HIPPOCAMPUS 6:685–692 (1996)

Modeling the Spontaneous Reactivation of Experience-Specific Hippocampal Cell Assembles During Sleep

Bin Shen and Bruce L. McNaughton

Division of Neural Systems, Memory and Aging, University of Arizona, Tucson, Arizona

ABSTRACT: During slow-wave sleep (SWS) following periods of spatial activity, hippocampal place cells that were temporally correlated, by virtue of the overlap of their place fields, exhibit enhanced temporal correlations, even though the animal sleeps in a different location (Wilson and McNaughton [1994] Science 267:676–679). The discharge of cells with overlapped place fields is more correlated in subsequent sleep, particularly during sharp waves, than in sleep episodes prior to the behavior, or than cell pairs with non-overlapped place fields. The reactivated correlated states appear during hippocampal sharp waves (SPWs), and are weak or absent in the inter-SPW interval. A simple conceptual hypothesis for this phenomenon is developed, based on the idea that hippocampal place fields reflect a two-dimensional distribution of continuously overlapping dynamic attractors in which each location is represented by the self-sustaining activity of a small subset of neurons with overlapping place fields. A numerical simulation of this hypothesis, based on a simplified representation of the CA3 recurrent network, accounts qualitatively for the main observations, including SPW-like dynamics. It is shown that, under conditions in which the connection patterns have been previously established, either associative or nonassociative mechanisms might underlie the reactivation of recently experienced states. These two alternatives appear, under at least some conditions (e.g., sparse coding), to be indistinguishable. © 1997 Wiley-Liss, Inc.

KEY WORDS: place cells, memory reactivation, associative learning

INTRODUCTION

What Is the Goal of the Model?

In this work we present a simple series of simulations that are intended to shed light on the possible mechanisms underlying the reactivation of correlated states of neuronal activity in hippocampus during periods of "off-line" processing such as slow wave sleep or quiet wakefulness. The main questions addressed are as follows: 1) Is it plausible to conclude, from data showing that pairwise cross-correlations between cells during waking are strongly re-expressed during sleep (Wilson and McNaughton, 1994; Skaggs and McNaughton, 1996), that this reflects the reactivation of global states of hippocampal activity, rather than merely fragmentary correlations between small groups or pairs of cells? 2) What sorts of synaptic architectures and what sorts of dynamic changes in coupling parameters might support this reactivation? Specifically, is cooperative synaptic enhancement, such as postulated by Hebb (1949), necessary or sufficient? Can one gain insight into the physiological dynamics of the hippocampus during "off-line" periods from the dynamics of a simple recurrent network containing embedded attractors presentated with random inputs?

What Are the Data to be Explained?

The hippocampal formation is the highest level of association cortex and plays an essential role both in spatial navigation (O'Keefe and Nadel, 1978) and in the initial encoding of important components of memory (Squire et al., 1989, for review). For at least some of these components, the role of the hippocampus appears to be transitory, in that damage to the hippocampus, while preventing further acquisition, leaves well-established memories relatively intact. One prominent theory is that the hippocampus acts to consolidate memories into their long-term form (McGaugh and Herz, 1972; Squire et al., 1984; MacKinnon and Squire, 1989; Kim and Fanselow, 1992). According to this view, information stored in the hippocampal formation during acquisition is used to reactivate traces of the experience elsewhere in the brain (particularly in neocortex), and this reactivation leads to memory consolidation. A computational explanation for why the memory system should be organized in this way was suggested recently by McClelland et al. (1995). Marr (1970, 1971) suggested that this reactivation process should occur during sleep, among other reasons, because the neocortex would not be busy processing external input during this time (Fishbein and Gutwein 1977). Wilson and McNaughton (1994) provided evidence for Marr's suggestion by showing that the correlation structure in the activity of a population of hippocampal place cells that occurs during spatial behavior, as a consequence of the pattern of overlap of the place fields, is preserved in subsequent slow-wave sleep. More recently, Qin et al. (1995) have shown that the same re-expression, during sleep, of behaviorally induced neural activity correlations occurs both within the neocortex, as well as between neocortical and hippocampal neurons. Thus, recently experienced states of correlated neuronal activities are

Accepted for publication August 19, 1996.
Address correspondence and reprint requests to Dr. B.L. McNaughton, 344 Life Sciences North, University of Arizona, Tucson, AZ 85724.

re-expressed widely across the brain during subsequent sleep, although it remains to be determined whether the effect observed in neocortex depends on an intact hippocampus. The re-emergence of pairwise correlations of neural activity is consistent with the idea that whole ensemble patterns of activity that occur during behavior re-emerge spontaneously during sleep; however, whether whole ensembles or mere fragments of them are reactivated cannot be tested directly without recording from the entire ensemble (or a large part of it), which is currently not technically possible. Therefore, to explore the possible mechanisms of this phenomenon, we have developed a numerical simulation for how the reactivation of such ensembles might arise in a network whose architecture captures some of the basic features of the CA3 field of the hippocampus. Wilson and McNaughton suggested that Hebbian associative synaptic modification within the recurrent collateral connections of CA3 pyramidal cells might underlie such an effect. During the course of these studies, however, we found that, given the appropriate pre-existing architecture, non-associative processes that merely bias the excitability of cells that were highly active during behavior, as shown originally by Pavlides and Winson (1989), might also account for the existing observations.

How Should the Reader Evaluate the Model?

A model is not an experiment. It is really nothing more than a (more-or-less) well-formulated hypothesis in which the implications of a collection of postulates is made explicit. If one is lucky, the exercise of modeling can also lead to insights about possible real phenomena and mechanisms which were not previously apparent from a mere statement of the postulates. We believe we will have accomplished our goals if, as a result of this work we have 1) demonstrated the plausibility of the conclusion of Wilson and McNaughton (1994) that global activity states are reactivated during sleep, and 2) gained insights into the possible mechanisms and dynamics underlying this process, which will lead to new experiments.

BODY

Network Structure

In order to study the problem in a simple, tractable form, a one-layer recurrent network of 100-point "pyramidal" neurons with continuous, sigmoid activation functions was studied. A single inhibitory neuron was used for global threshold control (e.g., McNaughton and Morris, 1987). The dynamics of the network are governed by the following equations:

$$C\frac{\mathrm{d}u}{\mathrm{d}t} = -R^{-1}u + Wg(u) - i1 + I_e \qquad (1)$$

$$\text{and } i = a1^T g(u) + \mathrm{b}$$

where u is the vector representing the membrane potential of the pyramidal cells, i is a scalar representing the strength of inhibition, I_e is the vector representing external input, 1 is a vector whose elements are all ones, W is the matrix of connections among pyramidal cells, R and C are diagonal matrices representing the membrane resistance and capacitance, respectively, a is an inhibitory constant, b is a bias constant, and $g(u)$ is an activation function representing firing rates of pyramidal cells at membrane potential u. In the present case, R and C are normalized to unit values, $a = 3$ and $b = 0.3$.

Specification of the Synaptic Weight Matrix

In an ensemble of simultaneously recorded place cells with place fields in a given environment, these fields are distributed approximately uniformly throughout the environment (Wilson and McNaughton, 1993), and, to a first approximation, each cell has a two-dimensional Gaussian probability density function for its firing:

$$P_i(x, y) = \frac{1}{2\pi\sigma^2}\exp\left(-\frac{(x - x_i)^2 + (y - y_i)^2}{2\sigma^2}\right) \qquad (2)$$

where $P_i(x, y)$ is the probability density of discharge of cell i, and x and y are the coordinates of the center of the place field of cell i. The Gaussian assumption is used for convenience. In the limit of large numbers of units, the exact form of the place fields is probably irrelevant, provided they are relatively compact. Also, the possible existence of multiple place fields for one cell in the same environment does not pose serious problems (Samsonovich and McNaughton, unpublished simulations). Given these basic assumptions, a simple Hebbian learning rule (Hebb, 1949) leads to a weight matrix in which the strength of the connections between two cells is proportional to the overlap of their place fields. Although the idea that CA3 place cells with overlapping Gaussian place fields would undergo mutual enhancement of their common synapses has been with us for many years, the first formal expression of equation 2 was suggested by Muller et al. (1991). The latter authors, however, used this idea to explore the possibility that the hippocampus acts according to the principles of mathematical graph theory and did not consider the nature of the attractor dynamics that such a matrix might lead to from the point of free recall, which is the theme of the present work.

The learning rule assumes that synaptic strengthening is proportional to the probability that the pre- and postsynaptic cells are coactive:

$$\Delta W_{ij} \propto \iint P_i(x, y)P_j(x, y)\,dx\,dy. \qquad (3)$$

Substituting equation 2 into equation 3, integrating over space, and choosing a constant of proportionality K, we have

$$\Delta W_{ij} = K\exp\left(-\frac{(x_i - x_j)^2 + (y_i - y_j)^2}{4\sigma^2}\right) \qquad (4)$$

The Gaussian weight distribution, W_{ij}, emerges after a period of uniform "exploration" of the environment represented by the cells in question. In the simulations, 30% of the initial weights were set

at random to a small positive value (0.01), and the rest were set to zero; during exploration, K in equation 4 was set to 1.0.

Simulation 1

The first simulation addressed the fundamental question of whether a network of the sort just described would have the property of spontaneously reactivating coherent representations of locations recently visited. By spontaneously, we mean that these representations would emerge preferentially during random input. For lack of more information, we make the tentative assumption that random inputs are all that is available during slow-wave sleep (SWS). Random input represents a "worst case" scenario. The presence of relevant structure in the input would obviously assist the recall process. It is intuitively apparent that, in the absence of overwhelming external input, the dynamically favored activity configurations in such a network are those in which activity is focused in tight clusters of cells with overlapping place fields, i.e., coherent representations of locations. Due to the structure of the weight distribution, these are the configurations in which the activity of a given cell receives the maximum support from cells with neighboring place fields. To explore this in simulation, 36 of the 100 units were selected at random to have "place fields" in the environment. The field centers were distributed in a 6×6 lattice, and the field size (i.e., σ in equation 2) was set to 0.15. Gaussian distributed random inputs (mean $= 0$; SD $= 0.1$) were presented to the network and maintained constant while the network was allowed to settle to its equilibrium state for a given input. The equilibrium values were taken as the firing rate vector for the current cycle. The network was reset to zero following each cycle. The simulation included 5,000 cycles. As a consequence of the nonlinear activation function g, the network responded to random inputs in one of two ways: Either all cells settled to states of low activity, or a few cells settled to states of high activity and the rest settled to low states (Fig. 1A). Thus, the sequence of equilibrium activity levels over cycles could be characterized as intermittent bursts of activity separated by periods of relative silence. This pattern is reminiscent of hippocampal neuronal activity during SWS, in which silent periods are punctuated by sharp-wave bursts (Fig. 1B). To explore this further, we examined the interburst interval distribution for the output of the model, and compared it with the inter–sharp-wave interval distribution recorded from a rat during SWS (Fig. 2). Both distributions were approximately exponential, which is what is expected for a random (Poisson) process. Note that, if actual sharp-wave bursts reflect convergence on an attractor state as seen in the simulation, the elevated activity must be extinguished by some intrinsic mechanism such as, possibly, the calcium-dependent potassium conductance. We did not simulate burst termination explicitly.

Consistent with the results of Pavlides and Winson (1989), the cells in the simulation that had place fields in the environment (i.e., those that had strengthened interconnections) were more active in "SWS" than cells without place fields (mean activities 0.020 vs. 0.005, respectively, $P < 0.01$). More importantly, as observed experimentally by Wilson and McNaughton, cells with overlapping place fields exhibited correspondingly strong positive cross-correlations, whereas cells with little or no overlap were weakly negatively correlated (Fig. 3).

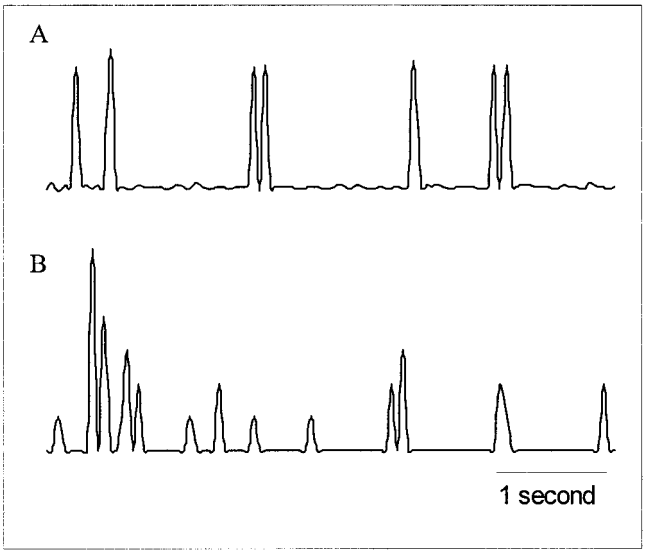

FIGURE 1. The simulation exhibits intermittent states of high activity separated by periods of low activity, which mimic the sharp wave activities during SWS. A: Temporal distribution of the network activity. The peaks represent high (attractor) states of network equilibrium activity in a single cycle of simulation, when the network is allowed to settle during a given random input. The horizontal axis is thus cycle number. Between cycles, the network was reset (see text), and random inputs were changed. Note that most random inputs result in a low-activity equilibrium state. B: Summed firing rate of 45 simultaneously recorded cells from a freely moving rat during SWS. As shown by Buzsaki (1989) these bursts of activity correspond to the sharp waves observed in hippocampal EEG during SWS and quiet wakefulness. In the simulation results the peaks of activity are all of nearly the same magnitude, whereas the peaks in the actual data are much more variable. We speculate that, in actual sharp-wave bursts, the net activity is also much more constant. The variability arises because different cells are involved in different groups during different sharp waves. Under such conditions, the exceedingly tiny sample size of simultaneously recorded cells (only 45 out of about 300,000) would result in highly variable measures of net activity during the bursts.

The foregoing results are consistent with the idea that, when the network settles to a high activity state, these states reflect coherent representations of locations. The alternative is that the reactivation may be fragmentary, involving pairs or small subsets of cells with place fields distributed over the environment. To explore this, we examined a group of nine cells, centered on the cell whose correlations are illustrated in Figure 3. Because the place field centers in the model were distributed in a unit rectangular grid, at the level of inhibition used in the simulations there are four dynamically preferred configurations involving a given cell, as illustrated in Figure 4. These correspond to location representations of the middles of the four grid squares of which the cell in question is a member. We compared the frequency with which the cell in question was in a high-activity state to the frequency with which one of its preferred configurations was active (i.e., four rectangularly adjacent cells were in a high state). As shown in Figure 4, the activation probability of the cell in

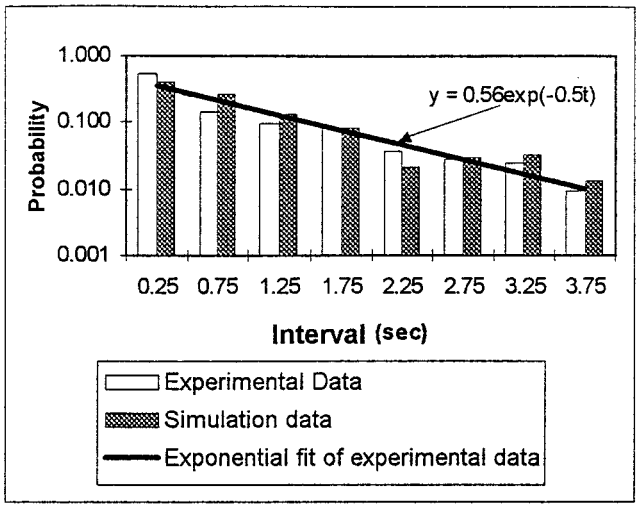

FIGURE 2. Spontaneous memory reactivation can be modeled by a Poisson process (in this case, a random input signal) which predicts that the interval distribution between successful reactivations ("high states of network activity") and failures will be exponentially distributed. It is interesting in this context that in an ensemble recording of 45 hippocampal neurons during sleep the interburst intervals also have an exponential distribution. While more recordings need to be analyzed, this illustrates the sort of insight and ideas for new analyses and experiments that even such a simple model may suggest.

question was almost equal to the sum of the probabilities of the individual clusters of which it was a member. Thus, in the model, the enhancement during "SWS" of pairwise cross-correlations of cells with overlapping place fields in the environment is almost entirely due to the reactivation of whole ensembles representing a given location.

Simulation 2

Place cells have place fields in more than one environment (O'Keefe and Conway, 1978; Kubie and Ranck, 1983), and there is little or no relationship among the sets of neighbors a cell may have in different environments (Hill, 1978; O'Keefe and Conway, 1978). It is reasonable to suppose that the synaptic weight distribution in the hippocampus simultaneously encodes the neighborhood relations of place fields in many environments. This could be simulated, for example, by randomly rearranging the place fields of the cells in the model and performing the learning procedure again multiple times. Indeed, much of the data on hippocampal place fields suggests that the synaptic matrix comes preconfigured in this way (possibly through early experience or through some ontogenetic process) to enable the possibility of many two-dimensional place field distributions (McNaughton et al., 1996).

These considerations raise the question of why the most recently visited environment dominates the correlation pattern of neural activity during SWS, at least for the first 30 min or so. The effect declines gradually over 30–60 min (Wilson and McNaughton, 1994; Kudrimoti et al., 1995). We hypothese that this involves two components, a long-term associative component that lasts for days and a short-term component that decays to zero over 30–60 min. The long-lasting component could be the result of long-term potentiation (LTP) during the original and subsequent exploration of the environment, or alternatively, could be prewired in the CA3 recurrent network before the exploration (see McNaughton et al., 1996). The short-term component can be either an associative synaptic modification as presented in simulation 1 or a nonassociative bias on cells that have recently been highly active.

To simulate these possibilities, five different place field configurations, corresponding to five different environments, were stored in the 100-neuron network according to equation 4. For each environment, 36 cells were chosen to have place fields. The selection process was arranged such that each pair of environments shared eight common place cells, but the locations of their fields were uncorrelated. Three conditions were simulated. In the first, no other factors were added, so that all environments had equal weight and inhibition. In the second condition, for one environ-

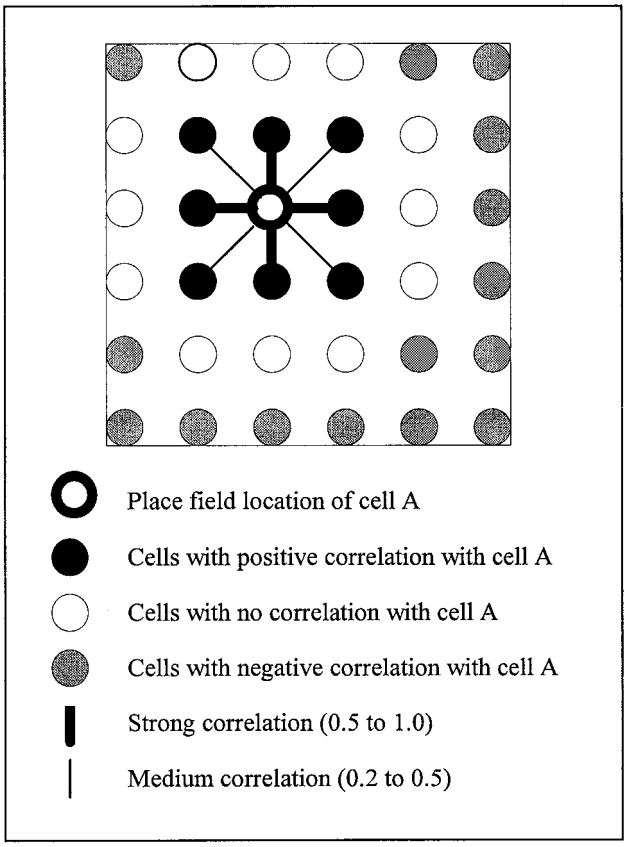

FIGURE 3. Correlation between cell A and the other place cells. The eight cells (marked with solid circles) that have place fields surrounding that of cell A have positive correlations with cell A. The cells (marked with gray circles) that have distant place fields from that of cell A are weakly negatively connected with cell A. Notice that the grid of cells surrounding and including cell A can be broken down into four separate sets of four cells each, of which cell A is a member of all four. Further analysis of the activations of these subsets is shown in Figure 4.

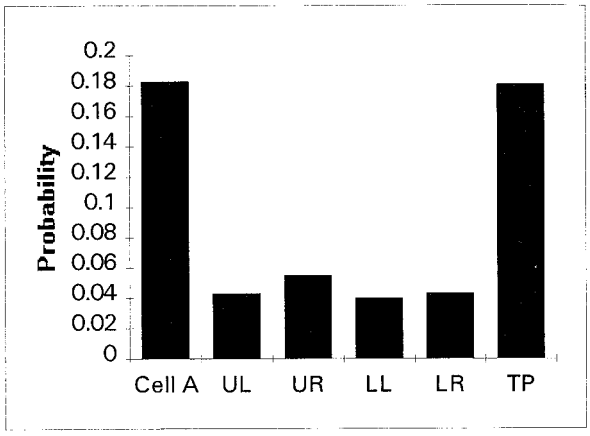

FIGURE 4. The activation probability of cell A and its surrounding spatial representations. The activation probability is defined as the relative frequency that a cell (or a spatial representation, a square grid of four cells in this simple model) was highly active, i.e., the firing rate of the cell (or of all the cells in the spatial representation) is above 0.7. Cell A and its eight neighbors form four spatial representations (see Fig. 3 for details): UL, the upper left four cells; UR, the upper right four cells; LL, the lower left four cells; LR, the lower right four cells; TP, the summed activation probability of the above four spatial representations. The fact that the activation probability of cell A is virtually identical with that of the sum of the four spatial representations to which it belongs is an indication of the coherence of the reactivation phenomenon. The probability of reactivation of partial representations, for example two pairs of neighboring cells, with each pair representing widely different regions of the environment, is exceedingly small.

ment (#5), the weights were increased by 50% over the others by setting K in equation 4 to 1.5; K = 1.0 for the other environments. In the third condition, the weights for all environments were equal (K = 1.0), but a nonassociative bias was added to the cells with place fields in environment 5, by setting the inhibitory constant b in equation 1 to 0.27; b = 0.30 for the other environments. This is equivalent to reducing the discharge threshold by 10%. For each condition, the network was presented with random inputs, as in simulation 1.

The simulation results of the three conditions are presented in Figures 5 and 6. When there was no bias, the network activity was distributed among all the environments (Fig. 5), and the correlation of cell discharge reflected the overlaps of place fields in different environments; therefore, a predominant correlation of cell discharge reflecting environment 5 would not be observed (Fig. 6D). When bias (either associative or nonassociative) was presented, the network activity was concentrated in environment 5 (Fig. 5), and the correlation of cell discharge dominantly reflected the overlaps within environment 5 (Fig. 6F). Figure 5 suggests that either associative or nonassociative bias may account for Pavlides and Winson's observation (1989). In Figure 6, the comparison of simulation results and experimental data suggests that a nonassociative excitability bias may also account for the spontaneous reactivation of recently experienced states (Wilson and McNaughton, 1994).

We further assessed the degree of coherence of reactivation for specific environments. The coherence index, which is the sum of correlations of cells with overlapped place fields, provides a measure of the clustering of correlations of each cell with its nearest neighbors in the environment.

As shown in Figure 7, whereas the control condition (#1) did not lead to any bias across environments, both the associative and nonassociative manipulations led to very large biases toward coherently reactivating the "most recently visited environment." Indeed, the two mechanisms led to essentially indistinguishable results. The quantitative difference was due to the different bias levels used in these two conditions, i.e., 50% bias in the associative case vs. 10% bias in the nonassociative case.

DISCUSSION

This simple simulation reveals several theoretically interesting results. First, a simple associative learning rule, implemented in a population of cells with the response properties of hippocampal place cells and intrinsic modifiable connections resembling those of field CA3, leads to a situation in which each location is represented by a quasi-stable dynamic attractor involving a small group of cells. Each cell is involved in numerous such attractors for points surrounding its preferred firing location (see also Tsodyks and Sejnowski, 1995). The distribution of attractors for a given environment may be continuous, in the sense that there is very little "energy" (Hopfield, 1982) difference between neighboring locations, and intrinsically two-dimensional, at least for cases in which exploration is confined to two dimensions. Second, numerous such planar configurations can be encoded in such a network with minimal interference. Third, once the long-term

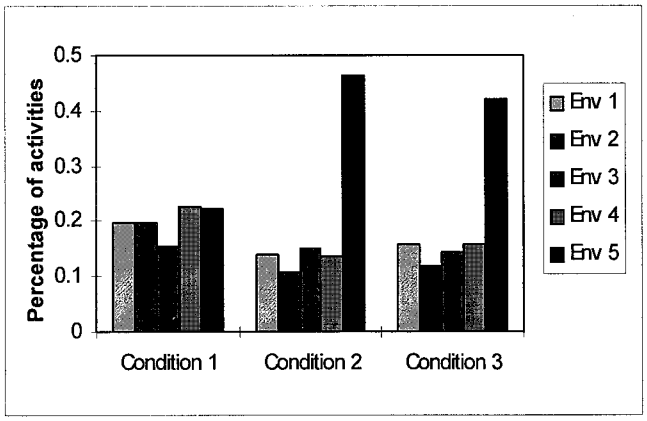

FIGURE 5. Activity distribution among five environments. Five different environments were stored in a network of 100 cells, with 36 place cells designated for each environment. The conditions are 1) no bias, 2) associative bias on cells of environment 5, and (3) nonassociative bias on cells of environment 5.

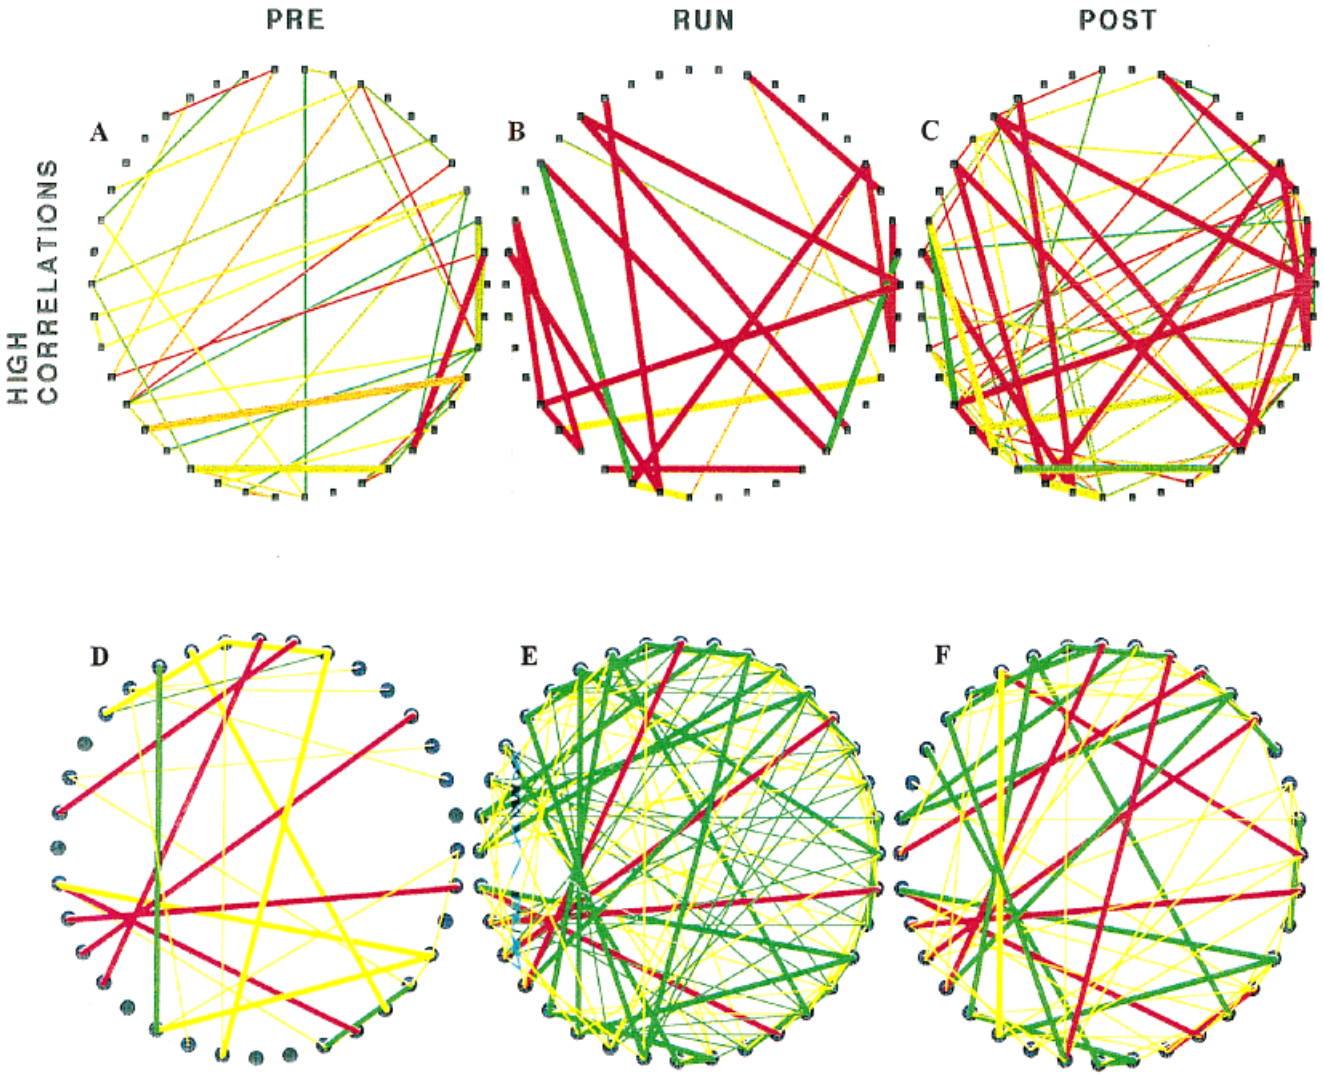

FIGURE 6. Nonassociative bias may result in reactivation of associative memories. Top: Data from parallel recording experiment reproduced from Wilson and McNaughton (1994). Each node represents a cell with a place field in the environment. Strength of correlation of activity between cell pairs is indicated by color (red = high correlation, yellow = weak). The wide lines indicate correlations present during behavior that were also present either during sleep prior to the spatial experience, or during sleep following the experience. Note that almost all correlations seen during behavior (RUN) are seen during subsequent sleep (POST), whereas very few of the same correlations were present during prior sleep (PRE). Bottom: In this simulation, five environments were stored in the network using overlapping sets of 36 of the 100 cells. The nodes correspond to the 36 neurons with place fields in environment 5. The legends are otherwise the same as those of the top figures. D: To simulate the "PRE" condition, the network was run under condition 1 (i.e., the five environments were equally weighted. Only a subset of the correlations seen in environment 5 emerged during "sleep." E: The correlation pattern of cells in environment 5 during uniform sampling of it in the simulated RUN phase (this was simply computed from the place field overlaps). F: When the network was allowed to "sleep" under condition 3 (nonassociative bias on units with fields in environment 5), the pattern of correlations was very similar to the run phase. Thus, in principle, given a pre-existing set of two-dimensional weight distributions, in which the synaptic weights between cells are functions of the overlap of the place fields in multiple environments, no further associative synaptic modification is necessary in order to elicit selective, spontaneous (i.e., with only random inputs) recall of representations of the most recently visited environment. A mere excitability bias is sufficient, at least under some conditions.

weight distributions for several environments have been established, if environmental exploration leads either to a transient elevation of the associative coupling among neighboring place cells, or merely to an excitability bias on the most recently active cells, then completely random inputs lead to the spontaneous recall of coherent representations of locations in the most recently visited environment. This result captures the increase of pairwise cross-correlations between behaviorally coactive place cells during sleep reported by Wilson and McNaughton (1994). It also provides possible insight into the stochastic nature of hippocampal sharp-waves (Buzsaki, 1989), which are the carriers of the memory trace reactivation effect.

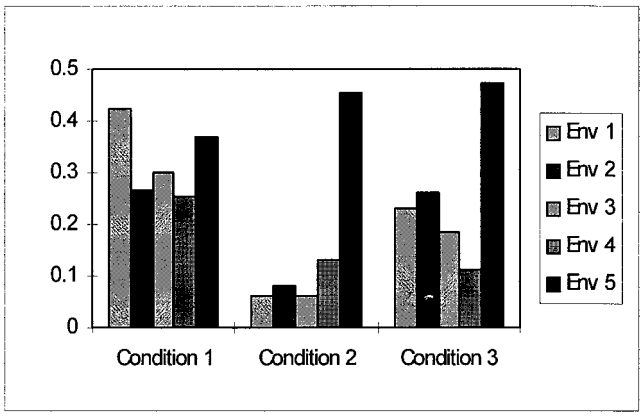

FIGURE 7. Coherent reactivation of environments during "sleep." Five environments are encoded in the network, using overlapping subsets of 36 neurons each, as in Figures 5 and 6. Coherence is defined in the text. This figure shows that, at least under the conditions simulated, either of the two possible mechanisms, associative (condition 2) or nonassociative (condition 3) are about equally able to select out the most recently visited environment. With no bias, (condition 1) the high states that emerge during "sleep" are all coherent (i.e., focused on a single environment), but the distribution of high-activity states over multiple cycles is about equally distributed across environments.

The fact that a nonassociative mechanism could subserve spontaneous memory reactivation in a system exhibiting attractor dynamics (i.e., one in which the appropriate synaptic matrix is already specified) is of considerable interest. For example, it could easily explain such phenomena as the differential priming effects for familiar items that is observed in hippocampally damaged patients (e.g., Shimamura and Squire, 1989). It could also account for some kinds of short-term or working memory (Baddeley, 1986) effects. It is well known that working memory is much superior for items for which internal representations have already been formed. The effect in our simulation occurs because the different environments, and locations within them, are encoded relatively sparsely (Marr, 1971; Tsodyks and Feigelman, 1988; McNaughton, 1989; Amit and Tsodyks, 1991; Gibson and Robinson, 1992). That is, they each involve only a limited fraction of the total population (in our case 36%). Clearly, if different configurations involving all cells were used for each environment, there would be no basis for this nonassociative memory trace. It is intriguing to speculate that the nonassociative bias might well be provided by persistent changes in gK_{Ca} such as have been observed in both hippocampus and neocortex in certain conditions (Disterhoft et al., 1991; Woody et al., 1991).

Acknowledgments

This work was supported by MH46823, NSF, and the McDonnell and Pew Foundations. We are grateful to D. Amit, A. Samsonovich, W.E. Skaggs, and M. Tsodyks for helpful discussions and suggestions.

REFERENCES

Amit DJ, Tsodyks MV (1991) Quantitative study of attractor neural network retrieving at low spike rates: I. Substrate-spikes, rates and neuronal gain. Network 2:259–274.

Baddeley AD (1986) Working memory. Oxford: Oxford University Press.

Buzsaki G (1989) Two-stage model of memory trace formation: a role for "noisy" brain states. Neuroscience 31:551–570.

Disterhoft JF, Black J, Moyer JR, Thompson LT (1991) Calcium-mediated changes in hippocampal neurons and learning. Brain Res Rev 16:196–198.

Fishbein W, Gutwein BM (1977) Paradoxical sleep and memory storage processes. Behav Biol 19:425–464.

Gibson WG, Robinson J (1992) Statistical analysis of the dynamics of a sparse associative memory. Neural Networks 5:645–661.

Hebb DO (1949) The organization of behavior. New York: John Wiley and Sons.

Hill AJ (1978) First occurrence of hippocampal spatial firing in a new environment. Exp Neurol 62:282–297.

Hopfield JJ (1982) Neural networks and physical systems with emergent collective computational abilities. Proc Natl Acad Sci USA 79:2554–2558.

Kim JJ, Fanselow MS (1992) Modality specific retrograde amnesia of fear. Science 256:675–677.

Kubie JL, Ranck JBJ (1983) Sensory-behavioral correlates in individual hippocampus neurons in three situations: space and context. In: Neurobiology of the hippocampus Seifert W, (ed), pp 433–447. New York: Academic Press.

Kudrimoti HS, McNaughton BL, Barnes CA, Skaggs WE (1995) Recent experience strengthens pre-existing correlations between hippocampal neurons during sleep. Society for Neuroscience Abstracts 21:941.

MacKinnon D, Squire LR (1989) Autobiographical memory in amnesia. Psychobiology 17:247–256.

Marr D (1970) A theory of cerebral neocortex. Proc R Soc Lond [Biol] 176:161–234.

Marr D (1971) Simple memory: a theory for archicortex. Philos Trans R Soc [Biol] 262:23–81.

McClelland JL, McNaughton BL, O'Reilly RC (1995) Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. Psychol Rev 102:419–457.

McGaugh JL, Herz MJ (1972) Memory consolidation. San Francisco: Albion.

McNaughton BL (1989) Neuronal mechanisms for spatial computation and information storage. In: Neural connections, mental computations (Nadel L, Cooper LA, Culicover P, Harnish RM, eds), pp. 285–350. Cambridge, MA: The MIT Press.

McNaughton BL, Morris RGM (1987) Hippocampal synaptic enhancement and information storage within a distributed memory system. Trends Neurosci 10:408–415.

McNaughton BL, Barnes CA, Gothard KM, Jung MW, Knierim JJ, Kudrimoti HK, Qin Y-L, Skaggs WE, Gerrard JL, Suster M, Weaver KL (1996) Deciphering the hippocampal polyglot: The hippocampus as a path integration system. Journal of Experimental Biology 199:173–185.

Muller RU, Kubie JL, Saypoff R (1991) The hippocampus as a cognitive graph (abriged version). Hippocampus 1:243–249.

Pavlides 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(8):2907–2918.

Qin Y-L, McNaughton BL, Skaggs WE, Barnes CA (1995) Reactivation during sleep of cortico-cortical and hippocampo-cortical correlation states from preceding behavior. Society for Neuroscience Abstracts 21:941.

Shimamura AP, Squire LR (1989) Impaired priming of new associations in amnesia. J Exp Psychol [Learn Mem Cogn] 15:721–728.

O'Keefe J, Conway D (1978) Hippocampal place units in the freely moving rat: why they fire where they fire. Exp Brain Res 31:573–590.

Skaggs WE, McNaughton BL (1996) Replay of neuronal firing sequences in rat hippocampus during sleep following spatial experience. Science 271:1870–1872.

Squire LR, Cohen NJ, Nadel L (1984) The medial temporal region and memory consolidation: a new hypothesis. In: Memory consolidation (Weingartner EG, eds), pp 185–210. Hillsdale, NJ: Erlbaum.

Squire LR, Shimamura AP, Amaral DG (1989) Memory and the hippocampus. In: Neural models of plasticity: Experimental and theoretical approaches (Byrne JH, Berry WO, eds), pp.208–239. New York: Academic Press.

Tsodyks MV, Feigelman MV (1988) The enhanced storage capacity of neural networks with low activity level. Europhysics Lett 46:101–105.

Tsodyks MV, Sejnowski TJ (1995) Associative memory and hippocampal place cells. Int J Neural Systems 6:81–86.

Wilson MA, McNaughton BL (1993) Dynamics of the hippocampal ensemble code for space. Science 261:1055–1058.

Wilson MA, McNaughton BL (1994) Reactivation of hippocampal ensemble memories during sleep. Science 267:676–679.

Woody CD, Gruen E, Birt D (1991) Changes in membrane currents during Pavlovian conditioning of single cortical neurons. Brain Res 539:76–84.