

Memory reprocessing in corticocortical and hippocampocortical neuronal ensembles

YU-LIN QIN⁴, BRUCE L. MCNAUGHTON^{1,4*}, WILLIAM E. SKAGGS⁴
AND CAROL A. BARNES^{1,2,4}

Departments of ¹Psychology, ²Neurology and ³Physiology, and ⁴ARL Division of Neural Systems, Memory and Aging, University of Arizona, Tucson, AZ 85724, USA

SUMMARY

Hippocampal cells that fire together during behaviour exhibit enhanced activity correlations during subsequent sleep, with some preservation of temporal order information. Thus, information reflecting experiences during behaviour is re-expressed in hippocampal circuits during subsequent 'offline' periods, as postulated by some theories of memory consolidation. If the hippocampus orchestrates the reinstatement of experience-specific activity patterns in the neocortex, as also postulated by such theories, then correlation patterns both within the neocortex and between hippocampus and neocortex should also re-emerge during sleep. Ensemble recordings were made in the posterior parietal neocortex, in CA1, and simultaneously in both areas, in seven rats. Each session involved an initial sleep episode (S1), behaviour on a simple maze (M), and subsequent sleep (S2). The ensemble activity-correlation structure within and between areas during S2 resembled that of M more closely than did the correlation pattern of S1. Temporal order (i.e. the asymmetry of the cross-correlogram) was also preserved within, but not between, structures. Thus, traces of recent experience are re-expressed in both hippocampal and neocortical circuits during sleep, and the representations in the two areas tend to correspond to the same experience. The poorer preservation of temporal firing biases between neurons in the different regions may reflect the less direct synaptic coupling between regions than within them. Alternatively, it could result from a shift, between behavioural states, in the relative dominance relations in the corticohippocampal dialogue. Between-structure order will be disrupted, for example, if, during behaviour, neocortical patterns tend to drive corresponding hippocampal patterns, whereas during sleep the reverse occurs. This possibility remains to be investigated.

1. INTRODUCTION

The phenomenon of retrograde amnesia for certain kinds of information after temporal-lobe damage (Scoville & Milner 1957; Squire 1992) has led to the general hypothesis that the hippocampus somehow facilitates the reactivation of neocortical activity patterns during 'offline' periods such as sleep or quiet wakefulness, when the neocortex is not actively engaged in processing incoming data (Marr 1971). It is presumed that this reactivation somehow permits recently acquired information to become appropriately integrated into long-term memory. This could occur either if the hippocampus generated explicit, compact representations of the events themselves (McClelland *et al.* 1995; Squire 1992; Marr 1971), or if it merely generated a cognitive map (O'Keefe & Nadel 1978) that could also serve as a contextual code (Teyler & Discenna 1986; Nadel *et al.* 1985; McNaughton *et al.* 1996) for each event. Either of these sorts of codes could be associated with the pattern in the

active neocortical modules at the time of the original experience. Because of the divergent, direct and indirect, return projections from hippocampus to neocortex (Amaral & Witter 1995) reactivation of the hippocampal component of the experience conceivably could reinstate the entire pattern of experience across the weakly interconnected neocortical network. The presumed attractor dynamics of the CA3 region of the hippocampus implies that reactivation of recent patterns in the hippocampus could arise either spontaneously (Shen & McNaughton 1997) or as a consequence of the reactivation of a subcomponent of the experience in the neocortex (Marr 1971) and subsequent pattern completion in the hippocampus. This general theory provides a framework for understanding the phenomenon of temporally graded retrograde amnesia after hippocampal damage. A minimal prediction of the theory is that, during offline periods, activity patterns corresponding to recent experiences should appear in both hippocampus and neocortex, and, moreover, that at any given time, the patterns in the two regions should be correlated with the same experience.

*Author for correspondence (bruce@nsma.arizona.edu).

Three studies have provided neurophysiological evidence for memory reactivation within the hippocampus during sleep. If a rat is confined to the place field of a particular hippocampal place cell, the firing rate of that cell is increased during subsequent sleep (Pavlidis & Winson 1989). More compellingly, if a rat is permitted to explore a larger environment during behaviour, so that many different place cells are activated, there is little change in the net firing rates of the corresponding place cells during subsequent slow-wave sleep, but a significant increase in the activity correlation between cell pairs whose activity had been correlated (by virtue of their place-field overlap) during behaviour (Wilson & McNaughton 1994). This could only occur if engrams, closely resembling the neuronal discharge patterns of prior experiences, were reactivated (Shen & McNaughton 1997). In addition to simple correlations, Skaggs & McNaughton (1996) have shown that the temporal biases in hippocampal cross-correlation functions persist from a period of behaviour into subsequent slow-wave sleep. Thus, information about the temporal sequence of events is also preserved and re-expressed in hippocampus.

Motivated by the theoretical considerations discussed above, the present study was designed to follow up on these observations by examining whether patterns of activity in the neocortex, and in hippocampal–neocortical relations, are replicated during sleep in a manner similar to that of intrahippocampal patterns. Some of the present results have previously been reported in abstract form (Qin *et al.* 1995, 1996).

2. MATERIALS AND METHODS

(a) *Subjects and behavioural procedures*

Seven male F-344 rats, ten months old (Harlan Sprague–Dawley, Indianapolis), were used. The rats were maintained at 85–90% of their *ad libitum* body mass, under illumination from 0800 to 2000 h, and all recordings were conducted between 0800 h and 1400 h. During pretraining, the animals learned to shuttle back and forth in a linear alleyway for a food reward. After recovery from surgery, the rats were adapted to the recording headstage and learned to run for food reinforcement on one of two similar elevated-track mazes, 7.5 cm wide. One track (3 rats) was triangular (61 cm per side), and contained reward sites at the midpoint of each side. The other track (4 rats) was rectangular (87 cm × 36 cm), and contained two reward sites located at adjacent corners on the long axis. The long side opposite the reward locations was bordered by a plywood wall. Both tracks were located in the same central position in a moderately illuminated room, 3.7 m × 3.7 m, containing several prominent visual landmarks.

Recording sessions consisted of three phases: an initial period (S1) of sleep or quiet wakefulness lasting at least 20 min; a maze-running phase (M), during which the rats ran a minimum of 15 circuits of the maze without changing direction, and a second sleep phase (S2), also lasting at least 20 min. During the two

sleep phases, the rats rested in a small bowl, which was placed on a pedestal about 20 cm above the maze. In these studies, the ‘sleep’ periods were defined by the animal’s relatively motionless behaviour, and hippocampal electroencephalogram (EEG) exhibiting predominantly large irregular activity, typically preceded by periods of neocortical spindle activity. There was a small amount of rapid eye-movement (REM) sleep, as indicated by theta activity. No further effort was made to characterize the behavioural states during these episodes, which thus can be considered to comprise mostly slow-wave sleep plus a small amount of quiet wakefulness and REM sleep.

(b) *Surgery*

All surgical procedures were carried out under deep sodium pentobarbital anaesthesia according to NIH guidelines. The stereotaxic surgery involved opening a craniotomy 2 mm in diameter in the right hemisphere over either the hindlimb sensorimotor region (Zilles (1985) area HL, 3.5 mm posterior and 2.0 mm lateral to bregma) or the posterior parietal area (Zilles (1985) area OC2ml, 5.0 mm posterior and 3.0 mm lateral to bregma). The dura was retracted and the tip of the microdrive assembly (see below) was placed on the exposed cortex so that the probes penetrated about 100 µm into the cortex. The drive assembly was anchored in place by using dental acrylic supported by seven small jeweller’s screws mounted in the skull.

(c) *Electrophysiological procedures*

Recordings were made with a microdrive assembly containing 14 independently movable tetrodes. The construction of tetrode recording probes (McNaughton *et al.* 1983; Recce & O’Keefe 1989; Wilson & McNaughton 1993) and the multielectrode ‘hyperdrive’ assembly have been described in detail elsewhere (Gothard *et al.* 1996). Briefly, four HML-insulated Nichrome wires (HP Reid Co.), 14 µm in diameter, were twisted together and the end of the bundle was cut off at right angles with sharp scissors. Each tetrode was mounted in the guide cannula of one of the 14 microdrives on the assembly, and each wire was connected to a separate pin on a multichannel connector. The microdrives permitted each tetrode to be moved up or down independently of the others by turning a screw, with a full turn moving the electrode tip vertically by about 320 µm. The entire drive assembly weighed 12–14 g.

A 50 channel field effect transmitter (FET) headstage assembly was mounted on top of the hyperdrive during recording sessions. The headstage assembly included two arrays of infrared light-emitting diodes (LEDs) whose positions were monitored and sampled at 20 Hz to provide information on both head location and head orientation in the horizontal plane. The FET outputs were transmitted to the main amplifiers via a lightweight cable.

Electrophysiological data were collected by means of an array of eight 80486-based microcomputers with synchronized, programmable, event-timing clocks

running at 10 kHz (Datawave Corporation). One of the probes was used as an EEG reference, placed near the hippocampal fissure. Another served as the indifferent reference electrode, located in the corpus callosum. The four channels of each of the remaining tetrodes were each amplified 2000–10 000 times (depending on the signal size), filtered between 600 Hz and 6 kHz, and digitized at 32 kHz per channel. An entire 1 ms sample of the waveforms of each action potential on the four tetrode channels was stored on disk, along with a time stamp code (0.1 ms resolution). In addition, the signals from one channel of each of several tetrodes was split off to a separate amplifier for EEG monitoring. EEG output on different channels was filtered with a band-pass of either 1–100 Hz or 1–3 kHz to permit simultaneous monitoring of hippocampal theta rhythm, sharp-waves and ripple activity as well as neocortical sleep spindles. Behavioural states were characterized as awake-theta, still-alert, slow-wave sleep, and REM sleep according to established behavioural and EEG criteria (Vanderwolf 1969; Winson 1977).

After surgery, the tetrode probes were advanced gradually, normally not more than 50–200 $\mu\text{m d}^{-1}$ depending on electrode location and the signal quality. During the first few recording sessions, all tetrodes were typically in the neocortex. About half of the probes were then gradually advanced to the CA1 layer of hippocampus to record hippocampal–neocortical interactions. Finally the remaining probes were advanced to CA1.

After recording, the animals were perfused transcardially with formalin, and the electrode tracks were verified histologically by using combined Nissl and myelin staining (Zilles 1985).

(d) Unit isolation

Single units were isolated from the stored data by means of a visually guided, multidimensional cluster analysis method described by McNaughton *et al.* (1989). Only units that were considered well isolated were accepted for analysis. Although it is never possible to define precisely the degree of unit isolation, the tetrode method has been demonstrated to provide isolation that is substantially superior to that of other methods (McNaughton *et al.* 1983; Recce & O'Keefe 1989; Gray *et al.* 1995).

(e) Analysis of spike-train interactions

The simple correlation between the spike trains of each pair of cells was used to describe the structure of the firing pattern of the recorded cells. For each pair of cells, the spike trains were divided into n 100 ms bins. If x_i and y_i denote the number of spikes fired by each cell in the i th bin, and \bar{x} and \bar{y} denote the mean numbers of spikes per bin, their rate-independent correlation coefficient is defined as

$$r_{x,y} = \frac{\sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^n (x_i - \bar{x})^2 \sum_{i=1}^n (y_i - \bar{y})^2}}. \quad (1)$$

Within the hippocampus, in an animal performing a spatial task, the extent of positive correlation between the spike trains of two pyramidal cells is typically related to the overlap of their place fields. In the neocortex, for the types of task used in this study, pyramidal cells also tend to show spatially heterogeneous patterns of firing, because different locations are associated with different types of behaviour or reward contingency (McNaughton *et al.* 1994). Between hippocampal and neocortical cells, relatively strong correlations would be observed if the hippocampal place field was located at one of the regions where the behavioural correlate of the neocortical cell was strongly expressed. Examples of these effects are illustrated in figure 1.

(f) Effects of experience on the activity correlations within and between anatomical structures

For each pair of cells in a data set, let $\text{Corr}(s1)$, $\text{Corr}(m)$ and $\text{Corr}(s2)$ be their activity correlations in S1, M and S2, respectively. Thus, each pair of cells gives rise to a single point in a three-dimensional space. A data set containing N pairs of cells produces a cloud of N such points. The regression slope over all cell pairs between $\text{Corr}(s1)$ and $\text{Corr}(m)$ measures the similarity of neuron-ensemble firing patterns between phase S1 and phase M. A significant positive slope may reflect the rats' prior experience or the patterns of connections imposed by early developmental processes. The regression over all cell pairs between $\text{Corr}(s2)$ and $\text{Corr}(m)$ measures the similarity of neurone-ensemble firing patterns between phase S2 and phase M. The difference between $\text{Corr}(s1)$ and $\text{Corr}(s2)$ is taken to reflect the influence of the experience in phase M. If this is the case, then the magnitude and sign of this difference should be linearly related to the magnitude and sign of $\text{Corr}(m)$.

Between-phase correlation coefficients were calculated on a within-session basis. If the numbers of neocortical and hippocampal cells recorded were n_c and n_h , respectively, then the numbers of pairwise activity correlations were $n_c(n_c - 1)/2$ within the neocortex, $n_h(n_h - 1)/2$ within the hippocampus, and $n_c n_h$ between regions. These pairwise correlations were calculated for each of phases S1, M, and S2 and were used to calculate $\text{Corr}(s1,m)$ and $\text{Corr}(s2,m)$. If a data set contained fewer than five cells in either hippocampus or neocortex, the corresponding within- and between-region correlations were excluded from the analysis.

Another way to determine whether there is an effect of experience on the activity pattern during sleep, one that cannot be predicted from the sleep before the experience, is to measure the strength of the linear relationship between variables $\text{Corr}(m)$ and $\text{Corr}(s2)$ after

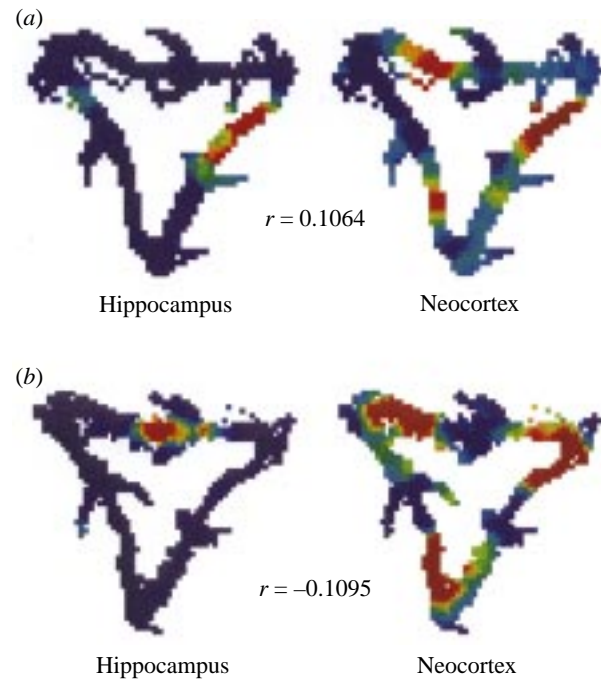


Figure 1. Examples of spatial firing patterns for hippocampal and neocortical neurons, on the triangular track used for some of the experiments. (a) Activity patterns for two cells whose spike trains were positively correlated; (b) activity patterns for two cells with negative correlation. As this figure suggests, spatial firing patterns for cells from the neocortical regions examined here were always nearly symmetrical across the three sides of the triangle, whereas hippocampal pyramidal cells always showed spatially asymmetrical firing patterns. Firing rates are scaled independently in each plot, with red corresponding to the highest rates, and dark blue corresponding to zero.

controlling for the effect of the variable $\text{Corr}(s1)$. This is given by the partial correlation coefficient between $\text{Corr}(m)$ and $\text{Corr}(s2)$ controlling for $\text{Corr}(s1)$, which is denoted as $\text{PCorr}(m,s2|s1)$ (Kleinbaum *et al.* 1988):

$$\text{PCorr}(m,s2|s1) = \frac{\text{Corr}(m,s2) - \text{Corr}(s1,m)\text{Corr}(s1,s2)}{\sqrt{(1 - \text{Corr}(s1,m)^2)(1 - \text{Corr}(s1,s2)^2)}} \quad (2)$$

(g) Temporal bias of cross-correlation

For a pair of cells, whose binned firing rates are given by x_i and y_i , a measure of temporal ordering was calculated on the basis of their cross-correlation histogram (Skaggs & McNaughton 1996). As illustrated in figure 2, the temporal bias B for the cells was defined to be the difference between the cross-correlation integrated over a T ms window after zero, and the cross-correlation integrated over a T ms window preceding zero. That is

$$B_{x,y}(T) = \int_0^T r_{x,y}(t) dt - \int_{-T}^0 r_{x,y}(t) dt. \quad (3)$$

For $t > 0$, $r_{x,y}(t)$ is the cross-correlation at lag t , defined as




Figure 2. Example of temporal bias for two neocortical cells. The three plots show cross-correlograms computed from data taken while the rat ran on a triangular track (Maze, *b*), during sleep beforehand (Sleep1, *a*), and during sleep afterwards (Sleep2, *c*). The measure of temporal bias is the difference between the light- and dark-shaded regions, i.e. the difference between the areas under the cross-correlograms for 200 ms after zero and 200 ms before zero. This example illustrates that the overall pattern of bias in the correlogram during Sleep2 resembles the bias during the Maze session more strongly than does the bias during Sleep1. As observed in the hippocampus (Skaggs & McNaughton 1996), there appears to be a compression of the peak in the correlogram during Sleep2 compared with Maze, suggesting an accelerated replay of sequence information; however, this aspect is not considered in detail here. Note that the count of spike pairs (ordinate) used 10 ms bins.

$$r_{x,y}(t) = \frac{\sum_{i=1}^{n-t} (x_i - \bar{x})(y_{i+t} - \bar{y})}{\sqrt{\sum_{i=1}^{n-t} (x_i - \bar{x})^2 \sum_{i=1+t}^n (y_i - \bar{y})^2}} \quad (4)$$

in which

$$\bar{x} = \sum_{i=1}^{n-t} x_i / (n-t), \quad \bar{y} = \sum_{i=1+t}^n y_i / (n-t).$$

When $t < 0$, we have $r_{x,y}(t) = r_{y,x}(-t)$.

$B_{x,y}(t)$ is a normalized measure of the difference between the number of events in which a spike from cell x was followed within T ms by a spike from cell y , and the number of events in which a spike from cell y was followed within T ms by a spike from cell x .

Following Skaggs & McNaughton (1996), for most analyses, T was taken as 200 ms.

3. RESULTS

Neurons in the hippocampus were classified as pyramidal cells or interneurons according to standard criteria (Ranck 1973). In the present study, a fraction of neocortical cells (about 10%), diffusely scattered in depth, showed qualities similar to those of hippocampal interneurons, including short spike width, amplitude symmetry between negative and positive peaks, high overall firing rates (typically above 10 Hz), and regular discharge patterns. In accordance with several intracellular recording experiments (Connors *et al.* 1982; Connors & Gutnick 1990), these cells are considered likely to be interneurons, and are not considered further in the present report. It may be noted that none of the cells recorded from the neocortex in the present study was modulated by the hippocampal EEG theta rhythm.

Hippocampal pyramidal neurons typically exhibited spatially selective firing on the apparatus, as expected from many previous studies. As described previously (McNaughton *et al.* 1994), many neocortical cells in the regions recorded from here exhibited differential firing related to some component of the rat's behaviour or to the maze geometry itself. This inference is based on the fact that the spatial firing patterns exhibited symmetries that were correlated with the shape of the apparatus; no attempt was made in the present studies to characterize the basis of these correlates. Examples of behavioural correlates (based on the spatial distribution of firing rates) are illustrated in figure 1.

Based on 241 hippocampal and 568 neocortical cells, 5659 pairwise correlations were calculated in each of the S1, M and S2 phases, including 2762 neocortical cell pairs, 1982 hippocampal cell-pairs, and 915 hippocampal–neocortical cell pairs. When the data from all recording sessions were pooled, $\text{Corr}(s2,m)$ was higher than $\text{Corr}(s1,m)$, for all types of interaction. These data are illustrated in figure 3. Thus, for each type of interaction, the larger the correlation between any two cells on the maze, the greater was the increase in the correlation between the same two cells from S1 to S2 ($p < 0.0001$ for all three comparisons, paired t -test).

To confirm that this effect was not confined to a small number of data sets, partial regression analysis was conducted on each data set independently. Figure 4 shows the average correlations (over data sets) between $\text{Corr}(s1)$ and $\text{Corr}(m)$ and partial correlations between $\text{Corr}(m)$ and $\text{Corr}(s2)$ after controlling for $\text{Corr}(s1)$.

For all comparisons, there was a significant ($p < 0.05$) overall effect of M on S2 after controlling for S1. For within-structure comparisons (NC and HC in figure 4), there were also significant relations between the cell–cell interactions on the maze and the corresponding interactions during S1 ($p < 0.05$); this result indicates that, within each structure, the interactions on the maze are partly predictable from the interactions that occur before the animal experiences the maze on the day in question. This was not true for the hippocampal–neocortical interaction (NC–HC, figure 3), for which the coherence between S1 and M

was not significantly different from zero, but was significantly different from the other comparisons.

As reported previously by Skaggs & McNaughton (1996), the temporal bias in the interaction between hippocampal cell pairs is preserved between M and S2 more strongly than between M and S1. This finding was replicated in the current experiment, as shown in the left column of figure 5 ($p < 0.001$, paired t -test). Temporal bias for neocortical cell pairs significantly reflected the maze bias during both S1 and S2, but the degree of coherence was significantly enhanced during S2 relative to S1 ($p < 0.01$, paired t -test). For hippocampal–neocortical cell pairs, there was no significant correlation of bias during either S1 or S2 with bias during M and no significant difference between S1 and S2 in this correlation.

4. DISCUSSION

The main finding of the present study is that, during periods of slow-wave sleep after an episode of spatially extended behaviour, patterns of neuronal correlation that were manifest during the behaviour reemerge in neocortical and hippocampal circuits as well as in the interactions between these two areas. For all three comparisons, the distribution of pairwise firing-rate correlations during sleep after the behaviour (S2) more strongly reflected the patterns observed during behaviour (M) than did the corresponding distributions recorded during sleep preceding behaviour (S1). The reactivation of patterns of correlation between hippocampus and neocortex is a necessary (but not sufficient) requirement of the theory that the hippocampus initiates and orchestrates the reactivation of traces of recent experience in neocortical circuits. Whether or not this is the case, this result indicates that, at a given moment, the patterns of activity in these two regions during sleep tend to reflect components of the same experience.

For within-area comparisons, there was also some correlation between S1 and M; this correlation indicates that patterns of network activity on the maze are related to patterns that occurred during sleep beforehand. This is to be expected for the posterior cortex, because many of these cells exhibit clear, mutually exclusive categories of behavioural correlates (e.g. left vs right turns (McNaughton *et al.* 1994)). This is presumably a reflection of previously established synaptic connectivity. It is also to be expected for within-hippocampus comparisons, because place fields are highly consistent from day to day in a given environment (Thompson & Best 1990), and the rats in the present study had had substantial previous experience in the recording apparatus. Moreover, there is evidence to suggest that the possible patterns of place-field overlap are to some extent preconfigured within the hippocampal synaptic matrix (Hill 1978; Quirk *et al.* 1992; Kudrimoti *et al.* 1997; Gothard *et al.* 1996; McNaughton *et al.* 1996). On the other hand, if, as the evidence suggests, the hippocampus is primarily involved in providing a spatial coordinate representation, whereas the neocortical areas studied here are primarily involved in sensory–motor integration, there should be no strongly constrained set of preconfigured

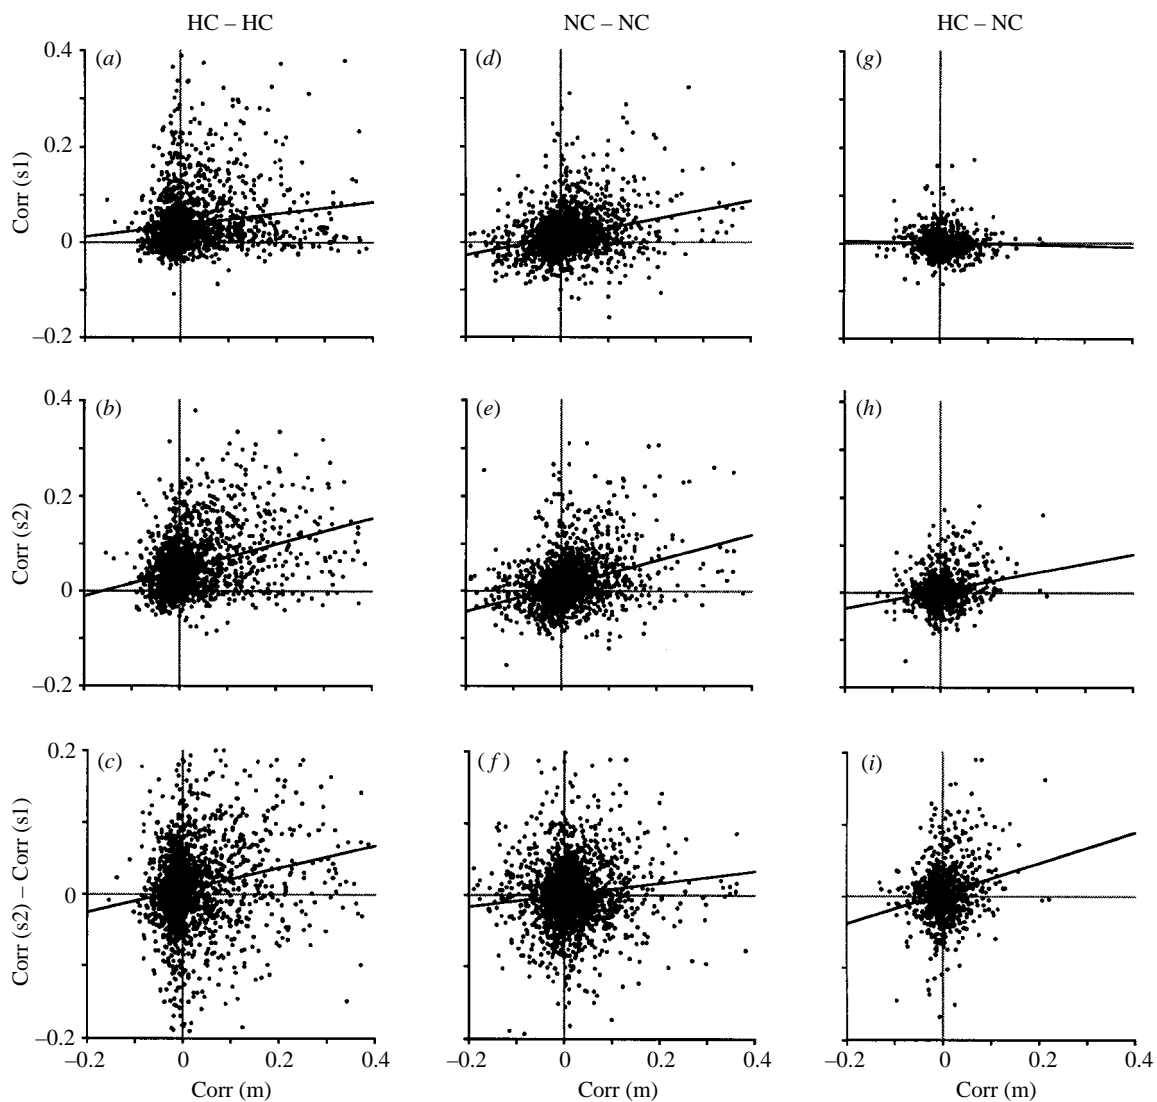


Figure 3. Coherence between correlation on the maze and correlation during sleep, for different types of cell pair. Each point in a scatter plot represents one pair of simultaneously recorded neurons. The abscissa is the correlation of these neurons on the maze, $\text{Corr}(m)$; the ordinate is their correlation during Sleep1 ($\text{Corr}(s1)$; a, d, g), their correlation during Sleep2 ($\text{Corr}(s2)$; b, e, h), or the difference between these correlations ($\text{Corr}(s2) - \text{Corr}(s1)$; c, f, i). The left column (a, b, c) includes only pairs where both cells were recorded from the CA1 layer of the hippocampus; the centre column (d, e, f) only pairs of neurons both from the neocortex, and the right column (g, h, i) pairs with one cell each from the hippocampus and the neocortex. Each plot demonstrates a statistically significant positive correlation ($p < 0.001$), with the exception of (g) (hippocampus–neocortex for Sleep1 versus maze).

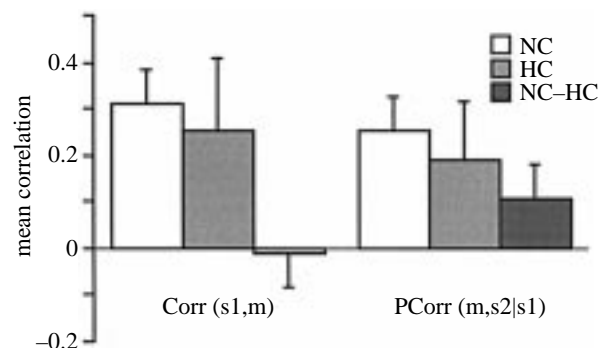


Figure 4. Histograms of correlations and partial correlations between Maze, Sleep1 and Sleep2 for different types of cell pair, across data sets. (a) Mean correlations between Maze and Sleep1, for correlations between HC, NC, and HC–NC cell pairs. (b) Partial correlations of Maze with Sleep2, after the expected effect of correlation in Sleep1 has been subtracted. The error bars represent 95% confidence intervals. All correlations and partial correlations are significantly greater than zero ($p < 0.05$), except the correlation between Maze and Sleep1 for HC–NC cell pairs.




Figure 5. Coherence between temporal bias on the maze and temporal bias during sleep, for different types of cell pair. Each point in a scatter plot represents one pair of simultaneously recorded neurons. The abscissa is the temporal bias of these neurons on the maze ($\text{Bias}(m)$); the ordinate is their bias during Sleep1 ($\text{Bias}(s1)$; a,d,g), their bias during Sleep2 ($\text{Bias}(s2)$; b,e,h), or the difference of these bias values (c,f,i). The left column (a,b,c) includes pairs where both cells were recorded from the CA1 layer of the hippocampus; the centre column (d,e,f) pairs of neurons both from the neocortex, and the right column (g,h,i) pairs with one cell each from the hippocampus and the neocortex. Each plot in the first two columns ($a-f$) demonstrates a statistically significant positive correlation ($p < 0.01$). None of the correlations for HC–NC pairs (g,h,i) is significant.

relations between the two regions. For example, it is *a priori* equally likely that the animal will turn right or left at location X. Thus, the fact that there was no strong relation between S1 and M in the patterns of correlation of hippocampal–neocortical pairs, but a significant relation between S2 and M, can be taken as evidence for a learning-related change in connection strengths in the direct or indirect pathways connecting these two regions. These changes apparently decayed within the *ca.* 24 h time periods that separated recording sessions.

The decay of hippocampal–neocortical correlations appears to present some difficulty for the simple form of the consolidation theory, because the theory assumes that the enhanced connections should persist

for a time period of weeks at least (in the rat). The effect might be explainable by a nonlinearity between the extent of reactivation in the hippocampus and the corresponding pattern in the neocortex, such that the observed reduction in intrahippocampal correlations from one day to the next (Wilson & McNaughton 1994; Skaggs & McNaughton 1996) leads to a greater reduction in hippocampal–neocortical correlations. This explanation assumes that the coherence seen in neocortex during S1 is due largely to stable patterns of connections in neocortex.

The present results also replicate the previous finding that some aspects of the temporal structure of hippocampal population activity in a spatial task are

preserved during sleep afterwards (Skaggs & McNaughton 1996), and demonstrate that a similar effect occurs for neocortical cell pairs, but not for mixed hippocampal–neocortical pairs. This pattern of results is consistent with the prediction of the consolidation theory that information is transferred from the neocortex to the hippocampus during the waking state, and then replayed back to the neocortex from the hippocampus during sleep. Thus, temporal structure intrinsic to the hippocampus or neocortex should be replicated, but, because of the postulated change in the direction of information flow, the temporal relations between the two structures should be altered. For example, suppose that in waking, a pattern A–B–C of activity in neocortex gives rise, after a small delay, to a pattern A'–B'–C' in the hippocampus; and that during sleep the same hippocampal pattern precedes the same neocortical pattern. Then we have the following scenario:

	Awake	Sleeping
NC:	A–B–C	A–B–C
HC:	A'–B'–C'	A'–B'–C'

The temporal relations are quite different in the two states, but a simple time-shift of the hippocampal pattern would bring them back into their original alignment. In reality, the patterns during sleep might also be subjected to a considerable temporal compression as well as to a time shift (Skaggs & McNaughton 1996; Skaggs *et al.* 1996; August & Levy 1996). These questions await further study.

An essential function of long-term memory is to construct, from past experience, internal representations of the world that permit adaptive generalization and appropriate responses to novel, or partly novel, situations. The development of such internal representations appears to require repeated interleaved exposure to exemplars of the categories to be formed, with small adjustments to the synapses on each trial (McClelland *et al.* 1995). Attempts to store new information all at once in such a memory lead to 'catastrophic interference' with items already stored (McClosky & Cohen 1989). Survival, however, often requires adding new items to the existing categorical structure, with only one or a few experiences. This leads to the idea of storage in a temporary buffer, followed by quasi-random reactivation during periods when the system is offline. The present findings, that memory traces are indeed reactivated during sleep both in the hippocampus and the neocortex, and that the patterns of correlation between these structures during sleep are consistent with prior behaviour, provide a key piece of evidence in support of this theory. It still remains to be demonstrated, however, that the coherence of pattern reactivation across widely separated neocortical modules is dependent on an orchestrating flow of stored information from the hippocampal formation.

Supported by grants MH46823, MH01227 and the McDonnell and Pew Foundations. We thank M. Suster, K. Moore, L. Frank, R. D'Monte, J. Gerrard, K. Stengel, V. Pawlowski and Jie Wang for technical assistance with this project.

REFERENCES

- Amaral, D. G. & Witter, M. P. 1995 Hippocampal formation. In *The rat nervous system*, 2nd edn (ed. G. Paxinos), pp. 443–493. San Diego, CA: Academic Press.
- August, D. A. & Levy, W. B. 1996 Temporal sequence compression by a hippocampal network model. In *INNS World Congress on Neural Networks*, pp. 1299–1304. Hillsdale, NJ: Lawrence Erlbaum. (In the press.)
- Connors, B. W., Gutnick, M. J. & Prince, D. A. 1982 Electrophysiological properties of neocortical neurons *in vitro*. *J. Neurophys.* **48**, 1302–1320.
- Connors, B. W. & Gutnick, M. J. 1990 Intrinsic firing patterns of diverse neocortical neurons. *Trends Neurosci.* **13**, 99–104.
- Gothard, K. M., Skaggs, W. E., Moore, K. M. & McNaughton, B. L. 1996 Binding of hippocampal CA1 neural activity to multiple reference frames in a landmark-based navigation task. *J. Neurosci.* **16**, 823–835.
- Gray, C. M., Maldonado, P. E., Wilson, M. A. & McNaughton, B. L. 1995 Tetrodes markedly improve the yield and reliability of multiple single-unit isolation from multi-unit recordings in cat striate cortex. *J. Neurosci. Meth.* **63**, 43–54.
- Hill, A. J. 1978 First occurrence of hippocampal spatial firing in a new environment. *Expl. Neurol.* **62**, 282–297.
- Kleinbaum, D. G., Kupper, L. L. & Muller, K. E. 1988 *Applied regression analysis and other multivariable methods*. Belmont, CA: Duxbury Press.
- Kudrimoti, H. S., McNaughton, B. L., Barnes, C. A. & Skaggs, W. E. 1995 Recent experience strengthens pre-existing correlations between hippocampal neurons during sleep. *Soc. Neurosci. Abstr.* **21**, 941.
- Marr, D. 1971 Simple memory: a theory for archicortex. *Phil. Trans. R. Soc. Lond. B* **262**, 23–81.
- McClelland, J. L., McNaughton, B. L. & O'Reilly, R. C. 1995 Why are there complementary learning systems in hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychol. Rev.* **5**, 245–286.
- McClosky, M. & Cohen, N. J. 1989 Catastrophic interference in connectionist networks: the sequential learning problem. In *The psychology of learning and motivation* (ed. G. H. Bower), pp. 109–165. New York: Academic Press.
- McNaughton, B. L., Barnes, C. A., Meltzer, J. & Sutherland, R. J. 1989 Hippocampal granule cells are necessary for normal spatial learning but not for spatially-selective pyramidal cell discharge. *Expl. Brain Res.* **76**, 485–496.
- McNaughton, B. L., Mizumori, S. J. Y., Barnes, C. A., Leonard, B. J., Marquis, M. & Green, E. J. 1994 Cortical representation of motion during spatial navigation in the rat. *Cerebr. Cortex* **4**, 27–39.
- McNaughton, B. L., O'Keefe, J. & Barnes, C. A. 1983 The stereotrode: a new technique for simultaneous isolation of several single units in the central nervous system from multiple unit records. *J. Neurosci. Meth.* **8**, 391–397.
- Nadel, L., Wilner, J. & Kurz, E. M. 1985 Cognitive maps and environmental context. In *Context and learning* (ed. P. D. B. A. Tomie), pp. 385–406. Hillsdale, NJ: Lawrence Erlbaum.
- O'Keefe, J. & Nadel, L. 1978 *The hippocampus as a cognitive map*. Oxford: Clarendon Press.
- Pavlidis, C. & Winson, J. 1989 Influences of hippocampal place cell firing in the awake state on the activity of these cells during subsequent sleep episodes. *J. Neurosci.* **9**, 2907–2918.
- Qin, Y.-L., McNaughton, B. L., Skaggs, W. E. & Barnes, C. A. 1995 Reactivation during sleep of cortico-cortical and hippocampo-cortical correlation states from preceding behaviour. *Soc. Neurosci. Abstr.* **21**, 941.
- Qin, Y.-L., McNaughton, B. L., Skaggs, W. E., Barnes, C. A., Suster, M. S., Weaver, K. L. & Gerrard, J. L. 1996 Interaction between hippocampus and neocortex in the

- replay of temporal sequences during sleep. *Soc. Neurosci. Abstr.* **22**, 1872.
- Quirk, G. J., Muller, R. U., Kubie, J. L. & Ranck, J. B. 1992 The positional firing properties of medial entorhinal neurons: description and comparison with hippocampal place cells. *J. Neurosci.* **12**, 1945–1963.
- Ranck, J. B. Jr 1973 Studies on single neurons in dorsal hippocampal formation and septum in unrestrained rats. I. Behavioural correlates and firing repertoires. *Expl Neurol.* **41**, 461–531.
- Recce, M. & O'Keefe, J. 1989 The tetrode: a new technique for multi-unit extracellular recording. *Soc. Neurosci. Abstr.* **19**, 1250.
- Scoville, W. B. & Milner, B. 1957 Loss of recent memory after bilateral hippocampal lesions. *J. Neurol. Neurosurg. Psychiatr.* **20**, 11–21.
- Shen, B. & McNaughton, B. L. 1997 Modeling the spontaneous reactivation of experience-specific hippocampal cell assemblies during sleep. *Hippocampus* **6**, 1–8.
- Skaggs, W. E. & McNaughton, B. L. 1996 Replay of neuronal firing sequences in rat hippocampus during sleep following spatial experience. *Science* **271**, 1870–1873.
- Skaggs, W. E., McNaughton, B. L., Wilson, M. A. & Barnes, C. A. 1996 Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus* **6**, 149–172.
- Squire, L. R. 1992 Memory and the hippocampus: a synthesis from findings with rats, monkeys, and humans. *Psychol. Rev.* **99**, 195–231.
- Teyler, T. J. & Discenna, P. 1986 The hippocampal memory indexing theory. *Behav. Neurosci.* **100**, 147–154.
- Thompson, L. T. & Best, P. J. 1990 Long-term stability of place-field activity of single units recorded from the dorsal hippocampus of freely behaving rats. *Brain Res.* **509**, 299–308.
- Vanderwolf, C. H. 1969 Hippocampal electrical activity and voluntary movement in the rat. *EEG Clin. Neurophys.* **26**, 407–418.
- Wilson, M. A. & McNaughton, B. L. 1993 Dynamics of the hippocampal ensemble code for space. *Science* **261**, 1055–1058.
- Wilson, M. A. & McNaughton, B. L. 1994 Reactivation of hippocampal ensemble memories during sleep. *Science* **265**, 577–708.
- Winson, J. & Abzug, C. 1977 Gating of neuronal transmission in the hippocampus: efficacy of transmission varies with behavioural state. *Science* **196**, 1223–1225.
- Zilles, K. 1985 *The cortex of the rat*. Berlin: Springer.

