doi:10.1038/nature08860

LETTERS

Human memory strength is predicted by theta-frequency phase-locking of single neurons

Ueli Rutishauser[1], Ian B. Ross[3], Adam N. Mamelak[3,4]* & Erin M. Schuman[1,2]†*

Learning from novel experiences is a major task of the central nervous system. In mammals, the medial temporal lobe is crucial for this rapid form of learning[1]. The modification of synapses and neuronal circuits through plasticity is thought to underlie memory formation[2]. The induction of synaptic plasticity is favoured by coordinated action-potential timing across populations of neurons[3]. Such coordinated activity of neural populations can give rise to oscillations of different frequencies, recorded in local field potentials. Brain oscillations in the theta frequency range (3–8 Hz) are often associated with the favourable induction of synaptic plasticity as well as behavioural memory[4]. Here we report the activity of single neurons recorded together with the local field potential in humans engaged in a learning task. We show that successful memory formation in humans is predicted by a tight coordination of spike timing with the local theta oscillation. More stereotyped spiking predicts better memory, as indicated by higher retrieval confidence reported by subjects. These findings provide a link between the known modulation of theta oscillations by many memory-modulating behaviours and circuit mechanisms of plasticity.

Many factors such as novelty, attention and arousal modulate the efficacy of memory formation[5]. The activity of many neurons is modulated by the theta rhythm[6–8]. In cellular studies, whether a given stimulus triggers synaptic changes can depend on when the stimulus arrives relative to the ongoing theta oscillation[9]. Theta oscillations are modulated by behaviour and brain state and have been implicated in memory formation[10–12]. For example, animals that exhibit more theta activity learn a new task faster than animals that have less pronounced theta oscillations[13]. Although they are most extensively studied in rodents, theta oscillations are also prominent in humans[14–16]. Theta oscillations are clearly observed in many structures, among them the hippocampus and the amygdala[4,8,14,15,17]. The amplitude of theta oscillations during learning is correlated with later retrieval success[18,19]. How the activity of individual neurons of the medial temporal lobe (MTL) relates to theta oscillations during learning, however, has not been addressed. In this work, we tested the idea that the timing of neuronal spikes relative to the ongoing theta rhythm during learning influences subsequent memory performance in humans.

Subjects performed a memory test that had two parts, learning and recognition (Fig. 1a). During learning, subjects viewed a set of 100 previously unseen images including cars, people, animals and tools. Each image was presented for 1 s. During recognition (at least 15 min later), subjects again viewed a set of 100 images, 50 of which were new and 50 of which had been presented during the learning trials. Subjects indicated whether they had seen the image before (response 'old') or not (response 'new') on a six-point confidence scale (Fig. 1a). Subjects remembered well: the average sensitivity, d', was 1.26 ± 0.08 (mean $\pm$ s.e.m., $P < 0.001$; 14 sessions from nine patients; see Supplementary Table 1). Subjects' reports of confidence mapped well onto their performance (Fig. 1b–d), indicating that they had a good sense of the quality of their memories. The receiver operating characteristics were asymmetrical (Fig. 1c), which is a signature of MTL-dependent recognition memory.

While subjects performed the above task, we recorded single-unit neuronal activity as well as the local field potential (LFP) using microwires implanted in the amygdala and hippocampus (Supplementary Fig. 5). We isolated 305 single neurons (120 from the hippocampus and 185 from the amygdala; see Supplementary Fig. 2 for spike sorting). Neurons fired, on average, 1.67 ± 0.13 spikes per second (Supplementary Fig. 3). Many neurons ($n = 65$, 21%) responded to the presentation of the image with a change in firing rate (Fig. 1e–h). This did not, however, predict whether subjects later remembered the stimulus (Fig. 1h), nor did a change in the population firing rate (Supplementary Fig. 4e, f). In the LFP, we observed prominent low-frequency (<10-Hz) oscillations (Fig. 2a). We found that a substantial fraction of all neurons (21%, 51 of the 246 that fired enough spikes; Methods) produced spikes that were phase-locked to an oscillation in the 3–8-Hz range (an example neuron is shown in Fig. 2b–f). Neurons were most phase-locked to LFP oscillations at a frequency of 5.14 ± 0.72 Hz. Neurons had a range of phase preferences (Fig. 2g), but the majority (61%) fired close to either the peak or the trough of the theta-frequency oscillation ($\pm 45^\circ$ around peak or trough).

Next we considered whether there is an LFP/spike timing relationship during learning that predicts whether the subject will successfully store a memory of the image and remember it later. Using the behavioural performance from the recognition memory trials, we categorized each learning trial according to whether the image presented was later remembered (a true positive) or forgotten (a false negative). We then compared the neuronal activity between the two types of trials. We hypothesized that successful memory formation is more likely when neurons fire synchronously with their neighbours.

We examined the timing relationship between the spikes of single neurons and the ongoing theta oscillation by quantifying the spike-field coherence (SFC)[20]. The SFC varies as a function of frequency and yields values between 0% and 100%. The larger the SFC, the more accurately the spikes follow a particular phase of this frequency (Supplementary Fig. 1). We calculated the SFC separately for subsequently remembered or forgotten learning trials (for each neuron). All spikes of phase-locked neurons for all true-positive and false-negative trials were included ($n = 33$ neurons with at least 50 spikes for each trial type; see Methods). We found a striking difference between the two kinds of trials: the theta-range SFC was ~50% higher in trials that were later remembered relative to trials that were later forgotten (Fig. 3a). The average SFC (across all neurons) in the 3–7-Hz range

[1]Computation and Neural Systems and Division of Biology, [2]Howard Hughes Medical Institute, California Institute of Technology, Pasadena, California 91125, USA. [3]Department of Neurosurgery, Epilepsy & Brain Mapping Unit, Huntington Memorial Hospital, Pasadena, California 91105, USA. [4]Department of Neurosurgery, Cedars-Sinai Medical Center, Los Angeles, California 90048, USA. †Present address: Max Planck Institute for Brain Research, Max-von-Laue Strasse 3, 60438 Frankfurt am Main, Germany.
*These authors contributed equally to this work.

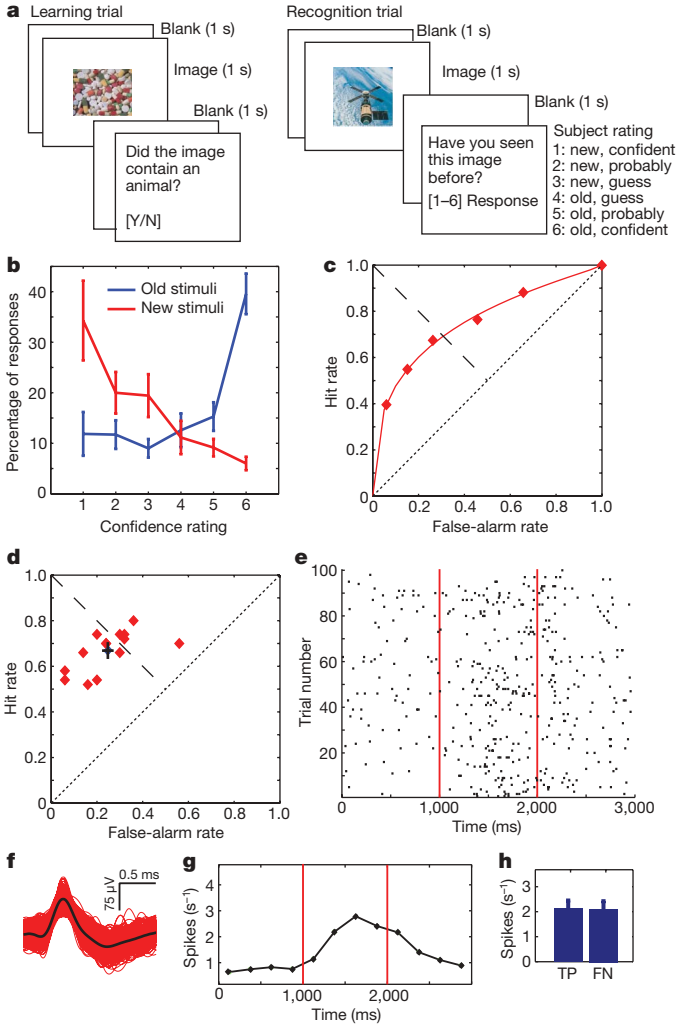

Figure 1 | Task, behavioural results and single-neuron example. **a**, Structure of the task. **b**, Percentage of responses 'new' and 'old' as a function of confidence. The highest probabilities for responses new and old were 1 and 6, respectively. Error bars, s.e.m. ($n = 14$ sessions). **c**, The average receiver operating characteristic of all subjects (average d' value, 1.26 ± 0.08; average slope of the z-score transform of the receiver operating characteristic, 0.76 ± 0.02 (significantly less than 1, $P = 0.003$)). Each red data point represents a confidence level, starting with the highest (6: old, confident) in the lower left quadrant. The dotted diagonal line indicates chance performance. **d**, The confidence level closest to the dashed diagonal line was chosen as the threshold that divides remembered images (hits) from forgotten images (misses) (red, individual sessions; blue, average). **e**, Single neuron in the left hippocampus during learning. The image was presented for 1 s (red lines). This neuron fired significantly more spikes following stimulus onset than in comparison with baseline ($P < 10^{-10}$). **f**, The raw waveforms for all spikes shown on the raster plot in **e** ($n = 516$). **g**, Average firing rate of the neuron as a function of time. **h**, The firing rate of this neuron did not differ significantly between stimuli that were later remembered and those that were forgotten (true positive (TP) versus false negative (FN), $P = 0.99$; $n = 35$ TP trials and 15 FN trials). Error bars, s.e.m.

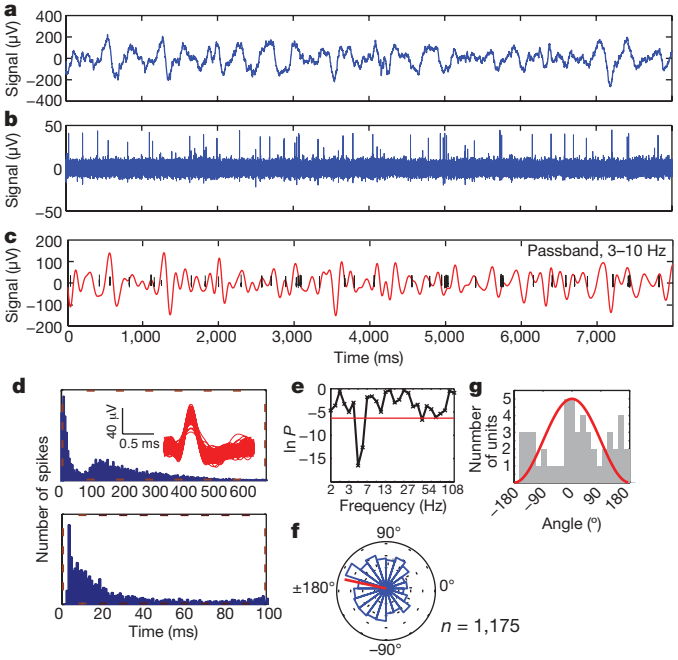

Figure 2 | Relationship between spikes and the theta oscillation (3–8 Hz) in the LFP. **a**, Raw signal recorded on a single microwire (2-Hz high-pass filtered). **b**, Spikes (300-Hz high-pass filtered), recorded on the same channel. All spikes (from multiple neurons) recorded on this channel are shown. **c**, The low-frequency component of the signal shown in **a** (passband, 3–10 Hz), together with a subset of the spikes (black) that were attributed to the single unit shown in **d**. **d**, Top: interspike interval frequency plot and waveforms from the single neuron shown in **c**. We note the two modes of the interspike interval. Bottom: enlargement of the 0–100-ms section of the top plot. **e**, Significance of phase-locking (Rayleigh test) as a function of frequency (1–180 Hz, logarithmically spaced). The threshold (red line) for significant phase-locking was set to $P = 0.0018$ (0.05/28, Bonferroni-corrected). The neuron shown exhibited maximal phase-locking at 4.8 Hz. **f**, Histogram of spike phase for the 4.8-Hz oscillation. The mean phase (167°) is near the trough (red; $R = 0.11$ (mean vector length), $P < 10^{-7}$, $\kappa = 0.22$ (circular variance)). **g**, Histogram of the preferred phase of all neurons that were phase-locked to an LFP oscillation in the 3–8-Hz range ($n = 51$ of 246, 21%). The red line indicates the phase notation used.

was approximately 43% higher in true-positive trials than in false-negative trials ($P = 0.00003$, two-tailed paired t-test; $2.02 \pm 0.19\%$ versus $1.41 \pm 0.13\%$). As a control, we randomly reassigned the label of true positive or false negative to the same set of learning trials and conducted the analysis again (Methods). The difference in the SFC was abolished (Fig. 3b). In the analysis described above, we pooled neurons from the hippocampus and amygdala, but the average theta-range SFC was also significantly different for amygdala ($n = 24$, $P = 0.0021$) and hippocampus ($n = 9$, $P = 0.0024$) neurons considered separately.

To further analyse the data, we used the spike-triggered average (STA). The STA is constructed by averaging LFP segments of ±480 ms centred on every spike. The resulting trace (Fig. 3c) deviates from zero if a systematic relationship exists between spike timing and the LFP. In our experiments, the STA of each significantly phase-locked neuron showed strong oscillations in the 3–8-Hz range (Fig. 3c, d; see Supplementary Fig. 9 for more examples). Comparing the STA (equalized for number of spikes to avoid bias) between remembered and forgotten trials illustrates the differences in phase-locking that predict the later memory performance (Fig. 3c, d). A second measure, the spike-triggered power (STP), quantifies the power of the oscillations that are present in the LFP. It is calculated by averaging the spectra of the individual LFP segments (each centred on a spike). The average STP of all theta-locked neurons indicated the strong presence of oscillations in the 4–8-Hz range (Fig. 3e). The power of the LFP at the time of spike occurrence (quantified by the STP) did not distinguish, however, between trials that were later remembered and those that were forgotten (Fig. 3f). We also compared the power of the entire LFP trace between remembered and forgotten trials and found no significant difference (Supplementary Fig. 7). Stimulus presentation can reset the phase of ongoing oscillations without changing their amplitudes[21]. In working memory tasks, such phase resets can differ between different trial types[21]. In our task, we found that the stimulus onset also reset ongoing theta oscillations (Supplementary Fig. 11a). We quantified phase resets for both true-positive and false-negative trials using the inter-trial coherence and found no significant difference (Supplementary Fig. 11c–g). Thus, the significant difference in the

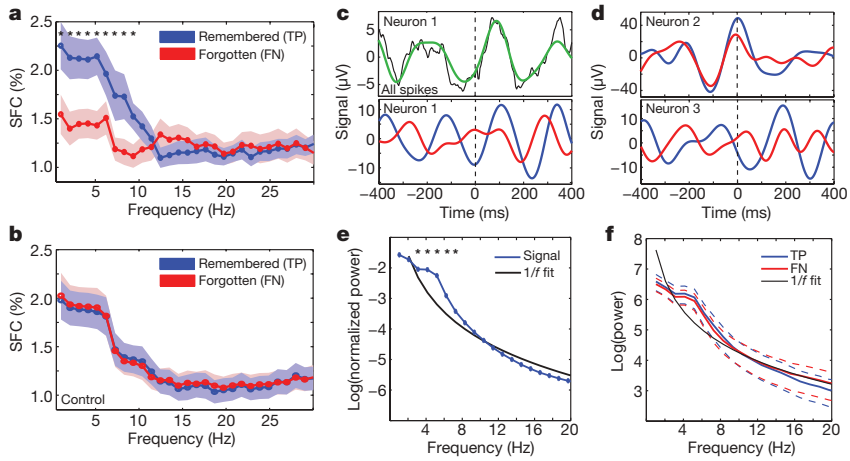

Figure 3 | The SFC distinguishes between learning trials that were later remembered and those that were forgotten. **a**, Strength of phase-locking as a function of frequency. Spikes fired during trials that were later remembered (blue) had a significantly higher SFC than spikes fired during trials that were not remembered (red). The SFCs were significantly different in the range of 2–10 Hz (false discovery rate (FDR) corrected for 100 multiple comparisons at 1–100 Hz, $^*q < 0.05$). The shaded area indicates s.e.m. for $n = 33$ cells. **b**, Control. The difference was abolished by the random assignment of true-positive and false-negative labels. **c**, Top: raw (black) and filtered (3–8 Hz; green) STAs of a single cell (recorded from the amygdala), constructed using $n = 578$ spikes. Bottom: STAs for spikes fired during trials that were later remembered (blue) and during trials that were later forgotten (red). **d**, STAs from two further single cells (top, hippocampus; bottom, amygdala), colour-coded as shown in **c**. All three STAs shown are from different patients. **e**, Average spectrum of the STP of all cells ($n = 33$) used in **a**. The presence of 3–8-Hz oscillations is clear. The black line indicates the power levels expected if the power were distributed as $1/f$ (fit of Af^{α}: $\alpha = -1.93 \pm 0.08$; A is a constant). The total power of each cell was normalized to 1. **f**, Same as in **e**, but calculated separately for both categories of trial. There was no significant difference. The areas inside the dotted lines indicate s.e.m.

SFC (Fig. 3a) is due to changes in phase-locking rather than changes in the amplitude or phase of the oscillations.

The SFC differed by 10–50% between stimuli that were later remembered or forgotten (Fig. 3). We used simulations to estimate how much spike-timing jitter is required to generate a difference of this magnitude (Supplementary Fig. 1h) and found that spike-timing noise of ±20 ms reduced the SFC in magnitude by 50%. Thus, from a computational perspective, an increase in spike-timing jitter of ±20 ms can account for the reduction in phase-locking we observed. Generally, the neurons we observed fired, on average, a few spikes

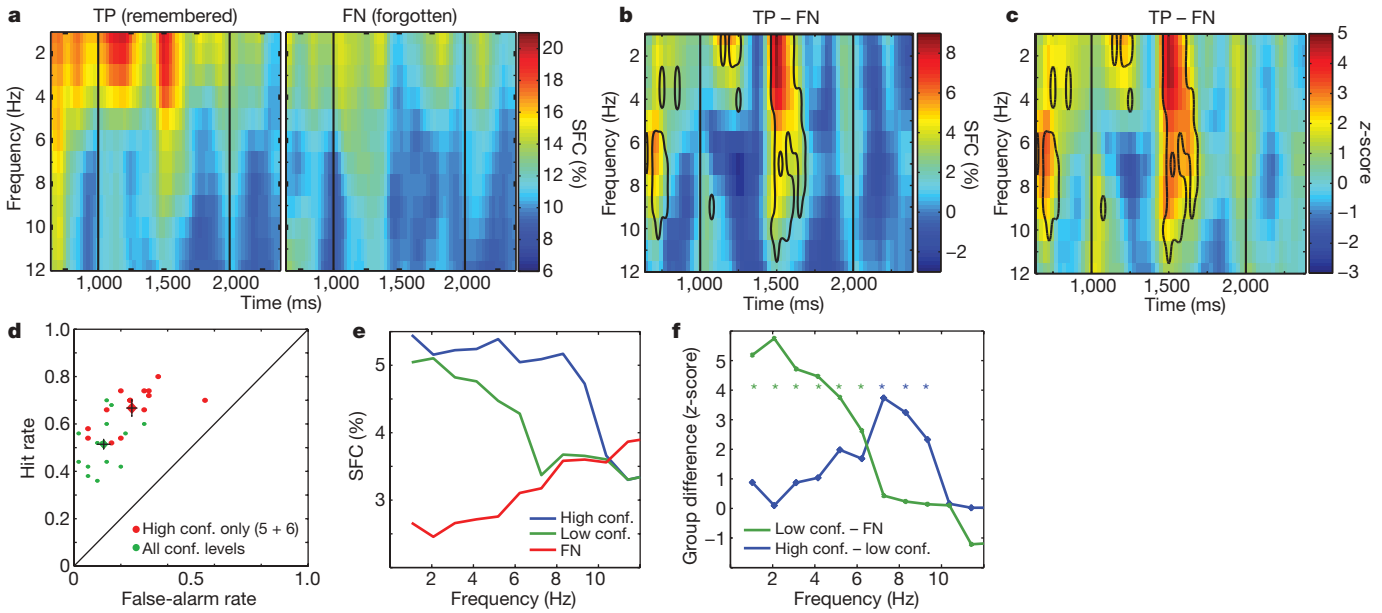

Figure 4 | SFC time course and relationship to subject confidence. **a**–**c**, Time courses of SFC differences. The time courses for true-positive (left) and false-negative (right) trials is plotted in **a**. The time course of the difference in SFC between true-positive and false-negative trials is plotted in **b**. The latency, relative to stimulus onset, of the most distinct difference was ~500 ms but there were also differences before stimulus onset. **c**, Significance of the difference, expressed as z-score (bootstrap statistic). The contours in **b** and **c** indicate areas of significantly higher SFC in true-positive trials (FDR corrected, $q < 0.05$; $z(q < 0.05) = 2.6$). A test for higher SFC in false-negative trials revealed no significant areas. **d**–**f**, Comparison of SFC (for the entire trial) between different levels of retrieval confidence. The behavioural performance of patients for only high-confidence (green) and all (red) memories is shown in **d**. The diamonds indicate the average performances (error bars, s.e.m.; $n = 14$ sessions). The SFC as a function of frequency and retrieval confidence is shown in **e**. Statistics of the pairwise difference of the SFCs (**e**) are shown in **f**. Asterisks indicate frequencies with significant differences (FDR corrected for 10 comparisons, $q < 0.05$), tested against the empirical null distribution (Methods). Differences are plotted in terms of z-scores relative to the same null distribution. We note that z-scores are shown only to illustrate the extent of the effect. The significances indicated (asterisks) are based directly on the P values. Note that the absolute values of the SFC are not comparable between panels. In **a**–**c**, all $n = 33$ cells were included; in **e**–**f**, $n = 26$ cells were included.

during every trial. Nevertheless, we observed a robust difference, which highlights the importance of single spikes.

In the analysis described above, we considered every spike fired during each learning trial. We next estimated the SFC as a function of time (sliding window of ±125 ms advanced in steps of 25 ms; Fig. 4a–c). We found that the accuracy of phase-locking was most powerful in two different phases of the learning trial: shortly before stimulus onset and ~500 ms after stimulus onset (Fig. 4a–c and Supplementary Fig. 6c). Thus, the increased SFC in later-remembered trials arises from two epochs of spike timing/LFP synchrony. This indicates the potential importance of the timing of to-be-learned stimuli, relative to the current dynamical brain state[22]. There are therefore two factors that can influence the accuracy of phase-locking: the onset (or the expectation thereof) of the stimulus, and the state before stimulus onset, which reflects whether the MTL is in a receptive or a non-receptive state.

Patients indicated a strong memory of some stimuli whereas for others they had only weak memories. We thus considered whether there was a difference (during learning) that indicated whether an image would be remembered with confidence or not. The confidence judgments made by subjects were well correlated with performance (the false-positive rate in high-confidence trials (12 ± 2%) was significantly lower than that observed in all trials (24 ± 3%); Fig. 4d). We thus next considered the SFC in forgotten (false negative), weakly remembered (true positive, confidence level 4) and strongly remembered (true positive, confidence levels 5 and 6) learning trials separately. We found that the SFC was significantly higher in true-positive trials with low confidence than in false-negative trials (significantly higher in the 3–6-Hz range; Fig. 4e, f). Furthermore, the SFC was also significantly higher (7–9-Hz range) in trials that were later remembered with high confidence (Fig. 4e, f) than it was in the weak-confidence trials. Thus, in addition to distinguishing remembered from forgotten stimuli, the accuracy of phase-locking during learning was also related to the subjective strength of memory (Fig. 4f). Different levels of subjective memory strength were associated with SFC in different theta ranges (3–6 Hz and 6–9 Hz), suggesting that two different encoding processes could be engaged simultaneously.

There are several possible mechanisms that could modulate spike-timing accuracy. Theta oscillations modulate excitability in areas projecting to and receiving information from the MTL[15,17,23]. One possibility is thus that upstream neurons are differentially coordinated with the MTL theta rhythm. Such coordination can occur without changing the local theta power. For example, activation of neurons of the locus coeruleus or dopaminergic neurons of the ventral tegmental area influences theta oscillations[24,25] and, therefore, phase-locking of neurons. Both of these types of neurons can be activated by novel stimuli. A second hypothesis is thus that noradrenaline or dopamine release controls spike-timing accuracy.

Episodic memories result from singular experiences[1]. The detection of novel stimuli is crucial for such learning and depends strongly on the integrity of the MTL[1,26]. Some neurons respond specifically to novel stimuli[27,28], and the firing rate of these neurons during memory retrieval correlates with memory strength[29]. We incorporated both features in our task. We showed that spike-timing accuracy was apparent in both the amygdala and the hippocampus. Both structures are important for the modulation of memory strength[5,30] and contain neurons phase-locked to theta oscillations[7,8]. The correlation between the accuracy of phase-locking during learning and later recognition establishes a direct relationship between a circuit-level physiological phenomenon (spike timing relative to the LFP) and its function (as expressed by human behaviour).

METHODS SUMMARY

Patients. The subjects were nine patients who were evaluated for possible surgical treatment of epilepsy using implantation of depth electrodes. The patients volunteered for the study and gave informed consent. The experiments were approved by the institutional review boards of the Huntington Memorial Hospital and the California Institute of Technology. We evaluated all patients using standard neuropsychological tests (Supplementary Table 1).

Electrophysiology. We recorded extracellular neural activity using 40-μm microwires inserted in a clinical depth electrode, which was implanted in the hippocampus and amygdala (bilaterally) as described previously[29]. Electrode locations were chosen according to clinical criteria alone. Target locations were verified using post-implantation structural magnetic resonance imaging (Supplementary Fig. 5). All recordings were locally grounded (bipolar recordings). Spikes were detected by the application of a threshold to an energy power signal computed from the raw trace. We sorted spikes using a template-matching method with careful manual post-processing (Methods).

Data analysis. The spike-field coherence, SFC(f), is a function of frequency, f, and takes values between 0% and 100%. The SFC is the ratio of the frequency spectrum of the STA and the average frequency spectrum of the LFP traces. The SFC is normalized both for firing rate as well as for LFP power changes. The average spectrum of the LFP traces is the STP. Frequency spectra were calculated using multitaper analysis. When comparing the SFC between two conditions, we always used the same number of spikes for each condition.

Full Methods and any associated references are available in the online version of the paper at www.nature.com/nature.

Received 19 August 2009; accepted 27 January 2010.
Published online 24 March 2010.

1. Squire, L. R., Stark, C. E. & Clark, R. E. The medial temporal lobe. *Annu. Rev. Neurosci.* **27,** 279–306 (2004).
2. Martin, S. J., Grimwood, P. D. & Morris, R. G. Synaptic plasticity and memory: an evaluation of the hypothesis. *Annu. Rev. Neurosci.* **23,** 649–711 (2000).
3. Markram, H., Lubke, J., Frotscher, M. & Sakmann, B. Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* **275,** 213–215 (1997).
4. Buzsáki, G. Theta oscillations in the hippocampus. *Neuron* **33,** 325–340 (2002).
5. Paller, K. A. & Wagner, A. D. Observing the transformation of experience into memory. *Trends Cogn. Sci.* **6,** 93–102 (2002).
6. Buzsáki, G., Leung, L. W. & Vanderwolf, C. H. Cellular bases of hippocampal EEG in the behaving rat. *Brain Res.* **287,** 139–171 (1983).
7. Jacobs, J., Kahana, M. J., Ekstrom, A. D. & Fried, I. Brain oscillations control timing of single-neuron activity in humans. *J. Neurosci.* **27,** 3839–3844 (2007).
8. Paré, D. & Gaudreau, H. Projection cells and interneurons of the lateral and basolateral amygdala: distinct firing patterns and differential relation to theta and delta rhythms in conscious cats. *J. Neurosci.* **16,** 3334–3350 (1996).
9. Pavlides, C., Greenstein, Y. J., Grudman, M. & Winson, J. Long-term potentiation in the dentate gyrus is induced preferentially on the positive phase of theta-rhythm. *Brain Res.* **439,** 383–387 (1988).
10. Winson, J. Loss of hippocampal theta rhythm results in spatial memory deficit in the rat. *Science* **201,** 160–163 (1978).
11. Hasselmo, M. E., Bodelon, C. & Wyble, B. P. A proposed function for hippocampal theta rhythm: separate phases of encoding and retrieval enhance reversal of prior learning. *Neural Comput.* **14,** 793–817 (2002).
12. Lisman, J. E. & Idiart, M. A. Storage of 7 ± 2 short-term memories in oscillatory subcycles. *Science* **267,** 1512–1515 (1995).
13. Berry, S. D. & Thompson, R. F. Prediction of learning rate from the hippocampal electroencephalogram. *Science* **200,** 1298–1300 (1978).
14. Cantero, J. L. *et al.* Sleep-dependent theta oscillations in the human hippocampus and neocortex. *J. Neurosci.* **23,** 10897–10903 (2003).
15. Kahana, M. J., Sekuler, R., Caplan, J. B., Kirschen, M. & Madsen, J. R. Human theta oscillations exhibit task dependence during virtual maze navigation. *Nature* **399,** 781–784 (1999).
16. Huh, K. *et al.* Human hippocampal EEG: effects of behavioral activation. *Neurology* **40,** 1177–1181 (1990).
17. Lee, H., Simpson, G. V., Logothetis, N. K. & Rainer, G. Phase locking of single neuron activity to theta oscillations during working memory in monkey extrastriate visual cortex. *Neuron* **45,** 147–156 (2005).
18. Klimesch, W., Doppelmayr, M., Russegger, H. & Pachinger, T. Theta band power in the human scalp EEG and the encoding of new information. *Neuroreport* **7,** 1235–1240 (1996).
19. Sederberg, P. B. *et al.* Hippocampal and neocortical gamma oscillations predict memory formation in humans. *Cereb. Cortex* **17,** 1190–1196 (2007).
20. Fries, P., Reynolds, J. H., Rorie, A. E. & Desimone, R. Modulation of oscillatory neuronal synchronization by selective visual attention. *Science* **291,** 1560–1563 (2001).
21. Rizzuto, D. S., Madsen, J. R., Bromfield, E. B., Schulze-Bonhage, A. & Kahana, M. J. Human neocortical oscillations exhibit theta phase differences between encoding and retrieval. *Neuroimage* **31,** 1352–1358 (2006).
22. Seager, M. A., Johnson, L. D., Chabot, E. S., Asaka, Y. & Berry, S. D. Oscillatory brain states and learning: impact of hippocampal theta-contingent training. *Proc. Natl Acad. Sci. USA* **99,** 1616–1620 (2002).
23. Siapas, A. G., Lubenov, E. V. & Wilson, M. A. Prefrontal phase locking to hippocampal theta oscillations. *Neuron* **46,** 141–151 (2005).

24. Brown, R. A., Walling, S. G., Milway, J. S. & Harley, C. W. Locus ceruleus activation suppresses feedforward interneurons and reduces β-γ electroencephalogram frequencies while it enhances θ frequencies in rat dentate gyrus. *J. Neurosci.* **25,** 1985–1991 (2005).
25. Orzeł-Gryglewska, J., Jurkowlaniec, E. & Trojniar, W. Microinjection of procaine and electrolytic lesion in the ventral tegmental area suppresses hippocampal theta rhythm in urethane-anesthetized rats. *Brain Res. Bull.* **68,** 295–309 (2006).
26. Knight, R. Contribution of human hippocampal region to novelty detection. *Nature* **383,** 256–259 (1996).
27. Rutishauser, U., Mamelak, A. N. & Schuman, E. M. Single-trial learning of novel stimuli by individual neurons of the human hippocampus-amygdala complex. *Neuron* **49,** 805–813 (2006).
28. Viskontas, I. V., Knowlton, B. J., Steinmetz, P. N. & Fried, I. Differences in mnemonic processing by neurons in the human hippocampus and parahippocampal regions. *J. Cogn. Neurosci.* **18,** 1654–1662 (2006).
29. Rutishauser, U., Schuman, E. M. & Mamelak, A. N. Activity of human hippocampal and amygdala neurons during retrieval of declarative memories. *Proc. Natl Acad. Sci. USA* **105,** 329–334 (2008).
30. Adolphs, R., Denburg, N. L. & Tranel, D. The amygdala's role in long-term declarative memory for gist and detail. *Behav. Neurosci.* **115,** 983–992 (2001).

Supplementary Information is linked to the online version of the paper at www.nature.com/nature.

Acknowledgements We thank all patients for their participation; the staff of the Huntington Memorial Hospital epilepsy unit; C. Heller, L. Philpot and W. Sutherling for their support; and F. Mormann, G. Laurent and R. Adolphs for discussion. Funding was provided by the Gordon and Betty Moore Foundation, the William T. Gimbel Discovery Fund and the Howard Hughes Medical Institute.

Author Contributions U.R. designed and performed experiments, wrote analysis methods, analysed data and wrote the paper; A.N.M. and I.B.R. performed surgery; A.N.M. designed experiments; E.M.S. designed experiments and wrote the paper. All authors discussed the results.

Author Information Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Correspondence and requests for materials should be addressed to E.M.S. (schumane@brain.mpg.de).

METHODS

Electrophysiology. We recorded spiking and LFP activity using a 64-channel Neuralynx system (Digital Lynx 10S, Neuralynx). Data was sampled at 32 kHz and stored continuously in raw form for later processing and analysis. The head stages used (HS-27 and HS-36, Neuralynx) had unity gain, very high input impedances ($\sim$1 TΩ) and no phase shift. Thus, there were no significant phase distortions at the LFP frequencies of interest[31].

Patients and electrodes. Each macroelectrode (four per patient) contained eight microwires, one of which was used as ground (possible number of channels, 28). In practice, we could successfully isolate at least one putative single unit on 9 ± 3 (s.d.) wires per patient. Target locations were verified using a human hippocampal atlas (fig. 115 in ref. 32). Magnetics resonance images were acquired on a Toshiba 1.5T system using a T1 FLAIR sequence. Only electrodes that could be localized to the hippocampus or the amygdala were included. All patients (Supplementary Table 1) had good recognition memory as well as clearly distinguishable spiking activity on at least one electrode. Four additional patients were excluded owing to having at-chance behaviour ($n = 2$) or no recordable spiking activity ($n = 2$).

Task. During learning trials, 100 novel images (each unique and shown only once) were shown for 1 s each. For some patients, only 50 novel images were shown owing to relatively poor memory. To encourage focus, patients were asked whether there was an animal in the image after every learning trial. Twenty per cent of the images contained an animal, and patients successfully indicated the presence or absence of an animal in >98% of the trials. During recognition trials, 100 images were shown: 50 were identical to images presented during learning trials (old) and 50 were novel (new). Patients had to identify each picture as old or new on a six-point scale. Performance feedback was only provided at the end of the experiment. We made the task difficult enough to ensure that patients forgot about 30% of previously shown images (thus responding 'new' when shown an image that was previously viewed). Between the learning and recognition block there was a >15-min delay, during which time a distraction task was administered (Stroop) to prevent active rehearsal.

All images were shown at the centre of the screen of a notebook computer placed in front of the patient. Patients responded by pressing marked buttons on a keyboard. The distance to the screen was approximately 50 cm and the visual angle of the screen was approximately 30° by 23°. Stimuli were 9° by 9°. A trial consisted of the following displays (in this order): delay (1 s), stimulus (1 s), delay (1 s) and question. During delay periods, the screen was blank. A normally distributed random delay of 90 ± 30 ms (s.d.) was added to the pre-stimulus delay period to make the exact appearance time of the stimulus unpredictable. After the delay, the question was displayed until an answer was provided. The answer could only be provided when the question was on the screen, to avoid artefacts. Stimuli were photographs of natural scenes (see Supplementary Fig. 12 for examples) of five different visual categories (animals, people, cars/vehicles, outdoor scenes/houses and flowers/food items). There were the same numbers of images presented in each category. The categories were balanced during retrieval to avoid bias in memory for individual subjects towards certain categories. All stimuli were novel and had never been seen by the patient. Each stimulus was presented at most twice (once during learning and once during retrieval). The task was implemented in MATLAB (version 2008b, Mathworks) using the Psychophysics Toolbox (version PTB-2)[33].

Data analysis: behaviour. All data analysis was performed using custom written MATLAB programs (Mathworks). We quantified retrieval performance using a receiver operating characteristic (ROC) analysis and d'. Each point in the ROC (Fig. 1c, d) is a function of the proportion of correctly remembered items (hit rate, y axis) and the proportion of falsely remembered items (false-alarm rate, x axis) at a given confidence level. The point in the lower left corner (lowest false-positive as well as true-positive rate) corresponds to the highest confidence level (old, confident). The degree of asymmetry of the ROC (individual subject as well as average) was assessed by fitting the z-transformed ROC with a straight line (using least-squares regression). The slope of this line is one if the ROC is symmetric and significantly less than one if it is asymmetric. As a summary measure of the entire ROC, we used d'. To divide trials into those remembered and those forgotten, we dynamically determined the threshold for each patient who localized the decision point most closely to the dashed diagonal in Fig. 1d. Points lying exactly on this line represent zero bias. This procedure allowed us to correct for individual biases in confidence judgment based on behavioural performance.

Data analysis: spike sorting. Signals were filtered with a band-pass filter of 300–3,000 Hz . All filters used were non-causal zero-phase-shift filters (fourth-order Butterworth). Spikes were sorted using a template-matching method with careful manual post-processing/evaluation[34]. We identified at least one single unit from 126 microwires, 26 of which had one unit and 100 of which had more than one unit (on average 2.8 ± 0.1 units on wires with more than one). We quantified the quality of unit isolation by the percentage of all interspike intervals that were shorter than 3 ms. We found that 0.25 ± 0.02% of all interspike intervals were shorter than 3 ms (Supplementary Fig. 2a). The signal-to-noise ratio (as defined in ref. 34) of the mean waveforms of each cluster relative to the background noise was 2.0 ± 0.05 (Supplementary Fig. 2b). For wires that had several clusters, we additionally quantified the goodness of separation by applying the projection test[34] to each possible pair of neurons. The projection test measures the number of standard deviations by which the two clusters are separated after normalizing the data, so that each cluster is normally distributed with a standard deviation of one. The average distance between all possible pairs ($n = 328$) was 18.0 ± 0.7 (Supplementary Fig. 2c).

Data analysis: estimation of phase-locking. The LFP was filtered with a low-pass filter <300 Hz and reduced to a 1,000-Hz sampling rate. It was recorded from the same microwires as the spiking activity. Before downsampling and low-pass-filtering, spikes were replaced by a cubic spline interpolation between 1 ms before and 2 ms after the peak. This is a precaution to avoid influences of the spike waveform itself. Our results, however, are qualitatively very similar both with and without this procedure. For each channel, we estimated the instantaneous phase as a function of time for 28 different frequencies, logarithmically spaced: $f = 2^x$ with $x \in \{6/8, 8/8, 10/8, 12/8, \ldots, 60/8\}$. Thus, the range of frequencies was 1.6–180 Hz. Channels with 60-Hz noise were filtered using a fourth-order Butterworth notch filter. The phase was estimated using the continuous wavelet transform. We used a complex Morlet wavelet with a length of four cycles. The phase of each spike at each frequency was determined using the result of above wavelet transformation. The phase was measured (in radians) in the range $-\pi$ to π ($-180°$ to $180°$) with phase zero equal to the peak and phases $-\pi$ and π equal to the trough of the oscillation (Supplementary Fig. 10). To test whether a neuron was significantly phase-locked, the sample of all phase angles was compared against uniformity using a Rayleigh test. For each neuron, a Rayleigh test was performed at each frequency (Fig. 2e shows an example). The threshold was 0.05, Bonferroni-corrected for multiple comparisons. A neuron was designated as theta phase-locked if the neuron passed this test for at least one frequency in the 3–8-Hz range (we tested six frequencies; thus, the threshold was 0.0083). To guarantee sufficient statistical power, only neurons that fired at least 50 spikes were evaluated for phase-locking.

We estimated whether significant theta power was present in the LFP by fitting a function Af^{α} to the power spectrum (of the raw or spike-triggered LFP; see Fig. 3 and Supplementary Fig. 7). We found that this function fitted the data well, which is indicative of $1/f$ behaviour (except in the theta range, where there was a significant deviation, indicating the presence of theta oscillations). This was accomplished by converting the spectrum to log–log coordinates and fitting a line (the slope of which corresponds to α) using a least-squares estimate. We found values of α of approximately two (see figure legends for exact values), which is in agreement with previous studies[35].

Data analysis: SFC. The spike-field coherence, SFC(f), is a function of frequency, f, and takes values between 0% and 100%. The SFC is the ratio of the frequency spectrum of the STA and the average frequency spectrum of the LFP traces (that were used to construct the STA). The average spectrum of the LFP traces is the STP. Formally, $\mathrm{SFC}(f) = [f\mathrm{STA}(f)/\mathrm{STP}(f)]100\%$. The STA was constructed by extracting, for every spike, a piece of LFP $\pm$480 ms long and centred on the spike. We used such relatively long traces because we wanted to examine low frequencies (theta). Averaging these traces of LFPs results in the STA. The frequency spectrum of the STA (fSTA) and the spectra of the individual traces were calculated using multitaper analysis. We used the Chronux Toolbox for this purpose[36]. Multitaper analysis is a powerful method to estimate robust single-trial frequency spectra[37]. We used 240 data points at 250 Hz, a time–bandwidth product of TW = 4 and seven tapers, resulting in a half-width (frequency resolution) of 4.2 Hz. The LFP traces used to calculate the STA and STP were unfiltered (except for the <300-Hz low-pass filter mentioned). The filtered STAs shown in Fig. 3 are only for presentation purposes.

Because the SFC varies as a function of both frequency as well as the number of spikes used to calculate it (Supplementary Fig. 1), we equalized the number of spikes in the two groups (true positive or false negative) by randomly selecting a subsample of spikes from the bigger group. This procedure was also applied for all STAs and STPs shown in the figures. The relative differences between the numerical values of the SFC, however, remained constant regardless of the numbers of spikes used (Supplementary Fig. 1f). We note also that the bias is larger for smaller numbers of spikes. Failure to correct for the bias would thus increase the numerical values of the SFC for the group with the smaller number of trials (the false-negative condition). However, we observed the opposite effect (Fig. 3a). Thus, even if there were bias (which there is not; see Fig. 3b), it would work in the opposite direction of the effect we observed.

We also performed bootstrap testing to verify that the analysis was unbiased. We randomly reordered the labels assigned to each trial (true positive or false negative). The total number of true positive and false negative labels thus remained constant. We performed this procedure 1,000 times and estimated the SFC for each (Fig. 3b). There was again no significant difference. The SFC is a population measure and can not be calculated for single spikes. To ensure sufficient numbers of spikes, we included in the analysis only cells that had at least 50 spikes for all true-positive as well as false-negative trials (100 spikes in total).

For the time-course analysis, the SFC was estimated in a sliding window 250 ms wide and advanced by 25 ms at each step. For each window, all spikes that fell within the 250-ms window were used to calculate the SFC. The same number of spikes from each cell was considered for both conditions to avoid bias. Cells that contributed fewer than three spikes to a bin were excluded (for that particular bin only). Bins that contributed fewer cells could be considered for this analysis because the time course is an average across all cells rather than a single cell (as for the previous SFC comparison, where a minimum of 50 spikes was required). In Fig. 3a (true-positive condition), a total of $N = 32{,}326$ spikes were considered (average number of spikes per bin, 14.2; average number of cells per bin, 32.57 ± 0.50). In Fig. 3b (false-negative condition), a total of $N = 32{,}648$ spikes were considered (average number of spikes per bin, 14.3; average number of cells per bin, 32.66 ± 0.51).

For statistical comparison between the two time courses, we empirically estimated the null distribution[38] of the difference $TP(f, t) - FN(f, t)$. For this purpose, we generated 2,000 bootstrap samples at every frequency and time by randomly permuting the labels of spikes (labels are true positive or false negative). This was accomplished as follows (for each cell and each bootstrap): pick the same number of spikes for each subgroup (true positive and false negative); randomly permute the labels between the two subgroups; calculate the statistic $TP(f, t) - FN(f, t)$. This procedure kept constant the total number of spikes assigned to each condition, time point and neuron identity. Also, it ensured that the null distribution had no bias (as a result of the first step). We then used this null distribution to estimate the statistical significance of the difference using *z*-scores (Fig. 4b, c). The null distributions estimated in this way were very close to normal, with a mean of zero, indicating no bias (Supplementary Fig. 6). Nevertheless all *P* values reported in Fig. 4 were empirically estimated using the estimated null distribution, which does not assume normality.

Statistics of the time course (Fig. 4) were FDR correct on the basis of these *P* values. All calculations were conducted on raw, unsmoothed data (some data were smoothed for display purposes; see Fig. 4). The same procedure was used to assess the significance of the difference between low and high confidence (black line in Fig. 4f). The SFCs shown in Fig. 4a–c were smoothed using a two-dimensional Gaussian kernel with s.d. of 25 ms and 0.8 Hz, for plotting purposes only. To compare weak memories with strong memories (Fig. 4e, f), we considered only those phase-locked cells from sessions where patients used all three confidence levels ($n = 26$ cells). For correlating the pre- and post-stimulus SFCs (Fig. 4), we used the average SFC from 0–800 ms and 1,200–2,000 ms (stimulus onset is at 100 ms) for each cell ($n = 33$), calculated using all trials (regardless of whether they are true positive or false negative). Other choices of bin showed similar results (no significant correlation).

All between-group comparisons of $SFC(f)$ were corrected for multiple comparisons (one comparison for each frequency) using the FDR procedure at $q < 0.05$. Between-group comparisons of SFC values were performed using paired *t*-tests (Fig. 3a, b).

Data analysis: Phase-reset analysis. The inter-trial coherence, $ITC(f, t)$, was estimated separately at each time *t* and frequency *f*. We estimated the instantaneous phase, $\varphi(f, t)$, and power, $P(f, t)$, of ongoing oscillations for every trial using the continuous wavelet transform (as described above). The ITC[39] was then calculated as

$$\mathrm{ITC}(f,t) = \frac{1}{n}\sum_{k=1}^{n}\frac{C(f,t,k)}{|C(f,t,k)|}$$

where *n* is the number of trials, $C(f, t, k)$ is the wavelet coefficient (a function of time, frequency and trial) and modulus indicates the complex norm. The ITC is a complex number; thus, all plots show the modulus, $|\mathrm{ITC}(f, t)|$. The ITC varies between zero (no phase-locking) and one (perfect phase-locking). The ITC is equivalent to other measures of phase concentration such as the value *R* in circular statistics. Only the instantaneous phase is used for calculating the ITC (regardless of power). Equal numbers of trials from each channel were used for group comparisons (true positive versus false negative). This avoids any bias. The bigger group was reduced by randomly selecting a subset equal to the smaller group. For statistical comparisons between the ITCs of different groups, an empirical null distribution was estimated (1,000 bootstraps). For each bootstrap and channel, trials were randomly reassigned the label true positive or false negative (using a permutation procedure, leaving all biases intact). The ITC was then calculated using the procedure used for the real data. This procedure was used to estimate a null distribution for each time and frequency. We then estimated *z*-scores using this non-parametric unbiased procedure, which does not assume normality.

The stimulus-onset triggered field coherence (SOFC) was calculated the same way as the SFC (see above), except that the LFP segments were aligned by the stimulus onset. Each trial thus contributed one LFP segment to the SOFC. The SOFC is more sensitive than the ITC because it uses the entire trace rather than just a single point to estimate phase resets. The stimulus-aligned average trace is expected to be flat if the phase of the oscillations is not related to the stimulus onset. The power spectra were estimated using multitaper analysis with six tapers (resulting in a half-width of 4.6 Hz, which is the frequency resolution of this procedure). Equal numbers of trials from each channel were used for group comparisons (true positive versus false negative).

Simulations. An extracellular signal was simulated by summing multiple sine waves of different frequencies, amplitudes and phases. The amplitudes of the sine waves were inversely proportional to the frequency. Cumulative Gaussian noise with a s.d. of five was added. This procedure generated a realistic artificial LFP signal that followed the expected $1/f$ power distribution. The activity of individual neurons was simulated by generating spikes either randomly (with a certain rate) or every time the underlying oscillation was at a certain phase (distributed normally with a s.d. of 2 ms around the preferred phase). The STA for simulated data was calculated by averaging segments of length 240 ms centred on every spike. We calculated the SFC for simulated LFP traces where each spike was phase-locked (Supplementary Fig. 1) to a constant phase with variable spike-timing accuracy. We varied noise levels using normally distributed spike-timing jitter, relative to perfect phase-locking for a simulated theta-phase-locked single unit (Supplementary Fig. 1h).

31. Nelson, M. J., Pouget, P., Nilsen, E. A., Patten, C. D. & Schall, J. D. Review of signal distortion through metal microelectrode recording circuits and filters. *J. Neurosci. Methods* **169,** 141–157 (2008).
32. Duvernoy, H. M. *The Human Hippocampus* 3rd edn, 165–217 (Springer, 2005).
33. Brainard, D. H. The Psychophysics Toolbox. *Spat. Vis.* **10,** 433–436 (1997).
34. Rutishauser, U., Schuman, E. M. & Mamelak, A. N. Online detection and sorting of extracellularly recorded action potentials in human medial temporal lobe recordings, in vivo. *J. Neurosci. Methods* **154,** 204–224 (2006).
35. Milstein, J., Mormann, F., Fried, I. & Koch, C. Neuronal shot noise and Brownian $1/f^2$ behavior in the local field potential. *PLoS ONE* **4,** e4338 (2009).
36. Mitra, P. P. & Bokil, H. *Observed Brain Dynamics* (Oxford Univ. Press, 2008).
37. Jarvis, M. R. & Mitra, P. P. Sampling properties of the spectrum and coherency of sequences of action potentials. *Neural Comput.* **13,** 717–749 (2001).
38. Efron, B. & Tibshirani, R. J. *An Introduction to the Bootstrap* 202–236 (Chapman & Hall, 1993).
39. Tallon-Baudry, C., Bertrand, O., Delpuech, C. & Pernier, J. Stimulus specificity of phase-locked and non-phase-locked 40 Hz visual responses in human. *J. Neurosci.* **16,** 4240–4249 (1996).