Supporting Text

Supporting Materials and Methods

Construction of 96-Channel Headstage and Surgery. The 96-channel headstage system contains two independently movable arrays (targeting bilateral CA1), each comprised of 48 channels, in a stereotrode format (1-2). Our ultralight headstage has three 36-pin connector array positioned in parallel (see Fig. 6). A bundle of 24 pieces of polyimide tubing was glued to each of two independently movable screw nuts located on the microdrive base. Each stereotrode was constructed by twisting a folded piece of 25 μm H-Formvar wire and was thread through one of the polyimide tubes. After every stereotrode has been inserted into separate tubes, the twisted ends of the wires were cut to a length that extended three-four millimeter beyond the end of the polyimide bundle. The free end of each stereotrode (insulation has been removed) was wrapped around adjacent connect pins. To enhance conduction, each wrapped connector pin was individually coated with silver paint. A reference wire (magnet wire, 0.01 mm^2, Belden Electronic Division, Richmond, IN) was soldered to the four pins on ends of each connector array.

Wild-type B6BCA/J mice were given continuous access to food and water in their cages. Mice were handled for several days before surgery to minimize the potential stress of human interaction. On the day of surgery, the mouse was anesthetized with i.p. injection of 60 mg/kg ketamine (Bedford Laboratories, Bedford, OH) and 4 mg·kg Domitor (Pfizer Animal Health, New York). The mouse head was immobilized in a stereotaxic frame, and his eyes were coated with sterile ocular lubricant to maintain moisture. After the hair removal, Betadine solution was applied to the skin surface, and an incision was made along the midline of the skull. The positions for the two bundles (2.0 mm lateral to bregma and 2.3 mm posterior to bregma on the both right and left sides, respectively) were then measured and marked. After four holes were drilled in the corners of a rectangular array, small screws were secured in each of these holes and fixed with dental cement. Holes for the two stereotrode bundles were then drilled, and dura was removed carefully. The stereotaxic apparatus was then used to lower the stereotrode bundles into the mouse's cortex. The gap surrounding the stereotrodes was filled with softened paraffin, and the headstage was stabilized with dental cement. The reference wire attached to the two posterior head screws was then soldered to the reference wire affixed to the connector pin arrays of the headstage, and copper mesh was wrapped around the entire headstage to protect the wires from potential damage. The mouse was then aroused with an injection of 2.5 mg/kg Antisedan and returned to his home cage.

***In Vivo* Recording and Spike Sorting.** The mouse was allowed to recover for several days before advancing the electrodes. The connector pin arrays on the headstage were then attached to preamplifiers with extended cables to allow for the monitoring of neuronal signals by using the 96-channel Plexon system. A helium-filled mylar balloon was tied to the cables to completely relieve the weight of the 96-channel cables, thereby enabling the mouse to move freely (see Fig. 6). The stereotrode bundles were advanced slowly toward the hippocampal CA1 area, in daily increments of about 0.07 mm, until the

tips of the stereotrodes had reached the CA1 region, as deduced from an assessment of field potential and neuronal activity patterns.

We have recorded the ensemble activity of a large number of individual neurons during several behavioral states. The recorded spike activities from those neurons were processed as follows: first, the spike waveforms and their associated time stamps for each of 96 channels were stored in data files by using Plexon (Dallas) system format (*.plx). The artifact waveforms were removed, and the spike waveforms were aligned at their minimum values in using OFFLINE SORTER 2.0 software, because the waveforms are more tightly clustered in principal component space when the waveforms are aligned around their minimum values. Then, the Plexon system data files (*.plx) were converted to Neuralynx system format (*.nst) and spike sorted with the MCLUST3.0 program (David Redish, University of Minnesota). This program permits classification of multidimensional continuous data; its cluster splitting feature (KLUSTAKWIK 1.5, Gyorgy Buzsaki, Rutgers, the State University of New Jersey) yields superior accuracy in comparison to the other available spike-sorting software and is therefore particularly suitable for spike sorting of hippocampal signals.

Principal component analysis was used to extract defining features from the spike wave shapes that are used as part of the input for the MCLUST3.0 spike-sorting program. The first two principal components, as well as the peak height, valley value, FFT, and total energy of spike waveform parameters, were calculated for each channel, then units were identified and isolated in high-dimensional space by using an autoclustering (KLUSTAKWIK 1.5) method (3-4). After autoclustering, the clusters containing nonspike waveforms were deleted by using the KLUSTAKWIK Selection function, and then the units were further isolated by using a manual cluster-cutting method in MCLUST. Only units with clear boundaries and <0.5% of spike intervals within a 1-ms refractory period are included in the present analysis. If the clusters cannot remain stable for 6 h, the entire units are excluded from further analysis. An example of the spike sorting of six stable individual units recorded in one of our experiments is shown in Fig. 7. On occasions, up to 14 stable units were recorded, as illustrated in Fig. 8. The good separation of these units is evident from the measurements of L-ratio and isolation distance (5) (see Figs. 7*C* and 8*B*). At the end of experiments, the mouse was anesthetized, and a small amount of current was applied to four channels in the headstage to mark the positioning of the electrode bundle. Histological staining with 1% cresyl echt violet was used to confirm the electrode positions.

Startle Episodes. We used three startle protocols: (*i*) A sudden air blow to the animal's back [termed air blow, 200 ms, 10 psi (1 psi = 6.89 kPa)]; (*ii*) A sudden drop of the animal inside a small elevator (termed elevator drop, vertical freefall height: 40 cm); and (*iii*) A sudden shake-like cage oscillation (termed shake, 200 ms; 300 rpm, Thermolyne Vortex CM37615, Fisher Scientific). The startle stimuli were delivered unexpectedly to freely behaving mice by using computer programs that provide precise mechanical control of an air valve, suspension rope, and a vortex machine. The air tube was attached to the recording cable in such a way the air blow was consistently delivered to the animal's neck. To assess startle-response properties of individual neurons, each type of

startle was delivered seven times at random intervals varying between 1 and 5 min. Environmental context-specificity was investigated by delivering air blows to mice in either a red bucket (30 cm in diameter) or a blue open box (30 × 35 cm) and also by dropping the mice inside of either a plastic bottle (15 cm in diameter) or a small black box (15 × 15 cm). The activity of large numbers of single CA1 neurons was monitored and recorded throughout the startle sessions (4-6 h in total), which were typically preceded by at least 10-20 min of a baseline resting period.

Although a significant proportion of the simultaneously recorded CA1 cells did not respond to any of the startle stimuli, a subset of these CA1 cells exhibited significant changes in firing rates. In general, based on their temporal response duration, the startles triggered dynamical changes that can be generally divided into four major firing modes: transient increase, transient decrease, prolonged increase, and prolonged decrease (see Fig. 10). The transient changes were as short as 250 ms or less, whereas the prolonged increases lasted up to 40 s in duration.

Measurement of Startle Memories. We used the conditioned place preference test (6) to measure the formation and retention of startling episodic memories in mice. The mice were handled and habituated before experiments. On the training day, the mice were allowed to freely explore the environment of a two-compartment conditioned place preference apparatus (20 × 40 × 20 cm) for 3 min (pretraining). Both proximal and distal cues were available in the room. The amount of time spent in each chambers were recorded. During the 10-min conditioning period, a sudden air blow was delivered whenever the mouse visited one of the randomly assigned compartment (termed air-blow compartment). The other unconditioned compartment is designated as the safe compartment. The retention of place conditioning memory was measured 3 h later, after training, by returning the mice to the place conditioning environment for 3 min. The time spent (sec) in the safe compartment was measured and used as a place preference index before and after training. An equal time spent in both compartments (90 s) would therefore indicate no preference toward either compartment. The formation of startle memories in mice, as indicated as its tendency to spend more time in the safe compartment, is shown in Fig. 9.

Data Analysis. *Preprocessing*. In our data analysis, we only include single units that were both clearly separated and remained stable for at least 6 h; units that failed to meet these criteria were excluded from further analysis. Firing frequencies (*f*) were evaluated in two 500-ms time bins immediately after the event start. This evaluation yields a sampling of seven points for each startle type. The 35 time intervals preceding the complete set of startle events were used to compute the population responses that define rest state. Each neuron response was defined by the formula:
$R_n = \left(f_{post-startle,\,n} - f_{pre}\right) / \left(f_0 + f_{pre}\right)$, where f_0 (2-3 Hz) is the global mean response frequency of putative excitatory neurons across the recorded neural population in an animal, excluding high firing-rate interneurons (cut off at 20 Hz) and f_{pre} is the base firing rate (computed from samples of firing rates in the 500-ms time bins before startle stimuli). This transformation emphasizes significant changes in firing patterns for units with both low- and high-baseline firing rates. Effectively, changes in responses of low-

firing units are proportional to absolute firing rate changes, whereas changes in responses of high-baseline units are proportional to the relative changes.

Multiple Discriminant Analysis (MDA). For each startle type, we use six of its seven repetitions as training data (30 data points were also selected from 35 rest epochs), and we use the remaining events to create a test set. A matrix of mean responses during rest and startle states are then computed, and a threshold criterion is used to exclude the neural features of unresponsive neurons. This procedure leads to a sparse between-class scatter matrix (7): $S_B = \sum_{i=1}^{N} n_i (m_i - m)(m_i - m)^t$. Here, n_i is the number of elements in each class, N is the number of classes, m_i is the mean vector for each class and m is the global mean vector. The discriminant projection vectors were calculated by performing an eigenvalue decomposition of the matrix product $S_W^{-1} \cdot S_B$, where the within-class scatter matrix S_W is defined as: $S_w = \sum_{i=1}^{N} \sum_{x \in D_i} (x - m_i)(x - m_i)^t$. Here D_i represents the set of population responses triggered by the ith startle type. In the cases analyzed here, the eigenvalue decomposition always produces at most $N - 1$ nonzero eigenvalues, hence it follows that the dimension of encoding subspace is also $N - 1$. For example, the first step of the discriminant analysis uses four classes (rest, air blow, drop, and shake), yielding a three-dimensional encoding subspace, whereas the subsequent context MDAs produced two-dimensional subspaces for discriminating their corresponding three categories (rest, air blow A, air blow B, and rest, drop A, drop B, respectively).

Encoding subspace. We fit the low-dimensional startle representations with multivariate Gaussian distributions: $P(x) = \frac{1}{(2\pi)^{N/2} |\Sigma|^{1/2}} \exp(-(x - m)^t \Sigma^{-1} (x - m)/2)$. Here, x is a point in the low-dimensional space, m is the center of startle cluster, $|\Sigma|$ and Σ^{-1} are the determinant and inverse of the covariance (scatter) matrix Σ. In some cases (encountered in data sets from mouse C and D), the first element on the main diagonal of the covariance matrix Σ was disproportionately larger than the subsequent ones. This difference was compensated by slightly incrementing all diagonal elements in Σ (Tikhonov regularization, see ref. 8). The use of Gaussian probability distributions allows us to compute the probability P_{ij} that startle event i belongs to category j**.** After computing all such probabilities, soft class memberships were assigned based on normalized probabilities: $P_{ij}' = P_{ij} / \sum_{j=1}^{N} P_{ij}$. To test the predictive power of the startle representations (see Table 1), a collection of 10 test data points (5 random startle points for air blow A, air blow B, drop A, drop B and shake, as well as their corresponding 5 rest samples) was used to cross-validate the predictive power of models constructed by using all other points (training). Average performances for class prediction were obtained by repeating MDA 1,000 times with random partitioning into training and test points. We noted that, in general, the prediction performance is correlated with the number of startle-responsive cells in the data sets.

Hierarchical clustering. To examine the structure of the neuronal population representation, the thresholded mean responses ($R < 0.3$ in first 500-ms time bin) of all neurons to the various startles were clustered by using a standard hierarchical clustering algorithm. The steps used by this agglomerative procedure are the following: initially each response vector is defined as a single cluster. The closest clusters are merged into a new cluster, and its mean is recomputed. This step is then repeated and the nearest-neighboring groups (based on the position of their centers) are successively merged until they eventually form a single group. At each intermediate step, two clusters are aligned and linked at their best matching endpoints forming a larger group, preserving their internal ordering while allowing for flipping. The hierarchical tree structure thus obtained can be displayed by using the dendrogram MATLAB (Mathworks, Natick, MA) function.

Principal Component Analysis (PCA). To further investigate the nature of the startle representation, we looked at the low-dimension encoding subspace generated by PCA. For our data, this space is determined by the first few eigenvectors of the total scatter matrix $S_{w'} = \sum_{i=1}^{N} (x_i - m)(x_i - m)^t$, which capture most of the variance observed in the original data set (this scatter matrix is different from the one created by MDA, because in this case, the training data are not partitioned into different classes). By definition, the resulting eigenvectors are orthogonal, therefore, they account for an independent source of variance. Because of its unsupervised nature, PCA can be used to identify and extract the information embedded in the network firing patterns. We use hierarchical clustering to determine the four startle clusters (step 1) and three context clusters (two additional steps) to be fitted with Gaussian distribution. Performance comparison with the MDA-based projections is then made under the same set of assumptions (see *Encoding Subspace* section).

The prediction power of the PCA method is assessed by using the same classification method and cross-validation used with MDA, thus allowing direct comparison between these two statistical techniques (see Fig. 11 and Table 2). It is noteworthy that the MDA method generally yields superior pattern separation and more accurate prediction than PCA method. This finding is consistent with the fact that MDA is better suited to maximally identify features for discriminating between classes, whereas PCA is more oriented to finding the defining features of each class (potentially finding features that are common across classes, consequently reducing the discriminating power of the low-dimensional encoding subspace) (7). Therefore, we choose to use MDA as the primary method for analyzing our data.

Remapped Neural Clique Responses by Using a Matrix Inversion Step. To translate the population responses into a startle-selective encoding coordinate system, we mapped the cluster centers onto corners of a hypercube in a "clique-space," where each axis corresponds to the expected responses of a particular clique. This geometrical operation involves a reorientation of the main axes of the low-dimensional encoding subspace. The mapping matrix M was calculated by using: $M = C^{-1} * B$, where C is the matrix containing the cluster centers, C^{-1} is its inverse, and B contains the new coordinates of these centers (e.g., [1,1,0,0], [1,0,1,0], and [1,0,0,1] for the air blow, drop, and shake,

respectively). More explicitly, the mapping *M* that converts the centers of air-blow, drop and shake representations along the *x*, *y*, and *z* axes (see Fig. 3*A*) into the general-startle, air, drop and shake codes (see Fig. 5*A*) is given by:

$$M = \begin{pmatrix} A_{1x} & A_{1y} & A_{1z} \\ A_{2x} & A_{2y} & A_{2z} \\ D_{1x} & D_{1y} & D_{1z} \\ D_{2x} & D_{2y} & D_{2z} \\ S_{1x} & S_{1y} & S_{1z} \end{pmatrix}^{-1} * \begin{pmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 \\ 1 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 \end{pmatrix}$$

Context codes were computed in a similar way.

To investigate in detail the temporal dynamics of the instantaneous projections in the clique-based subspace, we smoothed the firing rates of all units with Gaussian filters ($G(t) = 1/\sqrt{2\pi}\, e^{-(t-t_{spike})^2/2\sigma^2}$, width σ = 100 ms). Integration of different unit contributions is done by using the weight distribution from multiple discriminant analysis (MDA) times the remapping *M*, obtaining the responses shown in Fig. 5*A*. However, simple summation of these responses yields similar results, as illustrated in Fig. 13. Furthermore, the temporal scales of this summation process indicate that encoding robustness can still be achieved at time scales as low as 20-30 ms, indicating that the cospiking within neural cliques can serve as a network mechanism to overcome the response variability of individual neurons.

Mathematical Description of the Neural Clique Activation Function. The responses of the individual member neurons $R_n(t)$ are being linearly summed by using the remapped MDA weighting factors w_n to yield the response of the neural clique as $\sum_{n=1}^{N} w_n R_n(t)$. Event classification, by the neural clique activation codes, can now be achieved by means of evaluating a simple threshold function *h*. Consequently, the dynamical state of the neural clique is being described by $f(t) = h(\sum_{n=1}^{N} w_n R_n(t))$. Examples of threshold activation functions, such as the step threshold $h(x) = \begin{cases} 0, x < 0 \\ 1, x > 0 \end{cases}$, and the sigmoid threshold $h(x) = \dfrac{1}{1+\exp(-\beta x)}$, are presented in Fig. 14.

Extended Discussion

A central issue in the study of neural coding in the brain is the response variability of individual neurons (9-11). This variability at the level of individual neurons has posed a major challenge for understanding how the brain achieves its real-time encoding and decoding of behavioral experiences. A traditional way to deal with this issue is to average

the response of an individual neuron over many repetitions and trials. Although it allows the identification of event-related neural response, this practice of data averaging unfortunately loses crucial information regarding real-time neural coding functions. Here, we have combined the high-density *in vivo* recording technique with experimental designs (by using a set of startling episodes as a means) and investigated how the brain achieves its the network-level real-time encoding of behavioral experiences. Particularly, we asked the following three questions: (*i*) how are a variety of startling episodes represented at the CA1 network level? (*ii*) How does the CA1 network-level encoding overcome the response variability of individual neurons and achieve its robust real-time neural representation? and (*iii*) What are internal structures underlying the real-time CA1 network-level encoding?

Consistent with the knowledge that neurons in the hippocampus can respond to various external inputs (e.g., place field running, odor stimulation, or to corneal air puffs during nictitating membrane responses conditioning, etc.) (12-19), we find that a given type of startling episode triggers distinct and collective firing changes in many CA1 neurons. Because our study only included animals in which a significant portion of recorded cells (in dorsal hippocampus) were responsive to startling episodes, and also considering that the majority of hippocampal cells remained silent (20), we think that the actual percentage of responsive cells in the hippocampus are likely to be much lower than the one estimated from the four mice presented here. Nonetheless, based on the temporal dynamics, spike discharges of these neurons can be generally classified into four distinct firing modes: transient increase or decrease and prolonged increase or decrease. Moreover, CA1 cells exhibit a variety of response-tuning properties: some responded to several or all types of startling stimuli, whereas others showed extremely high specificity toward one specific type of startling stimulus. Interestingly, we also find that the firing patterns of some CA1 cells are critically dependent on the spatial environments in which the startling event occurs, thereby demonstrating that CA1 cells indeed encode startling episodes with both “what” and “where” information.

To seek out the network-level encoding patterns underlying the neural representation of startling episodes, we applied a set of statistical pattern classification tools to the data sets (7). It is evident that approaches such as peri-event histograms and cross-correlation methods are not suitable for dealing with the high-dimensionality of the large data sets. Thus, it is necessary to apply statistical pattern classification tools to reduce high-dimensionality of the data and to extract the underlying neural information from the ensemble activities. We focused on distinguishing intrinsically meaningful factors in the hippocampal network encoding, by identifying a highly informative low-dimensional subspace among the firing patterns of responsive neurons by using both MDA and PCA methods. These two methods not only provide for a mathematical classification of neural ensemble patterns but also allows their direct visualization. More importantly, it allowed us to analyze the underlying structure of episodic representations from the complex high-dimensional activity patterns of simultaneously recorded neurons.

Furthermore, the use of the MDA dimensionality-reduction method allows not only reliable classification but also permits the dynamical monitoring of the temporal

evolution of firing patterns in the low-dimensional encoding subspace. Although the instantaneous projections spent most of the time within the rest cluster, we have found that the dynamical trajectory also exhibited brief excursions in the immediate aftermath (seconds to minutes) of startle episodes (see Fig. 12 for a reactivation after a drop episode). Previous studies with correlation-based analysis report the heightened firing correlation between overlapping place cells during sleep (17, 21). Our dynamical decoding technique described here allows direct visualization of spontaneous, immediate reactivations in the awake-behaving state whose trajectory were directed toward the cluster representing the preceding startle. The observed intervals between a spontaneous reactivation and the initial activation (or between reactivations) were up to minutes, thereby ruling out the possibility that they were due to oscillatory sampling artifacts within the sliding window (of only 1-s width). The consistently smaller magnitude of the reactivation trajectories has also allowed a reliable distinction between "internally" vs. "externally" driven patterns. We hypothesize that these spontaneous reactivations may reflect the immediate postlearning fixation of memory traces (22).

Our combined applications of hierarchical clustering and sequential MDA enabled us to further identify discrete sets of functional coding units in the CA1 network. We find that the individual neurons within the same neural clique exhibit tight "collective cospiking" temporal dynamics that allow the neural clique to overcome the variability of single neurons for achieving real-time encoding robustness. Moreover, it is conceivable the dynamical state of the downstream neural clique activation can be assessed by applying biological threshold function (23) (see also Fig. 14). In addition to their increased robustness, neural cliques as coding units are less vulnerable to the death of one or a few neurons and may exhibit graceful degradation should such conditions occur. Currently we do not know how the individual neurons that comprise a functional coding clique are precisely connected at the anatomical level or how their memberships and firing properties are modulated by NMDA receptor-mediated synaptic plasticity (24-27). We also do not know to what extent they reflect a "hard-wired" response to the external sensory and/or emotional aspects of the events. The occurrence of postevent reactivations of these ensemble patterns (in the absence of the external sensory inputs) suggests that at least part of these ensemble activities could be associated with certain mnemonic aspects of episodic experiences.

Our additional mapping of the encoding subspace into one where the startle representations can directly correspond to neural clique activity patterns has further translated the ensemble activity into a set of useful network-level activation codes. This mathematical transformation allows us to convert the collective activity patterns of neural clique assembly into unique and efficient digital codes. These activation codes are also consistent with the idea that the activity state of a neural clique can be monitored by a downstream neuron or group of neurons based on a biologically plausible activation function (23). More importantly, the real-time activation codes of clique assemblies can conceivably form universal internal representations across different individuals (and even different animal species) to categorize similar cognitive events. We note that although the classification of startle events can be reliably achieved by using a minimal number of response-specific cliques, the prevalence of the multiresponsive neural cliques (e.g.,

general startle, drop/shake, etc.) suggests that CA1 cells play a crucial role in binding multimodal cortical information and, perhaps, even in processing cognitive abstraction and generalization (28-29).

The existence of hippocampal neural cliques suggests that the information representation and processing in the brain are achieved in real time by coordinated activities of neural cliques within multilayered neural circuits that span from sensory and motor control areas up to higher association and cognition areas. In this sense, the functions implemented by neural cliques at each layer depend on the specializations and functionalities of corresponding anatomical components of the nervous systems.

In conclusion, we have identified and visualized the network-level encoding patterns and postevent immediate processing of startling episodic experience in the CA1 region of the hippocampus. We have identified network-level functional coding units in the CA1 region, and we found that the individual members within each clique fired tightly together, cospiking in close temporal proximity during startling episodes. The cospiking of the neural clique members allows the cliques to overcome the response variability of individual neurons and thereby to achieve real-time network encoding of startling episodic experiences. The activation patterns of these coding units can be translated into a set of concise digital codes for universal representation and categorization of discrete behavioral episodes across different animals. Therefore, the "neural clique cospiking" principle provides a plausible network-level basis by which the nervous systems are built on to achieve real-time neural coding and processing of behavioral information.

1. McNaughton, B. L., O'Keefe, J. & Barnes, C. A. (1983) *J. Neurosci. Methods* **8,** 391–397.

2. Moita, M. A., Rosis, S., Zhou, Y., LeDoux, J. E. & Blair, H. T. (2004). *J. Neurosci.* **24,** 7015–7023.

3. Csicsvari, J., Hirase, H., Czurko, A. & Buzsaki, G. (1998) *Neuron* **21,** 179–189.

4. Harris, K. D., Henze, D. A., Csicsvari, J., Hirase, H. & Buzsaki, G. (2000) *J. Neurophysiol.* **84,** 401–414.

5. Schmitzer-Torbent, N., Jackson, J., Henze, D., Harris, K. & Redish, A. D. (2005) *Neuroscience*, in press.

6. Meyers, R. A., Zavala, A. R. & Neisewander, J. L. (2003) *NeuroReport* **14,** 2127–2131.

7. Duda, R.O., Hart, E. P. & Stork, D.G. (2001) *Pattern Classification* (Wiley, New York), 2nd Ed.

8. Hansen, P. C. (1999). *Rank-Deficient and Discrete Ill-Posed Problems* (Soc. Indus. and Appl. Math., Philadelphia).

9. Lestienne, R. (2001) *Prog. Neurobiol.* **65,** 545–591.

10. Abbott, L. F. & Dayan, P. (1999) *Neural Comput.* **11,** 91–101.

11. Fenton, A. A. & Muller, R. U. (1998) *Proc. Natl. Acad. Sci. USA* **95,** 3182–3187.

12. Cohen, N. J. & Eichenbaum, H. (1993) *Memory, Amnesia, and the Hippocampal System* (The MIT Press, Cambridge).

13. Berger, T. W., Alger, B. & Thompson, R. F. (1976) *Science* **192,** 483–485.

14. O'Keefe, J. & Nadal, L. (1978) *The Hippocampus as a Cognitive Map* (Oxford Univ. Press, Oxford).

15. Weiss, C., Kronforst-Collins, M. A. & Disterhoft, J. F. (1996) *Hippocampus* **6,** 192–209.

16. Deadwyler, S. A., Bunn, T. & Hampson, R. E. (1996) *J. Neurosci.* **16,** 354–372.

17. Wilson, M. A. & McNaughton, B. L. (1994) *Science* **265,** 676–679.

18. Wirth, S., Yanike, M., Frank, L. M., Smith, A. C., Brown, E. N. & Suzuki, W. A. (2003) *Science* **300,** 1578–1581.

19. Berger, T. W., Rinaldi, D. J., Weisz, D. J. & Thompson, R. F. (1983) *J. Neurophysiology* **50,** 1197–1219.

20. Henze, D. A., Borhegyi, Z., Csicsvari, J., Mamiya, A., Harris, K. D. & Buzsaki, G. (2000) *J. Neurophysiol.* **84,** 390–400.

21. Louie, K. & Wilson, M. A. (2001) *Neuron* **29,** 145–156.

22. Wittenberg, G. M. & Tsien, J. Z. (2002) *Trends Neurosci.* **25,** 501–505.

23. McCulloch, W. S. & Pitts, W. (1943) *Bull. Math. Biol.* **52,** 99–115.

24. Wigstrom, H. & Gustafsson, B. (1985) *Acta Physiol. Scand.* **123,** 519–522.

25. Malenka, R. C. & Nicoll, R. A. (1999) *Science* **285,** 1870–1874.

26. Tang, Y. P., Shimizu, E., Dube, G. R., Rampon, C., Kerchner, G. A., Zhuo, M., Liu, G. & Tsien, J. Z. (1999) *Nature* **401,** 63–69.

27. Tsien, J. Z. (2000) *Curr. Opin. Neurobiol.* **10,** 266–273.

28. Freiwald, W. A., Kreiter, A. K. & Singer, W. (2001) *Prog. Brain Res.* **130,** 111–140.

29. Crick, F. & Koch, C. (2003) *Nat. Neurosci.* **6,** 119–126.