undisputed that the hippocampus plays an integral role in episodic memory processing (Lisman et al., 2017; Marr, 1971; Squire, 1992) and the binding of multimodal information (Cooper and Ritchey, 2020). However, how it codes episodic memories remains controversial.

One important open question is whether neurons in the hippocampus code for specific elements or an entire episode. Concept Neurons in the hippocampus fire in response to specific invariant elements independent of the context in which they are presented (Gelbard-Sagiv et al., 2008; Mormann et al., 2008; Mormann et al., 2011; Quiroga et al., 2005). One contemporary idea is that the diverse elements that make up an episode are coded by the simultaneous activity of a set of these Concept Neurons (Quiroga, 2012; 2020) or by expanding the selectivity of existing Concept Neurons (Ison et al., 2015). According to this framework when you are sitting in your favourite coffee shop with your best friend, one set of Concept Neurons might code for the coffee shop and a separate set for your friend (Figure 1A).

Alternatively, single units in the hippocampus might sparsely encode a specific set of elements within an individual episode and act as pointers to cortical modules during memory

reinstatement. According to this so-called Indexing Theory (Teyler and DiScenna, 1986; Teyler and Rudy, 2007), the entire episode with your friend in the coffee shop is represented by a set of hippocampal neurons (Figure 1A). Unlike Concept Neurons, these Episode Specific Neurons (ESNs) would fire in response to the conjunction of all the diverse information within an episode and not in response to individual content elements. Despite computational models pointing towards the existence of ESNs (Bowman and Wyble, 2007; Krotov and Hopfield, 2020; Parish et al., 2021; Whittington et al., 2022), to this day there is no evidence for such a sparse conjunctive code in humans.

In the present work, we provide support for the existence of this content-agnostic episodic memory code implemented through Episode Specific Neurons. We leveraged intracranial microwire recordings to investigate the firing patterns of neurons in the human hippocampus and hypothesized that a significant number of hippocampal neurons reinstate their firing rate within a specific episode (i.e., fire during encoding and retrieval).

Importantly, these ESNs would code for the conjunctive elements present within an episode and are not tuned to individual elements within the episode. The existence of ESNs does not preclude Concept Neurons from participating in episodic memory processing. However, investigating the role of Concept Neurons in episodic memories goes beyond the scope of this work. As control analyses, we investigated whether this firing activity can be explained by a firing response to specific invariant elements, as occurs in Concept Neurons (Quiroga et al., 2005), or by a time preference, as occurs in Time Cells (TC; Reddy et al., 2021; Umbach et al., 2020).

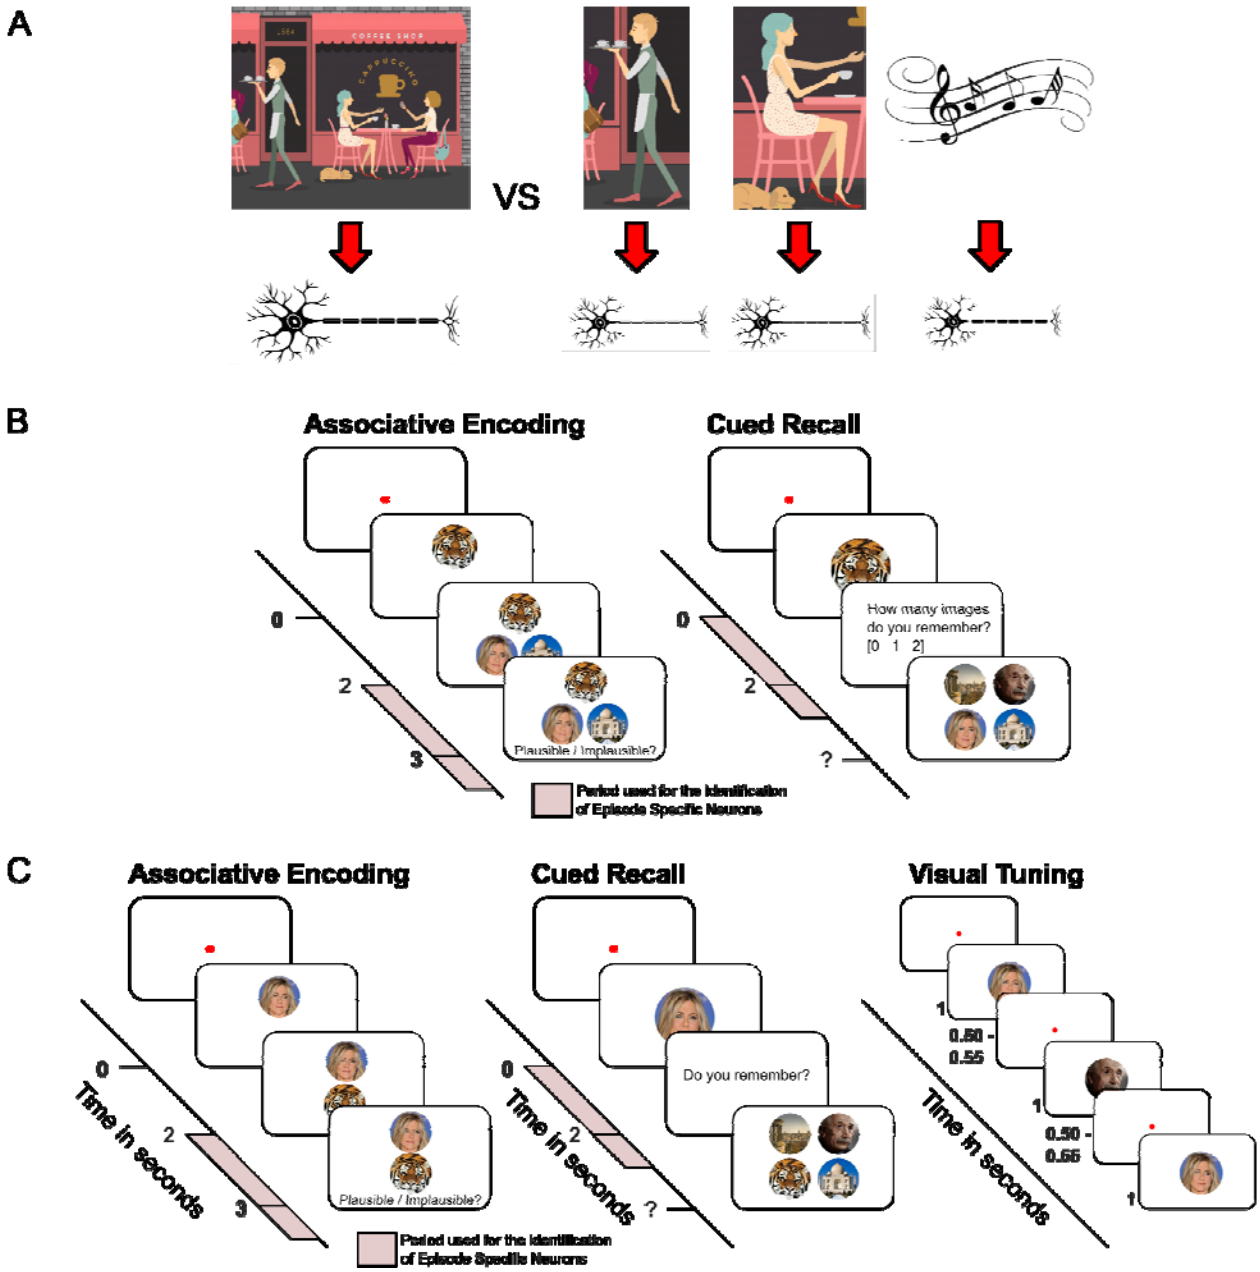

Figure 1. Difference between Indexing Theory and Concept Neuron based hippocampal coding of episodic memories, experiment procedure for experiments 1 & 2.

(A) Left: The classic Indexing Theory (Teyler and DiScenna, 1986) proposes that neurons in the hippocampus represent a conjunctive code that binds together all the elements that make up the episode in the form of an index. Within this framework, neurons do not directly code for the elements themselves (i.e., the smell of the coffee, your friend, the background music, the café, etc.), but rather act as pointers to these different elements which themselves are coded elsewhere (i.e., the neocortex). Right: Some hippocampal neurons are thought to code for specific elements or concepts, which is why they are called Concept Neurons (Mormann et al., 2008; Mormann et al., 2011;

Quiroga et al., 2005). Within this framework, a group of neurons collectively code an episodic memory, with each neuron representing a specific element involved in that episode (i.e., a neuron coding for the coffee, another neuron coding for your friend, etc. (Quiroga, 2012; 2020)). It is important to note that one index or one concept is likely to be coded by an assembly of neurons, not a single neuron.

(B) Outline of the procedure for Experiment 1. During encoding, all participants were instructed to imagine a vivid episode involving an animal cue and two associate images (two faces, two places or a face and a place) and rated its plausibility. This approach is suitable for investigating episodic memory as originally defined by Tulving in 1972 (Tulving, 1972). During recall, participants were asked to retrieve the associated images when cued with the animal cue. The experiment was self-paced and every episode was learned and tested only once. Following each encoding block of roughly 20 episodes, participants performed a short distractor task. The pink areas represent the time windows used for subsequent analyses (see Methods).

(C) Outline of the procedure for Experiment 2. Left: The memory task was largely the same as in Experiment 1 (see Figure 1B). However, events consisted of one cue (either an animal, a face or a place) and one associate image (either an animal, a face or a place). Right: After the memory task, patients performed a visual tuning task where the previously used stimuli were shown multiple times in quick succession without a memory component. This approach has been traditionally used to identify putative Concept Neurons.

Results:

We analyzed recordings from two separate experiments (Experiment 1: 585 neurons in the hippocampus, 16 participants, 7 female; age mean = 36.125 years, from 26-53 years; Experiment 2: 216 neurons in the hippocampus, 14 participants, 7 female; age mean = 33.857 years, from 19-58 years) where patients were implanted with stereotactic Behnke-Fried depth electrodes in the hippocampus (Figure S1), while they performed a memory association task (Figure 1B & Figure 1C).

During the encoding phase of Experiment 1 participants created a vivid mental story consisting of an animal cue and two associate images (two faces, two places or a face and a place). By contrast, Experiment 2 consisted of one cue and one associate image (both either an animal, face,

or place). The encoding and recall phase of the experiment was interleaved with a short distractor task where patients had to judge whether a series of 15 numbers was odd or even. The distractor phase lasted between 22.43s and 224.52s (median duration: 42.19s). During the recall phase, the animal cue was presented again and participants were asked to retrieve the associate image(s). The experiments were self-paced and every episode was learned and retrieved only once. Participants correctly recalled on average 68.38% (*SE* = 4.64%) episodes in the first experiment (see Supplements Table S1) and on average 65.63% (*SE* = 4.45%) episodes in the second experiment (see Supplements Table S2). This is substantially more than would be expected by chance (16.7% and 25% respectively).

Identifying Episode Specific Neurons (ESNs)

For every neuron, we determined the firing rate during each episode at encoding and retrieval. We then z-scored the firing rate across all encoding and retrieval episodes and excluded all later forgotten episodes. This was done independently for encoding and retrieval to account for general differences in firing rates. We measured episode-specific firing reinstatement as the product of the standardized firing rates at encoding and retrieval (Figure 2A).

Using a episode-shuffling procedure, we generated a distribution of reinstatement values expected by chance. A neuron was considered an ESN if (i) the empirical reinstatement value exceeded the 99th-percentile of the shuffled distribution for at least one episode and (ii) the standardized firing rate for encoding and retrieval of that episode each exceeded 1.645 ($\triangleq p_{\text{right-tailed}} < 0.05$). The second criterion prevented the identification of ESNs which would excessively fire at only one phase of the task (i.e., encoding or retrieval).

It could be argued that ESNs identified in this manner could reflect the firing of cells tuned to the image of the animal cue, rather than the conjunction of all elements since the cue is episode-

unique and presented during encoding and retrieval. To address this issue, in Experiment 1, we excluded neurons that showed a significant firing increase during the first second after the encoding of the animal cue for episodes that were later reinstated (see Methods). This procedure has traditionally been used to identify putative concept neurons (Quiroga et al., 2005; Mormann et al., 2008; Mormann et al., 2011). Using this approach, we identified a significant number of hippocampal ESNs in Experiment 1 (136 out of 585 neurons ≙ 23.25%; $p < 0.001$; permutation test; Figure 2B). Comparable results are obtained when (i) adding up the standardized firing rate between encoding and retrieval instead of multiplying them (E&R >= 1.645, reinstatement: E+R) (125 ESNs; $p < 0.001$), (ii) increasing the minimum standardized firing rate from z = 1.645 to z = 2.6 (E&R >= 2.6, reinstatement: E*R) (29 ESNs; $p < 0.001$) and (iii) using a different reinstatement measure that normalizes the encoding and retrieval product by their absolute difference (E&R >= 1.645, reinstatement: (E*R)/|E-R|) (53 ESNs; $p < 0.001$). This reinstatement measure has the important advantage of considering the similarity between the encoding and retrieval firing rate.

In Experiment 1, 117 out of 136 ESNs (≙ 86.03%) coded for a single episode, whereas the rest coded for multiple episodes. Two example ESNs are shown in Figure 3. These ESNs are unlikely to be concept cells tuned to the animal cue as the firing rate during encoding reaches its maximum only after the presentation of the associate stimulus (see Figure 4).

It is of note that the proportion of neurons that can be classified as ESNs is proportional to the number of events learned and retrieved (the same is the case for Concept Neurons). This is because we apply the threshold derived from the first permutation test to all episodes, without family-wise error correction. As such it is not suitable to determine the sparseness of the hippocampal code. However, the proportion of ESNs of all recorded neurons is useful as an estimation of how many ESNs we can expect in future analyses.

It is crucial to understand that this alpha-level inflation does not extend to the group-level permutation test, where the same number of tests are applied to randomly shuffled data. We have added a simulation using random values as spike rates and using circularly shuffled spike times to show that there is no inflation of the alpha error at the group-level at which we interpret our findings (see Methods; Figure S2).

ESNs are suggested to reflect a unique coding mechanism of the hippocampus (Teyler and DiScenna, 1986; Teyler and Rudy, 2007). In line with this, we did not find a significant number of ESNs in the parahippocampus (15 out of 104 neurons, $p = 0.5396$; permutation test). However, it should be noted that all 104 parahippocampal neurons originate from only five different microwire bundles over 14 sessions in five different patients, and therefore these results, albeit in line with the indexing theory, should be interpreted with caution. The second experiment did not contain a significant number of ESNs in the parahippocampus (3 out of 25 neurons, $p = 0.1199$).

To conclude, we find a significant number of ESNs in the hippocampus, but not in the parahippocampus. The analysis approach we use to identify ESNs is robust to deviations in the parameter space.

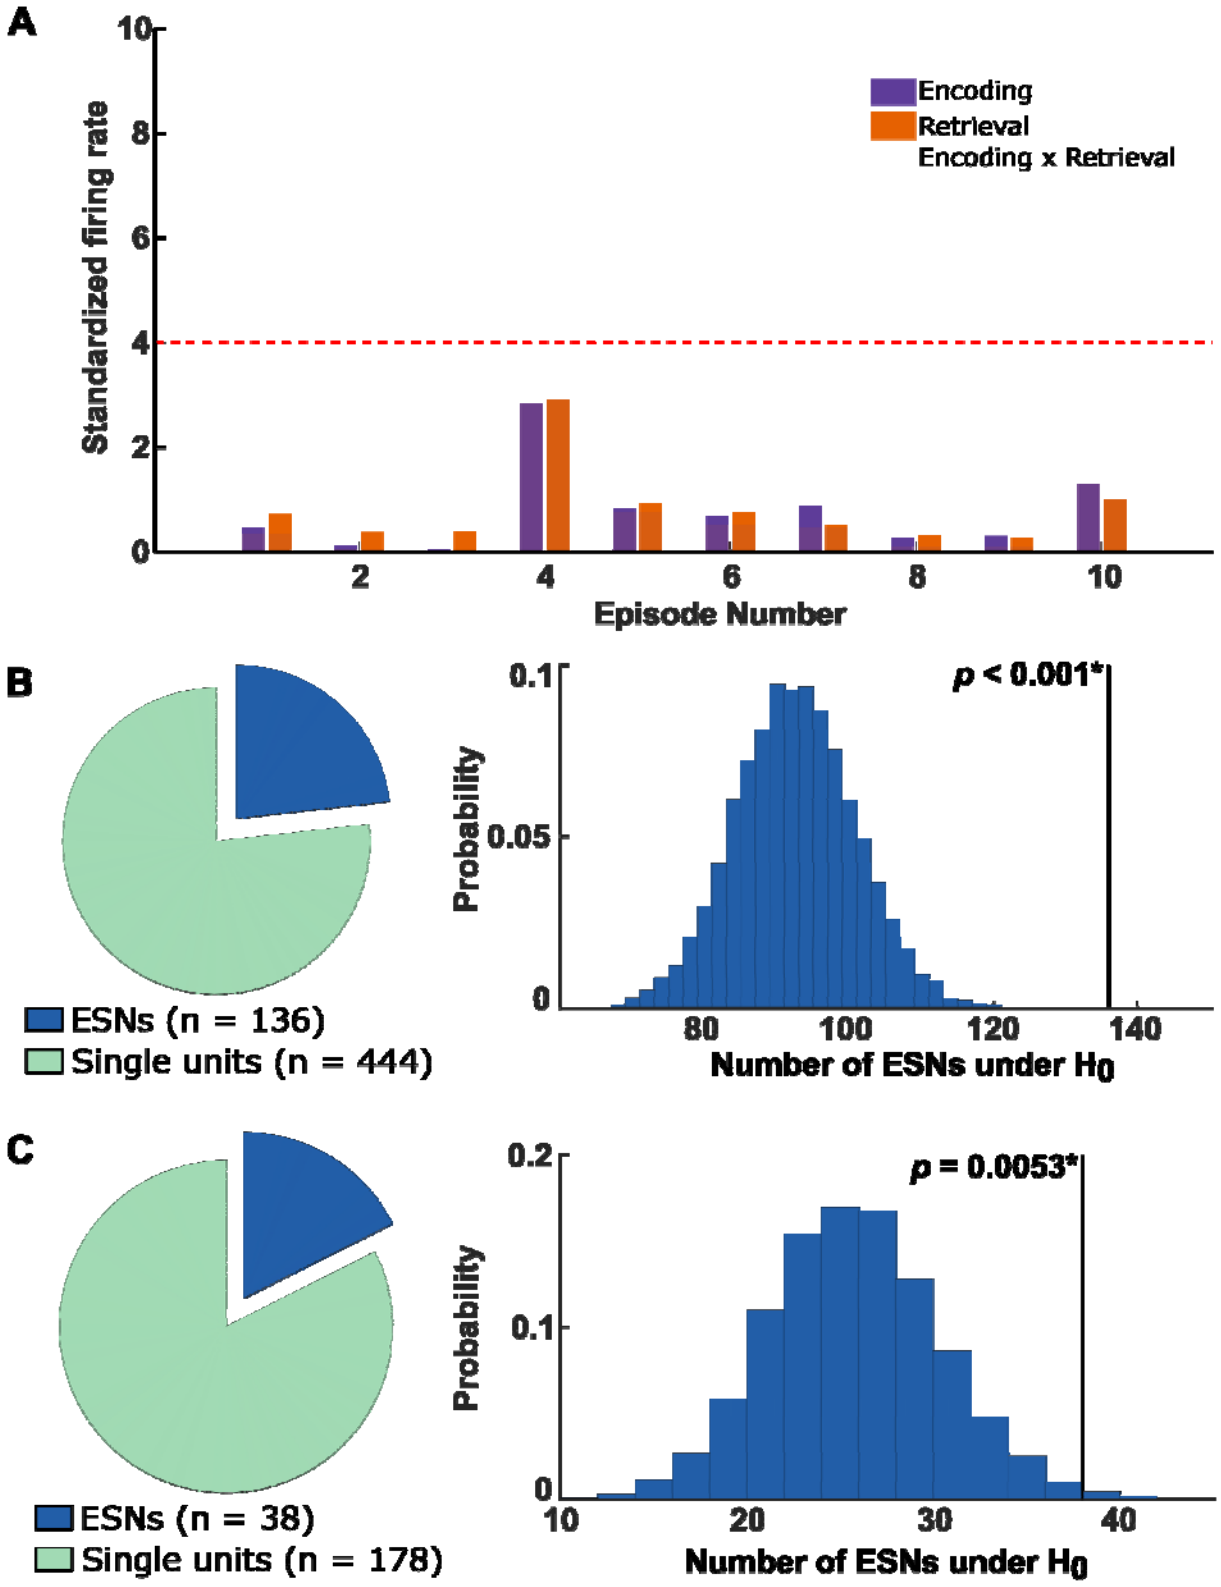

Figure 2. Analysis schematic and number of ESNs identified in experiments 1 & 2.

(A) A schematic for identifying Episode Specific Neurons (ESNs) is shown. The diagram shows the z-scored firing rate on the y-axis for ten simulated episodes on the x-axis colour-coded for encoding and retrieval (purple and orange, respectively). The transparent bars encompassing encoding and retrieval indicate the product of encoding and retrieval firing rates, which is used as the measure of episode-specific firing reinstatement. The dotted red line shows the threshold (derived from a shuffling procedure, see Methods). Because of the way ESNs are defined, they are required to fire substantially above their average firing rate during encoding and retrieval, which rules out neurons that generally show an increased firing rate during remembered episodes.

(B) Identified ESNs during Experiment 1. Left: The pie chart depicts the number of ESNs that show significant firing reinstatement to at least one episode (dark blue) and the number of neurons that showed no firing reinstatement (green). Right: The number of ESNs as expected by chance and the empirical number of ESNs (136 out of 585 neurons; $p < 0.001$; permutation test).

(C) Same as (B) but for Experiment 2. Out of a total of 216 hippocampal neurons, we identified 38 ESNs ($p = 0.0053$; permutation test).

ESNs do not code for the content/visual properties of the cue or associate image

Traditionally, visually responsive neurons have been identified using the repeated presentation of a stimulus. In the above analysis, we only present the animal cue once, which is suboptimal for ruling out Concept Neurons tuned to the animal cue. To ameliorate this shortcoming, in Experiment 2 we added a visual tuning task (Figure 1C) after the memory association task. During the visual tuning task, images from the memory task were repeatedly shown in quick succession. This approach is widely used to identify putative Concept Neurons that respond to one of the images independently of any memory processes (for example Concept Neurons see Figure S3; Mormann et al., 2008; Mormann et al., 2011; Quiroga et al., 2005). When excluding Concept Neuron activity in this independent dataset, we replicated our previous results and identified a significant number of ESNs (38 out of 216 neurons ≙ 17.59%; $p = 0.0053$;

permutation test; Figure 2C). In Experiment 2, 34 out of 38 ESNs (≙ 89.47%%) coded for a single episode, whereas the rest coded for multiple episodes.

However, traditional Concept Neuron detection methods might be too conservative to identify weakly tuned Concept Neurons. To address this concern, we drastically reduced the threshold of what constitutes a Concept Neuron, i.e., lowering the uncorrected threshold from $p = 0.0005$ to $p = 0.05$, which increased the number of Concept Neurons from 58 to 155 (out of 216 neurons). During a typical tuning task an average of 108.7 (min: 80; max: 156) different images are shown and each image is tested for visual tuning. There is no correction for multiple testing rendering a threshold at $p < 0.05$ very liberal.

Remarkably, incorporating this liberal threshold to exclude potential Concept Neurons, had little effect on the number of ESNs which remained almost unchanged (36 out of 216 neurons ≙ 16.67%; $p = 0.0025$; permutation test). It is conceivable that some images that are presented during the visual tuning task act as cues that reactivate some ESNs. These reactivated ESNs would then be erroneously rejected as Concept Neurons. However, in practice, only four potential ESNs were excluded based on the visual tuning task (six when lowering the Concept Neuron threshold). We suspect that ESNs were not reactivated during the visual tuning task because patients were not instructed to actively retrieve memories and thus may not have been in a "retrieval mode" (Lepage et al., 2000; Addante et al., 2011). Tulving first proposed this concept in 1983 (Tulving, 1983; Nyberg et al., 1995), referring to the cognitive state that occurs when we actively attempt to remember something. Being in a retrieval mode increases the likelihood that a memory cue successfully triggers memory retrieval. Nonetheless, we cannot rule out that in some cases ESNs were reactivated during the visual tuning task and subsequently rejected. However, this would only make our analysis for identifying ESNs more conservative.

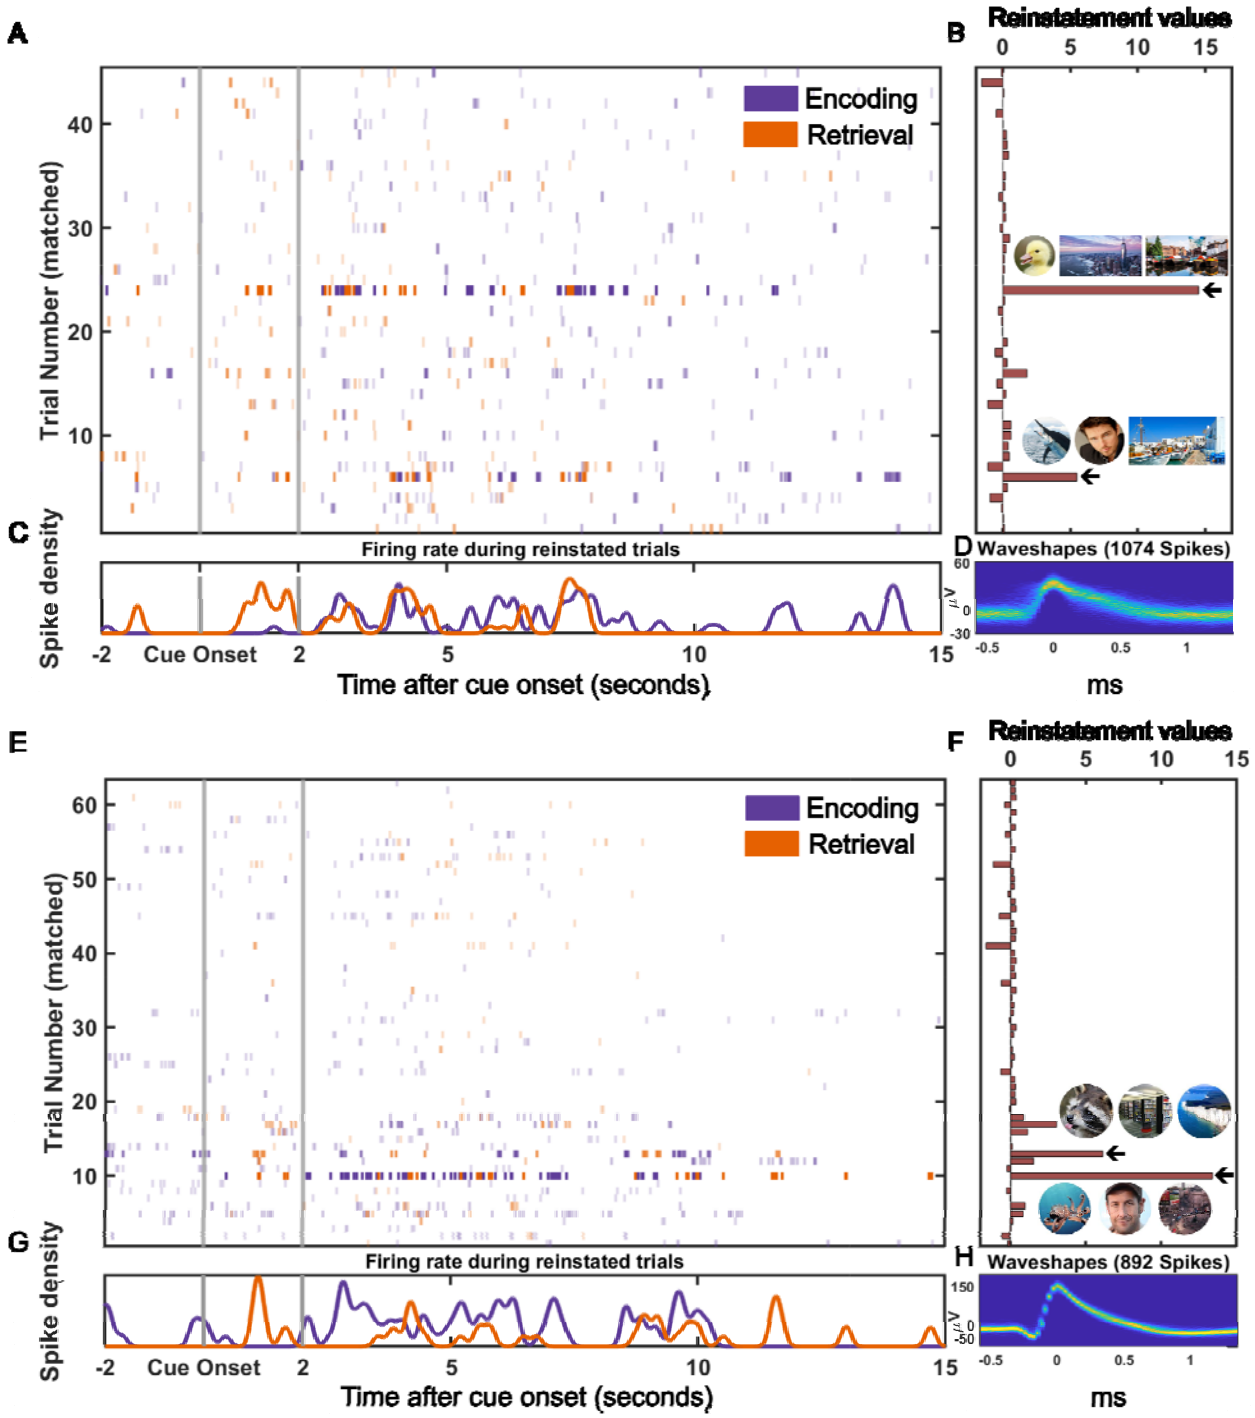

Figure 3. Firing patterns for two example Episode Specific Neurons.

(A) Spike raster plot. Each line indicates a spike. On the x-axis is time and on the y-axis episodes. Colour-coded in purple for encoding and orange for retrieval. The transparency is adjusted according to the reinstatement values in that specific episode.

(B) Reinstatement values and the animal cues with the respective associate images for reinstated episodes (indicated by the black arrows).

(C) Spike density plot for reinstated episodes. Note that the experiment is self-paced and episode length varies.

(D) 2D histogram of the waveshape of that particular unit (Niediek et al., 2016).

(E-H) same as (A-D) but for a different example ESN.

ESNs are limited to later remembered episodes

We have so far demonstrated that ESNs reinstate their firing rate when remembering a unique episode. This reinstatement cannot be explained by the semantic content or visual properties of the used image, which strengthens the notion that ESNs code for memories. In line with this, we did not find a significant number of miss-ESNs when limiting our analysis to later forgotten episodes (15 out of 585 neurons ≙ 2.56%; p = 0.4229; permutation test). However, this result could stem from a lower number of forgotten events (see Table S1). To counter this bias, we equalized event numbers between later remembered and later forgotten events for every neuron by randomly sampling (with replacement) later remembered events as many times as participants forgot an event. If any of the sampled events were later reinstated, we considered this neuron a miss-ESNs under the null hypothesis. By repeating this procedure 10,000 times we generated a distribution of how many miss-ESNs were expected if the number of later remembered and later forgotten events were equal. This analysis did not result in a significantly lower empirical number of miss-ESNs compared to hit-ESNs (p = 0.7032, bootstrapping test). To conclude, we did not find a significant number of ESNs when restricting our analysis to episodes that were forgotten. However, when considering that fewer episodes were forgotten than remembered there was no difference in the number of hit-ESNs and miss-ESNs.

Identification of temporal Episode Specific Neurons (tESNs)

The previous identification of ESNs relied on a rate code, i.e., the standardized mean firing rate during one episode at encoding and retrieval. We have adapted this analysis to identify neurons that reinstate a temporal pattern of firing. For every neuron, we considered the spiking activity six seconds before until one second after the response during encoding and retrieval (the first and last second was later excluded to avoid edge artefacts).

By convolving each spike with a gaussian kernel (standard deviation: 100ms, length: length: ±300ms, peak normalized to 1) we created a measure of instantaneous firing rate. Because we do not know the exact times when an episode is encoded or retrieved, we cross-correlated this episode-specific instantaneous firing rate during encoding and retrieval and considered the maximum value as the reinstatement value. We repeated this process after shuffling the encoding and retrieval episode order 1,000 times and took the 99th percentile as a threshold for the empirical reinstatement value. If the empirical reinstatement value reached this threshold, we considered the neuron a temporal Episode Specific Neuron (tESN; Figure S4). In the next step, we randomly drew for each neuron one of the previously calculated permutations. If these permuted values reached or exceeded the threshold the neuron was considered a tESNs under the null hypothesis. We repeated this process 1,000 times to build a null distribution against which we compared our empirical number of tESNs. We found a significant number of empirical tESNs in Experiment 1 (100 out of 585 neurons ≙ 17.09%; $p = 0.016$). For Experiment 2, we further excluded all episodes in which the given neuron showed a significant visual tuning (see Methods). With this additional constraint, we found a significant number of tESN in Experiment 2 (40 out of 216 neurons ≙ 18.52%; $p < 0.01$).

We employed a permutation test to assess the degree of overlap between episodes reinstated by a rate code (ESNs) and those reinstated by a temporal code (tESNs). Specifically, for each neuron we shuffled the identity of whether an episode was reinstated or not and compared the overlap in

the shuffled dataset with the empirical overlap (i.e., is an episode reinstated using a rate code also reinstated using a temporal code and vice versa). Our analysis revealed a significant overlap in Experiment 1 (Experiment 2), with 20.25% (26.19%) of all episodes reinstated by ESNs also being reinstated by tESNs, and 25.81% (24.44%) of all episodes reinstated by tESNs also being reinstated by ESNs (both $p < 0.001$).

We then tested the validity of this analysis using random spike times. We generated these random spike times by first rounding the empirical spike times to the nearest integer and then drawing an equal number of pseudorandom integer values from a discrete uniform distribution between the first and last empirical spikes times. We did not find a significant number of tESN in either experiment (both $p > 0.2$). We next repeated the analysis using 500 surrogate datasets. These datasets were created by segmenting all spike times into episodes in the order they occurred, and then circularly shuffling them. The results of this analysis indicated that the percentage of significant tESN identifications was below the 5% threshold (4.18%), providing further evidence of the credibility and reliability of our analysis.

In conclusion, we show in two separate experiments a significant number of neurons that reinstate an event specific temporal firing pattern during successful memory retrieval.

ESNs do not code for time

Recent studies in humans show that some hippocampal neurons code specific time points invariant across repetitions, which are referred to as Time Cells (Reddy et al., 2021; Umbach et al., 2020). We investigated whether our dataset contains such Time Cells (TC) using a similar method as employed by (Umbach et al., 2020). Due to the self-paced nature of our experiment, each encoding block varied in length. To accommodate this, we used both the unaltered block length, as well as a normalized block length within one recording session (see Methods). Of all

585 recorded cells, 12 (normalized) and 10 (non-normalized) fulfilled the criteria of TCs, which is below chance level (p values > 0.9; permutation test). Critically, there was no significant overlap between neurons that behaved like TCs and ESNs (p values > 0.3; permutation test) indicating that ESNs cannot be construed as TCs.

ESNs show a wider waveshape than other neurons

We found some evidence that spike waveshapes of ESNs are wider than those of other units (Figure S5A; p = 0.0563; with data from Experiment 1 and p = 0.0121 with data from both experiments combined; both independent samples t-test), possibly indicating that ESNs are physiologically different from other neurons. In the hippocampus, a wider waveshape has previously been associated with excitatory cells (Prestigio et al., 2019), therefore suggesting the ESNs are predominantly excitatory neurons. There was no significant difference in the spike height or Fano factor between ESNs and other neurons (unpaired t-tests; all p values > 0.3; Figures S5B and S6).

Recorded neurons are mostly single neurons and not multi-units

Although we tried to separate multi-units into single neurons as best as possible during the spike sorting procedure (see Methods), some units might still represent activity from multiple neurons. We thus employed the method outlined by Tankus and colleagues (Tankus et al., 2009) to classify units into single units and multi-units, using the inter-spike interval and spike waveshape variability as objective criteria. In the first experiment, 373 out of 585 units (≙ 63.76%) were classified as single units (95/136 ESNs ≙ 69.85%), while in the second experiment 132 out of 216 units (≙ 61.11%) were classified as single units (20/38 ESNs ≙ 52.63%). If we limit our analysis to neurons that satisfy these stringent criteria for putative single neurons, we still find a

significant number of ESNs in the first experiment (95 out of 373 single neurons; $p < 0.001$), and a number of ESNs trending towards significance in the second experiment (21 out of 132 single neurons; $p = 0.0714$).

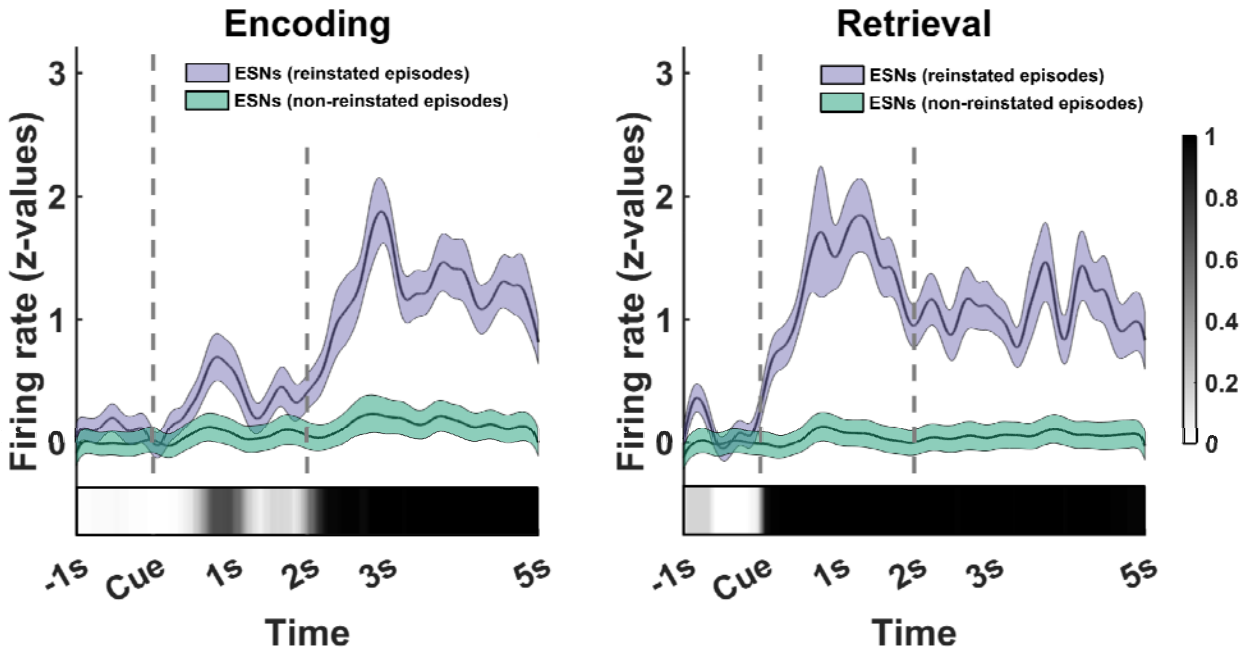

Figure 4. Firing rate of ESNs during reinstated (purple) and non-reinstated (green) episodes. Firing rate of ESNs (n = 136) from cue onset until five seconds later during memory encoding (left) and retrieval (right). We utilized a bootstrapping method to ensure an equal number of reinstated and non-reinstated episodes for each ESN, followed by a computation of the cluster-based permutation test (Maris and Oostenveld, 2007). The proportion of iterations that contained a significant cluster at a specific timepoint is represented by the grayscale inset at the bottom of the figure. The coloured shaded areas represent the standard error of the mean (SEM) that was calculated across all ESNs. We ensured an equal number of reinstated and non-reinstated episodes per ESN while calculating the SEM using a bootstrapping method.

Discussion

Using an associative episodic memory paradigm in human epilepsy patients, we identified hippocampal neurons that are active during the initial encoding of a unique episode and later reinstate their firing rate when successfully remembering the same episode. Therefore, we term these neurons Episode Specific Neurons (ESNs). The activity of these neurons could not be

explained by a firing rate increase towards specific images or time points. These results were replicated across two independent experiments and were robust against variations in how reinstatement was measured as well as changes in the hyperparameter space.

Previous studies have demonstrated that Concept Neurons increase their firing rate during memory retrieval when the image they are tuned to is part of the memory (Gelbard-Sagiv et al., 2008; Ison et al., 2015). We used two approaches to ensure that the ESNs we identified are not Concept Neurons that selectively respond to visual elements or semantic concepts: (i) in Experiment 1 we excluded ESNs that were visually responsive to the presentation of the animal cue at encoding. (ii) Following the episodic memory task in Experiment 2 patients completed a visual tuning task using all previously presented stimuli. This is a standard method to identify putative Concept Neurons (Ison et al., 2015; Mormann et al., 2008; Mormann et al., 2011; Quiroga et al., 2005) and allowed us to exclude episodes where a neuron showed a visual tuning to either the cue or the associate image. Using this approach, we replicated our results from Experiment 1 in a new sample of patients and found a significant number of ESNs while also verifying that these neurons do not selectively respond to visual elements or semantic concepts. Importantly, this finding was robust even when dramatically reducing the threshold of what constitutes a Concept Neuron. Taken together these analyses reinforce the argument that ESNs are memory related.

The existence of ESNs does not exclude Concept Neurons from playing a role in episodic memory processes. Concept Neurons might code the semantic aspect of an episode (i.e., the general concept of "coffee shop"). However, according to the Indexing Theory (Teyler and DiScenna, 1986; Teyler and Rudy, 2007), hippocampal neurons that perform this indexing function should have no initial tuning and are allocated to a specific episode during memory formation (i.e., the coffee shop in a specific setting). The behaviour of ESNs would be consistent

with such an indexing function and may add crucial event-specific information to an episode, that Concept Neurons cannot encode themselves.

We found a significant number of ESNs when excluding potential multi-units in the first experiment. However, in the second experiment we only obtained a statistical trend for a significant number of ESNs. This was likely because this restriction resulted in too few single neurons in the second experiment, thus reducing statistical power.

Because we do not know the exact time points when episodes are encoded or retrieved, we used a rate code approach in the first instance for these analyses (i.e., averaging the number of spikes over a time of interest and encoding and retrieval). In addition, we present evidence for a reinstatement of a temporal firing code which we uncovered by shifting the instantaneous firing rate (i.e., the spike times convolved with a gaussian kernel) using a cross-correlation.

Interestingly, we found a significant overlap of episodes that were coded both by tESNs and ESNs. This suggests that in many cases, a temporal firing code can still be identified through a rate code analysis. In some cases we detected a reinstated episode using either a rate code or a temporal code alone. It is possible that hippocampal neurons employ two distinct coding mechanisms, or in some instances, we may have missed certain spikes, resulting in the inability to identify either firing code, which is an interesting question for future studies.

The Indexing Theory proposes that this coding mechanism is unique to the hippocampus. In line with this, we did not find a significant number of ESNs in the parahippocampus. However, these findings are based on a relatively small sample size and should be considered preliminary. Future studies are needed to ascertain the regional specificity of ESNs to the hippocampus.

We did not find a significant number of ESNs when restricting our analysis to later forgotten episodes. However, there was no significant difference between the number of ESNs when considering later remembered and later forgotten events. This suggests that hippocampal neural

reinstatement might occur without behavioural memory retrieval. This could be due to downstream processing being disrupted (i.e., due to interference, selective attention). Alternatively, it is possible that in some cases during memory encoding patients created an episodic memory that did not incorporate the presented associate stimuli. While retrieval would lead to neural reinstatement, the patients would not be able to choose the correct associate images. Another possible explanation for this finding is that a relatively fewer number of forgotten episodes led to an insufficient power to detect a significant number of ESNs that code for later forgotten episodes.

Time Cells (TC) are neurons that invariantly fire at specific, reoccurring time points (Reddy et al., 2021; Umbach et al., 2020). We did not find a significant number of TCs in our study and there was no significant overlap between TCs and ESNs. This might be because the self-paced nature of the task introduced too much time variation between too few learning blocks to uncover TC dynamics. However, the absence of TCs in our paradigm corroborates ESNs as independent from TCs.

We found that ESNs have a wider waveshape than other neurons, which suggests that they are physiologically different from other single neurons. Specifically, the broader waveshape of ESNs suggests that they are likely excitatory cells (Prestigio et al., 2019). Alternatively, it is possible that ESNs and neurons with a narrower waveshape are located in different hippocampal subfields. Unfortunately, with the current methods, we lack the precision to designate neurons to individual subfields (Quiroga, 2019).

One limitation of the current study is that every event was encoded and retrieved only once. However, the very nature of episodic memories is one-shot learning and the ability to subsequently perform mental time travel. Any neural substrate that supports this function must occur after a single bout of learning and subsequent retrieval of a single episode. Our method

honours this fundamental characteristic which is the defining feature of episodic memory as originally stated by Endel Tulving (Tulving, 1972). Arguably, a repeated design would have allowed for a more reliable ESN identification. However, each memory reactivation leads to a transient plasticity of the memory trace until it is reconsolidated again. During this time window profound changes in the neurons that code for the initial memory trace might occur (Nader & Hardt, 2009). To avoid this potential confound, every episode is learned and retrieved only once in the present experiments. The stability of ESNs over repeated reactivations and extended periods, therefore, remains an interesting topic of research for future studies.

Our results are consistent with previous studies using fMRI that have shown item-specific activity reinstatement in the hippocampus (Chadwick et al., 2010; Mack and Preston, 2016) where similar representations are associated with distinct activity patterns (Bakker et al., 2008; Berron et al., 2016). These findings are suggestive of an episode-specific neural code, which is consistent with our results. However, due to the coarse resolution of fMRI, these previous results cannot disambiguate whether this event-specific code is driven by a population of event-specific concept neurons, or whether it is driven by a population of event-specific indexing neurons. We here provide evidence for the latter.

Previous intracranial work has identified a multitude of different neurons that detect episode boundaries and event onsets (Zheng et al., 2022) as well as novelty or familiarity (Rutishauser et al., 2006; Rutishauser et al., 2008; Rutishauser et al., 2015, Urgolites et al., 2022). Recent work showed that this is a generic signal that can be observed across the brain and is not unique to the hippocampus (Urgolites et al., 2022). These cell types generally fired not specific to a particular episode, but instead across many episodes (Rutishauser et al., 2006; Rutishauser et al., 2008; Rutishauser et al., 2015). Previous research quantified neural firing reinstatement during scene encoding and recognition by considering the activity of all recorded neurons (i.e., population

activity) (Zheng et al., 2022). In contrast, the vast majority of ESNs identified in our experiments coded for a single episode, and we demonstrated that neural reinstatement takes place at the single-neuron level. In line with our findings, recent work identified a sparse and item-specific memory recognition signal that was unique to the hippocampus (Urgolites et al., 2022). Importantly, we expect that an episode is coded by an assembly of ESNs from which we sampled only one due to the limited number of neurons that can be recorded with the currently available methods.

Our findings are in line with previous work showing that episodic memories in the hippocampus are coded in a sparse distributed way (Wixted et al., 2014; Wixted et al. 2018). However, there are various reasons why we refrain from making any claims regarding sparsity in the present study. A neural code can be sparse in two ways (Wixted et al., 2018). A neural code can be population sparse, which is the case when a low percentage of neurons respond to a given stimulus. It can also be lifetime sparse, which refers to a low percentage of stimuli that a given neuron responds to. On the one hand, we artificially induce lifetime sparsity (and by extension population sparsity) because we (i) standardize the firing rate during encoding and retrieval and then (ii) multiply these two values. On the other hand, we drastically reduce the sparsity because we test for reinstatement at each episode without correcting for multiple comparisons. It is very important to understand that while this leads to alpha-level inflation at the level of the neuron, this does not extend to the group-level at which we interpret our findings. We have confirmed that our analysis does not have a bias towards positive findings using a simulation (see Figure S2). Unfortunately, that also means that in the present study we have to refrain from making any claims regarding lifetime sparsity. Although most ESNs reinstate only a single episode, some code more than one. We expect that episodic memories are represented in the hippocampus as neural assemblies of single neurons and not individual neurons. Thus, it is plausible, that a subset

of neurons within a neural assembly coding for memory A are also part of an assembly that codes for memory B. This coding mechanism is more efficient as a partial overlap between neural assemblies reduces the number of required assemblies.

In conclusion, we found neurons in the hippocampus that show firing reinstatement in response to a specific conjunction of elements within a unique episode. These Episode Specific Neurons did not fire in response to individual concepts (Concept Neurons) or to specific, re-occurring time points (Time Cells). We propose that during memory formation an assembly of ESNs acts as a pointer or index that initially binds the elements of an episode together, in line with the Indexing Theory (Bowman and Wyble, 2007; Teyler and DiScenna, 1986; Teyler and Rudy, 2007). Reactivation of this pointer allows ESNs to reinstate the episodic memory previously encoded. Importantly, because ESNs reinstate unique episodes, they contain a time and content component. However, rather than reflecting the underlying coding mechanism, this time and content aspect necessarily emerges from the conjunctive code of an episode that is unique in content and time.

Methods

Procedure of memory experiment 1

During the encoding phase of the experiment the participant associated a cue with two other stimuli. For each episode, the cue was a new picture of an animal. The stimuli could be pictures of either places, faces or both. Every picture was only shown once. Two seconds after the animal cue was presented, the associate stimuli were shown, while the animal cue remained on the screen. The participant was asked to create a vivid imaginary story involving the cue and the two

stimuli. This part of the experiment was self-paced. The task continued once the participant rated the plausibility of the imaginary story (plausible/implausible).

After the encoding phase, the participant performed a distractor task to rule out working memory effects. During the distractor task, participants had to indicate whether a random number (up to two digits) that appeared serially on the screen was odd or even. After each response, the participant received feedback indicating a correct or incorrect response. This task consisted of 15 trials.

During the retrieval phase, all cues from the previous encoding phase were presented sequentially in pseudorandom order. Each animal cue was presented for two seconds and subjects were tasked to retrieve the corresponding images. The participant was then asked how many associated images they remembered (none, one, or two). Participants had as much time to respond as they required. If the participant indicated that they remembered one or two images, they then were asked to select two pictures from an array of four pictures (two targets and two distractors that consisted of pictures from the previous encoding block which were associated with a different cue).

The experiment ended after the retrieval phase if the total runtime exceeded 40 minutes, or if the patient asked to abort the experiment. Otherwise, the experiment continued with the next encoding block. The encoding block initially consisted of 20 episodes but could be adjusted depending on the cognitive abilities of the patient. If the hit rate fell below 66.25%, fewer episodes were shown for the next block and vice versa if the hit rate surpassed 73.75%.

The patients performed the memory task on a laptop computer (experiment 1: Toshiba Tecra W50, 60 Hz refresh rate; experiment 2: Lenovo L390 Yoga, 60.01 Hz refresh rate), while either seated in a chair next to their bed or their hospital bed.

Procedure of memory experiment 2

The second experiment is based on the first experiment with the following adaptations: participants are presented with one cue image (depicting an animal/place/face) and only one associate image (depicting an animal/place/face). During retrieval, participants were asked whether they remembered the associate image and the participants had to choose the correct associate from an array of four pictures (one target and three distractors that consisted of pictures from the previous encoding block which were associated with a different cue). The experiment was terminated upon request or when the runtime at the end of a retrieval block exceeded 30 minutes.

Visual tuning task procedure

For experiment 2, the memory task was followed by a visual tuning task. During this tuning task, every image that was shown during the preceding memory task was displayed. Each image was shown six times in pseudorandom order on the screen for a duration of one second. The inter-image interval was jittered between 500ms and 550ms. To ensure attention, patients had to categorize the image as an animal, a place, or a face using the arrow keys on the keyboard.

Participants

For experiment 1, eight patients were recorded in the Queen Elizabeth Hospital Birmingham (Birmingham, UK) (4 female; mean age: 36.25 years, from 26-49 years) and eight patients in the Universitätsklinikum Erlangen (Erlangen, Germany) (3 female; mean age: 36.125 years, from 26-53 years). For experiment 2, 14 patients were recorded in the Universitätsklinikum Erlangen (Erlangen, Germany) (7 female; mean age: 33.857, from 19-58 years).

Ethical approval

Ethical approval was granted by the National Health Service Health Research Authority (15/WM/0219) and the Ethik-Kommission of the Friedrich-Alexander Universität Erlangen-Nürnberg (142_12 B). Informed consent was obtained in accordance with the Declaration of Helsinki.

Behavioural analysis

For the analysis of the first experiment, we considered an episode a hit if the participant correctly identified both stimuli. We considered an episode a miss if the participant either indicated not to remember any stimuli or did not remember either stimulus correctly. Participants correctly recalled on average 68.38% (*SE* = 4.64%) episodes in the first experiment (see Table S1) and on average 65.63% (*SE* = 4.45%) episodes in the second experiment (see Table S2). This is substantially more than would be expected by chance (16.7% and 25% respectively).

Statistical analysis

All statistical analyses were conducted using MATLAB R2020a on a computer running Windows 10 Enterprise. The significance threshold for all statistical tests was set at 0.05. Unless specified otherwise, all permutation tests were implemented with N = 10,000 random draws.

Co-Registering

For all but one patient, a pre-operational T1-weighted MRI scan was co-registered with a post-operational scan and normalized in MNI space using SPM12. For one patient, a post-operational CT scan was used instead of a post-operational MRI scan. Each microelectrode was localized either within the hippocampus, within the parahippocampus, or outside of both brain structures

through visual inspection of an (see Figure S1). Only activity from microwires in Behnke-Fried electrodes assigned to the hippocampus was analysed in the main analysis of the current study. Neurons in the parahippocampus were analysed in an independent follow-up analysis.

Recording System and Electrodes

Patients were implanted with one to eight (see Table S1 and S2 for an overview) depth electrodes of the Behnke Fried type with microwire bundles (Ad-Tech Medical Instrument Corporation, USA) to localize epileptic foci. The electrode location was determined by clinical need. These single-use electrodes are made from platinum, have a diameter of 1.3mm and allow for simultaneous macro- and microcontact recordings. Platinum has a high impedance for lower frequency and a low impedance for higher frequency bands. As such it is suitable to pick up local extra-cellular action potentials. The micro contacts extended radially past the endpoint of the macro depth electrode, and each contained eight high-impedance microwires (38-micron diameter) and one low-impedance microwire that is typically used for referencing.

The electrodes were connected to an ATLAS system (Neuralynx Inc, USA) consisting of CHET-10-A pre-amplifiers and a Digital Lynx NX amplifier and recorded with a sampling rate of either 32,000 Hz (Location: Birmingham) or 32,768 Hz (Location: Erlangen). Upon acquisition, an analogue bandpass filter from 0.1 Hz to 9,000 Hz was applied.

Spike detection and spike sorting

In the following paragraph, we will outline the process used to filter the raw data, detect spike timestamps, extract features of the waveshape and cluster spike waveshapes into putative single neurons using the wave_clus toolbox. For a more in-depth description of the wave_clus algorithm, the reader is referred to (Chaure et al., 2018).

The unfiltered signal included both the local field potential and the action potentials of individual neurons. Action potentials are characterized by a very steep and transient amplitude in the signal. To extract these spikes, we first applied zero-phase filtering using a second-order bandpass elliptic filter in the range of 300-3,000 Hz. The resulting signal contained the information of the so-called spike band.

Next, we segmented the continuous filtered data into epochs of five minutes. Segmenting the continuous data into smaller epochs had the advantage that noise in the signal did not increase the detection threshold for the whole recording and instead was limited to the segment in which it occurred (Chaure et al., 2018).

Spike detection was performed separately for positive and negative deflections. Once a spike was identified, 64 data points around the spike maximum were extracted. This corresponds to a 2 ms window at a sampling rate of 32,000 Hz. The spike peak was aligned to the 20th sampling point. To avoid misalignment of the spike, the waveshape was first up-sampled to 320 data points using cubic spline-interpolated waveforms and then down-sampled again (Chaure et al., 2018).

Based on the extracted spike waveform, features were computed using a four-scale multiresolution decomposition with a Haar wavelet. This results in 64-wavelet coefficients for each spike. The 10 most significant coefficients were identified using a Lilliefors test and used for the clustering procedure (Chaure et al., 2018).

Nonparametric clustering in the feature space was performed using superparamagnetic clustering (SPC). SPC grouped spike waves into clusters based on nearest-neighbour interactions (Blatt et al., 1996). Template-matching in Euclidian space was performed to assign unclassified waveforms to one of the identified clusters. The resulting clustering solution was then manually inspected and further optimized by rejecting artefact clusters, splitting clusters that represent multi-unit activity and merging clusters that likely stem from the same neural source. See

Figures S5 to S7 for an overview of the spike width, spike height, the Fano factor and the firing rate separately for ESNs and all other single units.

Identification of Episode Specific Neurons (ESNs)

For every single unit, we determined the number of spikes within each episode. During encoding, spikes from the onset of the associate images (two seconds after the cue onset i.e., when the whole information of the episode was present) until the end of the episode were considered. During the retrieval phase, spikes from cue onset until the time point at which participants indicated how many images they remembered were considered. We chose this time window because an episode could be reinstated following cue presentation, while after the response patients were presented with an array of images that could have potentially induced single-unit firing. Because the experiment was self-paced and longer episodes trivially contained more spikes, the firing rate (in hertz) was computed for each episode and single unit. In the next step, we z-scored this firing rate per single unit within all encoding episodes and retrieval episodes separately. Afterwards, we excluded all episodes that were later forgotten (for hit-ESNs) or that were later remembered (for miss-ESNs). Only sessions with at least eight episodes after this restriction were considered for further analysis. We then multiplied this standardized firing rate for encoding and retrieval episodes elementwise to gain an indicator for the reinstatement of firing for each episode (Figure 2A).

To estimate a threshold at which episode-specific firing reinstatement occurs on a single-unit level, we permuted the order of the encoding episodes and recomputed the elementwise product of the shuffled episode series. We repeated this permutation step 10,000 times and stored all output values. The 99th percentile of these pooled values was then used as a threshold for firing reinstatement. As an additional constraint, z-scored firing during encoding and retrieval each had

to exceed 1.645 (≙ $p_{right\text{-}tailed} < 0.05$) to make sure the elementwise product was not predominantly driven by a high firing rate in one of the two phases alone (i.e., either encoding or retrieval). This procedure is allowing us to threshold, but we do not have family-wise error corrected statistical significance at the single-unit level (there is no alpha inflation at the group level, see #Simulation of ESN identification). Furthermore, we assume that single units fire independently. To ensure Concept Neurons tuned to the animal cue were not falsely interpreted as ESN activity, we excluded ESNs that showed a significant firing increase in response to the animal cue at encoding using the method described below under *Identification of putative Concept Cells.*

Alternative reinstatement measures are explored in the results section under #Identifying Episode Specific Neurons (ESNs) and include (i) adding up the standardized firing rate between encoding and retrieval instead of multiplying them (E&R >= 1.645, reinstatement: E+R), (ii) increasing the minimum standardized firing rate from z = 1.645 to z = 2.6 (E&R >= 2.6, reinstatement: E*R) and (iii) using a different reinstatement measure that normalizes the encoding and retrieval product by their absolute difference (E&R >= 1.645, reinstatement: (E*R)/|E-R|), thereby taking into account the similarity in the standardized firing rate between encoding and retrieval.

In the second step, we calculated whether the number of ESNs (as identified in the above procedure) was above chance level. We did this by randomly choosing one of the permutations calculated in the first step for every single unit and checking whether it would be classified as an ESN under the same criteria outlined above. This approach is similar to a set-level effect in SPM (Penny et al., 2011). This process was repeated 10,000 times and the total number of single units which would be classified as an ESN in every single iteration of this process was used to build a distribution against which we compared our empirically discovered number of ESNs.

Simulation of ESN identification

We created a simulation using random pseudo-spike rates to determine whether our ESN analysis pipeline contains a bias towards false positive results. To create this simulation, we simulated the firing rate of 585 single neurons during 40 encoding and 40 retrieval trials by randomly drawing from a standard uniform distribution in the open interval of 0 to 1. These values were first multiplied by a variance factor that cycled from 2 to 5 and then z-scored independently for encoding and retrieval. Just as in the main ESN analysis we computed a reinstatement value for each trial by multiplying the two standardized synthetic firing rates. Next, we created a threshold by permuting the encoding and retrieval trial order 10.000 times while recomputing the shuffled reinstatement value. The 99th percentile was used as a threshold while the empirical standardized pseudo-firing rate had to be at least 1.645 during encoding and retrieval. If these criteria were met, we considered the neuron an ESN.

Then we computed the second-order (group level) permutation test by drawing a random first-order permutation for every single neuron and contrasted these values with the single neuron specific threshold. If the shuffled values satisfied the criteria for ESNs (i.e., encoding and retrieval standardized pseudo-firing rate at or above 1.645 and a reinstatement value above the neuron specific threshold) we considered the single neuron an ESN under the null distribution. By repeating this step 10,000 times we created a distribution under the H0 against which we could compare our initial random values. We repeated this entire process 1,000 times for each level of variance (2 to 5).

Because our initial pseudo-spikes were just random values, we expected 5% of all repetitions to yield a significant number of ESNs at any level of variance. If there was a bias, then more than 5% of all repetitions would contain a significant number of ESNs. As evidenced by Figure S2 this was not the case for any level of variance.

To ensure the robustness of our analysis approach, we repeated the ESN identification analysis using 500 surrogate datasets. These datasets were generated by segmenting all spike times into all available episodes in the order they occurred, and then circularly shuffling them. The results of this analysis revealed that the percentage of significant ESN identifications was 1.8%, which is well below the 5% threshold. These findings provide additional evidence for the credibility and reliability of our analysis.

Identification of putative Concept Cells

We have followed the method outlined in Mormann et al. (Mormann et al., 2011; 2008) to detect significant single-unit responses towards images. To this end, the 1000ms period after the stimulus onset was divided into 19 overlapping 100ms bins. The spike counts of each bin over all presentations of an image were compared to the 500ms baseline periods before stimulus onset for all images in the session using a two-tailed Mann-Whitney U test. We used the Simes' procedure to correct for multiple comparisons (Rødland, 2006). We performed this test twice, once with the commonly used threshold of $p < 0.0005$ and again with a liberal threshold of $p = 0.05$.

Identification of temporal Episode Specific Neurons (tESNs)

The analysis to identify neurons that showed a temporal firing reinstatement for specific episodes closely followed the outline described in #Identification of Episode Specific Neurons (ESNs). For every neuron, we considered the spiking activity six seconds before until one second after the response during encoding and retrieval (the first and last second was later excluded to avoid edge artefacts). We set a minimum threshold of ten spikes per trial and ten trials per neuron to enter the analysis. These thresholds were chosen to avoid artificially high cross correlations due to low numbers of spikes, and for having enough trials for the randomization procedure.

We then convolved each spike with a gaussian kernel (standard deviation: 100ms, length: ±300ms, peak normalized to one) creating a measure of instantaneous firing rate.

A main problem with comparing neural time courses between encoding and retrieval is that we do not know the time point at which an episode was encoded or retrieved. Therefore, we cross-correlated the instantaneous firing rate during encoding with the instantaneous firing rate during the corresponding retrieval trial (maximum lag of ±2.5s). The maximum value of this sequence served as our empirical reinstatement value. We then shuffled the encoding and retrieval order and recomputed this reinstatement value 1,000 times. The 99th percentile of these values was used as a threshold. If the empirical reinstatement value reached this threshold, we considered the neuron a temporal Episode Specific Neuron (tESN). In the next step, for each neuron we randomly drew one of the permutations we calculated previously. Neurons whose permuted values reached or exceeded the threshold were considered tESNs under the null hypothesis. We repeated this process 1,000 times to build a null distribution against which we compared our empirical number of tESNs.

For experiment 2, we further excluded all trials in which the given neuron showed a significant visual tuning using the methodology outlined under #Identification of putative Concept Neurons. To evaluate the extent of trial overlap between rate code (ESNs) and temporal code (tESNs) reinstated trials, we used a permutation test. The trial identity (reinstated vs. non-reinstated) was shuffled within each neuron, and the resulting overlap values were compared with the empirical overlap. The analysis showed a significant overlap between rate code and temporally reinstated episodes in Experiment 1 and Experiment 2 ($p < 0.001$). We tested the validity of this analysis by repeating the same analysis using random spike times. We generated these random spike times by first rounding the empirical spike times to the nearest integer and then drawing an equal number of pseudorandom integer values from a discrete uniform distribution between the first

and last empirical spike times. We next followed a similar approach to ensure the robustness of our analysis in the tESN identification as outlined in #Simulation of ESN identification, whereby we repeated the analysis using 500 surrogate datasets. These datasets were created by segmenting all spike times into episodes in the order they occurred, and then circularly shuffling them. Upon analysis, the results indicated that the percentage of significant tESN identifications was below the 5% threshold (4.18%), providing further evidence of the credibility and reliability of our analysis.

Spike density calculation

To produce the visualisations in Figure 4, we extracted spikes from one second before the cue onset until five seconds after cue onset for each episode. Binary spike times were convolved with a 100 ms Gaussian kernel (length: ±300ms, peak normalized to one) to create a time-resolved signal of spike activity. We computed the average firing rate over time for all episodes (*ep*) during the baseline (*BL*) period 1,000 ms preceding the animal cue ($\overline{x}_{BL}$). We then z-scored the spike activity during the episode ($x_{ep,t}$) using the standard deviation ($std(\overline{x}_{BL})$) and mean ($\overline{\overline{x}}_{BL}$) across all pre-cue baseline periods (see equation (1)). To account for instances where no spiking activity occurred during the baseline period, 0.1 (see Ison et al., 2015) was added to the standard deviation ($std(\overline{x}_{BL})$). Episodes were then split into reinstated and non-reinstated episodes. Firing rates for each episode type (reinstated/non-reinstated) were then averaged over ESNs.

$$z_{ep,t} = \frac{x_{ep,t} - \overline{\overline{x}}_{BL}}{std(\overline{x}_{BL}) + 0.1} \quad (1)$$

We employed a bootstrapping technique (*N* = 100 random draws) to ensure that the number of reinstated and non-reinstated episodes was the same for each ESN. Next, we performed a computation of the standard error of the mean (SEM) and the cluster-based permutation test (Maris and Oostenveld, 2007). The grayscale inset at the bottom of the figure shows the proportion of iterations that had a significant cluster at a specific time point.

Identification of Time Cells

We defined the beginning of an encoding block as the most salient event. Based on Umbach and colleagues (Umbach et al., 2020), we then extracted all spikes within each block and convolved them with a 251 ms Gaussian kernel (width factor: 2.5). This created a block number x time points matrix. For our first analysis, we cut each encoding block into 40 equally sized bins, thereby normalizing block duration. We then used a Kruskal-Wallis test to determine whether any of the 40 bins significantly differed from each other.

We then performed a circular shifting permutation test to calculate whether we found a significant number of Time Cells. This is done by shifting a random number of values from the beginning of the vector to the end. This shifting was imposed on each block separately and repeated N = 10,000 times for every single unit.

In a second test, the block length was determined by the longest block and shorter blocks were filled up with NaN values. This resulted in no normalization of time between blocks. The rest of the procedure is the same as described in the above paragraph.

References

Addante, R. J., Watrous, A. J., Yonelinas, A. P., Ekstrom, A. D., & Ranganath, C. (2011). Prestimulus theta activity predicts correct source memory retrieval. Proceedings of the National Academy of Sciences *108*(26), 10702-10707.

Bakker, A., Kirwan, C.B., Miller, M., and Stark, C.E. (2008). Pattern separation in the human hippocampal CA3 and dentate gyrus. Science *319*, 1640-1642.

Berron, D., Schütze, H., Maass, A., Cardenas-Blanco, A., Kuijf, H.J., Kumaran, D., and Düzel, E. (2016). Strong evidence for pattern separation in human dentate gyrus. Journal of Neuroscience *36*, 7569-7579.
Bowman, H., and Wyble, B. (2007). The simultaneous type, serial token model of temporal attention and working memory. Psychological review *114*, 38.
Chadwick, M.J., Hassabis, D., Weiskopf, N., and Maguire, E.A. (2010). Decoding individual episodic memory traces in the human hippocampus. Curr Biol *20*, 544-547. 10.1016/j.cub.2010.01.053.
Cooper, R.A., and Ritchey, M. (2020). Progression from Feature-Specific Brain Activity to Hippocampal Binding during Episodic Encoding. J Neurosci *40*, 1701-1709. 10.1523/JNEUROSCI.1971-19.2019.
Gelbard-Sagiv, H., Mukamel, R., Harel, M., Malach, R., and Fried, I. (2008). Internally generated reactivation of single neurons in human hippocampus during free recall. Science *322*, 96-101.
Ison, M.J., Quiroga, R.Q., and Fried, I. (2015). Rapid encoding of new memories by individual neurons in the human brain. Neuron *87*, 220-230.
Krotov, D., & Hopfield, J. (2020). Large associative memory problem in neurobiology and machine learning. *arXiv preprint arXiv:2008.06996*.
Lepage, M., Ghaffar, O., Nyberg, L., & Tulving, E. (2000). Prefrontal cortex and episodic memory retrieval mode. Proceedings of the National Academy of Sciences *97*(1), 506-511.
Lisman, J., Buzsáki, G., Eichenbaum, H., Nadel, L., Ranganath, C., and Redish, A.D. (2017). Viewpoints: how the hippocampus contributes to memory, navigation and cognition. Nature neuroscience *20*, 1434-1447.
Mack, M.L., and Preston, A.R. (2016). Decisions about the past are guided by reinstatement of specific memories in the hippocampus and perirhinal cortex. Neuroimage *127*, 144-157.
Maris, E., and Oostenveld, R. (2007). Nonparametric statistical testing of EEG-and MEG-data. Journal of neuroscience methods *164*, 177-190.
Marr, D. (1971). Simple memory: a theory for archicortex. Philos Trans R Soc Lond B Biol Sci *262*, 23-81. 10.1098/rstb.1971.0078.
Mormann, F., Dubois, J., Kornblith, S., Milosavljevic, M., Cerf, M., Ison, M., Tsuchiya, N., Kraskov, A., Quiroga, R.Q., and Adolphs, R. (2011). A category-specific response to animals in the right human amygdala. Nature neuroscience *14*, 1247-1249.
Mormann, F., Kornblith, S., Quiroga, R.Q., Kraskov, A., Cerf, M., Fried, I., and Koch, C. (2008). Latency and selectivity of single neurons indicate hierarchical processing in the human medial temporal lobe. Journal of Neuroscience *28*, 8865-8872.
Nader, K., & Hardt, O. (2009). A single standard for memory: the case for reconsolidation. Nature Reviews Neuroscience, *10*(3), 224-234.
Niediek, J., Boström, J., Elger, C.E., and Mormann, F. (2016). Reliable analysis of single-unit recordings from the human brain under noisy conditions: tracking neurons over hours. PloS one *11*, e0166598.
Nyberg, L., Tulving, E., Kapur, S., Houle, S., & Mclntosh, A. R. (1995). Functional brain maps of retrieval mode and recovery of episodic information. NeuroReport *7*, 249-252.
Pacheco Estefan, D., Sánchez-Fibla, M., Duff, A., Principe, A., Rocamora, R., Zhang, H., ... & Verschure, P. F. (2019). Coordinated representational reinstatement in the human hippocampus and lateral temporal cortex during episodic memory retrieval. Nature communications, *10*(1), 1-13.

Parish, G., Michelmann, S., Hanslmayr, S., & Bowman, H. (2021). The Sync-Fire/deSync model: Modelling the reactivation of dynamic memories from cortical alpha oscillations. Neuropsychologia, *158*, 107867.
Prestigio, C., Ferrante, D., Valente, P., Casagrande, S., Albanesi, E., Yanagawa, Y., ... & Baldelli, P. (2019). Spike-related electrophysiological identification of cultured hippocampal excitatory and inhibitory neurons. Molecular Neurobiology, *56*(9), 6276-6292.
Quiroga, R.Q. (2012). Concept cells: the building blocks of declarative memory functions. Nature Reviews Neuroscience *13*, 587-597.
Quiroga, R. Q. (2019). Plugging in to human memory: advantages, challenges, and insights from human single-neuron recordings. Cell, *179*(5), 1015-1032.
Quiroga, R.Q. (2020). No pattern separation in the human hippocampus. Trends in Cognitive Sciences *24*, 994-1007.
Quiroga, R.Q., Reddy, L., Kreiman, G., Koch, C., and Fried, I. (2005). Invariant visual representation by single neurons in the human brain. Nature *435*, 1102-1107.
Reddy, L., Zoefel, B., Possel, J.K., Peters, J., Dijksterhuis, D.E., Poncet, M., van Straaten, E.C., Baayen, J.C., Idema, S., and Self, M.W. (2021). Human hippocampal neurons track moments in a sequence of events. Journal of Neuroscience *41*, 6714-6725.
Rutishauser, U., Mamelak, A.N., and Schuman, E.M. (2006). Single-trial learning of novel stimuli by individual neurons of the human hippocampus-amygdala complex. Neuron *49*, 805-813.
Rutishauser, U., Schuman, E.M., and Mamelak, A.N. (2008). Activity of human hippocampal and amygdala neurons during retrieval of declarative memories. Proceedings of the National Academy of Sciences *105*, 329-334.
Rutishauser, U., Ye, S., Koroma, M., Tudusciuc, O., Ross, I.B., Chung, J.M., and Mamelak, A.N. (2015). Representation of retrieval confidence by single neurons in the human medial temporal lobe. Nature neuroscience *18*, 1041-1050.
Squire, L.R. (1992). Memory and the hippocampus: a synthesis from findings with rats, monkeys, and humans. Psychological review *99*, 195.
Tankus, A., Yeshurun, Y., and Fried, I. (2009). An automatic measure for classifying clusters of suspected spikes into single cells versus multiunits. Journal of neural engineering *6*, 056001.
Teyler, T.J., and DiScenna, P. (1986). The hippocampal memory indexing theory. Behavioral neuroscience *100*, 147.
Teyler, T.J., and Rudy, J.W. (2007). The hippocampal indexing theory and episodic memory: updating the index. Hippocampus *17*, 1158-1169.
Tulving, E. (1983). *Elements of Episodic Memory*. Oxford: Clarendon Press.
Tulving, E. (2002). Episodic memory: From mind to brain. Annual review of psychology *53*, 1-25.
Tulving, E. (1972). Episodic and semantic memory. In E. Tulving & W. Donaldson (Eds.), *Organization of memory*. New York: Academic Press.
Umbach, G., Kantak, P., Jacobs, J., Kahana, M., Pfeiffer, B.E., Sperling, M., and Lega, B. (2020). Time cells in the human hippocampus and entorhinal cortex support episodic memory. Proceedings of the National Academy of Sciences *117*, 28463-28474.
Urgolites, Z. J., Wixted, J. T., Goldinger, S. D., Papesh, M. H., Treiman, D. M., Squire, L. R., & Steinmetz, P. N. (2022). Two kinds of memory signals in neurons of the human hippocampus. Proceedings of the National Academy of Sciences *119*, e2115128119.
Whittington, J. C., Warren, J., & Behrens, T. E. J. (2022). Relating transformers to models and neural representations of the hippocampal formation. *arXiv preprint arXiv:2112.04035*.

Wixted, J.T., Goldinger, S.D., Squire, L.R., Kuhn, J.R., Papesh, M.H., Smith, K.A., Treiman, D.M., and Steinmetz, P.N. (2018). Coding of episodic memory in the human hippocampus. Proceedings of the National Academy of Sciences *115*, 1093-1098.
Wixted, J.T., Squire, L.R., Jang, Y., Papesh, M.H., Goldinger, S.D., Kuhn, J.R., Smith, K.A., Treiman, D.M., and Steinmetz, P.N. (2014). Sparse and distributed coding of episodic memory in neurons of the human hippocampus. Proceedings of the National Academy of Sciences *111*, 9621-9626.
Zheng, J., Schjetnan, A. G., Yebra, M., Gomes, B. A., Mosher, C. P., Kalia, S. K., ... & Rutishauser, U. (2022). Neurons detect cognitive boundaries to structure episodic memories in humans. Nature Neuroscience, *25*(3), 358-368.

Acknowledgements:

We thank all the patients who have participated in our study.

Funding:

European Research Council grant 647954 (SH).

European Research Council grant 715714 (MW).

Wellcome Trust/Royal Society Sir Henry Dale Fellowship (107672/Z/15/Z) (BPS).

Author contributions:

Conceptualization: LDK, SH

Data curation: LDK

Formal analysis: LDK, SH

Funding acquisition: SH, BPS, MW

Investigation: LDK, FR, GP, MTW, MVDP, KW

Project administration: RC, VS, DTR, JL, SG, KW, RH, GK, HH

Software: LDK

Supervision: SH, HB

Visualization: LDK, SH, MTW

Writing – original draft: LDK, SH

Writing – review & editing: all authors

Competing interests:

The authors declare that they have no competing interests.

Data and materials availability:

Data and analysis code is currently available in a private repository on figshare (https://figshare.com/s/12fcaf4069e972a5c362) and will be made publicly available upon publication.

Supplementary Materials

Materials and Methods

Figures S1 to S7

Tables S1 and S2