Hippocampal synaptic enhancement and information storage within a distributed memory system

B. L. McNaughton and R. G. M. Morris

B. L. McNaughton is at the Department of Psychology, University of Colorado, Boulder, CO 80309, USA, and R. G. M. Morris is at the Laboratory for Cognitive Neuroscience, Department of Pharmacology, University of Edinburgh, Edinburgh EH8 9JZ, UK.

The hypothesis that the physical substrate of memory in the mammalian brain resides in alterations of synaptic efficacy has been proposed frequently in both neuroscience[1–5] and cognitive science[6–12], and has been widely investigated in behavioural, physiological and theoretical studies. Although this hypothesis remains unproven, considerable evidence suggests that a particular form of synaptic strengthening, induced by electrical stimulation of certain CNS fibre systems, may represent the activation of mechanisms that normally subserve associative memory. This phenomenon is known as long-term potentiation (LTP) or long-term enhancement (LTE). It has been most intensively investigated within the hippocampal formation, a brain structure that plays a crucial role in certain forms of associative memory. Physiological investigation has revealed that LTE exhibits most of the properties implicit in Hebb's original suggestion that associative memory results from a synaptic strengthening that is contingent upon the conjunction of activity in pre- and post-synaptic elements. In this article, we outline a simple neuronal model capable of superimposing multiple memory traces within the same matrix of connections, and consider the correspondence between such models and the properties of LTE in the context of the hippocampal circuitry in which it occurs. Certain predictions are derived from this framework concerning the behavioural consequences of experimental manipulation of LTE, and we conclude by describing experimental evidence that confirms these predictions and suggests that LTE is, in fact, fundamentally involved in memory.*

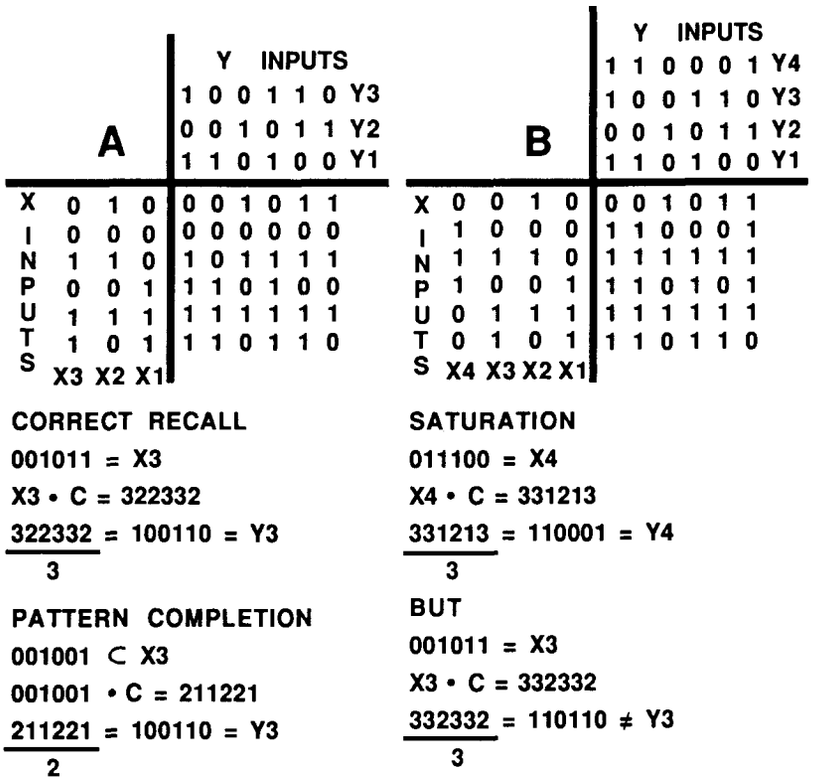

Fig. 1. *The fundamental principles of distributed associative memory can be easily illustrated by considering storage and recall of event pairs (X1:Y1, etc.) using a mechanism known as a correlation matrix. In this example, two sets of six binary information channels converge at a matrix C of 36 nodes. All nodes have an initial value of zero, but are permanently changed to a value of one whenever the corresponding X and Y input lines have the value one simultaneously. The matrix in* **(A)** *has experienced three sets of paired events. Any member of a pair can be recalled by multiplying the matrix by the other member, and dividing by the number of ones in the multiplier (decimals are ignored). Moreover, complete recall can also be accomplished by presenting only part of the corresponding pattern, provided that that part is unique (i.e. is not also a part of a different event). All three patterns stored in A can be correctly recalled, in spite of the fact that they overlap each other to some degree. Although the number of patterns such a matrix can store increases dramatically with the size of the matrix, there is always a limit above which recall of both new and older events become inaccurate. This is illustrated in* **(B)**, *where the storage of a fourth pair, which can itself be recalled correctly, has interfered with the recall of event 3. In general, this limit depends on the magnitudes and relative overlap of the input vectors. This saturation effect forms the basis for an experimental test of the relation between LTE and memory described in the text.*

**'Long term enhancement' (LTE) is used here in preference to the more commonly used 'long-term potentiation' (LTP) because the word 'potentiation' has a long prior history of specific application to a non-associative and short-lasting increase in the probability of transmitter release following tetanic stimulation[76]. Potentiation can be observed at a very large variety of both central and peripheral synapses (including those in the hippocampus), and has been shown to differ mechanistically from LTE[47]. While LTP and LTE refer to the same phenomenon, the indiscriminate use of the word 'potentiation' has, in our view, led to considerable confusion in the literature.*

To introduce the concept of synaptic associative memory, we begin by outlining the properties of a correlation matrix[6,13,14], a formalism that has been rediscovered several times over the last 30 years. Although not a neural model *per se*, it illustrates the fundamental principle that many memory traces can coexist as an overlaid distribution of connection strengths within a single network.

Consider a matrix *C* of six horizontal channels interacting with six vertical channels (Fig. 1A) at 36 nodes. Activity on a channel is represented by the value one, inactivity by zero. An input event (e.g. X1:Y1) is a pattern of activity on the set of input channels. The state space of the nodes is also binary and 'Hebbian' (i.e. if and only if both inputs to a node are simultaneously in a '1' state, the node undergoes an irreversible transi-

tion from 0 to 1). Such a matrix can store associations between pairs of binary stimulus events.

Recall of an input event is accomplished by multiplying the matrix $C^{\dagger}$, representing the connection strengths, by the corresponding paired associate, and performing an integer division (i.e. division with decimal truncation) on the result. The divisor is the number of ones in the cuing pattern. This division is analogous to setting a variable threshold and is essential for accurate recall when patterns that are not orthogonal (i.e. those sharing common active elements) are stored in the matrix.

Perfect recall can be achieved by this means provided not too many different patterns have been presented (Fig. 1A). Errors in recall will begin to occur as the matrix approaches saturation (Fig. 1B).

†Matrix multiplication, as used here, involves multiplying each row element of the matrix by the corresponding element of the input vector, and then summing the columns.

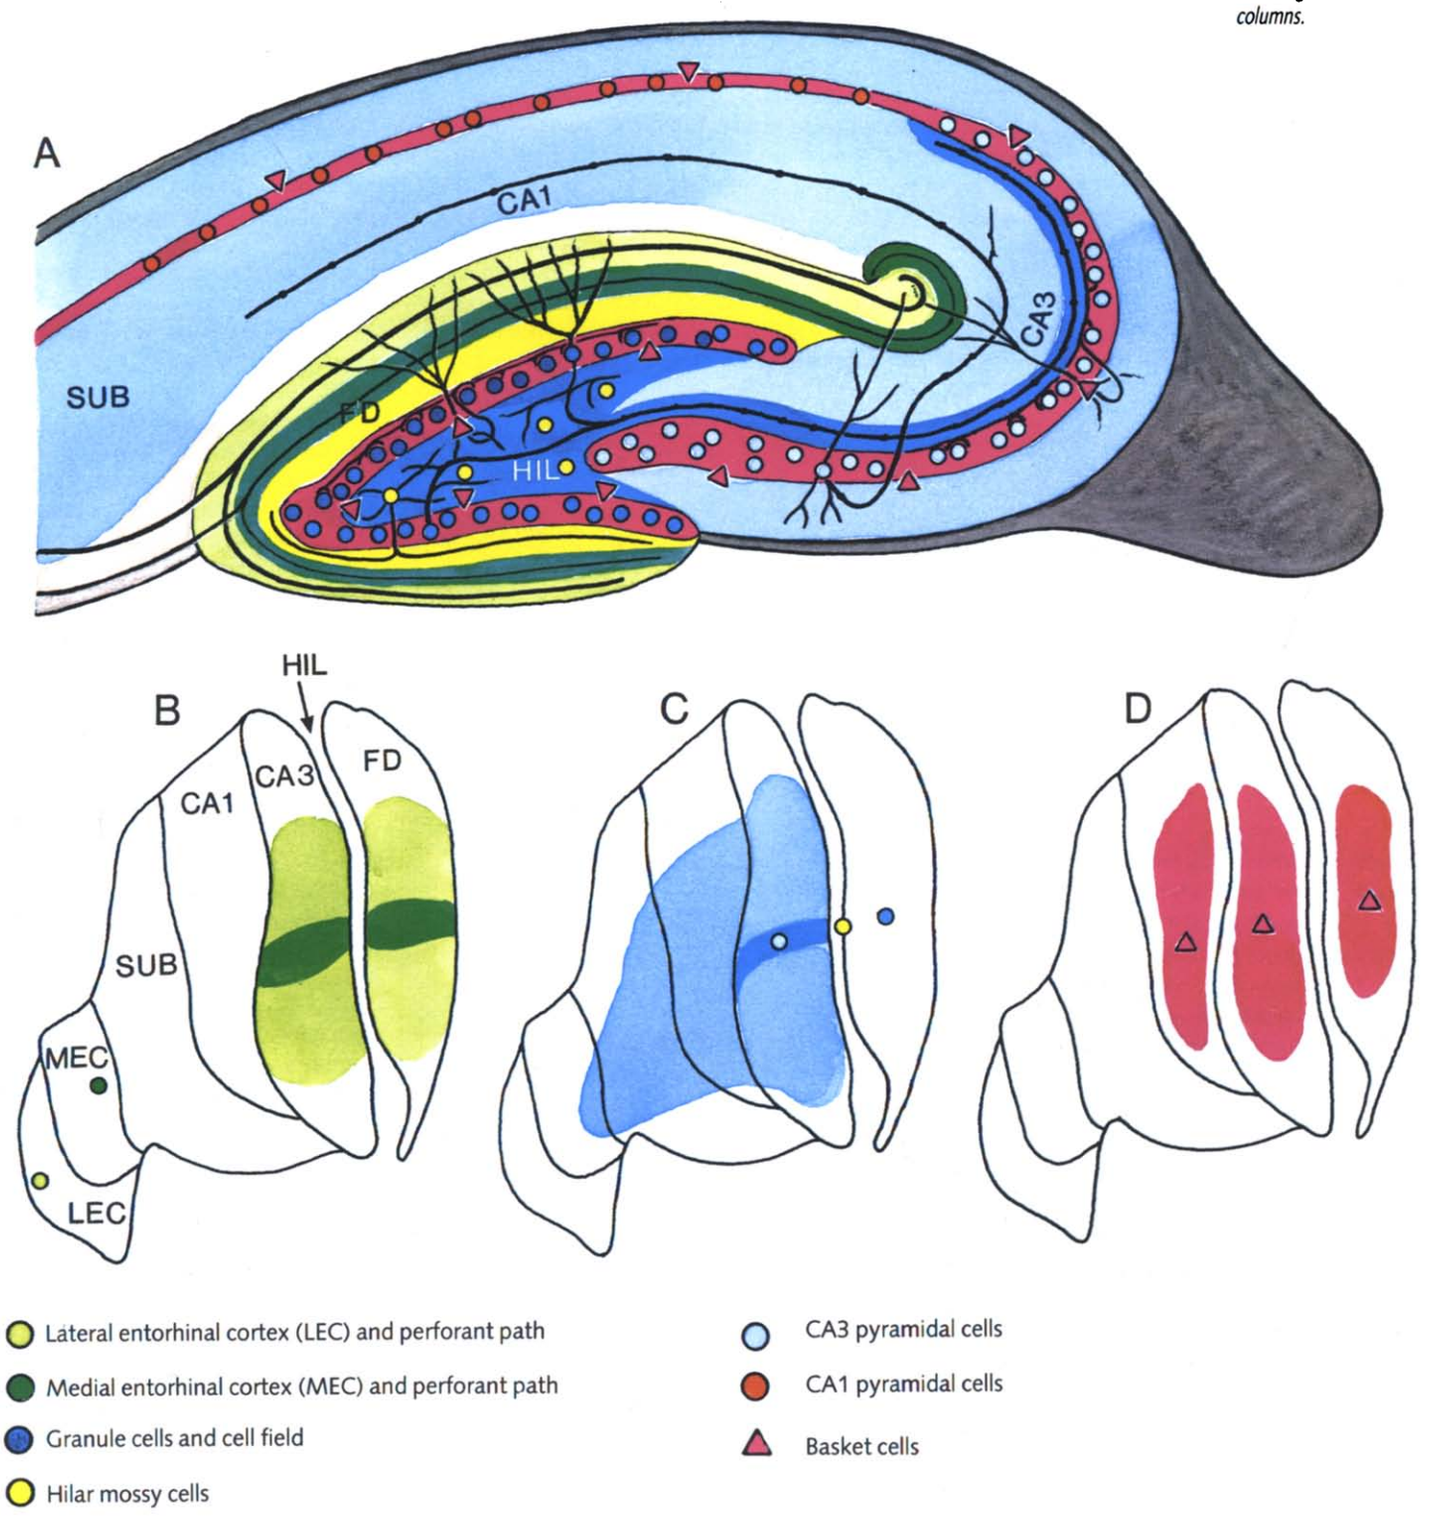

Fig. 2. *The major aspects of hippocampal connection topography are illustrated in a transverse section* **(A)** *and in unfolded two-dimensional representations* **(B, C, D)** *devised by Swanson* et al. *The circuitry of the hippocampal formation appears to manifest several of the formal properties of the simple nets considered in this article. These include the appropriate feedforward inhibitory (basket) cells, systems of relatively strong, anatomically specific inputs that converge with weak (but modifiable), anatomically diffuse inputs, and systems of modifiable excitatory recurrent collaterals. Moreover, long-term enhancement of synaptic transmission at many hippocampal synapses occurs by a mechanism possessing the necessary formal associative properties first proposed by Hebb.*

Conceptual neural implementation

How might the correlation matrix formalism relate to an actual nervous system. One way, first outlined by Marr[4,15], is illustrated in Fig. 3A. This network contains six principal neurons, one inhibitory interneuron and two types of inputs. The Y inputs each make a single powerful synapse with a principal neuron. This synapse is regarded as a detonator[16] in that it always produces intense depolarization and firing of the post-synaptic neuron. The X inputs each contact all principal neurons. These synapses have an initial strength of zero, but become enhanced to a strength of one according to Hebb's conjunction rule. The X inputs also make fixed excitatory connections with the inhibitory interneuron, which thus feeds forward an inhibitory signal proportional to the total number of active elements in the input pattern‡. The principal neurons have a resting threshold of one. The inhibitory signal performs the division operation. The system depicted in Fig. 3A has stored the same set of paired associates as that of Fig. 1A, and these can be recalled with the same accuracy.

‡*The inhibition can be made proportional to the excitatory input density in several ways. Firstly, the latency to discharge of the inhibitory cell will decrease as the excitatory input to it increases. This will lead to a proportional reduction in the latency of the resulting IPSP on the principal cells. In addition, both multiple interneurons, and a variable number of spikes per interneuron would contribute to the necessary proportionality.*

The learning performed by this network is heteroassociative in that the synaptic modification process in one pathway is governed by the activity of certain other fibres whose connection strengths are not themselves modified. One example of such learning is classical conditioning in vertebrates. Here, the Y inputs correspond to a representation of the unconditioned stimulus (US), the X inputs to the conditioned stimulus (CS), and the output pattern of the principal neurons to the memory representation of the US which comes to be evoked by activation of the CS[17]. Although this simple model fails to account for several important conditioning phenomena (e.g. blocking), it does explain how a CS comes to evoke a new response. This association of this internal representation of events is the defining characteristic of most higher forms of learning as well.

Heteroassociation is a special case of a more general paradigm known as autoassociation. This is required in situations such as spatial learning, where it may be necessary to recall a complete representation of a learned input from any partial subset of the original. Unlike the heteroassociative model, there are no explicit instructing inputs. Rather, elements of a neuronal population that have been active together become coupled in such a way that subsequent activation of a few members of the original set results in completion of the whole. This is the essence of Hebb's notion of a cell assembly. As an example, consider our ability, upon entering a familiar room in near darkness, to locate objects accurately on the basis of the few stimuli available. Autoassociation may be realized (Fig. 3B) by replacing the Y input to the network of Fig. 3A with specific detonator contacts from the X inputs. Each X fibre makes only one such contact, and each principal neuron receives only one (more complex examples might require conjunction of several detonators). The inhibitory division operation permits each of the stored input patterns to be correctly recalled even if the network is presented with only part of the original pattern. For example, pattern X2 may be completed by presentation of either 100000 or 001010. In general, this is true for any subset that is not also a subset of one of the other stored patterns. The performance of such networks improves dramatically with size.

A variant of this network also makes use of a single non-modifiable detonator synapse to drive each principal neuron externally, but in this case, the output pattern is fed back into the dendrites of the principal neurons via modifiable synapses (Fig. 3C).

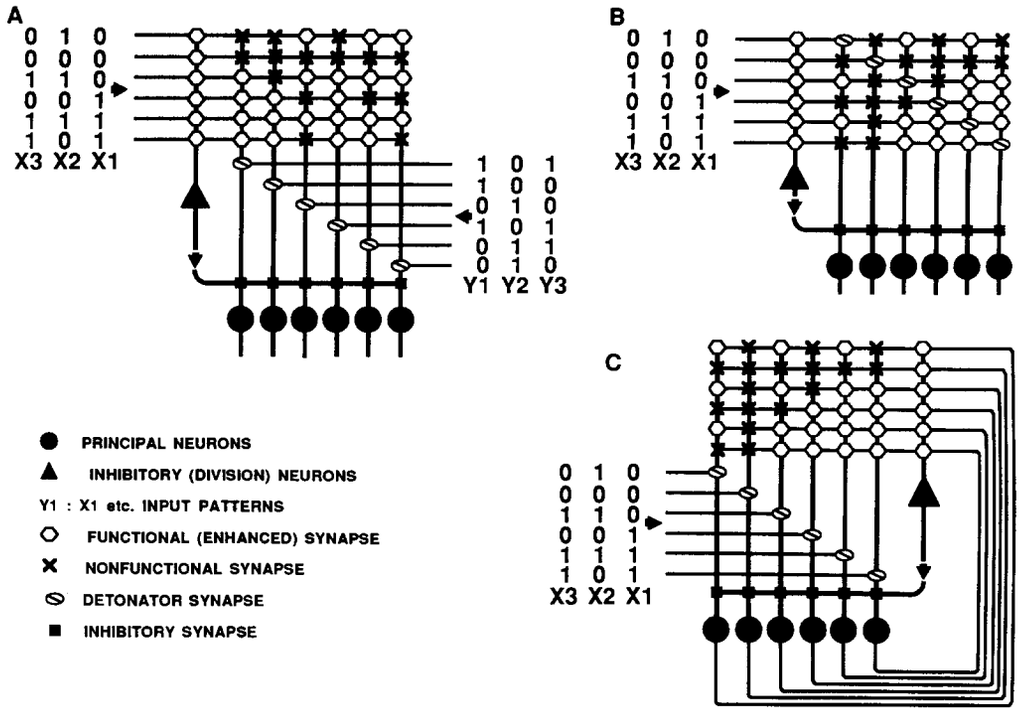

Fig. 3. (A) *A simple 'neuronal' implementation of the heteroassociative correlation matrix of Fig. 1A. The Y input fibres have exclusive control of the output cells by virtue of their terminating in 'detonator' synapses. The detonators also initiate 'Hebbian' synaptic enhancement of those terminals from the non-specific X input that were co-active with the particular Y input. Accurate recall is enabled by the inhibitory interneuron which is presumed to feed forward an inhibitory signal that divides the excitation of the principal neurons by the number of X fibres active in the current event. In real nervous systems, this division may be approximated by inhibitory synapses whose equilibrium potentials lie near the postsynaptic resting potential, and hence shunt the synaptic current arriving at the soma via the dendrites. While a single detonator input has been used here for simplicity, the idea should not be taken too literally. The same general principles apply when convergence of several or many inputs are required to initiate postsynaptic activation, with the exception that the output pattern will be some transform of the input, rather than identical to it.* **(B)** *A simple linear autoassociative network which has the ability to complete a stored representation when presented with only part of the original. This network is equivalent to making the X and Y inputs of Fig. 2A the same, and might be implemented simply by having each input fibre make a small number of detonator contacts, and a large number of weak, but modifiable, ones. Evidence for such a configuration has been detected within the fascia dentata.* **(C)** *A* recurrent *autoassociative network which associates the input to the principal cells with their own output. Depending on the details of its implementation this net can act as a pattern completion device with the additional capacity to keep patterns active by reverberation. If this reverberatory activity is not explicitly silenced before presentation of new information, this network will store event sequences and recall them from a fragment of the initial event in a sequence.*

This re-entrant organization was referred to as a 'collateral effect' by Marr[4]. Its properties differ from the previous network in two important ways. First, the output pattern of the re-entrant network persists after the input is terminated (unless it is explicitly turned off). Hebb first drew attention to such reverberatory activity as a possible mechanism for short-term storage of patterns. Second, for autoassociation, input to the re-entrant network must be held constant long enough during learning for the feedback to coincide with the input, and activity from the preceding cycle must be terminated before the arrival of new information. An interesting capacity to learn sequences arises when these conditions do not hold. In this case, the network associates its current input with the output of the preceding cycle. If trained with the sequence X_1, X_2, X_3, and then presented with the initial event X_1 only, the complete sequence is recalled.

The degradation of stored information through saturation of modifiable connections is a serious problem confronted by all distributed memory systems. As the number of stored events increases, so does the probability of spurious recall. We make use of predictions based on this saturation effect in the experiments described below. There are several ways of dealing with this problem, including efficient input coding (i.e. keeping the number of elements used to code a representation to a minimum), expansion of the input pathway onto a memory network containing a larger number of units, in such a way as to reduce the relative overlap of input patterns (e.g. Marr's codon hypothesis[4]) and making the memory impermanent (either by spontaneous decay or by competitive interactions). There may be particularly sound reasons for doing the latter. For example, once experience has established which elements of a particular environment are stable, more efficient representations of the stable elements might be stored elsewhere in a permanent memory while the large amounts of information stored initially can be dispensed with.

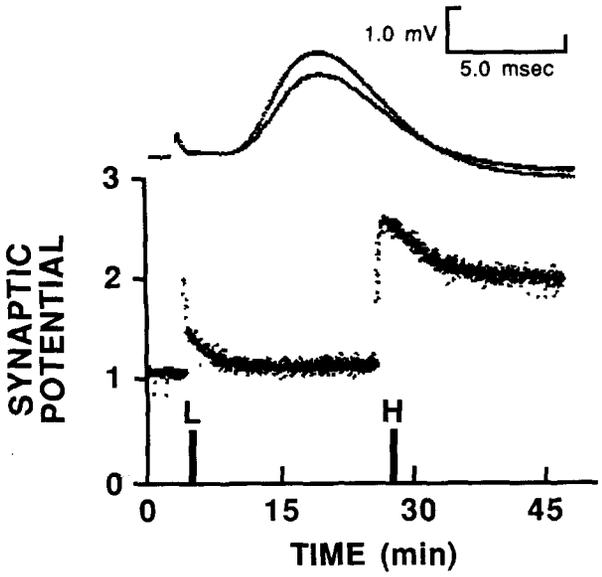

Fig. 4. *Data from an experiment illustrating non-associative potentiation, and associative LTE at terminals of the lateral perforant path in the fascia dentata. Data points represent the amplitude of electrically evoked synaptic field potential responses (normalized) to weak, fixed intensity stimuli delivered once every three seconds. At L, a stimulus train of 100 stimuli at 200 Hz was delivered at the same stimulus intensity (data points not shown). This resulted in transient potentiation of synaptic efficacy. At H, a second stimulus train was given which was identical to the first except that, during the train, the stimulus intensity was raised to co-activate many more convergent afferents. Long-term enhancement was superimposed on the non-associative potentiation. Similar studies have revealed cooperation (i.e. association) between different pathways, such as the medial and lateral, or the crossed and uncrossed, perforant pathways to the dentate gyrus, or the apical and basilar dendritic inputs to CA1 pyramids. The traces at the top of the figure illustrate the extracellularly recorded field EPSP evoked by test stimuli before, and two hours after induction of LTE.*

Instantiation in hippocampal circuitry?

Although the dynamics of activity within hippocampal circuitry are incompletely understood, several aspects of its anatomy and physiology correspond surprisingly well with the conceptual implementations just described (Fig. 2).

Cortical inputs projecting via the perforant path contact a considerably expanded number of cells (the granule cells of the denate gyrus) within transverse strips (lamellae)[20] along the longitudinal axis. Within this projection there are some clear analogies to the basic matrix arrangements of Fig. 3. The perforant path consists of two distinct fibre systems[21–24]. The medial pathway resembles the Y inputs of Fig. 3A in that it terminates close to the granule cell bodies, has synapses that are initially more powerful than the lateral pathway, and terminates topographically in a restricted transverse strip. The lateral pathway has synapses that are initially weak, and terminates in a topographically diffuse way – hence it resembles the X inputs. However, embedded in this system may also be a structure analogous to the autoassociative system of Fig. 3B. Most medial path synapses generate EPSPs of about 0.1 mV amplitude. However, about 3% of these terminals generate EPSPs 10–20 times larger[25]. These could act as the covert detonator synapses of the model.

The granule cells of the dentate gyrus project in a highly specific manner onto the pyramidal cells of CA3[26]. Each granule cell makes few but relatively powerful contacts close to the pyramidal cell bodies. The CA3 pyramidal cells project heavily back into their own dendrites via the 'longitudinal association' pathway. This configuration is clearly analogous to the re-entrant network of Fig. 3C. In addition, the same perforant path fibers that project to the dentate gyrus also terminate distally on the CA3 pyramidal cell dendrites. Thus, the granule cell output to CA3 represents a transform of the perforant path input to CA3. This arrangement thus resembles a heteroassociative configuration that might, in practice, be carrying out autoassociation by virtue of the detonator information being a transform of the information arriving distally.

We have emphasized above the importance of a division operation for accurate recall. Inhibitory interneurons, with what appear to be the necessary characteristics, are found throughout the hippocampus. For example, the basket cells of the dentate gyrus receive excitatory input from the afferents to the molecular layer, and feed inhibition forward onto the granule cells[27–29]. Similar configurations are found in CA1 and CA3, where the excitatory afferents terminate both on

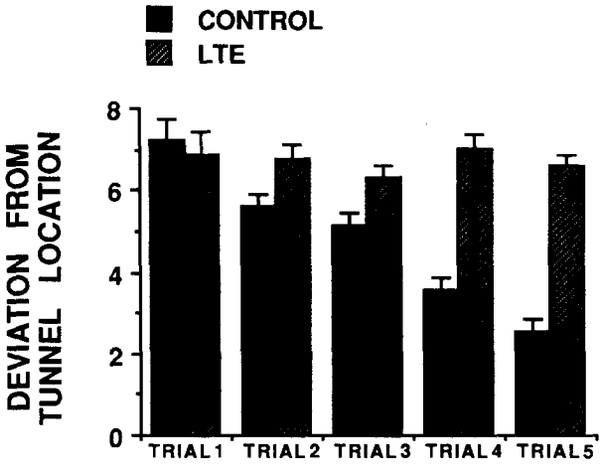

Fig. 5. *The hypothesis that the LTE mechanism underlies spatial information storage within the hippocampus predicts that bilateral saturation of this mechanism should produce a severe deficit in the acquisition of a spatial problem. This prediction is derived from the concept of matrix saturation illustrated in Fig. 1B. The prediction was confirmed using the Barnes circle maze task, in which animals learn to escape from bright light into a dark tunnel located under one of 18 holes equally spaced around the periphery of a circular platform. Once a day, the animals were released in a random orientation at the platform centre. The error score shown here is the distance of the hole that the animal first investigated from the correct one (in number of holes). The group receiving LTE induction had a severe impairment that was specific to learning of new locations and memory for recently learned locations. Spatial memory that was well established before the treatment was left intact, implying that the stored information in the hippocampus is ultimately transferred elsewhere.*

the pyramidal cells and on inhibitory interneurons[30]. The inhibitory interneurons are far fewer in number than the principal neurons, yet have an extensive, diffuse, axonal trajectory[31]. Physiological investigations have shown that the equilibrium potential for the inhibitory neurotransmitter GABA is close to the resting potential of the principal neurons[32]. Postsynaptic inhibition is thus primarily a current shunt. As a result of these factors, inhibition has the approximate effect of dividing the excitation by a variable related to the number of active afferents. The excitability properties and spontaneous activity of basket inhibitory cells is also consistent with such a role. These cells respond to stimulation of afferent fibers more quickly, and at a lower threshold, than the principal neurons[33,34]. This is necessary to insure that the division operation is already in effect before the principal cells begin to fire. Unlike the principal cells, the response of interneurons is graded in terms of the number of spikes elicited in response to a single input event. Single cell recording from freely moving animals has shown that putative basket cells (i.e. theta cells) show phasic activity, and convey little, if any, specific information[35]. This is in sharp contrast to recordings from pyramidal cells, which reveal a remarkable degree of spatial selectivity[36]. Thus, it appears that the basket cells are capable of signalling how many, but not which afferents are active.

There is another possible role for inhibition. Except for systems storing event sequences, effective storage and recall would appear to require that the input patterns be presented to quiescent networks. Otherwise, interference with preceding activity will occur. The whole hippocampal formation undergoes a cycle of excitation and inhibition of about 7Hz (the so called theta rhythm) while the animal is engaged in active movement[37] through its environment. Most hippocampal units are phase-locked[38] to this cycle. One function of the theta rhythm may be to clear the system before the arrival of new input information. The duration of one theta cycle is also long enough, in principle, for several cycles of the collateral effect to generate a progressively more accurate recall of stored patterns when the input is incomplete.

There are, of course, a multitude of other intrinsic and extrinsic connections to consider, and the evidence for the proposals just made is far from complete. Nevertheless, the hippocampal network seems to be organized in a manner compatible with the operation of a distributed, correlation-matrix memory.

Long-term enhancement of hippocampal connections

Occurrence

LTE is a persistent increase in the efficacy of synaptic transmission which results from high-frequency stimulation of certain CNS fiber systems[39,40]. It may be induced by brief episodes of electrical stimulation of the principal afferents to the hippocampus (perforant path–granule cell synapses; Fig. 4), and many, but not all of the intrinsic excitatory connections within and between the principal hippocampal subfields. Several other intrinsic and extrinsic connections have failed to exhibit LTE when tested by comparable techniques, although several of these inputs exhibit a modulatory role on LTE of the principal inputs. Persistent changes, similar to LTE in magnitude and timecourse, have been induced in electrically elicited responses in other cortical regions[41,42] (most prominently in other areas of limbic cortex).

Mechanisms of induction, expression and decay

In most cases, a few impulses delivered at intervals approaching the axonal refractory period have been most effective in eliciting LTE[43]. However, under conditions in which the postsynaptic neuron is artificially depolarized prior to the input, single stimuli are sufficient[44,45]. In addition to high frequency, it is clear that under normal circumstances the induction of LTE requires the convergence of a critical number of active afferents. In the dentate gyrus, both medial and lateral perforant paths can exhibit LTE. However, when activated in conjunction, the two pathways cooperate to generate considerably more LTE than either pathway alone. Indeed, at certain input strengths the association of the two inputs is required for any LTE[46]. Similar cooperativity has been demonstrated among other excitatory fibers within the hippocampus. While the 'threshold' for LTE in fascia dentata is apparently near the postsynaptic discharge threshold, spike discharge of the postsynaptic neuron is neither necessary nor sufficient for its induction[46]. Nevertheless, Hebb's principle is obeyed at these synapses in a statistical sense, if not in a physical one – the crucial regulatory factor is the depolarization of the dendritic tree to somewhere above the level normally required for cell discharge.

One is naturally led to ask why high frequency is

normally of such importance. There would appear to be at least two possible explanations. Firstly, most hippocampal excitatory synapses are both extremely weak, and rather unreliable. In the dentate gyrus, for example, the average EPSP generated by a perforant path synapse is a tenth of a millivolt, and about 400 convergent fibres are needed to discharge the granule cell[25]. Moreover, when activated at low frequency, a single terminal may fail to release any transmitter about 30% of the time. High frequency leads both to spatio-temporal summation, and to a dramatic increase in the probability of transmitter release, which may double or triple the effective depolarization induced by a group of convergent afferents[47]. Secondly, it is possible that some of the processes regulating the development of LTE may have somewhat long activation times, and may thus require that the depolarization be extended longer than the timecourse of a single EPSP. This might be of particular use in recurrent systems where the network output must be correlated with the appropriate input.

The locus of control over the induction of LTE resides with the postsynaptic neuron. This was originally inferred from the cooperativity effect. Subsequently, injections of the calcium chelator EGTA into single hippocampal neurons were found to block LTE in the impaled neuron without affecting either evoked synaptic transmission itself or LTE in the rest of the population[48]. More recently, it has been demonstrated that specific blockade of the NMDA-preferring glutamate receptor in the hippocampus blocks LTE[49,50]. The NMDA receptor complex is unusual because it is both chemical and voltage dependent[51,52], requiring prior depolarization of the postsynaptic cell in order for bound agonists to initiate a conductance increase to calcium ions. This appears to account for cooperativity, the input specificity of LTE, and the reduced requirement for high-frequency input to cells that are depolarized. The properties of the NMDA receptor may also account for the apparent modulatory effects of other inputs. For example, stimulation of the medial septum facilitates LTE of the perforant pathway[53]. On the other hand, high-frequency excitation of commissural afferents to the denate gyrus appears to reduce or prevent LTE of the perforant path[54].

While considerable progress has been made in determining the factors regulating the induction of LTE, the physical mechanism by which the increased synaptic efficacy is expressed is far from clear. Evidence has been marshalled for the following possibilities: formation of extra postsynaptic glutamate receptors[55], swelling of dendritic spines and reduction of their neck length[56,57], formation of new synaptic contacts[58], and increase in the amount of transmitter released[59]. In none of these cases has cause been adequately dissociated from secondary effects. A few possible mechanisms have been ruled out, including: changes in the postsynaptic membrane potential[60], increased probability of transmitter release such as underlies post-tetanic potentiation[47§], reduced inhibition[61], and reduced threshold in the afferent fibres themselves[60]. Also, there are two important constraints on any possible mechanism imposed by the pre- and post-synaptic specificity of LTE. The matrix models require that only those synapses on a given postsynaptic neuron that were coactive, and only those terminals of a given afferent that were involved in a convergence event of sufficient magnitude, should undergo LTE. These requirements have been born out by experimental observations[62–64].

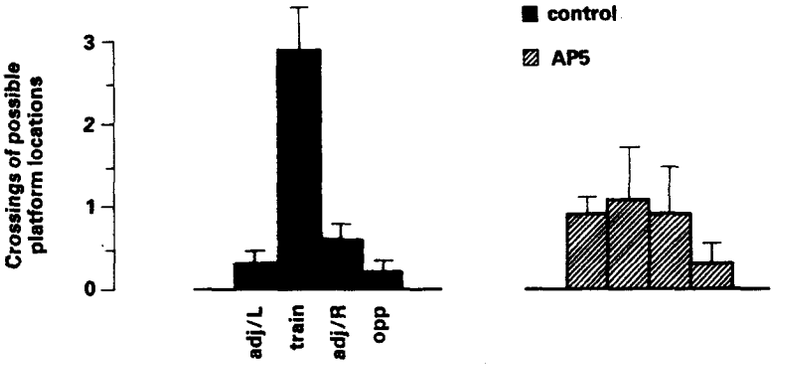

Fig. 6. *Intraventricular infusion of the drug AP5, which blocks LTE, caused a severe learning impairment for spatial information. Rats were trained to escape onto a hidden platform in a circular water-maze for 15 trials over 5 days. In a retention test conducted one day later, the platform was first removed. Control rats swam across the former location of the platform (train) far more often than other locations in adjacent quadrants of the pool to the left or right (adj/L, adj/R), or in the opposite quadrant (opp). Rats given AP5 showed little spatial bias. Data based on a replication of Ref. 73. Control rats N=22, AP5 rats N=14. The blockade of LTE affects learning of new locations, but not memory for old ones.*

The timecourse of the change in synaptic responsiveness following high-frequency activation is complex. After induction, LTE takes approximately 30 s to develop its full magnitude[65]. During this growth time, potentiation, a presynaptic increase in the probability of transmitter release (see footnote*, p. 408), is present but decays over about 5 min (Ref. 47). After reaching its peak, LTE also begins to decay. While LTE is a remarkably persistent phenomenon, whenever it has been studied in chronically prepared animals over adequate periods, it has been found to decay with a half-life of not more than several weeks. Although daily repetition does extend the half-life in much the same manner as memories are stabilized by repetition[66], hippocampal LTE may not be a mechanism for permanent memory. This decay may reflect the necessity to avoid saturation of the hippocampal network with information as discussed above. However, it is worth noting that in the visual cortex of kittens, irreversible changes in ocular dominance are controlled by a mechanism that bears a strong pharmacological and conceptual resemblance to LTE in the hippocampus[67].

The factors governing the decay of LTE have not been studied in detail. There are two possibilities, which are not necessarily exclusive. The commonly accepted view is that decay is spontaneous. The alternative is that it reflects a competitive process. While there is some evidence for heterosynaptic depression of this sort in the dentate gyrus[62,68], its properties are not yet well understood.

Behavioural evidence for a role of LTE in memory.

Evidence from both single unit recording and lesion studies[69,70] in rats indicates that the hippocampus plays a major role in spatial memory. The first indications that hippocampal LTE might underlie this type of memory came from the demonstration that the persistence of LTE was significantly correlated with performance of a spatial memory task[71]. More recently, experiments have addressed the behavioural effects of

§'Probability' is used here in the specific context of the quantal release equation m = n × p. Changes in the parameter 'n' have not been ruled out, and may underlie the increase in transmitter release observed by Bliss and his colleagues[59].

interfering with LTE. There are at least three ways in which such interference could be used to compromise the function of a distributed memory system. One way would be to saturate the network artificially, by increasing the efficacy of all (or most) of its synapses to near their maximum possible strength by the electrical induction of LTE. This should result in both retrograde and anterograde amnesia. The second way would be to block enhancement pharmacologically, in a manner that left baseline transmission unaffected. This should produce anterograde amnesia only. Finally, natural LTE could, in principle, be reversed. This should produce only retrograde amnesia. Our studies have addressed the first two of these possibilities. As yet, the technical means for selective reversal of LTE are not available.

Saturation

The first prediction has been tested[72] by examining the effects of high-frequency, LTE-inducing stimulation of the perforant path upon both recall and new learning in the circular platform spatial 'reference' memory paradigm devised by Barnes[71]. In this test, rats learn to escape from bright light into a tunnel accessible from only one point on the maze periphery, and invisible from the maze center.

In one experiment, rats with chronically implanted electrodes were first trained to their asymptotic performance on the task with the tunnel fixed in one location relative to the extramaze cues. Half of the rats were then subjected to bilateral induction and saturation of LTE, the remainder receiving low-frequency stimulation to the same pathway. The following day the animals were returned to the circular platform, but were now required to find the goal tunnel at a new location, 135° away from its original location. The LTE group failed completely to learn the new position of the goal tunnel over 5 days of training, whereas the low-frequency control rats learned rapidly (Fig. 5). However, on the first relocation trial, both groups were equally good at remembering the original tunnel location. High-frequency stimulation had thus produced anterograde, but not retrograde amnesia. In a second experiment, LTE-inducing stimulation delivered several minutes after the first relocation trial blocked the memory for the new tunnel location, demonstrating a limited retrograde amnesia. Finally, in a third experiment, rats trained on the general task demands in one room, and then subjected to LTE, were impaired on the acquisition of the tunnel location when the maze was moved to a different room.

These experiments show that driving the synaptic population to its maximal strength disrupts recently stored information and prevents new spatial learning[¶], as predicted by the hypothesis, but that hippocampal LTE may only be involved in the temporary storage of information destined for some other, more permanent memory system. The caveat, of course, is that high-frequency stimulation may have other effects apart from the induction of LTE.

Blockade

LTE can be blocked in rats implanted with an intraventricular cannula attached to an osmotic mini-pump containing the NMDA receptor antagonist AP5. Animals prepared in this way were trained in a spatial learning task in which they had to swim in a large pool of opaque water to escape onto a submerged platform located at a particular place. Over the course of five days of training, the group receiving AP5 took significantly longer and more circuitous routes to the platform than controls[73]. Following training, the platform was removed and the animals were required to swim for 60 s without the possibility of escape. While the controls swam directly to the former position of the platform and searched persistently in its vicinity, the AP5 animals showed little or no evidence of spatial memory (Fig. 6). In a second study, we confirmed that the same AP5 treatment was sufficient to block LTE without affecting normal evoked synaptic responses. Moreover, a similar experiment, involving acute microinfusions of AP5, demonstrated both that the spatial memory impairment is reversible, and that information acquired prior to the AP5 treatment is not disrupted.

Although the same caveat concerning possible secondary effects of the amnesic treatment is applicable to both sets of experiments, the demonstration that two radically different methods of interfering with LTE lead to the predicted learning and memory impairments shifts the burden of proof in favour of the hypothesis that LTE reflects a fundamental storage mechanism.

Concluding remarks

We hope to have illustrated in this article that full appreciation of the functional significance of LTE can only be achieved when its properties are considered in relation to both the formal structure of the hippocampal network in which it occurs, and to the type of information processed by this system. A number of important questions remain. For example, although there have been encouraging suggestions[74,75], it is still not known whether LTE actually occurs during the course of exposure to novel information. The likelihood of ever detecting such changes will, of course, depend on such parameters as how much new information is presented, how many synapses are involved in the storage, and whether there is active competition for synaptic efficacy, as alluded to above. These questions are unanswered. At present, there are no data on whether LTE is distributed within the population of afferent terminals as moderate changes in a large number of connections, or large changes at a few terminals (as assumed in the present models). Also, the detailed routes by which information (presumably) stored in the hippocampal matrix returns to the neocortex, and the use to which this information is put (e.g. consolidation), are largely unknown. Much more theoretical work needs to be done on the properties of sparsely connected networks, and the effects of topographically restricted connectivity, both of which are characteristics of the hippocampus. These sorts of uncertainty obviously leave a serious gap in our understanding of the exact nature of the role of LTE in memory. One prediction is clear. Further understanding of how the physical dynamics of this network contribute to a solution of the computational problems of learning and memory will require a considerably closer collaboration between the behavioural, cognitive, and neural sciences.

[¶]It is somewhat of a paradox that nictitating membrane conditioning is improved by prior saturation of LTE in the hippocampus[18]. However, information storage for this type of learning lies outside the hippocampus[19], and blockade of contextual conditioning by prior saturation of LTE could, perhaps, be responsible for the faster learning in the delayed conditioning paradigm used by Berger.

Selected references

1 Tanzi, E. (1893) *Rivista sperimental di Freniatria e Medicina Legale delle Alienazioni Mentali Societa Italiana de Psichiatria* 19, 419–472
2 Hebb, D. O. (1949) *The Organization of Behavior*, John Wiley & Sons
3 Brindley, G. S. (1969) *Proc. R. Soc. London Ser. B.* 184, 173–191
4 Marr, D. (1971) *Proc. R. Soc. London Ser. B.* 262, 23–81
5 Eccles, J. C. (1977) *Brain Res.* 127, 327–352
6 Kohonen, T. (1978) *Associative Memory*, Springer-Verlag
7 Palm, G. (1982) *Neural Assemblies: An Alternative Approach to Artificial Intelligence*, Springer-Verlag
8 Hinton, G. and Anderson, J. A. (1981) *Parallel Models of Associative Memory*, Erlbaum
9 Feldman, J. A. (1982) *Biol. Cybern.* 46, 27–39
10 Rumelhart, D. E. and McClelland, J. L. (1986) *Parallel Distributed Processing* (Vols 1 & 2), MIT Press
11 Little, W. and Shaw, G. (1978) *Math. Biosci.* 39, 281–290
12 Hopfield, J. J. (1982) *Proc. Natl Acad. Sci. USA* 79, 2554–2558
13 Steinbuch, K. (1961) *Kybernetik* 1, 36–45
14 Willshaw, D. J., Buneman, O. P. and Longuet-Higgins, H. C. (1969) *Nature* 222, 960–964
15 Marr, D. (1969) *J. Physiol. (London)* 202, 437–470
16 Andersen, P. and Loyning, Y. (1962) *Colloquium Internationale de CNRS* 107, 23–45
17 Dickinson, A. (1980) *Contemporary Animal Learning Theory*, Cambridge University Press
18 Berger, T. W. (1984) *Science* 224, 627–630
19 McCormick, D. A., Clark, E. A., Larond, D. G. and Thompson, R. F. (1982) *Proc. Natl Acad. Sci. USA* 79, 2731–2742
20 Andersen, P., Bliss, T. V. P. and Skrede, K. (1971) *Exp. Brain Res.* 13, 222–238
21 Hjorth–Simonsen, A. and Jeune, B. (1972) *J. Comp. Neurol.* 144, 215–232
22 Steward, O. (1976) *J. Comp. Neurol.* 167, 285–314
23 McNaughton, B. L. (1981) *Brain Res.* 199, 1–19
24 Wyss, J. M. (1981) *J. Comp. Neurol.* 199, 495–512
25 McNaughton, B. L., Barnes, C. A. and Anderson, P. (1981) *J. Neurophysiol.* 46, 952–966
26 Blackstad, T. W., Brink, K., Hem, J. and Jeune, B. (1970) *J. Comp. Neurol.* 138, 433–450
27 Buszaki, G. and Eidelberg, E. (1981) *Brain Res.* 230, 346–350
28 Douglas, R. M., McNaughton, B. L. and Goddard, G. V. (1983) *J. Comp. Neurol.* 219, 285–294
29 Abraham, W. C. and Bliss, T. V. P. (1985) *Brain Res.* 331, 303–313
30 Frotscher, M. (1985) *J. Neurocytol.* 14, 245–259
31 Struble, R. G., Desmond, N. L. and Levy, W. B. (1978) *Brain Res.* 152, 580–585
32 Alger, R. E. and Nicoll, R. A. (1982) *J. Physiol. (London)* 328, 105–123
33 Bland, B. H., Andersen, P., Ganes, T. and Sveen, O. (1980) *Exp. Brain. Res.* 38, 205–219
34 Fox, S. E. and Ranck, J. B., Jr (1981) *Exp. Brain Res.* 41, 399–410
35 Ranck, J. B. (1973) *Exp. Neurol.* 41, 461–555
36 O'Keefe, J. and Dostrovsky, J. (1971) *Brain Res.* 34, 171–175
37 Vanderwolf, C. H. (1969) *EEG Clin. Neurophysiol.* 26, 407–418
38 Fox, S. E., Wolfson, S. and Ranck, J. B. (1986) *Exp. Brain Res.* 62, 495–508
39 Lomo, T. (1966) *Acta Physiol. Scand.* 68 (Suppl. 277)
40 Bliss, T. V. P. and Lomo, T. (1973) *J. Physiol. (London)* 232, 331–356
41 Racine, R. J., Milgram, N. W. and Hafner, S. (1983) *Brain Res.* 260, 217–231
42 Lee, K. (1983) *Neurobiology of the Hippocampus* (Siefert, W., ed.), pp. 265–272, Academic Press
43 Douglas, R. M. (1977) *Brain Res.* 126, 361–365
44 Wigstrom, H. and Gustaffson, B. (1985) *Acta Physiol. Scand.* 123, 519–522
45 Wigstrom, H., Gustafsson, B., Huang, Y. Y. and Abraham, W. C. (1986) *Acta Physiol. Scand.* 126, 317–319
46 McNaughton, B. L., Douglas, R. M. and Goddard, G. V. (1978) *Brain Res.* 157, 277–293
47 McNaughton, B. L. (1982) *J. Physiol. (London)* 324, 249–262
48 Lynch, G. S., Larson, J., Kelso, S., Barrinuevo, G. and Schottler, F. (1983) *Nature* 305, 719–721
49 Collingridge, G. L., Kehl, S. J. and McLennan, H. (1983) *J. Physiol. (London)* 334, 33–46
50 Harris, E. W., Ganong, A. H. and Cotman, C. W. (1984) *Brain Res.* 323, 132–137
51 Mayer, M. L., Westbrook, G. L. and Guthrie, P. B. (1984) *Nature* 309, A281–A263
52 Nowak, L., Bregestovski, P., Ascher, P., Herbet, A. and Prochiantz, A. (1984) *Nature* 307, 462–465
53 Robinson, G. and Racine, R. J. (1982) *Brain Res.* 249, 162–166
54 Douglas, R. M., Goddard, G. V. and Riives, M. (1982) *Brain Res.* 240, 259–272
55 Lynch, G. S., Halpain, S. H. and Baudry, M. (1982) *Brain Res.* 244, 101–111
56 Fifkova, E. and Van Harreveld, A. (1975) *J. Neurocytol.* 6, 211–230
57 Lee, K., Schottler, F., Oliver, M. and Lynch, G. (1980) *J. Neurophysiol.* 44, 247–258
58 Chang, F. L. F. and Greenough, W. T. (1984) *Brain Res.* 309, 35–46
59 Dolphin, A. L., Errington, M. L. and Bliss, T. V. P. (1982) *Nature* 297, 496–498
60 Anderson, P., Sundberg, S. H., Sveen, O., Swann, J. W. and Wigstrom, H. (1980) *J. Physiol. (London)* 302, 463–482
61 Haas, H. L. and Rose, G. (1982) *J. Physiol. (London)* 329, 541–552
62 Levy, W. B. and Steward, O. (1979) *Brain Res.* 175, 233–245
63 McNaughton, B. L. and Barnes, C. A. (1977) *J. Comp. Neurol.* 175, 439–454
64 Andersen, P., Sundberg, S. H., Sveen, O. and Wigstorm, H. (1977) *Nature* 266, 737
65 McNaughton, B. L. (1983) in *Neurobiology of the Hippocampus* (Siefert, W., ed.), pp. 233–252, Academic Press
66 Barnes, C. A. and McNaughton, B. L. (1980) in *Psychobiology of Aging* (Stein, D. G. ed.), pp. 253–272, Elsevier
67 Singer, W. (1985) in *Synaptic Modification, Neuron Selectivity and Nervous System Organization* (Levy, W. B., Andersen, J. A. and Lehmkule, S., eds), pp. 1–38, Erlbaum
68 Abraham, W. C. and Goddard, G. V. (1983) *Nature* 305, 717–719
69 O'Keefe, J. and Nadel, L. (1978) *The Hippocampus as a Cognitive Map*, Oxford University Press
70 Morris, R. G. M. (1983) in *Neurobiology of the Hippocampus* (Siefert, W., ed.), pp. 405–432, Academic Press
71 Barnes, C. A. (1979) *J. Comp. Physiol. Psychol.* 93, 74–104
72 McNaughton, B. L., Barnes, C. A., Rao, G., Baldwin, J. and Rasmussen, M. (1986) *J. Neurosci.* 6, 565–571
73 Morris, R. G. M., Anderson, E., Lynch, G. S. and Baudry, M. (1986) *Nature* 319, 774–776
74 Sharp, P., McNaughton, B. L. and Barnes, C. A. (1985) *Brain Res.* 339, 361–365
75 Sharp, P., Barnes, C. A. and McNaughton, B. L. (1987) *Behav. Neurosci.* 101, 170–178
76 Magleby, K. L. (1979) *Prog. Brain Res.* 49, 175–182

Acknowledgements

Preparation of this paper was supported by PHS grant number NS20331 to BLMcN and a grant from the MRC to RGMM. We are grateful to C. A. Barnes, R. J. Sutherland, E. J. Green and D. J. Willshaw for their comments on earlier drafts of the manuscript and to G. Rao and S. Scott for assistance with figures.

New Neuroscience Journals

GLIA This new bimonthly journal, to be launched in 1988 by Alan R. Liss, will cover a broad range of experimental topics related to research on glia. Manuscripts for consideration may be submitted to the Editors-in-Chief: Bruce R. Ransom, Dept of Neurology, Yale University School of Medicine, 333 Cedar Street, New Haven, CT 06510, USA, or Helmut Kettenmann, Institut fur Neurobiologie der Universitat Heidelberg Im Nuenheimen Feld 504, D-6900 Heidelberg 1, FRG. Subscription prices for Vol. 1 are $65.00 for individuals and $150.00 for institutions. Subscription information and sample copy requests may be obtained from Alan R. Liss, 41 East 11th Street, New York, NY 10003, USA.

SYNAPSE *Synapse* was launched in 1987 as a major new journal for the presentation of all aspects of synaptic structure and function. It is bimonthly and is published by Alan R. Liss. Manuscripts may be submitted to its Editor-in-Chief, John E. Johnson Jr, Dept of Neurology, Eighth Floor–Meyer Building, Johns Hopkins School of Medicine, 600 North Wolfe Street, Baltimore, MD 21205, USA. Subscription prices: $210.00 within the USA and $229.50 outside. For information contact Alan R. Liss, 41 East 11th Street, New York, NY 10003, USA.

NEURON This is a new monthly journal to be launched in 1988 by Cell Press with a focus on cellular and molecular neurobiology. Its format and production schedule will be similar to those of *Cell*. Manuscripts may be submitted to the Editors: Zack W. Hall, A. J. Hudspeth and Louis F. Reichardt, Dept of Physiology, University of California School of Medicine, San Francisco, CA 94143, USA.