HIPPOCAMPUS 11:240–250 (2001)

Role of Long-Term Synaptic Modification in Short-Term Memory

Raymond P. Kesner[1]* and Edmund T. Rolls[2]

[1]*Department of Psychology, University of Utah, Salt Lake City, Utah*

[2]*Department of Experimental Psychology, University of Oxford, Oxford, UK*

ABSTRACT: One way that some types of short-term or working memory may be implemented in the brain is by using autoassociation networks that recirculate information to maintain the firing of a subset of neurons in what is termed an attractor state. We describe how long-term synaptic modification is necessary to set up the appropriate stable attractors, each one of which corresponds to a memory of a particular item. Once the synapses have been modified, any of the short-term memory states may be triggered by an appropriate input which starts the neurons firing in one of the attractors, and then the firing is maintained in that attractor by the already modified synapses, with no further synaptic modification necessary. This analysis leads to the prediction that if this type of implementation is used for working memory, then long-term synaptic modification may be necessary only during an acquisition phase of a task, and once the task has been acquired, the performance of the working memory task should be unimpaired if no further synaptic modification is allowed. We show that a considerable body of research findings on the effects of agents that block synaptic modification on working memory tasks can be understood in this way. Many of the findings are consistent with the hypothesis that blocking synaptic modification in the hippocampus impairs the acquisition, but not the later performance, of hippocampal-dependent working memory tasks. *Hippocampus 2001;11:240–250.*

KEY WORDS: working memory; short-term memory; long-term synaptic modification; AP5; learning; performance; spatial context; temporal context; auto associative memory; attractor neural networks; hippocampus; CA3 region of the hippocampus

INTRODUCTION

Many experiments have been performed to investigate whether short-term or working memory tasks are affected by agents that impair synaptic modification. In some studies, impairments in working memory tasks are reported (Caramanos and Shapiro, 1994; Kesner and Dakis, 1997; Tonkiss and Rawlins, 1991; Walker and Gold, 1994), but not in others (Caramanos and Shapiro, 1994; Tonkiss and Rawlins, 1991; Steele and Morris, 1999). In this paper we clarify how in one type of implementation of working memory in biologically plausible neuronal networks, synaptic modification is necessary only in an acquisition phase, and not later after the possible memory states have been learned. We then show that the results from many of the studies that have been performed can be understood, and at least must be considered, in the light of, these concepts.

The term working memory can be used in a number of ways. One definition is that it is a short-lasting form of memory that can be maintained even when other items are being processed by perceptual systems. However, it is likely that the brain has many working memory systems, each specialized for holding different types of information for potentially different periods. For example, it has been shown that hippocampal damage in rats does not result in short-term or working memory impairments for visual objects, motor responses, odors, and reward value information (Kesner, 1998). In contrast, hippocampal damage results in marked short-term or working memory deficits for spatial information based on memory for spatial location, allocentric spatial distance, egocentric spatial distance, or head direction (DeCoteau et al., 1999; Long and Kesner, 1996, 1998). Clearly, one of the aims of neuroscience is to understand the properties of the different working memory systems in the brain, and how each operates (see Rolls and Treves, 1998, Chapters 3, 6, 8, and 10). In this paper we consider experiments on the type or types of working memory in which hippocampal circuitry is implicated. It is further sometimes implied that working memory is implemented by ongoing neuronal activity, and the model for working memory considered in this paper does maintain the memory by ongoing neuronal activity. Some authors use the term working memory even when the memory must be maintained for long delays (minutes to hours), but in this paper the emphasis will be within a time frame (seconds to minutes) in which firing rates can realistically be maintained to implement a working memory. It should be noted that it is in principle possible to make a memory

Grant sponsor: Medical Research Council; Grant number: PG9826105; Grant sponsor: Human Frontier Science Program; Grant sponsor: NSF; Grant number: IBN 9817583.
*Correspondence to: R.P. Kesner, Department of Psychology, University of Utah, 380 South 1530 East, Room 502, Salt Lake City, UT 84112.
E-mail: rpkesner@behsci.utah.edu
Accepted for publication 22 October 2000.

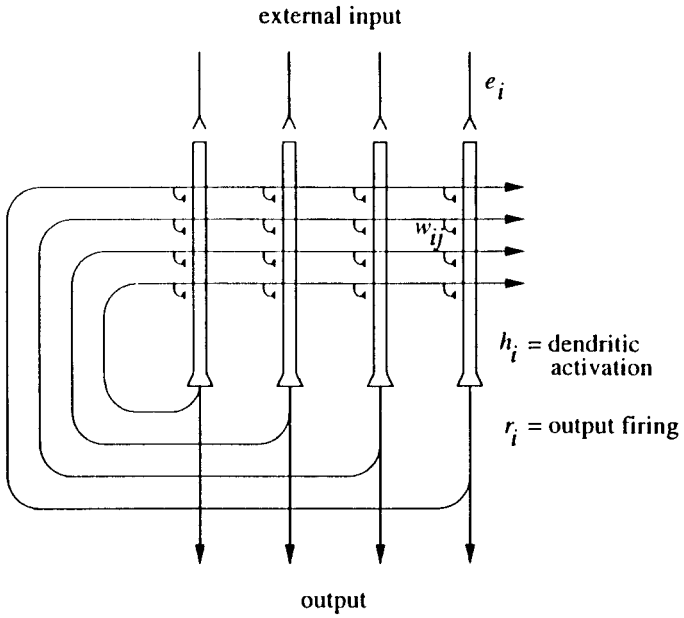

FIGURE 1. Architecture of an autoassociative neural network. External input e_i is applied to each neuron i by unmodifiable synapses. This produces firing r_i of each neuron. Each output neuron i is connected by a recurrent collateral synaptic connection to the other neurons in the network, via modifiable synaptic connection strengths or weights w_{ij}.

which responds differently to novel and familiar stimuli by having neurons that show habituation between the novel and familiar presentations, and in which there is no ongoing neuronal activity (Rolls and Treves, 1998, pp. 149–151). Such a system is limited because it is primarily useful for responding differentially to recently seen and novel stimuli, rather than recalling one or several items from a working memory store, e.g., where from among a large number of spatial locations a stimulus has just been presented. Thus, the novelty/familiarity memory system that could be implemented by nonassociative habituation is not considered further here.

NEURONAL NETWORK SYSTEMS FOR WORKING MEMORY

Autoassociative memories, or attractor neural networks, store memories, each one of which is represented by a pattern of neural activity. The functional architecture of an autoassociation memory is shown in Figure 1. Such networks can recall the appropriate memory when provided with a fragment of one of the memories. This is called completion. Many different memories can be stored in the network and retrieved correctly. The network can in principle learn each memory in one trial. The associative synaptic modification which implements the learning could involve long-term potentiation (LTP), and this can occur rapidly, within a trial. An autoassociation memory can be used as a short-term memory, in which feedback processing around the recurrent collateral connection loop keeps a representation active until another input cue is received. The operation of this type of network has been studied extensively, as it could be involved in a number of types of memory operation in the brain (Hebb, 1949; Marr, 1971; Kohonen, 1977; Hopfield, 1982; Amit, 1989, 1995; Rolls, 1989; McNaughton and Morris, 1989; Treves and Rolls, 1991, 1994; Hasselmo and Bower, 1993; Shapiro and Hetherington, 1993; Muller and Stead, 1996; Granger et al., 1994; Renart et al., 1999a, 1999b, 2000; August and Levy, 1999; Eichenbaum et al., 1999). A description of the operation and properties of this type of neuronal network, and its relevance to understanding the operation of different brain systems, is provided by Rolls and Treves (1998, Chapter 3).

The purpose of this paper is to make clear why long-term synaptic modification is necessary to set up a new short-term memory using this class of network, but is not necessary when the short-term memory is triggered and used later; and to consider whether the effects on short-term memory of agents that block long-term synaptic modification can be understood in terms of their effects on the setting up, but not the later use of, short-term memories in this class of neuronal network. This is the only class of network we know that can implement a short-term memory by continued firing of a subset of neurons, with a different subset for each memory. We do note that there is a proposal that a set of neurons could be kept active in a recurrent network with nonmodifiable recurrent synapses by keeping the active neurons above their NMDA receptor voltage-dependent thresholds so that this nonlinearity in the firing rate characteristic of the neurons can keep the initially activated neurons active (Lisman et al., 1998). However, we note that the arbitrary set of active neurons could not actually be used as a short-term memory without a readout mechanism from the network to learn which memory was represented by each set of firing neurons, and this would require associative synaptic modification during a setting-up phase, so that this proposal in no way obviates the need for associative synaptic modification during a setting-up phase, just as does the attractor network described here. A similar argument applies to the scheme proposed by Jensen and Lisman (1996), which also requires long-term synaptic modification for good performance. We also note that the system described by Lisman et al. (1998) does not have any of the desirable properties of short-term memory systems such as completion, generalization, and noise tolerance (Rolls and Treves, 1998, Chapter 3), and this is because these properties rely on associative synaptic modification. In contrast, nonassociative processes such as habituation of recently activated synapses or depolarizing after-potentials might implement altered neuronal responses to a repeated stimulus in a novel/familiar recognition memory task (Baylis and Rolls, 1987; Brown and Xiang, 1998), but would not be expected to require associative long-term synaptic modification in order for the memory to be set up. To the extent that blockers of long-term associative synaptic modification such as long-term potentiation (LTP) implemented using n-methyl-d-aspartate (NMDA) receptors do impair setting up short-term memories but not their later use, the indication is that the short-term memory is being implemented by autoassocation of the type described here rather than nonassociative novel/familiar processes. If NMDA receptor blockade turns out also to impair the actual use of the working memory even after

it has been set up, then the implication is that the NMDA receptors are involved in simply maintaining neuronal firing in a recurrent network. Although much of the maintenance of firing is likely to be implemented by AMPA receptor-mediated events, the longer synaptic time constant of NMDA receptors could be useful in the maintenance of ongoing neuronal recurrent activity (Wang, 1999). It turns out that many of the studies described here of working memory in the hippocampal system are consistent with the hypothesis that NMDA receptor activation is necessary for the setting up but not later use of hippocampus-dependent working memory systems. The studies described here are thus useful in understanding the details of operation of hippocampus-dependent working memory circuits, and imply that these may be set up by synaptic modification in autoassociation networks, and that the later maintenance of the firing is not dependent on NMDA-receptor mediated synaptic currents.

We now describe in a little more detail how an autoassociation memory operates as a short-term memory, and then discuss when the synaptic modification takes place in it in relation to how behavioral tasks are designed to measure short-term memory function. One of those who first investigated autoassociation memories was Marr (1971), in the context of hippocampal function, and with respect to completion. Kohonen (1997) studied autoassociation memories, also in the context of a content-addressable memory that could complete. Hopfield (1982) brought the theoretical physics approach to bear on such memories, and proved that if they had complete connectivity, and were trained with a Hebb rule, then there were a number of stable retrieval states in which neuronal firing could continue in what is termed a stable attractor. The term attractor refers to the fact that patterns similar to one of those stored in the network will lead to recall of the pattern stored in the network, and completion is an example of this attractor process.

The prototypical architecture of an autoassociation memory is shown in Figure 1. The external input e_i is applied to each neuron i by unmodifiable synapses. This produces firing r_i of each neuron. Each output neuron i is connected by a recurrent collateral synaptic connection to the other neurons in the network, via modifiable connection weights w_{ij}. This architecture effectively enables the output firing rates of all the neurons in the network to be associated with themselves, using the synaptic connections implemented via the recurrent collaterals. Later on, during recall, presentation of part of the external input will force some of the output neurons to fire, but through the recurrent collateral axons and the modified synapses, other output neurons can be brought into activity. This process can be repeated a number of times, and recall of a complete pattern may be perfect. Effectively, a pattern can be recalled or recognized because of associations formed between its parts. This requires distributed representations, i.e., representations in which each pattern to be stored consists of a set of active inputs.

A large number of different patterns can be stored in an autoassociation network by setting up the synaptic connection strengths (or weights) as follows. The firing of every output neuron i is forced to a value r_i determined by the external input e_i. Then a Hebb-like associative local learning rule is applied to the recurrent synapses in the network:

$$\delta w_{ij} = k r_i r_j \qquad \text{(Hebb rule)} \tag{1}$$

where r_j is the firing of the jth axon in Figure 1. It is notable that in a fully connected network, this will result in a symmetric matrix of synaptic weights, i.e., the strength of the connection from neuron 1 to neuron 2 will be the same as the strength of the connection from neuron 2 to neuron 1 (both implemented via recurrent collateral synapses).

During recall the external input e_i is applied, and produces output firing. The firing is fed back by the recurrent collateral axons shown in Figure 1 to produce activation of each output neuron through the modified synapses on each output neuron. The internal activation h_i produced by the recurrent collateral effect on the ith neuron is the sum of the activations produced in proportion to the firing rate of each axon r_j operating through each modified synapse w_{ij}, i.e.,

$$h_i = \Sigma_j r_j w_{ij} \tag{2}$$

where Σ_j indicates that the sum is over the C input axons to each neuron, indexed by j.

The output firing r_i is a function of the activation produced by the recurrent collateral effect (internal recall) and by the external input (e_i):

$$r_i = f(h_i + e_i) \tag{3}$$

The activation function should be nonlinear, and may be for example binary threshold (no firing below a threshold and maximum firing above a threshold), linear threshold (linear after the threshold has been reached), or sigmoid (see Fig. 1.3 of Rolls and Treves, 1998). A nonlinear activation function can minimize interference between the pattern being recalled and other patterns stored in the network, and can also be used to help a positive feedback system to remain stable. The network can be allowed to repeat this recurrent collateral loop a number of times. Each time the loop operates, the output firing becomes more like the originally stored pattern, and this progressive recall is usually complete within 5–15 iterations. Once the recall is complete, the neurons will keep firing in the stable attractor until another external input is applied to the network, and pushes the network towards another stored pattern.

With this type of network, the use of associative Hebbian learning in the recurrent collateral synaptic connections sets up the network to have a number of stable retrieval states, each one corresponding to the set of neurons that were active when the pattern was being learned by the network. An example is now provided to make this clear. If the firing applied in Figure 2 as the external input was 1010 (where 1 indicates high activity, and 0 indicates zero activity), then the firing on the axons would be as shown (1010 reading from the bottom), and the synaptic weights produced by the learning would be as shown, with strong connections (shown as a strength of 1 rather than 0) between neurons 1 and 3, 3 and 1, 1 and 1, and 3 and 3. At any time thereafter, if the input pattern, or a part of it, was applied as an external input, then the network would continue to reexcite itself through the recurrent collateral connections, and would continue to fire (indefinitely in principle) with the firing pattern 1010. This illustration should make it very clear that synaptic modification is needed during a

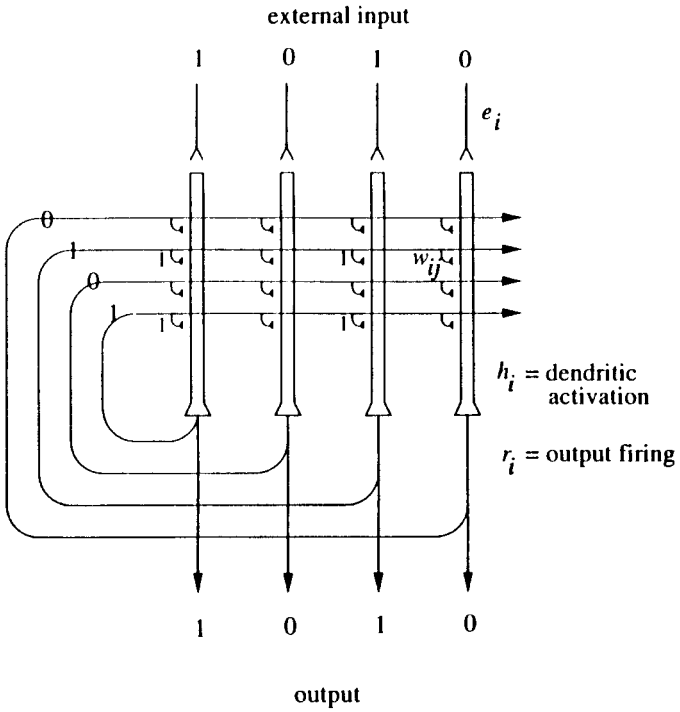

FIGURE 2. Autoassociative network during training on the external pattern 1010, showing how this is reflected in the firing of the recurrent collateral axons, and in the synapses that are modified during learning.

learning phase to learn all the patterns that might later need to be activated as short-term memories. But after the synapses have been modified by learning, it is possible to set the neurons firing in an attractor short-term memory state with no further synaptic modification. The particular attractor state that is selected depends simply on which external input retrieval cue pattern is applied. We note in addition that individual attractor nets are suitable for holding one item active in a delay period but cannot easily hold more than one item active, and that attractor nets are unlikely to bridge long delays because when the organism is distracted away from the memory task, other inputs are likely to trigger the attractor, and the originally active attractor state will be lost (see Rolls and Treves, 1998, Chapter 3).

This analysis then leads to the following predictions. In studies of the role of synaptic modification in short-term or working memory tasks, performance will be impaired if the application of the inhibitor of synaptic modification is made during the learning phase. In contrast, application of an inhibitor of synaptic modification should have no effect on the later performance of those particular short-term memory tasks (provided that the inhibitors do not produce generic impairments in the operation of the networks, such as reducing cell excitability). In addition, the most clear dissociation between the effects of synaptic modification inhibitors on the initial setting up and the later use of short-term memory networks is likely to be in tasks where there is a discrete set of short-term memory states, each of which can be learned. This would be the case where, e.g., a rat was trained to perform a short-term memory task with multiple different possible locations for a hidden platform. Once the places of the platforms had been learned, no further learning would be needed in a short term-memory task, using the first part of a trial to show which of the one or more places was correct to go to after a memory delay period. On the other hand, if the platforms were moved to a variable position (or new testing environment) on each trial and then the rat was required to remember its position, some new learning, and therefore synaptic modification, might be needed to store a representation of that particular trial-unique place. Similarly, if the delay period or intertrial interval is changed on each trial, some new learning, and therefore synaptic modification, might be needed to store a representation of that particular trial-unique temporal interval. We therefore predict that how the rat is trained, as well as when the inhibitor of synaptic modification is given in relation to training and the later short-term memory testing, will both be important in interpreting the effects of experimental investigations on the role of impairment of long-term synaptic modification on short-term memory tasks.

We note that single attractor networks can normally keep one item active in short-term memory. If several items must be kept active, one possibility is to have several separate, possibly weakly coupled, attractor networks, each implemented by recurrent collateral connections (Levy et al., 1999; Renart et al., 1999a, 1999b). Another possibility is to use a mechanism with discrete timesteps, and use the output of an attractor net at one timestep to retrieve the appropriate pattern after a suitable delay at the next timestep. This forms a sequence memory (see Rolls and Treves, 1998, section 3.3.10). However, in practice in the brain an external clock or timing device would be needed to keep the sequential stages discrete, because otherwise attractor networks with integrate-and-fire neurons with continuous dynamics fall into their attractor very fast, as shown by Treves (1993) (see also Rolls and Treves, 1998, Appendix 5) and would not easily inherently achieve successive discrete states without an external timing mechanism. However, in both cases, associative synaptic modification is still needed to set up, but not to use later, the short-term memory, and the predictions are the same as those described above.

EFFECTS OF BLOCKERS OF SYNAPTIC MODIFICATION ON THE ACQUISITION AND PERFORMANCE OF WORKING MEMORY TASKS

Here we address the issue of how inhibitors of synaptic modification affect hippocampal-dependent short-term or working memory tasks. The hippocampal formation may play a role in these working memory tasks by virtue of the recurrent collateral connections between cells at one or several of its stages. These stages might include the presubiculum and parasubiculum, lesions of which impair a spatial working memory task (Kesner and Giles, 1998), and/or the entorhinal cortex. In the context of working memory, the hypothesis would be that networks in some part of the hippocampal formation operate as attractor networks, with

activity being maintained in the working-memory period by recirculating activity in the recurrent collateral connections. This hypothesis would predict that setting up the attractor states in the first place would involve synaptic modification, during initial learning. Once the possible attractor states have been learned, it would be predicted that any one of these attractor states could be cued by a particular sensory input with no further synaptic modification. The initial learning would require synaptic modification; but the actual use of the working memory at any later time would not be dependent on further synaptic modification.

In considering actual experiments, a number of points that may affect the interpretation of the results will need to be borne in mind. First, there could only be a clear dissociation between the initial learning phase and the later performance phase of the spatial working memory task if the inhibitor of synaptic modification does not produce any generic impairment of synaptic transmission. To the extent that some synaptic transmission (in addition to synaptic modification) may be N-methyl-D-aspartate (NMDA)-receptor-dependent (Lisman et al., 1998), not only the learning, but also the performance phase, would be expected to be impaired. Second, only some experiments may set up a working memory performance testing phase which does not require any new learning. If a small set of platform locations on each of which pretraining has been given so that adequate spatial representations of each platform position have already been trained prior to the working memory performance testing phase, then no new learning may be needed in the performance testing phase. If, in contrast, the performance testing phase involves placing an escape platform in a particular place in which it has not been found previously, then some new learning of the spatial location of the platform would be required in the spatial working memory testing phase. If during pretraining in the spatial working memory task the delay and intertrial intervals (ITI) are maintained and are invariant, then no new learning may be needed in the performance testing phase. If, in contrast, the performance testing phase involves changes in the temporal structure of the task, i.e., a change in the delay interval between the study and test phase components of a trial setting up new temporal parameters, then some new learning would be required in the spatial working memory testing phase. Third, the particular hypothesis being addressed here is that the attractor network is implemented in some part of the hippocampal formation; therefore, to test the hypothesis that new learning is required to set up the possible attractor states, but not to use them in the performance of a spatial working memory task, the inhibitor of synaptic modification must reach the correct part of the hippocampal formation where the attractor is implemented.

Because the autoassociative recurrent collateral model described above is applied in this paper to the modification of synapses within the hippocampal formation, we limit this review to intracranial injections which are likely to affect the hippocampal formation, and to spatial tasks that have been shown to be dependent on hippocampal formation function. There are two major procedures to inject NMDA receptor blocking agents into the hippocampal formation. The first is an indirect method with a cannula implanted into the lateral ventricle and the use of a minipump to slowly release the NMDA blocking agent over a long time period. This method is likely to result in the injected agent affecting much of the hippocampal formation, including areas such as the presubiculum, parasubiculum, and entorhinal cortex. Furthermore, the drug is likely to affect all regions of the hippocampus, including CA1, CA3, and dentate gyrus. The second is a more direct method with a cannula implanted into the dentate gyrus, CA3, CA1, presubiculum, parasubiculum, or entorhinal cortex. In this case the drug injections are likely to target one or more subregions of the hippocampal formation. However, if the spatial working memory function were implemented in a part of the hippocampal formation other than the injection site, such intrahippocampal injections would not directly test the hypothesis considered here.

Since it usually takes many trials to learn a spatial working memory task, many injections are needed to test the effects of receptor blockade on the acquisition of a spatial working memory task. The number of studies that have reported the effects of NMDA receptor blockade in such a paradigm is therefore rather limited. In one study, Kesner and Dakis (1997) injected 10 min prior to each trial 36 nanomoles (nM) or 54 nM of phencyclidine (PCP), an NMDA antagonist, or saline into the dentate gyrus of the hippocampus during the acquisition of a spatial continuous recognition memory task on a 12-arm maze. The results indicate that relative to controls and rats with 36-nM PCP injections, rats with 54-nM PCP injections are impaired in the acquisition of the task across all lags (delays and intervening spatial locations), as measured by an increase in latency to respond to repeated items. These drug doses have been shown not to disrupt sensory-motor function. A similar pattern of results was reported following peripheral injections of 4 mg/kg PCP (Kesner and Dakis, 1993). In a different study, Walker and Gold (1994) showed that injections of 0.63 nM CPP or 50 nM D,L-AP5 (NMDA receptor antagonists) into the CA3 region of the ventral hippocampus, with diffusion into the CA1 and dentate gyrus region 15 min prior to testing, decreased learning of a spontaneous alternation task. Both studies suggest that there is a deficit in learning a spatial working memory task following intracranial injections of NMDA blocking agents into either the dorsal or ventral hippocampus. We note that a possible interpretation is that these parts of the hippocampus are needed in order to learn the spatial representation required, which is then sent to another part of the hippocampal formation for the actual implementation of the spatial working memory by an attractor network.

An alternative approach is to train rats in a spatial short-term or working memory task, and then inject an NMDA receptor antagonist to observe its effects on performance of the already trained task. In addition, it is then possible to test whether the NMDA receptor blockers impair the acquisition in the same animals of the same task, but in an unfamiliar spatial environment. In one such study, Caramanos and Shapiro (1994) trained rats for 65 days in an eight-arm maze with four arms that had one piece of food in each of these four arms. The rat could use working memory to avoid returning to an arm that had already been visited. In addition, four arms never had food, allowing spatial reference memory to be tested. At the end of training, 20- or 30-mM AP5 infusions into the lateral ventricle resulted in continued normal performance for

both working and reference memory throughout the 4 days of testing in the pretrained environment. Thus, AP5 did not impair the performance of a spatial working memory task when the places had been well learned previously. In contrast, when the rats were then tested for 6 days in a new and unfamiliar environment, there were profound deficits in the spatial working memory task, with no improvements across the 6 days of testing. Thus, AP5 disrupted the learning of a new spatial working memory task. A similar pattern of results was reported following peripheral injections of 0.0625 mg/kg of MK-801 (an NMDA receptor antagonist) (Shapiro and O'Conner, 1992).

Tonkiss and Rawlins (1991) first trained rats in a spatial alternation paradigm in a T-maze for 6 days at eight trials per day. Then after the training, during ventricular infusions of 15 mM AP5, there was very good performance of this spatial working memory task in which the delay was 1–2 s between trials (with just a mild deficit on the first day of infusions).

Robinson and Mao (1997) pretrained rats on a delayed nonmatching-to-position task with delays of 0, 4, 8, or 16 s. They then showed that there was little impairment of this spatial working memory task produced by the injection of 1.8 μg, 2.4 μg, 3.6 μg, or 4.8 μg MK-801 into the dentate gyrus region of the dorsal hippocampus and 4 μg or 12 μg MK-801 into the CA1 region of the ventral hippocampus 2 min before testing. In another study, Kawabe et al. (1998) trained rats for 12 days, using the standard eight-arm maze procedure with each arm baited once to provide a spatial working memory task. When they retested the rats following direct injections of 20 or 40 mM AP5 into the dentate gyrus region of the dorsal hippocampus made 20 min prior to testing, they found that there was a dose-dependent deficit in spatial working memory performance, as shown by an increase in the number of return errors. The hypothesis we are considering would not have predicted this, but we note that 2 μl of AP5 injections into the dorsal hippocampus on each side are quite high and can easily result in spread into other neural regions.

Hippocampal-dependent short-term or working memory tasks do not always involve the modification of synapses to provide for a representation of spatial information, but can also involve synaptic modification involving the representation of temporal information. For example, it has been shown that rats with hippocampal lesions are not able to perform a working memory task that requires memory for the temporal order of spatial location information (Chiba et al., 1994). In this task, rats are required to remember an event (e.g., spatial location) dependent on the temporal order of occurrence of events. More specifically, on an eight-arm maze during the study phase of each trial, rats were allowed to visit each of the eight arms once in an order that was randomly selected for that trial. The test phase required the rats to choose which of two arms occurred earlier in the sequence of arms visited during the study phase, i.e., a temporal order memory task. The arms selected as test arms varied according to temporal lag or distance (0–6) or the number of arms that occurred between the two test arms in the study phase. This task requires rats to encode and chunk spatial locations into discrete temporal events. Rats with electrolytic lesions of the hippocampus have difficulty separating one event (spatial location) from another, perhaps because of an inability to inhibit interference that is likely to accompany sequentially occurring events. In additional research, it was shown that ibotenic lesions of the CA3 or CA1 region disrupt temporal order working memory (Kesner et al., 2000). It is important to point out that the number of possible temporal order representations of eight spatial locations is very large, so that pretraining is not likely to modify all the critical set of NMDA receptors required for temporally dependent memory, so that it would be predicted that injections of NMDA blocking agents should still be capable of disrupting temporal order memory performance. Indeed, it was shown in one experiment that 3–4 mg/kg of PCP disrupted performance on the above-mentioned temporal order memory paradigm (Long and Kesner, 1995). In a different experiment, Ohno and Watanabe (1995, 1998) and Ohno et al. (1992) used a three-panel runway task that required temporal ordering of information. In this task there are four choice locations, and there are three doors per choice location. The rats receive six trials per day in which the patterns of three doors per choice were different from those used on other days. After the rat has made a set of choices on trial 1, the rat must remember and repeat the same choices on the next five trials. After 15–20 pretraining sessions on this task, rats received bilateral injections of NMDA antagonists into the CA1 region of the dorsal hippocampus 10 min before the presentation of the first trial of that day. The results showed that 3.2 ng, 10 ng, and 32 ng CGS 19755 (an NMDA competitive antagonist), and 3.2 μg, but not 0.32 or 1.0 μg AIDA (a selective antagonist of the group 1 metabotropic receptor) or 0.1 and 0.18 μg MK-801 (an NMDA receptor antagonist), produced deficits, suggesting that rats had difficulty in learning the task for that day.

An alternative way to conceptualize the importance of temporal processing is to examine the temporal components associated with the standard working memory paradigms. There are two critical time intervals that need to be processed. The first involves the time on an individual trial between the study phase and the test phase which has been labeled the delay period; the second involves the time between each trial, which includes a set of study and test phases which is labeled the intertrial interval (ITI). One might predict that inhibitors of synaptic modification should not affect working memory performance after pretraining as long as the delay and ITI intervals are constant, and the delay interval does not become too long or filled with interfering stimuli which disrupt the attractor state. However, any changes in the delay period or ITI period or the relative relationship between the delay period and ITI might require new synaptic modification in that, for example, the animal has to learn to wait for a very long time before being allowed to make the choice. Thus, altering the temporal components of these tasks might require new contextual learning, and make these tasks susceptible to NMDA inhibitors until new learning of the new temporal components of the task has been accomplished. The experiments we describe next suggest that altering the temporal components of the task does make the task sensitive to NMDA receptor blockers, and on this evidence does require new learning of the temporal contexts involved.

The evidence for this idea comes from the studies by Tonkiss and Rawlins (1991), Bolhuis and Reid (1992), and Steele and Morris (1999). Tonkiss and Rawlins (1991) first trained rats in a

spatial alternation paradigm in a T-maze for 6 days at eight trials per day. Then after training, during ventricular infusions of 15 mM AP5, there was very good performance of this spatial working memory task in which the delay was 1–2 s between trials. However, when the rats were exposed to a new 20-s delay period between trials, rats with AP-5 infusions did not learn the task relative to rats with vehicle injections. Bolhuis and Reid (1992) trained rats for 5 days on an eight- arm maze, using a four-arm study and four-arm test phase procedure with a 5-min delay period. Following infusions of AP5, new trials were presented with 20- or 60-min delays interspersed with 1- and 5-min delays with no improvement in performance compared to vehicle control rats. The deficits in the above-mentioned studies are probably due to the requirement of new learning associated with the new time delays requiring the need to modify new synapses, and the fact that an attractor network is not likely to be able to maintain its firing for such long periods, as other inputs may push the attractor into another memory state unless special training has been given to expect a long delay.

In a different experiment, Steele and Morris (1999) trained rats in a matching-to-place task in a water maze. The rats were trained to find a hidden platform that was placed in one of nine possible locations for each day with either a 15-s intertrial interval for all four trials or intertrial intervals that included 20 or 120 min between trials 1 and 2, followed by 15-s intertrial intervals between trials 2 and 3, and trials 3 and 4. The starting position was varied from trial to trial. Learning was observed as measured by a reduced latency on the second trial of each day. The results indicated that AP5 injections into the hippocampus after training did not disrupt performance in a working memory task in a water maze when trained with a 15-s intertrial interval for all four trials, but when the training intertrial interval included an unexpected 20 or 120 min between trials 1 and 2 followed by 15-s intertrial intervals between trials 2 and 3 and trials 3 and 4, disruptive effects appeared. These studies indicate that when the task is performed in a way that is likely to rely on a spatial and temporal working memory of the type that could be implemented if an attractor is used, blocking synaptic modification does not impair use of the pretrained attractor memory. Alternatively, NMDA inhibitors disrupt working memory performance at long, but not short delay periods because of the disruption of ongoing active attractor states. Finally, there is another possibility that short-term memory for spatial information is not dependent on NMDA receptor activation, but that NMDA receptor activation is important for rapid consolidation of spatial information into long-term memory.

To test the idea that blocking synaptic modification does not impair use of the pretrained attractor memory and at the same time to localize injections into the CA3 region of the hippocampus, rats were trained to remember one of two arms within an eight-arm maze (Kesner, in preparation). On any day there were four independent working memory trials with two trials at a 10-min delay and two trials at a 30-min delay. After training, different subjects received either saline or 36 nM PCP directly into the CA3 region of the dorsal hippocampus. It can be shown that PCP or AP5 injections into the CA3 region do not spread into the CA1 and dentate gyrus regions, suggesting localized action of the drug in the CA3 region. The order and dose of drug injections were counterbalanced with 2 days of drug injections followed by 2 days of saline injections. The results are shown in Figure 3 and indicate that there is good working memory for the 1-min and 30-min delays for the rats that received either saline or 36 nM PCP injection into the CA3 region. This suggests that after extensive training with a 1- or 30-min delay period, NMDA receptors in the CA3 region are not involved in maintaining spatial working memory for a 1- or 30-min delay period. Thus, involvement of NMDA receptors is not a function of the delay period, but a function of change relative to pretraining requiring new learning and new modification of NMDA receptors. The intracranial injections used in this study leave untested the possibility that some spatial working memory is being maintained in this task in an attractor network that is in another part of the hippocampal formation, but this possibility is made unlikely by the fact that peripheral injections of PCP (4 mg/kg) also did not impair the spatial working memory performance on the pretrained task at the 1-min and 30-min delays (see Fig. 3). In another study using the same short-term memory procedure with 10-s delays, AP5 injected directly into the CA3 region did not impair performance of the spatial working memory task when the places had been well-learned previously. In contrast, when the rats were then tested in a new and unfamiliar environment, AP5 injections into the CA3 region produced a profound deficit in the spatial working memory task (I. Lee and R.P. Kesner, unpublished observations). Similar deficits were found with AP5 injections into the CA3 region in the same working memory task when a delay of 10 s was changed to 5 min (I. Lee and R.P. Kesner, unpublished observations). Thus, AP5 injections into the CA3 region disrupted spatial working memory tasks whenever there was a change in the spatial or temporal context.

DISCUSSION

Many studies indicate that after initial training, the actual performance of spatial working memory tasks does not require synaptic modification, as assessed by the use of blockers of NMDA receptors. In some cases, as with local injections into the CA3 region, the experimental findings might arise because the injections did not actually reach whichever attractor may maintain the memory state, and this could be in some other part of the hippocampal formation. However, the fact that peripheral injections of NMDA receptor blockers do not impair the performance of spatial working memory tasks would also support the hypothesis that once attractor states have been trained into recurrent networks by synaptic modification during initial learning, the later use of these attractor states for holding a working memory state for a short period does not require further synaptic modification.

With respect to the acquisition or new learning of spatial or temporal working memory tasks, most studies do find an impairment produced by treatments, such as PCP, MK-801, and AP5, that block the formation of LTP in the hippocampal system. It appears that intrahippocampal injections of AP5 or peripheral in-

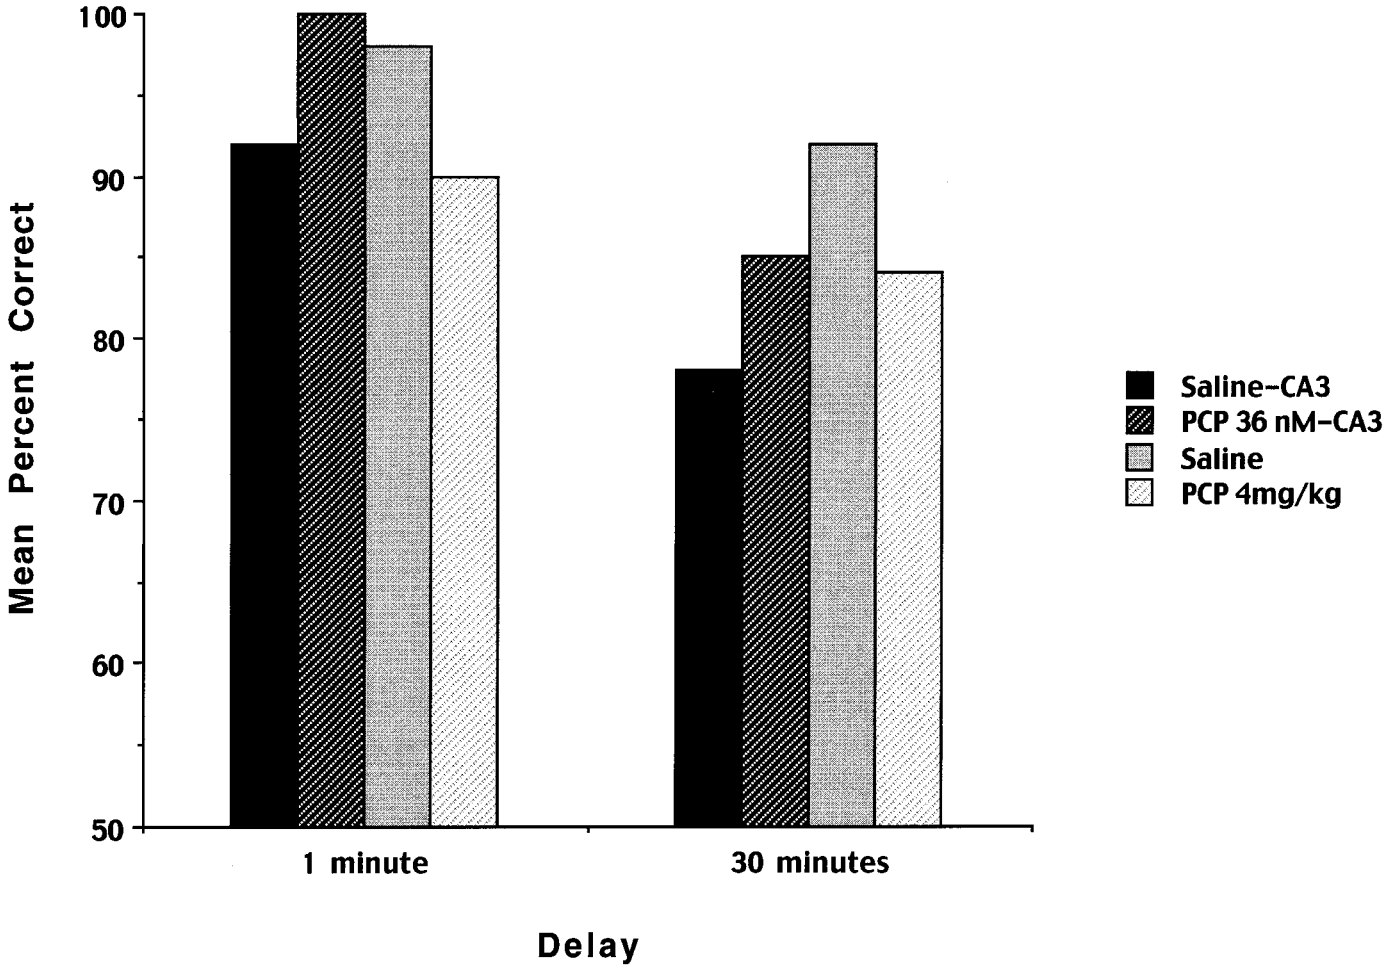

FIGURE 3. Mean percent correct performance following either PCP or vehicle injections directly into the CA3 region, or peripheral injections of phencyclidine (PCP) or vehicle as a function of a 1- or 30-min delay between study phase and test phase.

jections of PCP and MK-801 produce a disruption of LTP at dose ranges that also produce behavioral effects in learning of spatial or temporal working memory tasks (Abraham and Mason, 1988; Caramanos and Shapiro, 1994; Gilbert and Mack, 1990; Kesner and Dakis, 1993; Morris et al., 1989; Stringer and Guyenet, 1983; Stringer et al., 1983). The treatment effects on acquisition or new learning of spatial or temporal working memory tasks could arise for two reasons. One possibility is that the synaptic modification required to set up each of the attractor states could not be set up in a network that would later be used as an attractor. A second possibility is that some parts of the hippocampus must learn using synaptic modification in order to set up the spatial and temporal representations that will then be sent to the attractor network in another part of the hippocampal formation, or elsewhere in the brain. If the spatial and temporal representations to be later remembered cannot be set up in the first place, then of course later use of these long-term spatial or temporal representations to trigger an attractor into one of its short-term memory states would not be possible.

For future studies, a good paradigm to address the hypothesis described here would include testing the efficacy of inhibitors of synaptic modification on both new learning and the later performance of spatial and/or temporal short-term memory tasks. Good examples of such experiments are those described by Caramanos and Shapiro (1994) and Tonkiss and Rawlins (1991). However, such studies would be ideally appropriate if the injections were aimed to affect the actual network that supports the attractor states, evidence for which would include recording of continuing neuronal firing during the delay period. We next discuss evidence of this type, and other evidence, on where the actual attractor networks might be for the spatial and temporal working memory tasks analyzed here. However, we emphasize that the hypothesis described here is generic. If short-term memory is implemented by attractor networks, then we predict that the acquisition of new attractor states should depend on synaptic modification, but once learned, no further synaptic modification should be necessary in order to use these attractor states in short-term memory tasks.

What subregions and networks within the hippocampal formation could be considered as good candidates for the attractor network or networks that implement hippocampus-dependent spatial working memory? From a neuroanatomical point of view, one type of implementation of an attractor network would utilize a set of

neurons with recurrent collaterals which undergo Hebb-like synaptic modification (see Rolls and Treves, 1998, Chapter 3). Another possible implementation would involve feedforward connections from one set of neurons to a second set, which in turn have feedback connections to the first set, and with Hebb-like synaptic modification in the forward and backprojection synapses. Such an architecture is found in many parts of the cerebral cortex, and attractor states are one possible function that could be performed by this backprojection architecture, as described by Rolls and Treves (1998, section 10.2). The mossy/granule-cell system in the hippocampus might also implement an architecture of this type, as suggested by Buckmaster and Schwartzkroin (1994). On that basis, the CA3 recurrent collateral network (see Rolls and Treves, 1998, Chapter 6), and the entorhinal cortex with its recurrent collaterals which are a characteristic property of most neocortical areas, are potential candidates for mediating a spatial working memory attractor system. Evidence in favor of the CA3 region comes from the observation that injections of NMDA receptor blockers into the CA3 region, which has recurrent collaterals, disrupt the learning of a spatial working memory task, but do not disrupt later performance of a spatial working memory task (I. Lee and R.P. Kesner, unpublished observations; Walker and Gold, 1994). The hypothesis would be that the CA3 attractor network maintains activity during the spatial working memory period by recirculating activity in the recurrent collateral connections. Thus, one should be able to measure increased firing rates during the delay period. Hampson et al (1993) trained rats in a delayed matching to position task and recorded from CA3 hippocampal cells (and CA1 cells, which would be expected to reflect any continuing activity in the CA3 cells). They reported that one subset of CA3 and CA1 cells increased their firing rates during the 30-s delay interval. Similar increases in firing rate in hippocampal neurons have been reported for monkeys during the delay period in spatial working memory tasks (Cahusac et al., 1989; Colombo et al., 1998). It should be noted that others have not found sustained firing in rat hippocampal cells during a delay in a T-maze spatial working memory task (Wible et al., 1986). Finally, the fact that hippocampal place cells in rats (Samsonovitch and McNaughton, 1997) and spatial view cells in monkeys (Robertson et al., 1998; Rolls, 1999) continue to fire in the dark, with some drift without the view details visible, can also be taken as consistent with the hypothesis that there is an attractor in the CA3 region, which keeps the neurons firing in the absence of one of the major driving sensory inputs, the view of the environment. Consistently, kainic acid lesions of the CA3 region have been shown to disrupt spatial working memory (Handelmann and Olton,1981; Jarrard, 1983), although the CA3 effect in Jarrard (1983) was rather small. In addition, there is an impairment in memory for all items of a list of five spatial locations when an electrolytically induced dorsal hippocampal lesion includes the CA3 region, but memory for the last spatial location in the list is intact when the CA3 region is spared, suggesting that short-term or working memory for spatial location information is mediated by the CA3 region (Kesner et al., 1988). Furthermore, it has been shown that ibotenic acid lesions of the CA3 or CA1 region disrupt temporal order working memory (Kesner et al., 2000). Thus, based on injections of NMDA receptor inhibitors into CA3, sustained firing during the delay period of CA3 neurons, and lesions of CA3, there is reasonable evidence for a role of CA3 as an attractor network that is used in hippocampal-dependent spatial and temporal working memory tasks.

Support for the possibility that the entorhinal cortex might implement an attractor useful in some short-term memory tasks is that cells in the entorhinal cortex show sustained firing during the delay period in a spatial working memory task in monkeys (Suzuki et al., 1997) and in an odor working memory task in rats (Eichenbaum, 1997). In some studies, lesions of the entorhinal cortex in rats produce spatial working memory deficits (Holscher and Schmidt, 1994; Nagahara et al., 1995), but in other studies there are no deficits (Bouffard and Jarrard, 1988; Kesner and Giles, 1998; Rothblat et al., 1993).

In conclusion, the hypothesis described here is that some short-term memory tasks can be implemented by firing in attractor networks, with synaptic modification being necessary to set up the attractor states, but not to use them for later performance. We have reviewed evidence that for hippocampal-dependent spatial working memory, synaptic modification is indeed not necessary for performance, but is necessary for learning. Although we do not take a strong position on where in the hippocampal formation the relevant attractor network may be, we describe some evidence that it might be at least in part implemented by the well-developed hippocampal CA3 cell recurrent network, which has Hebb-modifiable synapses (Debanne et al., 1998). We note that the hypothesis applies just as much to attractor networks in other brain regions, and predict for example that for the spatial delayed response type of working memory task implemented in the dorsolateral prefrontal cortex (Goldman-Rakic, 1996; Rolls and Treves, 1998, Chapter 10) in monkey synaptic modification is required during the acquisition but not the later performance of the task.

REFERENCES

Abraham WC, Mason SE. 1988. Effects of the NMDA receptor/channel antagonists CPP and MK801 on hippocampal field potentials and long-term potentation in anesthetized rats. Brain Res 462:40–46.

Amit DJ. 1989. Modelling brain function. New York: Cambridge University Press.

Amit DJ. 1995. The Hebbian paradigm reintegrated: local reverberations as internal representations. Behav Brain Sci 18:617–657.

August DA, Levy WB. 1999. Temporal sequence compression by an integrate-and-fire model of hippocampal area CA3. J Comput Neurosci 6:71–90.

Baylis GC, Rolls ET. 1987. Responses of neurons in the inferior temporal cortex in short term and serial recognition memory tasks. Exp Brain Res 65:614–622.

Bolhuis JJ, Reid IC. 1992. Effects of intraventricular infusion of the N-methyl-D-aspartate (NMDA) receptor antagonist AP5 on spatial memory of rats in a radial arm maze. Behav Brain Res 47:151–157.

Bouffard JP, Jarrard LE. 1988. Acquisition of a complex place task in rats with selective ibotinate lesions of hippocampal formation: combined lesions of subiculum and entorhinal cortex versus hippocampus. Behav Neurosci 102:828–834.

Brown MW, Xiang JZ. 1998. Recognition memory: neuronal substrates of the judgement of prior occurrence. Prog Neurobiol 55:149–189.
Buckmaster PS, Schwartzkroin PA. 1994. Commentary: hippocampal mossy cell function: a speculative view. Hippocampus 4:393–402.
Cahusac PMB, Miyashita Y, Rolls ET. 1989. Responses of hippocampal formation neurons in the monkey related to delayed spatial response and object-place memory tasks. Behav Brain Res 33:229–240.
Caramanos Z, Shapiro ML. 1994. Spatial memory and N-methyl-D-aspartate receptor antagonists APV and MK-801: memory impairments depend on familiarity with the environment, drug dose, and training duration. Behav Neurosci 108:30–43.
Chiba AA, Kesner RP, Reynolds AM. 1994. Memory for spatial location as a function of temporal lag in rats: role of hippocampus and medial prefrontal cortex. Behav Neural Biol 61:123–131.
Colombo M, Fernandez T, Nakamura K, Gross CG. 1998. Functional differentiation along the anterior-posterior axis of the hippocampus in monkeys. J Neurophysiol 80:1002–1005.
Debanne D, Gahwiler BH, Thompson SM. 1998. Long-term synaptic plasticity between pairs of individual CA3 pyramidal cells in rat hippocampus slice cultures. J Physiol (Lond) 507:237–247.
DeCoteau WE, Stone AH, Huff LL, Kesner RP. 1999. Role of the hippocampus and medial caudate in memory for direction. Soc Neurosci Abstr 25:1624.
Eichenbaum H. 1997. Declarative memory: insights from cognitive neurobiology. Annu Rev Psychol 48:547–572.
Eichenbaum H, Dudchenko P, Wood E, Shapiro M, Tanila H. 1999. The hippocampus, memory, and place cells: is it spatial memory or a memory space? Neuron 23:209–226.
Gilbert ME, Mack CM. 1990. The NMDA antagonist, MK-801, suppresses long-term potentiation, kindling, and kindling-induced potentiation in the perforant path of the unanesthetized rat. Brain Res 519: 89–96.
Goldman-Rakic PS. 1996. The prefrontal landscape: implications of functional architecture for understanding human mentation and the central executive. Philos Trans R Soc Lond [Biol] 351:1445–1453.
Granger R, Whitson J, Larson J, Lynch G. 1994. Non-Hebbian properties of long-term potentiation enable high-capacity encoding of temporal sequences. Proc Natl Acad Sci USA 91:10104–10108.
Hampson RE, Heyser CJ, Deadwyler SA. 1993. Hippocampal cell firing correlates of delayed-match-to-sample performance in the rat. Behav Neurosci 107:715–739.
Handelmann GE, Olton DS. 1981. Spatial memory following damage to hippocampal CA3 pyramidal cells with kainic acid: impairment and recovery with preoperative training. Brain Res 217:41–58.
Hasselmo ME, Bower JM. 1993. Acetylcholine and memory. Trends Neurosci 16:218–222.
Hebb DO. 1949. The organization of behavior. New York: Wiley.
Holscher C, Schmidt WJ. 1994. Quinolinic acid lesion of the rat entorhinal cortex pars medialis produces selective amnesia in allocentric working memory (WM), but not in egocentric WM. Behav Brain Res 63:187–194.
Hopfield JJ. 1982. Neural networks and physical systems with emergent collective computational abilities. Proc Natl Acad Sci USA 79:2554–2558.
Jarrard LE. 1983. Selective hippocampal lesions and behavior: effects of kainic acid lesions on performance of place and cue tasks. Behav Neurosci 97:873–889.
Jensen O, Lisman JE. 1996. Novel lists of 7 +/− 2 known items can be reliably stored in an oscillatory short-term memory network: interaction with long-term memory. Learn Mem 3:257–263.
Kawabe K, Ichitani Y, Iwasaki T. 1998. Effects of intrahippocampal AP5 treatment on radial-arm maze performance in rats. Brain Res 781: 300–306.
Kesner RP. 1998. Neurobiological views of memory. In: Martinez JL, Kesner RP, editors. Neurobiology of learning and memory. San Diego: Academic Press.
Kesner RP, Dakis M. 1993. Phencyclidine disrupts acquisition and retention performance within a spatial continuous recognition memory task. Pharmacol Biochem Behav 44:419–424.
Kesner RP, Dakis M. 1997. Intrahippocampal injections of phencyclidine but not naloxone disrupt acquisition of a spatial continuous recognition memory task. Pharmacol Biochem Behav 56:97–101.
Kesner RP, Giles R. 1998. Neural circuit analysis of spatial working memory: role of pre- and parasubiculum, medial and lateral entorhinal cortex. Hippocampus 8:415–423.
Kesner RP, Crutcher K, Beers DR. 1988. Serial position curves for item (spatial location) information: role of the dorsal hippocampal formation and medial septum. Brain Res 454:219–226.
Kesner RP, Gilbert PE, Wallenstein GV. 2000. Testing neural network models of memory with behavioral experiments. Curr Opin Neurobiol 10:260–265.
Kohonen T. 1997. Associative memory: a system theoretical approach. New York: Springer.
Levy N, Horn D, Ruppin E. 1999. Associative memory in a multi-modular network. Neural Comput 11:1717–1737.
Lisman JE, Fellous JM, Wang XJ. 1998. A role for NMDA-receptor channels in working memory. Nat Neurosci 1:273–275.
Long JM, Kesner RP. 1995. Phencyclidine impairs temporal order memory for spatial locations in rats. Pharmacol Biochem Behav 52:645–648.
Long JM, Kesner RP. 1996. The effects of dorsal vs. ventral hippocampal, total hippocampal, and parietal cortex lesions on memory for allocentric distance in rats. Behav Neurosci 110:922–932.
Long JM, Kesner RP. 1998. The effects of hippocampal and parietal cortex lesions on memory for egocentric distance and spatial location information in rats. Behav Neurosci 112:480–495.
Marr D. 1971. Simple memory: a theory for archicortex. Philos Trans R Soc Lond [Biol] 262:23–81.
McNaughton BL, Morris RGM. 1989. Hippocampal synaptic enhancement and information storage within a distributed memory system. Trends Neurosci 10:408–415.
Morris RGM, Halliwell RF, Bowery N. 1989. Synaptic plasticity and learning II: do different kinds of plasticity underlie different kinds of learning? Neuropsychologia 27:41–59.
Muller RU, Stead M. 1996. Hippocampal place cells connected by Hebbian synapses can solve spatial problems. Hippocampus 6:709–719.
Nagahara AH, Otto T, Gallagher M. 1995. Entorhinal-perirhinal lesions impair performance of rats on two versions of place learning in the Morris water maze. Behav Neurosci 109:3–9.
Ohno M, Watanabe S. 1995. Intrahippocampal administration of (+)-SKF 10,047, a sigma ligand, reverses MK-801-induced impairment of working memory in rats. Brain Res 684:237–242.
Ohno M, Watanabe S. 1998. Enhanced N-methyl-D-aspartate function reverses working memory failure induced by blockade of group 1 metabotropic glutamate receptors in the rat hippocampus. Neurosci Lett 240:37–40.
Ohno M, Yamamoto T, Watanabe S. 1992. Effects of intrahippocampal injections of N-methyl-D-aspartate receptor antagonists and scopolamine on working and reference memory assessed in rats by a three-panel runway task. J Pharmacol Exp Ther 263:943–950.
Renart A, Parga N, Rolls ET. 1999a. Backprojections in the cerebral cortex: implications for memory storage. Neural Comput 11:1349–1388.
Renart A, Parga N, Rolls ET. 1999b. Associative memory properties of multiple cortical modules. Network 10:237–255.
Renart A, Parga N, Rolls ET. 2000. A recurrent model of the interaction between the prefrontal cortex and inferior temporal cortex in delay memory tasks. In: Solla SA, Leen TK, Mueller K-R, editors. Advances in neural information processing systems 12. Cambridge, MA: MIT Press. p 171–177 (NIPS 99).
Robertson RG, Rolls ET, Georges-François P. 1998. Spatial view cells in the primate hippocampus: effects of removal of view details. J Neurophysiol 79:1145–1156.

Robinson JK, Mao JB. 1997. Differential effects on delayed nonmatching-to-position in rats of microinjections of muscarinic receptor antagonist scopolamine or NMDA receptor antagonist MK-801 into the dorsal or ventral extent of the hippocampus. Brain Res 765:51–60.

Rolls ET. 1989. Functions of neuronal networks in the hippocampus and neocortex in memory. In: Byrne JH, Berry WO, editors. Neural models of plasticity: experimental and theoretical approaches. San Diego: Academic Press. p 240–265.

Rolls ET. 1999. Spatial view cells and the representation of place in the primate hippocampus. Hippocampus 9:467–480.

Rolls ET, Treves A. 1998. Neural networks and brain function. Oxford: Oxford University Press.

Rothblat LA, Vnek N, Gleason TC, Kromer LF. 1993. Role of the parahippocampal region in spatial and non-spatial memory; effects of parahippocampal lesions on rewarded alternation and concurrent object discrimination learning in the rat. Behav Brain Res 55:93–100.

Samsonovitch A, McNaughton BL. 1997. Path integration and cognitive mapping in a continuous attractor neural network model. J Neurosci 17:5900–5920.

Shapiro ML, Hetherington PA. 1993. A simple network model simulates hippocampal place fields: parametric analyses and physiological predictions. Behav Neurosci 107:34–50.

Shapiro ML, O'Connor C. 1992. N-methyl-D-aspartate receptor antagonist MK-801 and spatial memory representation: working memory is impaired in an unfamiliar environment but not in a familiar environment. Behav Neurosci 106:604–612.

Steele RJ, Morris RGM. 1999. Delay-dependent impairment of a matching-to-place task with chronic and intrahippocampal infusion of the NMDA-antagonist D-AP5. Hippocampus 9:118–136.

Stringer JL, Guyenet PG. 1983. Elimination of long-term potentiation in the hippocampus by phencyclidine and ketamine. Brain Res 258:159–164.

Stringer JL, Greenfield LJ, Hackett JT, Guyenet PG. 1983. Blockade of long-term potentiation by phencyclidine and σ opiates in the hippocampus in vivo and in vitro. Brain Res 280:127–138.

Suzuki WA, Miller EK, Desimone R. 1997. Object and place memory in the macaque entorhinal cortex. J Neurophysiol 78:106210–106281.

Tonkiss J, Rawlins JNP. 1991. The competitive NMDA antagonist AP5, but not the non-competitive antagonist MK801, induces a delay-related impairment in spatial working memory in rats. Exp Brain Res 85:349–358.

Treves A. 1993. Mean-field analysis of neuronal spike dynamics. Network 4:259–284.

Treves A, Rolls ET. 1991. What determines the capacity of autoassociative memories in the brain? Network 2:371–397.

Treves A, Rolls ET. 1994. A computational analysis of the role of the hippocampus in memory. Hippocampus 4:374–391.

Walker DL, Gold PE. 1994. Intrahippocampal administration of both the D- and the L-isomers of AP5 disrupt spontaneous alternation behavior and evoked potentials. Behav Neural Biol 62:151–162.

Wang XJ. 1999. Synaptic basis of cortical persistent activity: the importance of NMDA receptors to working memory. J Neurosci 19:9587–9603.

Wible CG, Findling RL, Shapiro M, Lang EJ, Crane S, Olton DS. 1986. Mnemonic correlates of unit activity in the hippocampus. Brain Res 399:97–110.