

An attractor network in the hippocampus: Theory and neurophysiology

Edmund T. Rolls¹

Department of Experimental Psychology, University of Oxford, Oxford OX1 3UD, England, United Kingdom

A quantitative computational theory of the operation of the CA3 system as an attractor or autoassociation network is described. Based on the proposal that CA3–CA3 autoassociative networks are important for episodic or event memory in which space is a component (place in rodents and spatial view in primates), it has been shown behaviorally that the CA3 supports spatial rapid one-trial learning and learning of arbitrary associations and pattern completion where space is a component. Consistent with the theory, single neurons in the primate CA3 respond to combinations of spatial view and object, and spatial view and reward. Furthermore, single CA3 neurons reflect the recall of a place from an object in a one-trial object-place event memory task. CA3 neurons also reflect in their firing a memory of spatial view that is retained and updated by idiothetic information to implement path integration when the spatial view is obscured. Based on the computational proposal that the dentate gyrus produces sparse representations by competitive learning and via the mossy fiber pathway forces new representations on the CA3 during learning (encoding), it has been shown behaviorally that the dentate gyrus supports spatial pattern separation during learning, and that the mossy fiber system to CA3 connections are involved in learning but not in recall. The perforant path input to CA3 is quantitatively appropriate to provide the cue for recall in CA3. The concept that the CA1 recodes information from CA3 and sets up associatively learned back-projections to neocortex to allow subsequent retrieval of information to neocortex provides a quantitative account of the large number of hippocampo–neocortical back-projections.

Part of the evidence that the CA3 is an autoassociation or attractor network comes from data that shows that the hippocampus is involved in spatial function and memory. Part of this evidence is based on the connectivity and associative synaptic plasticity of the hippocampus, and forms part of the foundation for a quantitative theory of how the CA3 operates as an attractor and in concert with other parts of the hippocampal system. Neurophysiological data (from primates, because of its relevance to understanding hippocampal function in humans) is provided that indicates that the hippocampus operates to implement event (e.g., object-and-place and reward-and-place) memory, the recall of event memories, and the holding on-line of neuronal activity when the spatial view details are obscured. These types of evidence are described and integrated in this paper.

Evidence from the effects of damage to the primate hippocampus

Damage to the hippocampus or to some of its connections such as the fornix in monkeys produces deficits in learning about the places of objects and about the places where responses should be made (Buckley and Gaffan 2000). For example, macaques and humans with damage to the hippocampal system or fornix are impaired in object-place memory tasks in which not only the objects seen, but where they were seen, must be remembered (Smith and Milner 1981; Gaffan and Saunders 1985; Parkinson et al. 1988; Gaffan 1994; Burgess et al. 2002; Crane and Milner 2005). Posterior parahippocampal lesions in macaques impair even a simple type of object-place learning in which the memory load is just one pair of trial-unique stimuli (Malkova and Mishkin 2003). (It is further predicted that a more difficult object-place learning task with non-trial-unique stimuli and with many object-place pairs would be impaired by neurotoxic hippocampal

lesions.) Furthermore, neurotoxic lesions that selectively damage the primate hippocampus impair spatial scene memory, tested by the ability to remember where in a scene to touch to obtain reward (Murray et al. 1998). Also, fornix lesions impair conditional left–right discrimination learning, in which the visual appearance of an object specifies whether a response is to be made to the left or the right (Rupniak and Gaffan 1987). A comparable deficit is found in humans (Petrides 1985). Fornix-sectioned monkeys are also impaired in learning on the basis of a spatial cue which object to choose (e.g., if two objects are on the left, choose object A, but if the two objects are on the right, choose object B) (Gaffan and Harrison 1989a). These tasks require associations to be made between objects and (spatial) responses. Monkeys with fornix damage are also impaired in using information about their place in an environment. For example, Gaffan and Harrison (1989b) found learning impairments when which of two or more objects the monkey had to choose depended on the position of the monkey in the room. Rats with hippocampal lesions are impaired in using environmental spatial cues to remember particular places (O'Keefe and Nadel 1978; Jarrod 1993; Cassaday and Rawlins 1997; Martin et al. 2000; Kesner et al. 2004), and to utilize spatial cues or bridge delays (Rawlins 1985; Kesner 1998; Kesner and Rolls 2001; Kesner et al. 2004; Rolls and Kesner 2006).

Many of these memory functions are important in event or episodic memory, in which the ability to remember what happened where on typically a single occasion is important. It is suggested below that an autoassociation memory implemented by the CA3 neurons enables event or episodic memories to be formed by enabling associations to be formed between spatial and other, including object, representations.

Information stored in the hippocampus will need to be retrieved and affect other parts of the brain in order to be used. The information about episodic events recalled from the hippocampus could be used to help form semantic memories (Rolls 1989b,c, 1990a; Treves and Rolls 1994). For example, remember-

¹Corresponding author.

E-mail Edmund.Rolls@psy.ox.ac.uk; fax 44-1865-310447.

Article is online at <http://www.learnmem.org/cgi/doi/10.1101/lm.631207>.

ing many particular journeys could help to build a geographic cognitive map in the neocortex. The hippocampus and neocortex would thus be complementary memory systems, with the hippocampus being used for rapid, "on the fly," unstructured storage of information involving activity potentially arriving from many areas of the neocortex, while the neocortex would gradually build and adjust the semantic representation on the basis of much accumulating information (Rolls 1989b; Treves and Rolls 1994; McClelland et al. 1995; Moscovitch et al. 2005). The theory described below shows how information could be retrieved within the hippocampus and how this retrieved information could enable the activity in neocortical areas that was present during the original storage of the episodic event to be reinstated, thus implementing recall, by using hippocampo-neocortical back-projections (see Fig. 1).

Hippocampal connectivity

Systems-level anatomy

The primate hippocampus receives inputs via the entorhinal cortex (area 28) and the highly developed parahippocampal gyrus (areas TF and TH) as well as the perirhinal cortex from the ends of many processing streams of the cerebral association cortex, including the visual and auditory temporal lobe association cortical areas, the prefrontal cortex, and the parietal cortex (Van Hoesen 1982; Amaral 1987; Amaral et al. 1992; Suzuki and Amaral 1994b; Witter et al. 2000b; Lavenex et al. 2004) (see Fig. 1).

The hippocampus is thus by its connections potentially able to associate together object and spatial representations. In addition, the entorhinal cortex receives inputs from the amygdala and the orbitofrontal cortex, which could provide reward-related information to the hippocampus (Suzuki and Amaral 1994a; Carmichael and Price 1995; Stefanacci et al. 1996; Pitkanen et al. 2002).

The primary output from the hippocampus to neocortex originates in CA1 and reaches there via multiple-stage back-projection pathways via the subiculum, entorhinal cortex, and parahippocampal structures (areas TF-TH) as well as prefrontal cortex (Van Hoesen 1982; Witter 1993; Delatour and Witter 2002; van Haeften et al. 2003) (see Fig. 1), although there are other outputs (Rolls and Kesner 2006).

Hippocampal circuitry

The internal connectivity of the hippocampus is illustrated in Figure 1 (see Amaral and Witter 1989; Storm-Mathiesen et al. 1990; Amaral 1993; Witter et al. 2000b; Naber et al. 2001; Lavenex et al. 2004).

Projections from the entorhinal cortex layer 2 reach the granule cells (of which there are 10^6 in the rat) in the dentate gyrus (DG), via the perforant path (pp) (Witter 1993). The granule cells project to CA3 cells via the mossy fibers (mf), which provide a sparse but possibly powerful connection to the 3×10^5 CA3 pyramidal cells in the rat. Each CA3 cell receives ~ 50 mossy fiber inputs, so that the sparseness of this connectivity is thus 0.005%. In contrast, there are many more—possibly weaker—

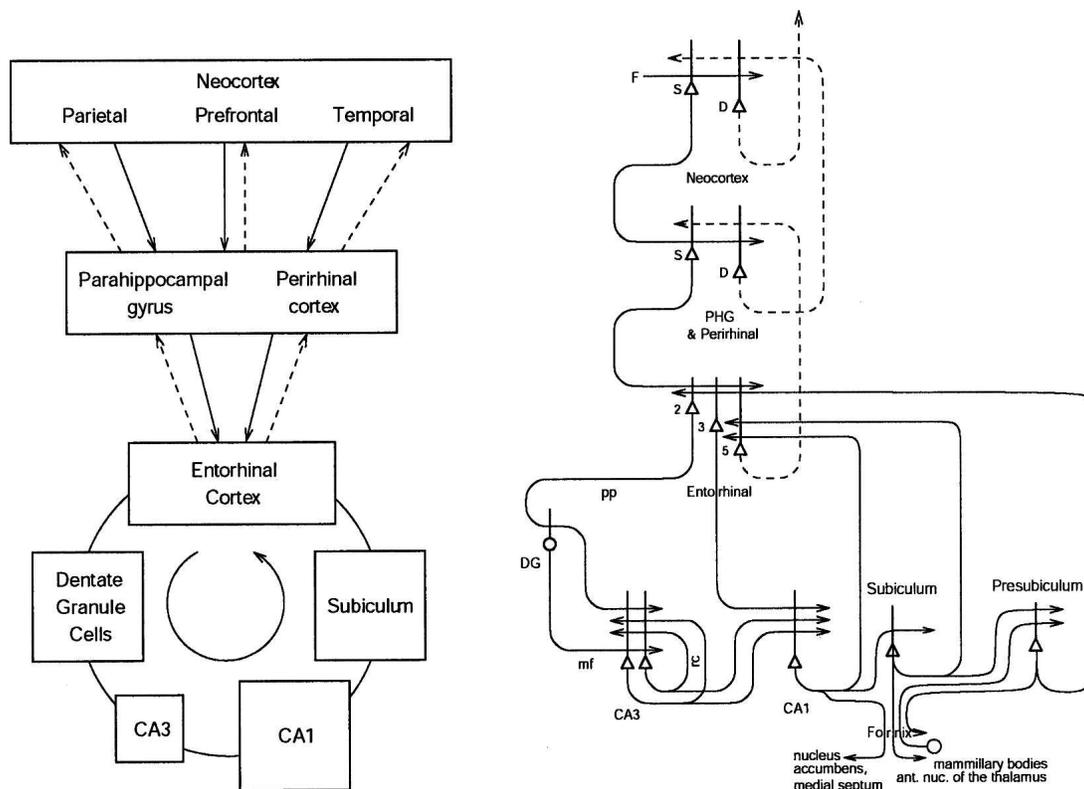


Figure 1. Forward connections (solid lines) from areas of cerebral association neocortex via the parahippocampal gyrus and perirhinal cortex, and entorhinal cortex, to the hippocampus; and back-projections (dashed lines) via the hippocampal CA1 pyramidal cells, subiculum, and parahippocampal gyrus to the neocortex. There is great convergence in the forward connections down to the single network implemented in the CA3 pyramidal cells and great divergence again in the back-projections. (Left) Block diagram; (right) more detailed representation of some of the principal excitatory neurons in the pathways. (D) Deep pyramidal cells; (DG) dentate granule cells; (F) forward inputs to areas of the association cortex from preceding cortical areas in the hierarchy; (mf) mossy fibers; (PHG) parahippocampal gyrus and perirhinal cortex; (pp) perforant path; (rc) recurrent collateral of the CA3 hippocampal pyramidal cells; (S) superficial pyramidal cells; (2) pyramidal cells in layer 2 of the entorhinal cortex; (3) pyramidal cells in layer 3 of the entorhinal cortex. The thick lines above the cell bodies represent the dendrites.

direct perforant path inputs also from layer 2 of the entorhinal cortex onto each CA3 cell, in the rat of the order of 4×10^3 . The largest number of synapses ($\sim 1.2 \times 10^4$ in the rat) on the dendrites of CA3 pyramidal cells is, however, provided by the (recurrent) axon collaterals of CA3 cells themselves (rc) (see Fig. 2). It is remarkable that the recurrent collaterals are distributed to other CA3 cells throughout the hippocampus (Amaral and Witter 1989, 1995; Amaral et al. 1990; Ishizuka et al. 1990), so that effectively the CA3 system provides a single network, with a connectivity of $\sim 2\%$ between the different CA3 neurons given that the connections are bilateral. The neurons that comprise CA3, in turn, project to CA1 neurons via the Schaffer collaterals. In addition, projections that terminate in the CA1 region originate in layer 3 of the entorhinal cortex (see Fig. 1).

CA3 as an autoassociation or attractor memory

Arbitrary associations formed in the memory, storage capacity, recall, and pattern completion

Many of the synapses in the hippocampus show associative modification as shown by long-term potentiation, and this synaptic modification appears to be involved in learning (see Morris 1989, 2003; Morris et al. 2003; Lynch 2004). On the basis of the evidence summarized above, Rolls (1987, 1989a,b,c, 1990a,b, 1991) and others (McNaughton and Morris 1987; Levy 1989; McNaughton 1991) have suggested that the CA3 stage acts as an autoassociation memory that enables episodic memories to be formed and stored in the CA3 network, and that subsequently the extensive recurrent collateral connectivity allows for the retrieval of a whole representation to be initiated by the activation of some small part of the same representation (the cue). The crucial synaptic modification for this is in the recurrent collateral synapses. (A description of the operation of autoassociative networks is provided by Hertz et al. 1991, Rolls and Treves 1998, Rolls and Deco 2002, and Rolls 2008.) Pioneering work on collateral connections and completion was performed by Marr (1971), although he did not discuss CA3 and did not produce a theory of how recall from the hippocampus to the neocortex could occur. The architecture of an autoassociation network is shown in Figure 3, and the learning rule for the change in the synaptic weight is as shown in Equation 1:

$$\delta w_{ij} = k \times r_i \times r_j' \quad (1)$$

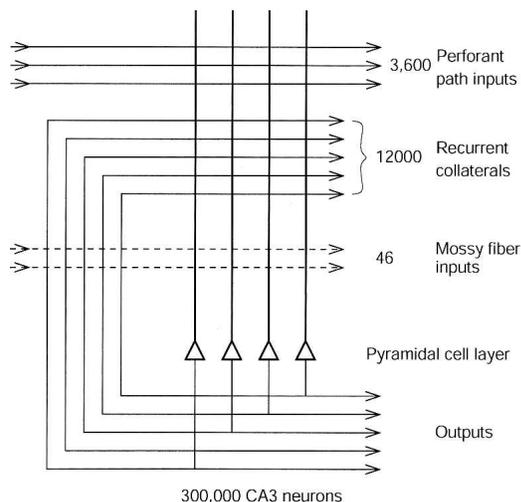


Figure 2. The numbers of connections from three different sources onto each CA3 cell from three different sources in the rat (after Treves and Rolls 1992; Rolls and Treves 1998) (Fig. 6).

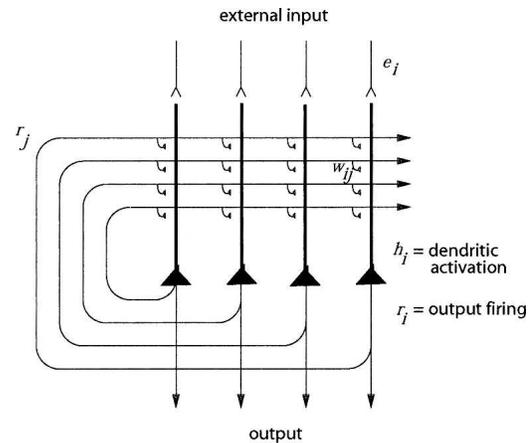


Figure 3. The architecture of an autoassociation or attractor neural network.

where k is a constant, r_i is the activation of the dendrite (the postsynaptic term), r_j' is the presynaptic firing rate, and w_{ij} is the synaptic weight.

The hypothesis is that because the CA3 operates effectively as a single network, it can allow arbitrary associations between inputs originating from very different parts of the cerebral cortex to be formed. These might involve associations between information originating in the temporal visual cortex about the presence of an object and information originating in the parietal cortex about where it is. I note that although there is some spatial gradient in the CA3 recurrent connections, so that the connectivity is not fully uniform (Ishizuka et al. 1990), nevertheless the network will still have the properties of a single interconnected autoassociation network allowing associations between arbitrary neurons to be formed, given the presence of many long-range connections that overlap from different CA3 cells.

Crucial issues include how many memories could be stored in this system (to determine whether the autoassociation hypothesis leads to a realistic estimate of the number of memories that the hippocampus could store); whether the whole of a memory could be completed from any part; whether the autoassociation memory can act as a short-term memory, for which the architecture is inherently suited; and whether the system could operate with spatial representations, which are essentially continuous because of the continuous nature of space. These and related issues are considered below and in more detail elsewhere (Rolls and Kesner 2006; Rolls 2008).

Storage capacity

We have performed quantitative analyses of the storage and retrieval processes in the CA3 network (Treves and Rolls 1991, 1992). We have extended previous formal models of autoassociative memory (see Amit 1989) by analyzing a network with graded response units, so as to represent more realistically the continuously variable rates at which neurons fire, and with incomplete connectivity (Treves 1990; Treves and Rolls 1991). We have found that, in general, the maximum number p_{\max} of firing patterns that can be (individually) retrieved is proportional to the number C^{RC} of (associatively) modifiable recurrent collateral synapses per cell, by a factor that increases roughly with the inverse of the sparseness a of the neuronal representation.² The neuronal

²Each memory is precisely defined in the theory: it is a set of firing rates of the population of neurons (which represent a memory) that can be stored and later retrieved, with retrieval being possible from a fraction of the originally stored set of neuronal firing rates.

population sparseness a of the representation can be measured by extending the binary notion of the proportion of neurons that are firing to any one stimulus or event as

$$a = \sum_{i=1,n} (r_i/N)^2 / \sum_{i=1,n} (r_i^2/N) \quad (2)$$

where r_i is the firing rate of the i th neuron in the set of N neurons. The sparseness ranges from $1/N$, when only one of the neurons responds to a particular stimulus (a local or grandmother cell representation), to a value of 1.0, attained when all the neurons are responding to a given stimulus. Approximately,

$$p_{\max} \cong \frac{C^{\text{RC}}}{a \ln(1/a)} k \quad (3)$$

where k is a factor that depends weakly on the detailed structure of the rate distribution, on the connectivity pattern, and so on, but is roughly in the order of 0.2–0.3 (Treves and Rolls 1991). For example, for $C^{\text{RC}} = 12,000$ and $a = 0.02$, p_{\max} is calculated to be $\sim 36,000$. This analysis emphasizes the utility of having a sparse representation in the hippocampus, for this enables many different memories to be stored.³

The sparseness estimates obtained by measuring immediate early gene changes (e.g., of arc; Chawla et al. 2005), which are effectively population sparsenesses, would be expected to depend greatly on the range of environments or stimuli in which this was measured. The population sparseness is measured by the firing rate distribution of the population of neurons while one stimulus, event, or place is being represented (Franco et al. 2007; Rolls 2008). If the environment was restricted to one stimulus, event, or place while the immediate early gene change was being produced, this could reflect the population sparseness. If the environment was changing, the measure from immediate early gene changes would be rather undefined and difficult to interpret, as all the populations of neurons activated in an undefined number of testing situations would be likely to be activated.

In order for most associative networks to store information efficiently, heterosynaptic Long Term Depression (as well as LTP) is required (Rolls and Treves 1990, 1998; Treves and Rolls 1991; Fazeli and Collingridge 1996; Rolls and Deco 2002; Rolls 2008). This type of LTD helps to remove the correlations between the training patterns that arise because the neurons have positive-only firing rates. The effect of the LTD can be to enable the effect of the mean presynaptic firing rate to be subtracted from the patterns (Rolls and Treves 1990, 1998; Treves and Rolls 1991; Rolls and Deco 2002; Rolls 2008). Simulations of attractor neural networks that are fully consistent with the analytic theory of the number of patterns that can be stored in attractor networks with sparse patterns (Treves and Rolls 1991; Rolls and Treves 1998) are provided by Simmen et al. (1996) and Rolls et al. (1997b).

We have also indicated how to estimate I , the total amount of information (in bits per synapse) that can be retrieved from the network. I is defined with respect to the information i_p (in bits per cell) contained in each stored firing pattern, by subtracting the amount i_i lost in retrieval and multiplying by p/C^{RC} :

$$I = \frac{p}{C^{\text{RC}}} (i_p - i_i) \quad (4)$$

³The sparseness a in Equation 3 is strictly the population sparseness (Treves and Rolls 1991; Franco et al. 2007). The population sparseness a^p would be measured by measuring the distribution of firing rates of all neurons to a single stimulus at a single time. The single neuron sparseness or selectivity a^s would be measured by the distribution of firing rates to a set of stimuli, which would take a long time. The selectivity or sparseness a^s of a single neuron measured across a set of stimuli often takes a similar value to the population sparseness a in the brain, and does so if the tuning profiles of the neurons to the set of stimuli are uncorrelated (Franco et al. 2007). These concepts are elucidated by Franco et al. (2007).

The maximal value I_{\max} of this quantity was found to be in several interesting cases ~ 0.2 – 0.3 bits per synapse, with only a mild dependency on parameters such as the sparseness of coding a (Treves and Rolls 1991).

We may then estimate (Treves and Rolls 1992) how much information has to be stored in each pattern for the network to efficiently exploit its information retrieval capacity I_{\max} . The estimate is expressed as a requirement on i_p :

$$i_p > a \ln(1/a) \quad (5)$$

As the information content of each stored pattern i_p depends on the storage process, we see how the retrieval capacity analysis, coupled with the notion that the system is organized so as to be an efficient memory device in a quantitative sense, leads to a constraint on the storage process. A more detailed analysis of how much information can be stored with different types of representation, and how many patterns can be correctly retrieved, leading to the conclusion that a sparse distributed representation is useful, especially for optimizing the number of memories that can be stored and correctly retrieved, is provided elsewhere (Treves and Rolls 1991, 1992; Rolls and Treves 1998; Rolls 2008).

Several points that arise are treated elsewhere (Rolls and Kesner 2006; Rolls 2008). Here I note that given that the memory capacity of the hippocampal CA3 system is limited, it is necessary to have some form of forgetting in this store, or other mechanism to ensure that its capacity is not exceeded. (Exceeding the capacity can lead to a loss of much of the information retrievable from the network.) Heterosynaptic LTD could help this “forgetting,” by enabling new memories to overwrite old memories (Rolls 1996a, 2008). The limited capacity of the CA3 system does also provide one of the arguments that some transfer of information from the hippocampus to neocortical memory stores may be useful (see Treves and Rolls 1994). Given its limited capacity, the hippocampus might be a useful store for only a limited period, which might be in the order of days, weeks, or months. This period may well depend on the acquisition rate of new episodic memories. If the animal were in a constant and limited environment, then as new information is not being added to the hippocampus, the representations in the hippocampus would remain stable and persistent. These hypotheses have clear experimental implications, both for recordings from single neurons and for the gradient of retrograde amnesia, both of which might be expected to depend on whether the environment is stable or frequently changing. They show that the conditions under which a gradient of retrograde amnesia might be demonstrable would be when large numbers of new memories are being acquired, not when only a few memories (few in the case of the hippocampus being less than a few hundred) are being learned.

Recall

A fundamental property of the autoassociation model of the CA3 recurrent collateral network is that the recall can be symmetric, that is, the whole of the memory can be retrieved from any part. For example, in an object-place autoassociation memory, an object could be recalled from a place retrieval cue, and vice versa. This is not the case with a pattern association network. If, for example, the CA3 activity represented a place/spatial view and perforant path inputs with associative synapses to CA3 neurons carried object information (consistent with evidence that the lateral perforant path [LPP] may reflect inputs from the perirhinal cortex connecting via the lateral entorhinal cortex; Hargreaves et al. 2005), then an object could recall a place, but a place could not recall an object.

A prediction of the theory is thus that the CA3 recurrent

collateral associative connections enable arbitrary associations to be formed between whatever is represented in the hippocampus, in that, for example, any place could be associated with any object, and in that the object could be recalled with a spatial recall cue, or the place with an object recall cue.

In one test of this hypothesis, Day et al. (2003) trained rats in a study phase to learn in one trial an association between two flavors of food and two spatial locations. During a recall test phase, they were presented with a flavor that served as a cue for the selection of the correct location. They found that injections of an NMDA blocker (AP5) or AMPA blocker (CNQX) to the dorsal hippocampus prior to the study phase impaired encoding, but injections of AP5 prior to the test phase did not impair the place recall, whereas injections of CNQX did impair the place recall. The interpretation is that somewhere in the hippocampus NMDA receptors are necessary for forming one-trial odor–place associations, and that recall can be performed without further involvement of NMDA receptors.

In a hippocampus subregion test of this, rats in a study phase are shown one object in one location and then a second object in another location. (There are 50 possible objects and 48 locations.) In the test phase, the rat is shown one object in the start box, and then after a 10-sec delay, must go to the correct location (choosing between two marked locations). CA3 lesions made after training in the task produced chance performance on this one-trial object–place recall task (Kesner 2007; Kesner et al. 2007). A control fixed visual conditional to place task with the same delay was not impaired, showing that it is recall after one-trial (or rapid) learning that is impaired. In the context of arbitrary associations between whatever is represented in CA3, the theory also predicts that cued place-object recall tasks and cued place-odor recall tasks should be impaired by CA3 lesions.

Evidence that the CA3 system is not necessarily required during recall in a reference memory spatial task, such as the water maze spatial navigation for a single spatial location task, is that CA3-lesioned rats are not impaired during recall of a previously learned water maze task (Brun et al. 2002; Florian and Roulet 2004). However, if completion from an incomplete cue is needed, then CA3 NMDA receptors are necessary (presumably to ensure satisfactory CA3–CA3 learning) even in a reference memory task (Nakazawa et al. 2002). Thus, the CA3 system appears to be especially needed in rapid, one-trial object-place recall, and when completion from an incomplete cue is required.

In a neurophysiological investigation described below of one-trial object-place learning followed by recall of the spatial position in which to respond when shown the object, Rolls and Xiang (2005) showed that some primate hippocampal (including CA3) neurons respond to an object cue with the spatial position in which the object had been shown earlier in the trial. Thus, some hippocampal neurons appear to reflect spatial recall given an object recall cue. It will be interesting in future research to investigate the recall of an object given a place cue to initiate the recall.

Completion

Another fundamental property of autoassociation memories is that the recall can be complete even from a small fragment. Thus, it is a prediction that when an incomplete retrieval cue is given, CA3 may be especially important in the retrieval process. Tests of this prediction of a role for CA3 in pattern completion have been performed, as follows.

Rats were tested on a cheese board with a black curtain with four extramaze cues surrounding the apparatus. (The cheese board is like a dry land water maze with 177 holes on a 119-cm diameter board.) Rats were trained to move a sample phase object covering a food well that could appear in one of five possible

spatial locations. During the test phase of the task, following a 30-sec delay, the animal needs to find the same food well in order to receive reinforcement with the object now removed. After reaching stable performance in terms of accuracy to find the correct location, rats received lesions in CA3. During post-surgery testing, four extramaze cues were always available during the sample phase. However, during the test phase zero, one, two, or three cues were removed in different combinations. The results indicate that controls performed well on the task regardless of the availability of one, two, three, or all cues, suggesting intact spatial pattern completion. Following the CA3 lesion, however, there was an impairment in accuracy compared to the controls especially when only one or two cues were available, suggesting impairment in spatial pattern completion in CA3-lesioned rats (Gold and Kesner 2005). A useful aspect of this task is that the test for the ability to remember a spatial location learned in one presentation can be tested with a varying number of available cues, and many times in which the locations vary, to allow for accurate measurement of pattern completion ability when the information stored on the single presentation must be recalled.

In another study, Nakazawa et al. (2002) trained CA3 NMDA receptor-knockout mice in an analogous task, using the water maze. When the animals were required to perform the task in an environment where some of the familiar cues were removed, they were impaired in performing the task. The result suggests that the NMDA receptor-dependent synaptic plasticity mechanisms in CA3 are critical to perform the pattern completion process in the hippocampus.

Continuous, spatial, patterns and CA3 representations

The fact that spatial patterns, which imply continuous representations of space, are represented in the hippocampus has led to the application of continuous attractor models to help understand hippocampal function. This has been necessary because space is inherently continuous, because the firing of place and spatial view cells is approximately Gaussian as a function of the distance away from the preferred spatial location, because these cells have spatially overlapping fields, and because the theory is that these cells in CA3 are connected by Hebb-modifiable synapses. This specification would inherently lead the system to operate as a continuous attractor network. Continuous attractor network models have been studied by Amari (1977), Zhang (1996), Samsonovich and McNaughton (1997), Battaglia and Treves (1998a), Taylor (1999), Stringer and Rolls (2002), Stringer et al. (2002a,b, 2004), and Rolls and Stringer (2005), and are described next, and elsewhere in more detail (Rolls and Deco 2002; Rolls 2008).

A “Continuous Attractor” neural network (CANN) can maintain the firing of its neurons to represent any location along a continuous physical dimension such as spatial position, head direction, and so on. It uses excitatory recurrent collateral connections between the neurons (as are present in CA3) to reflect the distance between the neurons in the state space of the animal (e.g., place or head direction). These networks can maintain the bubble of neural activity constant for long periods wherever it is started to represent the current state (head direction, position, spatial view, etc.) of the animal and are likely to be involved in many aspects of spatial processing and memory, including spatial vision. Global inhibition is used to keep the number of neurons in a bubble or packet of actively firing neurons relatively constant, and to help to ensure that there is only one activity packet. Continuous attractor networks can be thought of as very similar to autoassociation or discrete attractor networks (see Rolls and Deco 2002) and have the same architecture, as illustrated in Figure 3. The main difference is that the patterns stored in a

CANN are continuous patterns, with each neuron having broadly tuned firing that decreases with, for example, a Gaussian function as the distance from the optimal firing location of the cell is varied, and with different neurons having tuning that overlaps throughout the space. Such tuning is illustrated in Figure 4. For comparison, autoassociation networks normally have discrete (separate) patterns (each pattern implemented by the firing of a particular subset of the neurons), with no continuous distribution of the patterns throughout the space (see Fig. 4). A consequent difference is that the CANN can maintain its firing at any location in the trained continuous space, whereas a discrete attractor or autoassociation network moves its population of active neurons toward one of the previously learned attractor states and thus implements the recall of a particular previously learned pattern from an incomplete or noisy (distorted) version of one of the previously learned patterns.

The energy landscape of a discrete attractor network (Rolls and Deco 2002; Rolls 2008) has separate energy minima, each one of which corresponds to a learned pattern, whereas the energy landscape of a continuous attractor network is flat, so that the activity packet remains stable with continuous firing wherever it is started in the state space. (The “state space” refers in this case to the set of possible spatial states of the animal in its environment, e.g., the set of possible places in a room, or a set of spatial views of scenes.)

So far we have said that the neurons in the continuous attractor network are connected to each other by synaptic weights

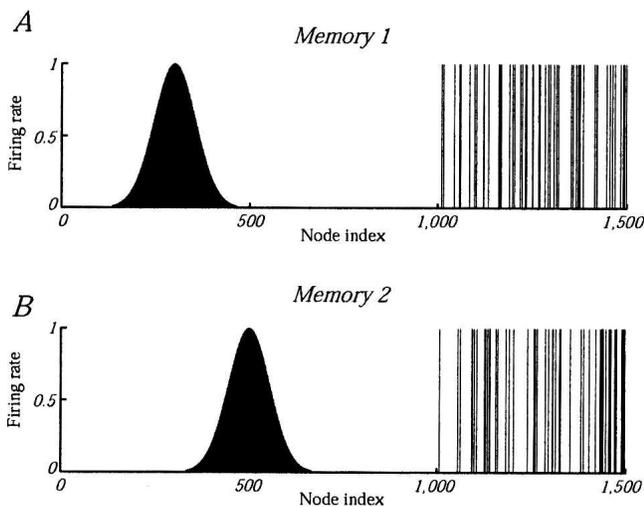


Figure 4. The types of firing patterns stored in continuous attractor networks are illustrated for the patterns stored on neurons 1–1000 for Memory 1 (when the firing is that produced when the spatial state represented is that for location 300) and for Memory 2 (when the firing is that produced when the spatial state represented is that for location 500). The continuous nature of the spatial representation results from the fact that each neuron has a Gaussian firing rate that peaks at its optimal location. This particular mixed network also contains discrete representations that consist of discrete subsets of active binary firing rate neurons in the range 1001–1500. The firing of these latter neurons can be thought of as representing the discrete events that occur at the location. Continuous attractor networks by definition contain only continuous representations, but this particular network can store mixed continuous and discrete representations and is illustrated to show the difference of the firing patterns normally stored in separate continuous attractor and discrete attractor networks. For this particular mixed network, during learning, Memory 1 is stored in the synaptic weights, then Memory 2, and so on, and each memory contains part that is continuously distributed to represent physical space and part that represents a discrete event or object.

w_{ij} that are a simple function, for example, Gaussian, of the distance between the states of the agent in the physical world (e.g., head directions, spatial views) represented by the neurons. In many simulations, the weights are set by formula to have weights with these appropriate Gaussian values. However, Stringer et al. (2002b) showed how the appropriate weights could be set up by learning. They started with the fact that since the neurons have broad tuning that may be Gaussian in shape, nearby neurons in the state space will have overlapping spatial fields, and will thus be coactive to a degree that depends on the distance between them. They postulated that therefore the synaptic weights could be set up by associative learning based on the coactivity of the Gaussian profile neurons produced by external stimuli as the animal moved in the state space. Stringer et al. (2002b) showed that after training at all positions in the state space, the synaptic connections develop strengths that are an almost Gaussian function of the distance between the cells in the state space, and this applies not only to head direction cells (Stringer et al. 2002b; Stringer and Rolls 2006), but also to rat place cells (Stringer et al. 2002a,b) and primate spatial view cells (Stringer et al. 2004, 2005; Rolls and Stringer 2005).

Combined continuous and discrete memory representations in the same (e.g., CA3) network, and episodic memory

Space is continuous, and object representations are discrete. If these representations are to be combined in, for example, an object-place memory, then we need to understand the operation of networks that combine these representations. It has now been shown that attractor networks can store both continuous patterns and discrete patterns (as illustrated in Fig. 4) and can thus be used to store, for example, the location in (continuous, physical) space (e.g., the place “out there” in a room represented by spatial view cells) where an object (a discrete item) is present (Rolls et al. 2002).

The capacity of a continuous attractor network, and multiple charts

If spatial representations are stored in the hippocampus, the important issue arises in terms of understanding memories that include a spatial component or context of how many such spatial representations could be stored in a continuous attractor network. The very interesting result is that because there are, in general, low correlations between the representations of places in different maps or charts (where each map or chart might be of one room or locale), very many different maps can be simultaneously stored in a continuous attractor network (Battaglia and Treves 1998a).

Idiothetic update by path integration

We have considered how spatial representations could be stored in continuous attractor networks, and how the activity can be maintained at any location in the state space in a form of short-term memory when the external (e.g., visual) input is removed. However, many networks with spatial representations in the brain can be updated by internal, self-motion (i.e., idiothetic) cues even when there is no external (e.g., visual) input. The way in which path integration could be implemented in recurrent networks such as the CA3 system in the hippocampus or in related systems is described next.

Single-cell recording studies have shown that some neurons represent the current position along a continuous physical di-

mension or space even when no inputs are available, for example, in darkness. Examples include neurons that represent the positions of the eyes (i.e., eye direction with respect to the head), the place where the animal is looking in space, head direction, and the place where the animal is located. In particular, examples of such classes of cells include head direction cells in rats (Ranck 1985; Taube et al. 1990, 1996; Muller et al. 1996) and primates (Robertson et al. 1999), which respond maximally when the animal's head is facing in a particular preferred direction; place cells in rats (O'Keefe and Dostrovsky 1971; McNaughton et al. 1983; O'Keefe 1984; Muller et al. 1991; Markus et al. 1995) that fire maximally when the animal is in a particular location; and spatial view cells in primates that respond when the monkey is looking toward a particular location in space (Rolls et al. 1997a; Robertson et al. 1998; Georges-François et al. 1999).

One approach to simulating the movement of an activity packet produced by idiothetic cues (which is a form of path integration whereby the current location is calculated from recent movements) is to use a look-up table that stores (taking head direction cells as an example), for every possible head direction and head rotational velocity input generated by the vestibular system, the corresponding new head direction (Samsonovich and McNaughton 1997). An analogous approach has been described for entorhinal cortex grid cells (McNaughton et al. 2006). Another approach involves modulating the strengths of the recurrent synaptic weights in the continuous attractor on one but not the other side of a currently represented position, so that the stable position of the packet of activity, which requires symmetric connections in different directions from each node, is lost, and the packet moves in the direction of the temporarily increased weights, although no possible biological implementation was proposed or how the appropriate dynamic synaptic weight changes might be achieved (Zhang 1996). Another mechanism (for head direction cells) (Skaggs et al. 1995) relies on a set of cells, termed (head velocity) rotation cells, which are coactivated by head direction cells and vestibular cells and drive the activity of the attractor network by anatomically distinct connections for clockwise and counterclockwise rotation cells, in what is effectively a look-up table. However, these proposals did not show how the synaptic weights for this path integration could be achieved by a biologically plausible learning process.

Stringer et al. (2002b) introduced a proposal with more biological plausibility about how the synaptic connections from idiothetic inputs to a continuous attractor network can be learned by a self-organizing learning process. The mechanism associates a short-term memory trace of the firing of the neurons in the attractor network reflecting recent movements in the state space (e.g., of places), with an idiothetic velocity of movement input (see Fig. 5). This has been applied to head direction cells (Stringer et al. 2002b; Stringer and Rolls 2006), rat place cells (Stringer et al. 2002a,b), and primate spatial view cells (Stringer et al. 2004, 2005; Rolls and Stringer 2005). These attractor networks provide a basis for understanding cognitive maps and how they are updated by learning and by self-motion. The implication is that to the extent that path integration of place or spatial view representations is performed within the hippocampus itself, then the CA3 system is the most likely part of the hippocampus to be involved in this because it has the appropriate recurrent collateral connections. Consistent with this, Whishaw and colleagues (Maaswinkel et al. 1999; Whishaw et al. 2001; Wallace and Whishaw 2003) have shown that path integration is impaired by hippocampal lesions. Path integration of head direction is reflected in the firing of neurons in the presubiculum, and mechanisms outside the hippocampus probably implement path integration for head direction (Stringer and Rolls 2006; Rolls 2008).

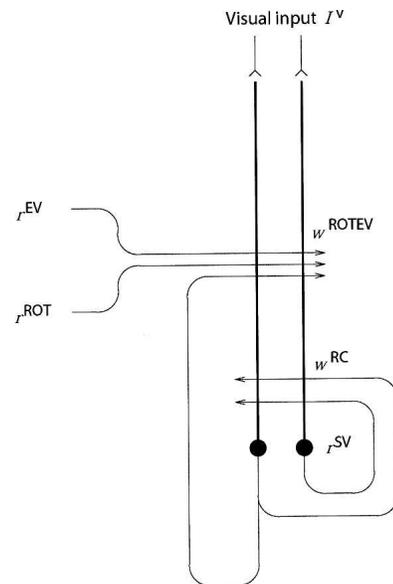


Figure 5. Neural network architecture for two-dimensional continuous attractor models of spatial view cells incorporating self-organizing idiothetic update to implement path integration. The recurrent network of spatial view cells with firing rates r^{SV} receives external inputs from three sources: (1) spatial view inputs derived from the visual scene with firing rates I^V ; (2) a population of clockwise and anticlockwise head rotation velocity cells with firing rates r^{ROT} ; and (3) a population of eye velocity cells with firing rates r^{EV} . The recurrent collateral synapses w^{RC} form a conventional continuous attractor network for spatial view and represent CA3 neurons in the model. The Sigma-Pi synapses w^{ROTEV} could be replaced by neurons that respond to combinations of these inputs (Stringer and Rolls 2006). The principle is that during training, the combination of head and eye velocity signals, and a memory trace delayed version of the spatial view firing, is associated with the current spatial view represented in the post-synaptic term of w^{ROTEV} . The difference between the current and the previous spatial view gives the change of spatial view, and this is included in the association with the two velocity signals and the current spatial view. After training, just the two velocity signals and the current spatial view allow path integration to the next spatial view to be produced (after Stringer et al. 2005).

The dynamics of the recurrent network

The analysis described above of the capacity of a recurrent network such as the CA3 considered steady-state conditions of the firing rates of the neurons. The question arises of how quickly the recurrent network would settle into its final state. With reference to the CA3 network, how long does it take before a pattern of activity, originally evoked in CA3 by afferent inputs, becomes influenced by the activation of recurrent collaterals? In a more general context, recurrent collaterals between the pyramidal cells are an important feature of the connectivity of the cerebral neocortex. How long would it take these collaterals to contribute fully to the activity of cortical cells? If these settling processes took in the order of hundreds of milliseconds, they would be much too slow to contribute usefully to cortical activity, whether in the hippocampus or the neocortex (Rolls 1992, 2003; Panzeri et al. 2001; Rolls and Deco 2002).

It has been shown that if the neurons are treated not as McCulloch-Pitts neurons, which are simply "updated" at each iteration, or cycles of time steps (and assume the active state if the threshold is exceeded), but instead are analyzed and modeled as "integrate-and-fire" neurons in real continuous time, then the network can effectively "relax" into its recall state very rapidly, in one or two time constants of the synapses (Treves 1993; Battaglia and Treves 1998b; Rolls and Treves 1998; Rolls and Deco 2002). This corresponds to perhaps 20 msec in the brain. One factor in

this rapid dynamics of autoassociative networks with brain-like “integrate-and-fire” membrane and synaptic properties is that with some spontaneous activity, some of the neurons in the network are close to threshold already before the recall cue is applied, and hence some of the neurons are very quickly pushed by the recall cue into firing, so that information starts to be exchanged very rapidly (within 1–2 msec of brain time) through the modified synapses by the neurons in the network. The progressive exchange of information starting early on within what would otherwise be thought of as an iteration period (of perhaps 20 msec, corresponding to a neuronal firing rate of 50 spikes/sec) is the mechanism accounting for rapid recall in an autoassociative neuronal network made biologically realistic in this way. Further analysis of the fast dynamics of these networks if they are implemented in a biologically plausible way with “integrate-and-fire” neurons is provided in Section 7.7 of Rolls and Deco (2002), in Appendix A5 of Rolls and Treves (1998), by Treves (1993), and by Panzeri et al. (2001).

Mossy fiber inputs to the CA3 cells

We hypothesize that the mossy fiber inputs force efficient information storage by virtue of their strong and sparse influence on the CA3 cell firing rates (Rolls 1987, Rolls 1989b,c; Treves and Rolls 1992). (The strong effects likely to be mediated by the mossy fibers were also emphasized by McNaughton and Morris 1987 and McNaughton and Nadel 1990.) We hypothesize that the mossy fiber input appears to be particularly appropriate in several ways. First of all, the fact that mossy fiber synapses are large and located very close to the soma makes them relatively powerful in activating the postsynaptic cell. (This should not be taken to imply that a CA3 cell can be fired by a single mossy fiber EPSP.) Second, the firing activity of dentate granule cells appears to be very sparse (Jung and McNaughton 1993), and this, together with the small number of connections on each CA3 cell, produces a sparse signal, which can then be transformed into an even sparser firing activity in CA3 by a threshold effect.⁴ Third, nonassociative plasticity of mossy fibers (see Brown et al. 1989, 1990) might have a useful effect in enhancing the signal-to-noise ratio, in that a consistently firing mossy fiber would produce nonlinearly amplified currents in the postsynaptic cell, which would not happen with an occasionally firing fiber (Treves and Rolls 1992). This plasticity, and also learning in the dentate, would also have the effect that similar fragments of each episode (e.g., the same environmental location) recurring on subsequent occasions would be more likely to activate the same population of CA3 cells, which would have potential advantages in terms of economy of use of the CA3 cells in different memories and in making some link between different episodic memories with a common feature, such as the same location in space. Fourth, with only a few, and powerful, active mossy fiber inputs to each CA3 cell, setting a given sparseness of the representation provided by CA3 cells would be simplified, for the EPSPs produced by the mossy fibers would be Poisson-distributed with large

membrane potential differences for each active mossy fiber. Setting the average firing rate of the dentate granule cells would effectively set the sparseness of the CA3 representation, without great precision being required in the threshold setting of the CA3 cells (Rolls et al. 1997b). Part of what is achieved by the mossy fiber input may be setting the sparseness of the CA3 cells correctly, which, as shown above, is very important in an autoassociative memory store. Fifth, the non-associative and sparse connectivity properties of the mossy fiber connections to CA3 cells may be appropriate for an episodic memory system that can learn very fast, in one trial. The hypothesis is that the sparse connectivity would help arbitrary relatively uncorrelated sets of CA3 neurons to be activated for even somewhat similar input patterns without the need for any learning of how best to separate the patterns, which in a self-organizing competitive network would take several repetitions (at least) of the set of patterns. The mossy fiber solution may thus be adaptive in a system that must learn in one trial, and for which the CA3 recurrent collateral learning requires uncorrelated sets of CA3 cells to be allocated for each (one-trial) episodic memory. The hypothesis is that the mossy fiber sparse connectivity solution performs the appropriate function without the mossy fiber system having to learn by repeated presentations of how best to separate a set of training patterns. The perforant path input would, the quantitative analysis shows, not produce a pattern of firing in CA3 that contains sufficient information for learning (Treves and Rolls 1992).

On the basis of these points, we predict that the mossy fibers may be necessary for new learning in the hippocampus but may not be necessary for recall of existing memories from the hippocampus. Experimental evidence consistent with this prediction about the role of the mossy fibers in learning has been found in rats with disruption of the dentate granule cells (Lassalle et al. 2000).

As acetylcholine turns down the efficacy of the recurrent collateral synapses between CA3 neurons (Hasselmo et al. 1995), then cholinergic activation also might help to allow external inputs rather than the internal recurrent collateral inputs to dominate the firing of the CA3 neurons during learning, as the current theory proposes. If cholinergic activation at the same time facilitated LTP in the recurrent collaterals (as it appears to in the neocortex), then cholinergic activation could have a useful double role in facilitating new learning at times of behavioral activation, when presumably it may be particularly relevant to allocate some of the limited memory capacity to new memories.

Perforant path inputs to CA3 cells

By calculating the amount of information that would end up being carried by a CA3 firing pattern produced solely by the perforant path input and by the effect of the recurrent connections, we have been able to show (Treves and Rolls 1992) that an input of the perforant path type, alone, is unable to direct efficient information storage. Such an input is too weak, it turns out, to drive the firing of the cells, as the “dynamics” of the network is dominated by the randomizing effect of the recurrent collaterals. This is the manifestation, in the CA3 network, of a general problem affecting storage (i.e., learning) in all autoassociative memories. The problem arises when the system is considered to be activated by a set of input axons making synaptic connections that have to compete with the recurrent connections, rather than having the firing rates of the neurons artificially clamped into a prescribed pattern.

An autoassociative memory network needs afferent inputs also in the other mode of operation, that is, when it retrieves a previously stored pattern of activity. We have shown (Treves and Rolls 1992) that the requirements on the organization of the

⁴For example, if only one granule cell in 100 were active in the dentate gyrus, and each CA3 cell received a connection from 50 randomly placed granule cells, then the number of active mossy fiber inputs received by CA3 cells would follow a Poisson distribution of average $50/100 = 1/2$, that is, 60% of the cells would not receive any active input, 30% would receive only one, 7.5% two, little more than 1% would receive three, and so on. (It is easy to show from the properties of the Poisson distribution and our definition of sparseness, that the sparseness of the mossy fiber signal as seen by a CA3 cell would be $x/(1+x)$, with $x = C^{MF} a_{DG}$, assuming equal strengths for all mossy fiber synapses.) If three mossy fiber inputs were required to fire a CA3 cell and these were the only inputs available, we see that the activity in CA3 would be roughly as sparse, in the example, as in the dentate gyrus. C^{MF} is the number of mossy fiber connections to a CA3 neuron, and a_{DG} is the sparseness of the representation in the dentate granule cells.

afferents are in this case very different, implying the necessity of a second, separate input system, which we have identified with the perforant path to CA3. In brief, the argument is based on the notion that the cue available to initiate retrieval might be rather small, that is, the distribution of activity on the afferent axons might carry a small correlation, $q \ll 1$, with the activity distribution present during learning. In order not to lose this small correlation altogether, but rather transform it into an input current in the CA3 cells that carries a sizable signal—which can then initiate the retrieval of the full pattern by the recurrent collaterals—one needs a large number of associatively modifiable synapses. This is expressed by the formulas that give the specific signal S produced by sets of associatively modifiable synapses, or by nonassociatively modifiable synapses: if C^{AFF} is the number of afferents per cell,

$$S_{\text{ASS}} \sim \frac{\sqrt{C^{\text{AFF}}}}{\sqrt{p}} q \quad S_{\text{NONASS}} \sim \frac{1}{\sqrt{C^{\text{AFF}}}} q. \quad (6)$$

Associatively modifiable synapses are therefore needed, and are needed in a number C^{AFF} of the same order as the number of concurrently stored patterns p , so that small cues can be effective; whereas nonassociatively modifiable synapses—or even more so, nonmodifiable ones—produce very small signals, which decrease in size the larger the number of synapses. In contrast with the storage process, the average strength of these synapses does not now play a crucial role. This suggests that the perforant path system is the one involved in relaying the cues that initiate retrieval in the CA3 autoassociation network.

Dentate granule cells

The theory is that the dentate granule cell stage of hippocampal processing that precedes the CA3 stage acts in several ways to produce during learning the sparse yet efficient (i.e., nonredundant) representation in CA3 neurons that is required for the autoassociation implemented by CA3 to perform well (Rolls 1989b; Treves and Rolls 1992; Rolls and Kesner 2006; Rolls et al. 2006).

The first way is that the perforant path–dentate granule cell system with its Hebb-like modifiability is suggested to act as a competitive learning network to remove redundancy from the inputs, producing a more orthogonal, sparse, and categorized set of outputs (Rolls 1987, 1989a,b,c, 1990a,b, 2008; Rolls and Treves 1998). (Competitive networks are described elsewhere, see Hertz et al. 1991; Rolls and Treves 1998; Rolls and Deco 2002; Rolls 2008.) The nonlinearity in the NMDA receptors may help the operation of such a competitive net, for it ensures that only the most active neurons left after the competitive feedback inhibition have synapses that become modified and thus learn to respond to that input (Rolls 1989a). Because of the feedback inhibition, the competitive process may result in a relatively constant number of strongly active dentate neurons relatively independently of the number of active perforant path inputs to the dentate cells. The operation of the dentate granule cell system as a competitive network may also be facilitated by a Hebb rule of the form:

$$\delta w_{ij} = k \times r_i(r_j' - w_{ij}) \quad (7)$$

where k is a constant, r_i is the activation of the dendrite (the postsynaptic term), r_j' is the presynaptic firing rate, w_{ij} is the synaptic weight, and r_j' and w_{ij} are in appropriate units (Rolls 1989a). Incorporation of a rule such as this, which implies heterosynaptic long-term depression as well as long-term potentiation (see Levy and Desmond 1985; Levy et al. 1990), makes the sum of the synaptic weights on each neuron remain roughly

constant during learning (cf. Oja 1982; see Rolls 1989a, 2008; Rolls and Treves 1998; Rolls and Deco 2002).

This functionality could be used to help build hippocampal place cells in rats from the grid cells present in the medial entorhinal cortex (Hafting et al. 2005). Each grid cell responds to a set of places in a spatial environment, with the places to which a cell responds set out in a regular grid. Different grid cells have different phases (positional offsets) and grid spacings (or frequencies) (Hafting et al. 2005). We (Rolls et al. 2006) have simulated the dentate granule cells as a system that receives as inputs the activity of a population of entorhinal cortex grid cells as the animal traverses a spatial environment and have shown that the competitive net builds dentate-like place cells from such entorhinal grid cell inputs. This occurs because the firing states of entorhinal cortex cells that are active at the same time when the animal is in one place become associated together by the learning in the competitive net, yet each dentate cell represents primarily one place because the dentate representation is kept sparse, thus helping to implement symmetry-breaking (Rolls et al. 2006).

The second way is also a result of the competitive learning hypothesized to be implemented by the dentate granule cells (Rolls 1987, 1989a,b,c, 1990a,b, 1994). It is proposed that this allows overlapping (or very similar) inputs to the hippocampus to be separated, in the following way (see also Rolls 1996b). Consider three patterns B, W, and BW, where BW is a linear combination of B and W. (To make the example very concrete, we could consider binary patterns where B = 10, W = 01, and BW = 11.) Then the memory system is required to associate B with reward, W with reward, but BW with punishment. Without the hippocampus, rats might have more difficulty in solving such problems, particularly when they are spatial, for the dentate/CA3 system in rodents is characterized by being implicated in spatial memory. However, it is a property of competitive neuronal networks that they can separate such overlapping patterns, as has been shown elsewhere (Rolls 1989a, 2008; Rolls and Treves 1998); normalization of synaptic weight vectors is required for this property. It is thus an important part of hippocampal neuronal network architecture that there is a competitive network that precedes the CA3 autoassociation system. Without the dentate gyrus, if a conventional autoassociation network were presented with the mixture BW having learned B and W separately, then the autoassociation network would produce a mixed output state and would therefore be incapable of storing separate memories for B, W, and BW. It is suggested therefore that operation of the dentate granule cells as a competitive network is one of the powerful computational features of the hippocampus, and that could enable it to help solve spatial pattern separation tasks (Rolls and Kesner 2006). Performing this spatial pattern separation before the CA3 stage allows the CA3 to operate as a different type of network, which associates together the spatial representation, made relatively non-overlapping with other spatial representations, with other representations, for example, of objects, and for which a different, autoassociative, architecture is needed (Rolls 2008).

This computational hypothesis and its predictions have been tested. Rats with dentate gyrus lesions are impaired at a metric spatial pattern separation task (Gilbert et al. 2001; Goodrich-Hunsaker et al. 2005; Kesner 2007). The recoding of grid cells in the entorhinal cortex (Hafting et al. 2005) into small place field cells in the dentate granule cells that has been modeled (Rolls et al. 2006) can also be considered to be a case in which overlapping inputs must be recoded so that different spatial components can be treated differently. I note that Sutherland and Rudy's configural learning hypothesis was similar but was not tested with spatial pattern separation. Instead, when tested with, for example, tone and light combinations, it was not con-

sistently found that the hippocampus was important (Sutherland and Rudy 1991; O'Reilly and Rudy 2001). I suggest that application of the configural concept, but applied to spatial pattern separation, may capture part of what the dentate gyrus acting as a competitive network could perform, particularly when a large number of such overlapping spatial memories must be stored and retrieved.

The third way in which the dentate gyrus is hypothesized to contribute to the sparse and relatively orthogonal representations in CA3 arises because of the very low contact probability in the mossy fiber–CA3 connections, and is described below and by Treves and Rolls (1992).

A fourth way is that as suggested and explained above, the dentate granule cell–mossy fiber input to the CA3 cells may be powerful and its use particularly during learning would be efficient in forcing a new pattern of firing onto the CA3 cells during learning.

In the ways just described, the dentate granule cells could be particularly important in helping to build and prepare spatial representations for the CA3 network. The actual representation of space in the primate hippocampus includes a representation of spatial view, whereas in the rat hippocampus it is of the place where the rat is. The representation in the rat may be related to the fact that with a much less developed visual system than the primate, the rat's representation of space may be defined more by the olfactory and tactile as well as distant visual cues present, and may thus tend to reflect the place where the rat is. However, the spatial representations in the rat and primate could arise from essentially the same computational process as follows (Rolls 1999; de Araujo et al. 2001). The starting assumption is that in both the rat and the primate, the dentate granule cells (and the CA3 and CA1 pyramidal cells) respond to combinations of the inputs received. In the case of the primate, a combination of visual features in the environment will, because of the fovea providing high spatial resolution over a typical viewing angle of perhaps 10°–20°, result in the formation of a spatial view cell, the effective trigger for which will thus be a combination of visual features within a relatively small part of space. In contrast, in the rat, given the very extensive visual field subtended by the rodent retina, which may extend over 180°–270°, a combination of visual features formed over such a wide visual angle would effectively define a position in space that is a place (de Araujo et al. 2001).

Although spatial view cells are present in the parahippocampal areas (Rolls et al. 1997a, 1998; Robertson et al. 1998; Georges-François et al. 1999), and neurons with place-like fields (though in some cases as a grid; Hafting et al. 2005) are found in the medial entorhinal cortex (Moser and Moser 1998; Brun et al. 2002; Fyhn et al. 2004; Moser 2004), there are back-projections from the hippocampus to the entorhinal cortex and thus to parahippocampal areas, and these back-projections could enable the hippocampus to influence the spatial representations found in the entorhinal cortex and parahippocampal gyrus. On the other hand, as described above, the grid-like place cells in the medial entorhinal cortex could, if transformed by the competitive net functionality of the dentate cells, result in the place cell activity (without a repeating grid) that is found in dentate and rat hippocampal neurons (Rolls et al. 2006).

CA1 cells

Associative retrieval at the CA3-to-CA1 (Schaffer collateral) synapses

The CA3 cells connect to the CA1 cells by the Schaeffer collateral synapses. The following arguments outline the advantage of this

connection being associatively modifiable and apply independently of the relative extent to which the CA3 or the direct entorhinal cortex inputs to CA1 drive the CA1 cells during the learning phase.

The amount of information about each episode retrievable from CA3 has to be balanced against the number of episodes that can be held concurrently in storage. The balance is regulated by the sparseness of the coding. Whatever the amount of information per episode in CA3, one may hypothesize that the organization of the structures that follow CA3 (i.e., CA1, the various subicular fields, and the return projections to neocortex) should be optimized so as to preserve and use this information content in its entirety. This would prevent further loss of information, after the massive but necessary reduction in information content that has taken place along the sensory pathways and before the autoassociation stage in CA3. We have proposed (Treves and Rolls 1994; Treves 1995) that the need to preserve the full information content present in the output of an autoassociative memory requires an intermediate recoding stage (CA1) with special characteristics. In fact, a calculation of the information present in the CA1 firing pattern, elicited by a pattern of activity retrieved from CA3, shows that a considerable fraction of the information is lost if the synapses are nonmodifiable, and that this loss can be prevented only if the CA3-to-CA1 synapses are associatively modifiable. Their modifiability should match the plasticity of the CA3 recurrent collaterals. The additional information that can be retrieved beyond that retrieved by CA3 because the CA3-to-CA1 synapses are associatively modifiable is strongly demonstrated by the hippocampal simulation described by Rolls (1995) and is quantitatively analyzed by Schultz and Rolls (1999).

Recoding in CA1 to facilitate retrieval to the neocortex

If the total amount of information carried by CA3 cells is redistributed over a larger number of CA1 cells, less information needs to be loaded onto each CA1 cell, rendering the code more robust to information loss in the next stages. For example, if each CA3 cell had to code for 2 bits of information, for example, by firing at one of four equiprobable activity levels, then each CA1 cell (if there were twice as many as there are CA3 cells) could code for just 1 bit, for example, by firing at one of only two equiprobable levels. Thus the same information content could be maintained in the overall representation while reducing the sensitivity to noise in the firing level of each cell. In fact, there are more CA1 cells than CA3 cells in rats (2.5×10^5). There are even more CA1 cells (4.6×10^6) in humans (and the ratio of CA1 to CA3 cells is greater). The CA1 cells may thus provide the first part of the expansion for the return projections to the enormous numbers of neocortical cells in primates, after the bottleneck of the single network in CA3, the number of neurons in which may be limited because it has to operate as a single network.

Another argument on the operation of the CA1 cells is also considered to be related to the CA3 autoassociation effect. In this, several arbitrary patterns of firing occur together on the CA3 neurons and become associated together to form an episodic or "whole scene" memory. It is essential for this CA3 operation that several different sparse representations are present conjunctively in order to form the association. Moreover, when completion operates in the CA3 autoassociation system, all the neurons firing in the original conjunction can be brought into activity by only a part of the original set of conjunctive events. For these reasons, a memory in the CA3 cells consists of several different simultaneously active ensembles of activity. To be explicit, the parts A, B, C, D, and E of a particular episode would each be represented, roughly speaking, by its own population of CA3

cells, and these five populations would be linked together by autoassociation. It is suggested that the CA1 cells, which receive these groups of simultaneously active ensembles, can detect the conjunctions of firing of the different ensembles that represent the episodic memory, and allocate by competitive learning neurons to represent at least larger parts of each episodic memory (Rolls 1987, 1989a,b,c, 1990a,b). In relation to the simple example above, some CA1 neurons might code for ABC, and others for BDE, rather than having to maintain independent representations in CA1 of A, B, C, D, and E. This implies a more efficient representation, in the sense that when eventually after many further stages, neocortical neuronal activity is recalled (as discussed below), each neocortical cell need not be accessed by all the axons carrying each component A, B, C, D, and E, but instead by fewer axons carrying larger fragments, such as ABC and BDE. This process is performed by competitive networks, which self-organize to find categories in the input space, where each category is represented by a set of simultaneously active inputs (Rolls and Treves 1998; Rolls 2000; Rolls and Deco 2002).

CA1 inputs from CA3 versus direct entorhinal inputs

Another feature of the CA1 network is its double set of afferents, with each of its cells receiving most synapses from the Schaeffer collaterals coming from CA3, but also a proportion (~1/6; Amaral et al. 1990) from direct perforant path projections from entorhinal cortex. Such projections appear to originate mainly in layer 3 of the entorhinal cortex (Witter et al. 1989), from a population of cells only partially overlapping with that (mainly in layer 2) giving rise to the perforant path projections to DG and CA3. This suggests that it is useful to include in CA1 not only what it is possible to recall from CA3, but also the detailed information present in the retrieval cue itself (see Treves and Rolls 1994).

Another possibility is that the perforant path input provides the strong forcing input to the CA1 neurons during learning and that the output of the CA3 system is associated with this forced CA1 firing during learning (McClelland et al. 1995). During recall, an incomplete cue could then be completed in CA3, and the CA3 output would then produce firing in CA1 that would correspond to that present during the learning. This suggestion is essentially identical to that of Treves and Rolls (1994) about the back-projection system and recall, except that McClelland et al. (1995) suggest that the output of CA3 is associated at the CA3-to-CA1 (Schaeffer collateral) synapses with the signal present during training in CA1, whereas in the theory of Treves and Rolls (1994), the output of the hippocampus consists of CA1 firing that is associated in the entorhinal cortex and earlier cortical stages with the firing present during learning, providing a theory of how the correct recall is implemented at every back-projection stage through the neocortex (see below).

Back-projections to the neocortex

The need for information to be retrieved from the hippocampus to affect other brain areas was noted in the Introduction. The way in which this could be implemented via back-projections to the neocortex is now considered.

It is suggested that the modifiable connections from the CA3 neurons to the CA1 neurons allow the whole episode in CA3 to be produced in CA1. This may be assisted as described above by the direct perforant path input to CA1. The CA1 neurons would then activate, via their termination in the deep layers of the entorhinal cortex, at least the pyramidal cells in the deep layers of the entorhinal cortex (see Fig. 1). These entorhinal cortex layer 5 neurons would then, by virtue of their back-projections (Lavenex and Amaral 2000; Witter et al. 2000a) to the parts of cerebral cortex that originally provided the inputs to the

hippocampus, terminate in the superficial layers (including layer 1) of those neocortical areas, where synapses would be made onto the distal parts of the dendrites of the (superficial and deep) cortical pyramidal cells (Rolls 1989a,b,c). The areas of cerebral neocortex in which this recall would be produced could include multimodal cortical areas (e.g., the cortex in the superior temporal sulcus that receives inputs from temporal, parietal, and occipital cortical areas, and from which it is thought that cortical areas such as 39 and 40 related to language developed), and also areas of unimodal association cortex (e.g., inferior temporal visual cortex). The back-projections, by recalling previous episodic events, could provide information useful to the neocortex in the building of new representations in the multimodal and unimodal association cortical areas, which by building new long-term representations can be considered as a form of memory consolidation (Rolls 1989a,b,c, 1990a,b), or in organizing actions.

The hypothesis of the architecture with which this would be achieved is shown in Figure 1. The feedforward connections from association areas of the cerebral neocortex (solid lines in Fig. 1) show major convergence as information is passed to CA3, with the CA3 autoassociation network having the smallest number of neurons at any stage of the processing. The back-projections allow for divergence back to neocortical areas. The way in which I suggest that the back-projection synapses are set up to have the appropriate strengths for recall is as follows (Rolls 1989a,b,c). During the setting up of a new episodic memory, there would be strong feedforward activity progressing toward the hippocampus. During the episode, the CA3 synapses would be modified, and via the CA1 neurons and the subiculum, a pattern of activity would be produced on the back-projecting synapses to the entorhinal cortex. Here the back-projecting synapses from active back-projection axons onto pyramidal cells being activated by the forward inputs to entorhinal cortex would be associatively modified. A similar process would be implemented at preceding stages of neocortex, that is, in the parahippocampal gyrus/perirhinal cortex stage and in association cortical areas, as shown in Figure 1.

The concept is that during the learning of an episodic memory, cortical pyramidal cells in at least one of the stages would be driven by forward inputs, but would simultaneously be receiving back-projected activity (indirectly) from the hippocampus, which would by pattern association from the back-projecting synapses to the cortical pyramidal cells become associated with whichever cortical cells were being made to fire by the forward inputs. Then later on, during recall, a recall cue from perhaps another part of cortex might reach CA3, where the firing during the original episode would be completed. The resulting back-projecting activity would then, as a result of the pattern association learned previously, bring back the firing in any cortical area that was present during the original episode. Thus retrieval involves reinstating the activity that was present in different cortical areas that was present during the learning of an episode. (The pattern association is also called "heteroassociation," to contrast it with autoassociation. The pattern association operates at multiple stages in the back-projection pathway, as made evident in Figure 1.) If the recall cue was an object, this might result in recall of the neocortical firing that represented the place in which that object had been seen previously. As noted elsewhere in this paper and by McClelland et al. (1995), that recall might be useful to the neocortex to help it build new semantic memories, which might inherently be a slow process and is not part of the theory of recall.

A plausible requirement for a successful hippocampally directed recall operation is that the signal generated from the hippocampally retrieved pattern of activity, and carried backward toward neocortex, remain undegraded when compared to

the noise due, at each stage, to the interference effects caused by the concurrent storage of other patterns of activity on the same back-projecting synaptic systems. That requirement is equivalent to that used in deriving the storage capacity of such a series of heteroassociative memories, and it was shown in Treves and Rolls (1991) that the maximum number of independently generated activity patterns that can be retrieved is given, essentially, by the same formula as Equation 3 above, where, however, a is now the sparseness of the representation at any given stage, and C is the average number of (back-)projections each cell of that stage receives from cells of the previous one (Treves and Rolls 1994). (k' is a similar slowly varying factor to that introduced above.) If p is equal to the number of memories held in the hippocampal memory, it is limited by the retrieval capacity of the CA3 network, p_{\max} . Putting together the formula for the CA3 recurrent collaterals (Equation 3) with that for the backprojections, one concludes that, roughly, the requirement implies that the number of afferents of (indirect) hippocampal origin to a given neocortical stage (C^{HBP}), must be $C^{\text{HBP}} = C^{\text{RC}} a_{\text{nc}} / a_{\text{CA3}}$, where C^{RC} is the number of recurrent collaterals to any given cell in CA3, the average sparseness of a representation is a_{nc} , and a_{CA3} is the sparseness of memory representations there in CA3 (Treves and Rolls 1994).

The above requirement is very strong: even if representations were to remain as sparse as they are in CA3, which is unlikely, to avoid degrading the signal, C^{HBP} should be as large as C^{RC} , that is, 12,000 in the rat. If, then, C^{HBP} has to be of the same order as C^{RC} , one is led to a very definite conclusion: a mechanism of the type envisaged here could not possibly rely on a set of monosynaptic CA3-to-neocortex back-projections. This would imply that, to make a sufficient number of synapses on each of the vast number of neocortical cells, each cell in CA3 has to generate a disproportionate number of synapses (i.e., C^{HBP} times the ratio between the number of neocortical and that of CA3 cells). The required divergence can be kept within reasonable limits only by assuming that the back-projecting system is polysynaptic, provided that the number of cells involved grows gradually at each stage, from CA3 back to neocortical association areas (Treves and Rolls 1994) (cf. Fig. 1).

The theory of recall by the back-projections thus provides a quantitative account of why the cerebral cortex has as many back-projections as forward projection connections (Treves and Rolls 1994). Further aspects of the operation of the back-projecting systems are described elsewhere (Rolls 2008).

Neurophysiology of the primate hippocampus and attractors

Hippocampal neurophysiology is essential to understand the information that could be stored in a hippocampal attractor network in the CA3 in primates. To understand how the hippocampus works, it is not sufficient to state just that it can store information—one needs to know what information. The systems-level neurophysiology of the primate hippocampus that is especially relevant to what is stored in hippocampal attractors, and whether there are hippocampal attractors, is described next, with fuller accounts available elsewhere (Rolls and Xiang 2006). The primate neurophysiology provides a perspective relevant to understanding the function of the human hippocampus that is somewhat different from that provided by the properties of place cells in rodents, which have been reviewed elsewhere (see McNaughton et al. 1983; O'Keefe 1984; Muller et al. 1991; Jeffery and Hayman 2004; Jeffery et al. 2004).

Spatial view neurons in the primate hippocampus

The primate hippocampus contains spatial cells that respond when the monkey looks at a certain part of space, for example, at

one quadrant of a video monitor while the monkey is performing an object-place memory task in which he must remember where on the monitor he has seen particular images (Rolls et al. 1989). Approximately 9% of the hippocampal neurons have such spatial view fields, and ~2.4% combine information about the position in space with information about the object that is in that position in space (Rolls et al. 1989). The representation of space is for the majority of hippocampal neurons in allocentric not egocentric coordinates (Feigenbaum and Rolls 1991). These spatial view cells can be recorded while monkeys move themselves round the test environment by walking (or running) on all fours (Rolls et al. 1997a, 1998; Robertson et al. 1998; Georges-François et al. 1999). These hippocampal "spatial view neurons" respond significantly differently for different allocentric spatial views and have information about spatial view in their firing rate, but do not respond differently just on the basis of eye position, head direction, or place. This type of spatial representation is ideal for an event or episodic memory in which a memory is formed in an autoassociation network between, for example, a spatial location being viewed and an object or reward at that location. This is typical of human episodic memory and could not be implemented by rodent place cells that represent the location of the animal, not the location in space at which something is seen.

Evidence that the representation in the CA pyramidal cells has a characteristic property of an attractor network is that spatial view neurons continue firing when the view details are obscured and the macaque is in the dark. In particular, if the view details are obscured by curtains and darkness, then some spatial view neurons continue to respond for up to a few minutes whenever the monkey looks toward the spatial view field, showing that these neurons can be updated for at least short periods by idiothetic (self-motion) cues including eye position and head direction signals (Rolls et al. 1997b; Robertson et al. 1998) (see example in Fig. 6). The fact that some drift of the location being represented with the curtains drawn to obscure the view and during darkness is evidence that this is a short-term memory function operating in the absence of the normal input, in this case visual, that provides the reference position for a spatial view neuron. These findings are evidence for the operation of an attractor that influences the activity of hippocampal pyramidal cells. The firing continues better in CA1 than CA3 neurons, and this is consistent with the hypothesis that the spatial view path integration is being performed in the CA3 neurons resulting in some firing, and that the associative retrieval in the CA3-to-CA1 connections allows a better representation to be produced in CA1, particularly under noisy low signal conditions, as described above.

Object-place neurons in the primate hippocampus

A fundamental question about the function of the primate including human hippocampus in relation to episodic memory is whether object as well as allocentric spatial information is represented, as would be required if it is to participate in, for example, object-place memory, by a process such as autoassociation in the CA3 network. To investigate this, Rolls et al. (2005) made recordings from single hippocampal formation neurons while macaques performed an object-place memory task that required the monkeys to learn associations between objects and where they were shown in a room. Some neurons (10%) responded differently to different objects independently of location; other neurons (13%) responded to the spatial view independently of which object was present at the location; and some neurons (12%) responded to a combination of a particular object and the place where it was shown in the room. These results show that there are separate as well as combined representations

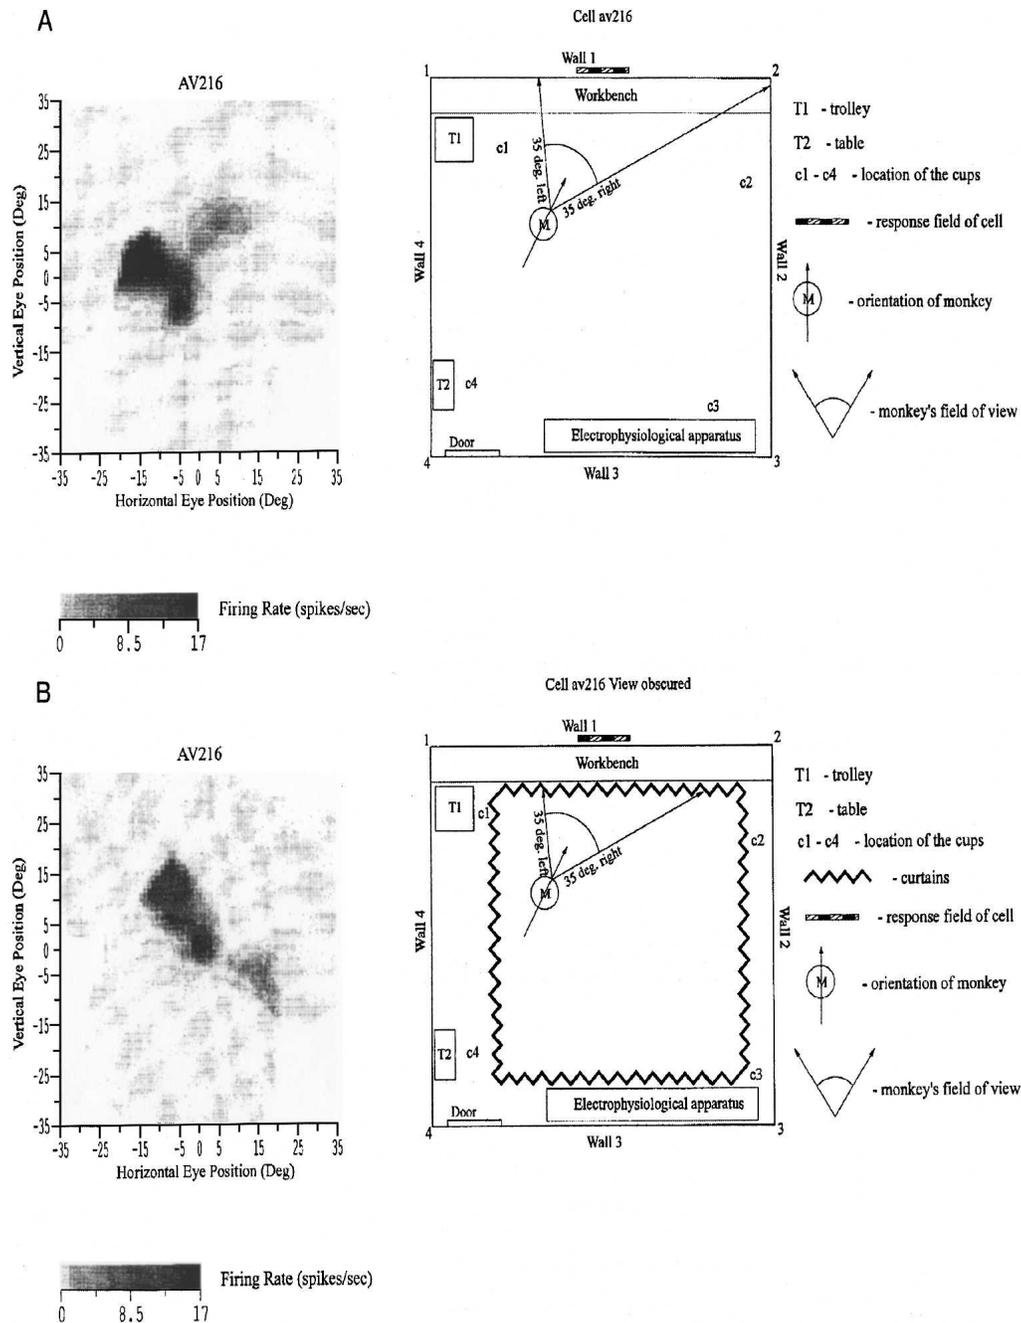


Figure 6. Idiothetic update of the firing of a hippocampal spatial view cell. This cell fired when the macaque looked toward the effective spatial view even when it was obscured by curtains. The firing is shown with the monkey stationary with his head facing in the direction indicated by the arrow when the curtains were (A) drawn open or were (B) drawn closed. (Left) Firing rate of the cell in spikes/second is indicated by the blackness (a calibration bar in spikes/second is shown below) projected onto the monkey's field of view. The two-dimensional firing rate profile of the cell was smoothed for clarity using a two-dimensional Gaussian spatial filter. The space adequately sampled by the eye movements of the monkey is indicated by shading. (Right) A plan view of the room to indicate the monkey's view of the wall is shown. (M) Position of the monkey (after Robertson et al. 1998).

of objects and their locations in space in the primate hippocampus. This is a property required in an episodic memory system, for which associations between objects and the places where they are seen are prototypical. The results thus show that a requirement for a human episodic memory system, separate and combined neuronal representations of objects and where they are seen "out there" in the environment, is present in the primate hippocampus (Rolls et al. 2005). What may be a corresponding finding in rats is that some rat hippocampal neurons respond on

the basis of the conjunction of location and odor (Wood et al. 1999).

Spatial view cells, and object-place cells, are also present in the parahippocampal areas (Rolls et al. 1997a, 1998, 2005; Robertson et al. 1998; Georges-François et al. 1999). There are back-projections from the hippocampus to the entorhinal cortex and thus to parahippocampal areas, and these back-projections could enable the hippocampus to influence the spatial representations found in the entorhinal cortex and parahippocampal gyrus. On

the other hand, some of the spatial functions may be provided for in these parahippocampal areas, which will, in turn, influence the hippocampus. However, the hippocampus itself may be able to make a special contribution to event or episodic memory, by enabling in the CA3 autoassociation network with its very widespread recurrent collateral connections an association between any one item (such as a particular place being viewed) with any other item (such as a particular object) to form an arbitrary association to represent an event.

Recall-related neurons in the primate hippocampus

It has now been possible to investigate directly, neurophysiologically, the hippocampal recall process in primates (Rolls and Xiang 2006). We used a visual object-place memory task because this is prototypical of episodic memory. It has been shown that a one-trial odor-place recall memory task is hippocampal-dependent in rodents (Day et al. 2003). We designed a one-trial object-place recall task in which the whole memory was recalled from a part of it. The task is illustrated in Figure 7. Images of new objects were used each day, and within a day the same objects were used, so that with non-trial unique objects within a day, the recall task is quite difficult.

Recordings were made from 347 neurons in the hippocampus of a macaque performing the object-place recall task. The following types of neurons were found in the task (Rolls and Xiang 2006). One type of neuron had responses that occurred to one of the objects used in the task. Several of these neurons had activity that was related to the recall process. An example of one of these neurons is shown in Figure 8. The neuron had activity that was greater to object 1 not only when it was shown in stages 1, 2, and 3 of the task, but also in the delay period following stage 3 when the object was no longer visible, and in stage 4, when also the object was no longer visible and the macaque was touching the remembered location of that object. Thus, while the location

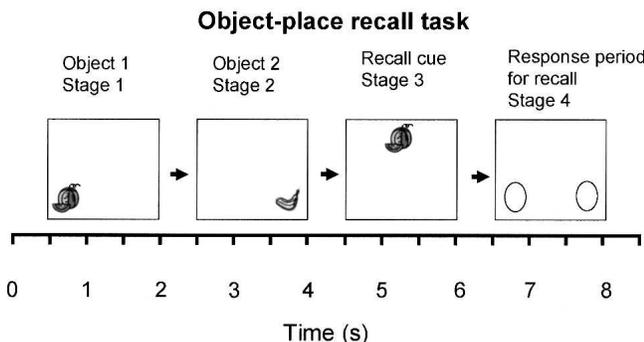


Figure 7. The object-place recall task. One trial is shown. After a 0.5-sec tone to indicate the start of a trial, in Stage 1, one of two objects (O1) is shown at one of the places (P1). (The object and the place are chosen randomly on each trial.) To ensure that the monkey sees the stimulus, the monkey can touch the screen at the place to obtain one drop of juice reward by licking. After a 0.5-sec delay, in Stage 2, the other of the two objects (O2) is shown at the other place (P2). (One drop of fruit juice was available as in Stage 1.) After a 0.5-sec delay, in Stage 3, the recall cue, one of the objects chosen at random, is shown at the top center of the screen. (One drop of fruit juice was available as in Stage 1.) After a 0.5-sec delay, in Stage 4, the macaque must then recall the place in which the object shown as the recall cue in Stage 3 was presented, and must then touch that place on the screen to obtain four licks of fruit juice, thus indicating that he has recalled the location correctly. In Stage 4 of the trials, the left and right positions (P1 and P2) have no image present, with the two possible locations for a response indicated by identical circles. The task requires the monkey to perform recall of the place from the object, within the period beginning at the presentation of the recall cue at the start of Stage 3 and ending when the response is made in Stage 4.

Object recall-related neuronal activity in the object-place recall task

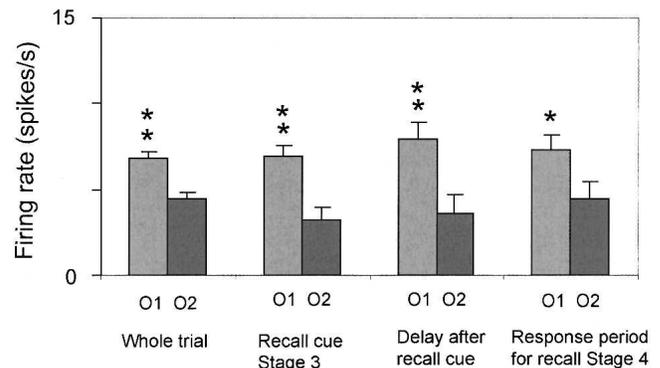


Figure 8. Activity of a neuron with responses related to one of the objects used in the object-place recall task. The firing rates to object 1 (O1) and object 2 (O2) are shown (mean firing rate in spikes/second across trials \pm SEM). The first histogram pair (on the left) shows the responses to the two objects measured throughout the trial whenever object 1 or object 2 was on the screen. The second histogram pair shows the neuronal responses when the objects were being shown in stage 3 as the recall cue. The third histogram pair shows the neuronal responses in the 0.5-sec delay period after one of the objects had been shown in stage 3 as the recall cue. The neuron continued to respond more after object 1 than after object 2 had been seen, in this recall period in which the place was being recalled from the object. The fourth histogram pair shows the neuronal responses in stage 4 when the macaque was recalling and touching the place at which the cue recall object had been shown. The responses of the neuron were object-related even when the object was not being seen, but was being used as a recall cue, in the delay after stage 3 of the task, and in stage 4. ** $P < 0.01$; * $P < 0.05$.

was being recalled from the object, this type of neuron continued to respond as if the object were present, that is, it kept the representation of the object active after the object was no longer visible, and the place to touch was being recalled. Sixteen of the neurons responded in this way, and an additional six had object-related firing that did not continue following stage 3 of the task in the recall period. The difference of the firing rates of these 22 neurons to the different objects were in many cases highly statistically significant (e.g., $P < 10^{-6}$). We performed a Fisher exact probability test to confirm that the set of statistically significant results in the 22 neurons could not have arisen by chance within the 347 tests performed and were able to reject this with $P < 5.4 \times 10^{-8}$. Thus, the population of 22 neurons had statistically very high significance in its object-related responses. None of these neurons had differential responses for the different places used in the object-place recall task.

A second type of neuron had responses related to the place (left or right) in which an object was shown in stages 1 or 2 of each trial. An example of one of these neurons is shown in Figure 9. The neuron responded more when an object was shown in the left position (P1) than in the right position (P2) on the screen. Interestingly, when the recall object was shown in stage 3 of the trial in the top center of the screen, the neuron also responded as if the left position (P1) was being processed on trials in which the left position had to be recalled. This firing continued in the delay period after the recall cue had been removed at the end of stage 3, and into stage 4. Thus, this type of neuron appeared to reflect the recall of the position on the screen at which the object had been represented. Analysis of trials on which errors were made indicated that the responses were not just motor response related, for if owing to some response bias the monkey touched the incorrect side, the neuron could still respond according to the correct recalled location. Thirteen neurons had differential re-

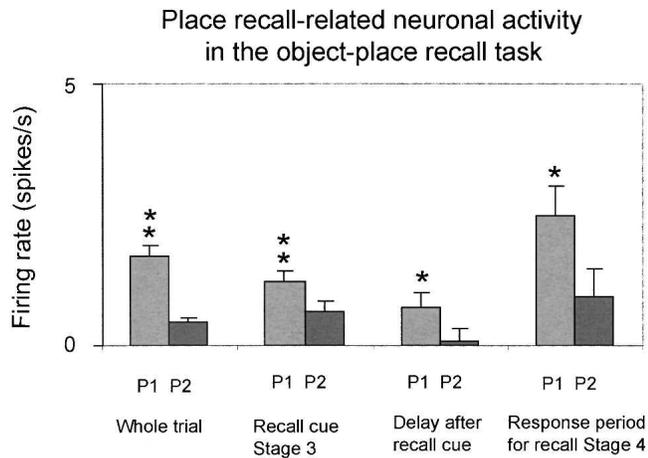


Figure 9. Activity of a neuron with responses related to the left place (P1) in the object-place recall task. The firing rates to place 1 (P1) and place 2 (P2) are shown (mean firing rate in spikes/second across trials \pm SEM). The first histogram pair (on the left) shows the responses to the two places measured when a stimulus was on the screen in stage 1 or stage 2. The second histogram pair shows the neuronal responses when the objects were being shown in stage 3 as the recall cue, and depending on whether the place to be recalled was place 1 or place 2. The third histogram pair shows the neuronal responses in the 0.5-sec delay period after one of the objects had been shown in stage 3 as the recall cue. The neuron responded more when place 1 was the correct place to be recalled on a trial. The fourth histogram pair shows the neuronal responses in stage 4 when the macaque was recalling and touching the place at which the cue recall object had been shown. The responses of the neuron were place-related even in stage 3 when the object being shown as a place recall cue was at the top of the screen, in the delay after stage 3 of the task, and in stage 4. ** $P < 0.01$; * $P < 0.05$.

sponses to the different places P1 and P2 and continued to show place-related activity in the recall part of the task, stage 3. Five other neurons had left-right place-related responses without a memory recall component, in that they did not respond in stage 3 of the task, when a nonspatial recall stimulus was being shown, and a place should be being recalled (see Table 1). We performed a Fisher exact probability test to confirm that the set of statistically significant results in the 18 neurons could not have arisen by chance within the 347 tests performed, and were able to reject this with $P < 0.05$. Thus, the population of 18 neurons as a population had statistically significant place-related responses. The new finding is that 13 of the neurons had place-related responses when a place was being recalled by an object cue.

The responses of the population of neurons recorded in one macaque are shown in Table 1. In addition to the neurons described above, three further neurons responded to particular combinations of objects and places, for example, to object 1 when it was shown in place 1, but not to other combinations.

The recording sites of the object and of the place neurons were within the hippocampus proper (Rolls and Xiang 2006). The mean firing rate of the population of responsive neurons (see

Table 1. Numbers of neurons in the hippocampus with different types of responses during the object-place recall task

Response	No. of neurons
Object with activity continuing after the recall cue	16
Object with activity not continuing after the recall cue	6
Place with activity during and after the recall cue	13
Place with activity during the recall cue	5
Object-place	3
Total	347

Table 1) to the most effective object or place was 7.2 ± 0.6 spikes/sec (\pm SEM), and their mean spontaneous rate was 3.2 ± 0.6 spikes/sec.

These findings (Rolls and Xiang 2006) are the first we know in the primate hippocampus of neuronal activity that is related to recall. It is particularly interesting that the neurons with continuing activity to the object after it had disappeared in the recall phase of the task could reflect the operation of a short-term memory implemented by the CA3 recurrent collateral attractor network. By continuing to respond to the object while the place is being recalled in the task, the object-related neurons could be part of the completion of the whole object-place combination memory from an autoassociation or attractor process in CA3 (Rolls and Kesner 2006). It is also very interesting that the appropriate set of spatial view (place) related neurons started to fire in the recall phase of the task when one of the objects was shown, for this is what is predicted for completion in a CA3 attractor network of the object-place memory based on an object retrieval cue.

The neurons with recall-related activity in the object-place recall task also provide neurophysiological evidence on the speed of association learning in the hippocampal formation. Given that this is a one-trial object-place recall task, with the association between the object and its place being made in stages 1 and 2 of each trial (see Fig. 7), it is clear that it takes just one trial for the object-place associations to be formed that are relevant to the later recall on that trial. This is the speed of learning that is required for episodic memory, and this neurophysiological evidence shows that this type of rapid, one-trial, object-place learning is represented in the primate hippocampus.

Reward-place neurons in the primate hippocampus

The primate anterior hippocampus (which corresponds to the rodent ventral hippocampus) receives inputs from brain regions involved in reward processing such as the amygdala and orbitofrontal cortex (Pitkanen et al. 2002). To investigate how this affective input may be incorporated into primate hippocampal function, Rolls and Xiang (2005) recorded neuronal activity while macaques performed a reward-place association task in which each spatial scene shown on a videomonitor had one location that if touched yielded a preferred fruit juice reward, and a second location that yielded a less preferred juice reward. Each scene had different locations for the different rewards. Of 312 hippocampal neurons analyzed, 18% responded more to the location of the preferred reward in different scenes, and 5% to the location of the less preferred reward (Rolls and Xiang 2005). When the locations of the preferred rewards in the scenes were reversed, 60% of 44 neurons tested reversed the location to which they responded, showing that the reward-place associations could be altered by new learning in a few trials. The majority (82%) of these 44 hippocampal reward-place neurons tested did not respond to object-reward associations in a visual discrimination object-reward association task. Thus the primate hippocampus contains a representation of the reward associations of places "out there" being viewed, and this is a way in which affective information can be stored as part of an episodic memory and how the current mood state may influence the retrieval of episodic memories. An autoassociation network in CA3 representing spatial views and rewards would enable particular spatial views to become associated with particular rewards, and later for the place at which a reward was available to be retrieved from a reward cue, and vice versa for the reward available at a location to be retrieved from the spatial view. There is consistent evidence that rewards available in a spatial environment can influence the responsiveness of rodent place neurons (Hölscher et al. 2003; Tabuchi et al. 2003).

In conclusion, in this paper, I have described many types of

evidence that the CA3 system acts as an autoassociation or attractor memory. The theory is developed further and compared to and contrasted with other theories and approaches elsewhere (Rolls and Kesner 2006; Rolls 2008).

Acknowledgments

Different parts of the research described here were supported by Programme Grants from the Medical Research Council, by a Human Frontier Science program grant, by an EEC BRAIN grant, by the MRC Oxford Interdisciplinary Research Centre in Cognitive Neuroscience, and by the Oxford McDonnell Centre in Cognitive Neuroscience. The author has performed the experimental and theoretical work that is incorporated in some of the ideas presented here on the hippocampus with many colleagues, including Alessandro Treves, Simon Stringer, Ray Kesner, Robert Robertson, Pierre Georges-François, and Shane O'Mara, and their contributions are sincerely acknowledged.

References

- Amaral, D.G. 1987. Memory: Anatomical organization of candidate brain regions. In *Handbook of physiology, Section 1: The nervous system* (ed. V.B. Mountcastle), pp. 211–294. American Physiological Society, Washington, DC.
- Amaral, D.G. 1993. Emerging principles of intrinsic hippocampal organisation. *Curr. Opin. Neurobiol.* **3**: 225–229.
- Amaral, D.G. and Witter, M.P. 1989. The three-dimensional organization of the hippocampal formation: A review of anatomical data. *Neuroscience* **31**: 571–591.
- Amaral, D.G. and Witter, M.P. 1995. The hippocampal formation. In *The Rat nervous system* (ed. G. Paxinos), pp. 443–493. Academic Press, San Diego.
- Amaral, D.G., Ishizuka, N., and Claiborne, B. 1990. Neurons, numbers and the hippocampal network. *Prog. Brain Res.* **83**: 1–11.
- Amaral, D.G., Price, J.L., Pitkanen, A., and Carmichael, S.T. 1992. Anatomical organization of the primate amygdaloid complex. In *The amygdala* (ed. J.P. Aggleton), pp. 1–66. Wiley-Liss, New York.
- Amari, S. 1977. Dynamics of pattern formation in lateral-inhibition type neural fields. *Biol. Cybern.* **27**: 77–87.
- Amit, D.J. 1989. *Modeling brain function*. Cambridge University Press, Cambridge.
- Battaglia, F.P. and Treves, A. 1998a. Attractor neural networks storing multiple space representations: A model for hippocampal place fields. *Phys. Rev. E* **58**: 7738–7753.
- Battaglia, F.P. and Treves, A. 1998b. Stable and rapid recurrent processing in realistic auto-associative memories. *Neural Comput.* **10**: 431–450.
- Brown, T.H., Ganong, A.H., Kairiss, E.W., Keenan, C.L., and Kelso, S.R., eds. 1989. *Long-term potentiation in two synaptic systems of the hippocampal brain slice*. Academic Press, San Diego.
- Brown, T.H., Kairiss, E.W., and Keenan, C.L. 1990. Hebbian synapses: Biophysical mechanisms and algorithms. *Annu. Rev. Neurosci.* **13**: 475–511.
- Brun, V.H., Otnass, M.K., Molden, S., Steffenach, H.A., Witter, M.P., Moser, M.B., and Moser, E.I. 2002. Place cells and place recognition maintained by direct entorhinal-hippocampal circuitry. *Science* **296**: 2243–2246.
- Buckley, M.J. and Gaffan, D. 2000. The hippocampus, perirhinal cortex, and memory in the monkey. In *Brain, perception, and memory: Advances in cognitive neuroscience* (ed. J.J. Bolhuis), pp. 279–298. Oxford University Press, Oxford.
- Burgess, N., Maguire, E.A., and O'Keefe, J. 2002. The human hippocampus and spatial and episodic memory. *Neuron* **35**: 625–641.
- Carmichael, S.T. and Price, J.L. 1995. Limbic connections of the orbital and medial prefrontal cortex in macaque monkeys. *J. Comp. Neurol.* **346**: 403–434.
- Cassaday, H.J. and Rawlins, J.N. 1997. The hippocampus, objects, and their contexts. *Behav. Neurosci.* **111**: 1228–1244.
- Chawla, M.K., Guzowski, J.F., Ramirez-Amaya, V., Lipa, P., Hoffman, K.L., Marriott, L.K., Worley, P.F., McNaughton, B.L., and Barnes, C.A. 2005. Sparse, environmentally selective expression of Arc RNA in the upper blade of the rodent fascia dentata by brief spatial experience. *Hippocampus* **15**: 579–586.
- Crane, J. and Milner, B. 2005. What went where? Impaired object-location learning in patients with right hippocampal lesions. *Hippocampus* **15**: 216–231.
- Day, M., Langston, R., and Morris, R.G. 2003. Glutamate-receptor-mediated encoding and retrieval of paired-associate learning. *Nature* **424**: 205–209.
- de Araujo, I.E.T., Rolls, E.T., and Stringer, S.M. 2001. A view model which accounts for the spatial fields of hippocampal primate view cells and rat place cells. *Hippocampus* **11**: 699–706.
- Delatour, B. and Witter, M.P. 2002. Projections from the parahippocampal region to the prefrontal cortex in the rat: evidence of multiple pathways. *Eur. J. Neurosci.* **15**: 1400–1407.
- Fazeli, M.S. and Collingridge, G.L., eds. 1996. *Cortical plasticity: LTP and LTD*. Bios, Oxford.
- Feigenbaum, J.D. and Rolls, E.T. 1991. Allocentric and egocentric spatial information processing in the hippocampal formation of the behaving primate. *Psychobiology* **19**: 21–40.
- Florian, C. and Roullet, P. 2004. Hippocampal CA3-region is crucial for acquisition and memory consolidation in Morris water maze task in mice. *Behav. Brain Res.* **154**: 365–374.
- Franco, L., Rolls, E.T., Aggelopoulos, N.C., and Jerez, J.M. 2007. Neuronal selectivity, population sparseness, and ergodicity in the inferior temporal visual cortex. *Biol. Cybern.* **96**: 547–560. doi: 10.1007/s00422-007-0149-1.
- Fyhn, M., Molden, S., Witter, M.P., Moser, E.I., and Moser, M.B. 2004. Spatial representation in the entorhinal cortex. *Science* **305**: 1258–1264.
- Gaffan, D. 1994. Scene-specific memory for objects: A model of episodic memory impairment in monkeys with fornix transection. *J. Cogn. Neurosci.* **6**: 305–320.
- Gaffan, D. and Harrison, S. 1989a. A comparison of the effects of fornix section and sulcus principalis ablation upon spatial learning by monkeys. *Behav. Brain Res.* **31**: 207–220.
- Gaffan, D. and Harrison, S. 1989b. Place memory and scene memory: Effects of fornix transection in the monkey. *Exp. Brain Res.* **74**: 202–212.
- Gaffan, D. and Saunders, R.C. 1985. Running recognition of configural stimuli by fornix transected monkeys. *Q. J. Exp. Psychol. B* **37**: 61–71.
- Georges-François, P., Rolls, E.T., and Robertson, R.G. 1999. Spatial view cells in the primate hippocampus: Allocentric view not head direction or eye position or place. *Cereb. Cortex* **9**: 197–212.
- Gilbert, P.E., Kesner, R.P., and Lee, I. 2001. Dissociating hippocampal subregions: Double dissociation between dentate gyrus and CA1. *Hippocampus* **11**: 626–636.
- Gold, A.E. and Kesner, R.P. 2005. The role of the CA3 subregion of the dorsal hippocampus in spatial pattern completion in the rat. *Hippocampus* **15**: 808–814.
- Goodrich-Hunsaker, N.J., Hunsaker, M.R., and Kesner, R.P. 2005. Effects of hippocampus sub-regional lesions for metric and topological spatial information processing. In *Abstract viewer/itinerary planner, program no. 647.1*. Society for Neuroscience, Washington, DC.
- Hafting, T., Fyhn, M., Molden, S., Moser, M.B., and Moser, E.I. 2005. Microstructure of a spatial map in the entorhinal cortex. *Nature* **436**: 801–806.
- Hargreaves, E.L., Rao, G., Lee, I., and Knierim, J.J. 2005. Major dissociation between medial and lateral entorhinal input to dorsal hippocampus. *Science* **308**: 1792–1794.
- Hasselmo, M.E., Schnell, E., and Barkai, E. 1995. Dynamics of learning and recall at excitatory recurrent synapses and cholinergic modulation in rat hippocampal region CA3. *J. Neurosci.* **15**: 5249–5262.
- Hertz, J., Krogh, A., and Palmer, R.G. 1991. *An introduction to the theory of neural computation*. Addison-Wesley, Wokingham, UK.
- Hölscher, C., Jacob, W., and Mallot, H.A. 2003. Reward modulates neuronal activity in the hippocampus of the rat. *Behav. Brain Res.* **142**: 181–191.
- Ishizuka, N., Weber, J., and Amaral, D.G. 1990. Organization of intrahippocampal projections originating from CA3 pyramidal cells in the rat. *J. Comp. Neurol.* **295**: 580–623.
- Jarrard, E.L. 1993. On the role of the hippocampus in learning and memory in the rat. *Behav. Neural Biol.* **60**: 9–26.
- Jeffery, K.J. and Hayman, R. 2004. Plasticity of the hippocampal place cell representation. *Rev. Neurosci.* **15**: 309–331.
- Jeffery, K.J., Anderson, M.I., Hayman, R., and Chakraborty, S. 2004. A proposed architecture for the neural representation of spatial context. *Neurosci. Biobehav. Rev.* **28**: 201–218.
- Jung, M.W. and McNaughton, B.L. 1993. Spatial selectivity of unit activity in the hippocampal granular layer. *Hippocampus* **3**: 165–182.
- Kesner, R.P. 1998. Neural mediation of memory for time: Role of hippocampus and medial prefrontal cortex. *Psychol. Bull. Rev.* **5**: 585–596.
- Kesner, R.P. 2007. Behavioral functions of the CA3 subregion of the hippocampus. *Learn. Mem.* (this issue). doi: 10.1101/lm.688207.
- Kesner, R.P. and Rolls, E.T. 2001. Role of long term synaptic modification in short term memory. *Hippocampus* **11**: 240–250.
- Kesner, R.P., Lee, I., and Gilbert, P. 2004. A behavioral assessment of hippocampal function based on a subregional analysis. *Rev. Neurosci.* **15**: 333–351.
- Kesner, R.P., Hunsaker, M.R., and Warthen, M.W. 2007. *CA3 mediates arbitrary associations and episodic memory processing based on*

- performance of reciprocal object place recall tasks. Society for Neuroscience Abstracts, Washington, DC.
- Lassalle, J.M., Bataille, T., and Halley, H. 2000. Reversible inactivation of the hippocampal mossy fiber synapses in mice impairs spatial learning, but neither consolidation nor memory retrieval, in the Morris navigation task. *Neurobiol. Learn. Mem.* **73**: 243–257.
- Lavenex, P. and Amaral, D.G. 2000. Hippocampal-neocortical interaction: A hierarchy of associativity. *Hippocampus* **10**: 420–430.
- Lavenex, P., Suzuki, W.A., and Amaral, D.G. 2004. Perirhinal and parahippocampal cortices of the macaque monkey: Intrinsic projections and interconnections. *J. Comp. Neurol.* **472**: 371–394.
- Levy, W.B. 1989. A computational approach to hippocampal function. In *Computational models of learning in simple neural systems* (eds. R.D. Hawkins and G.H. Bower), pp. 243–305. Academic Press, San Diego.
- Levy, W.B. and Desmond, N.L. 1985. The rules of elemental synaptic plasticity. In *Synaptic modification, neuron selectivity, and nervous system organization* (eds. W.B. Levy et al.), pp. 105–121. Erlbaum, Hillsdale, NJ.
- Levy, W.B., Colbert, C.M., and Desmond, N.L. 1990. Elemental adaptive processes of neurons and synapses: A statistical/computational perspective. In *Neuroscience and connectionist theory* (eds. M.A. Gluck and D.E. Rumelhart), pp. 187–235. Erlbaum, Hillsdale, NJ.
- Lynch, M.A. 2004. Long-term potentiation and memory. *Physiol. Rev.* **84**: 87–136.
- Maaswinkel, H., Jarrard, L.E., and Whishaw, I.Q. 1999. Hippocampal-ectomized rats are impaired in homing by path integration. *Hippocampus* **9**: 553–561.
- Malkova, L. and Mishkin, M. 2003. One-trial memory for object-place associations after separate lesions of hippocampus and posterior parahippocampal region in the monkey. *J. Neurosci.* **23**: 1956–1965.
- Markus, E.J., Qin, Y.L., Leonard, B., Skaggs, W., McNaughton, B.L., and Barnes, C.A. 1995. Interactions between location and task affect the spatial and directional firing of hippocampal neurons. *J. Neurosci.* **15**: 7079–7094.
- Marr, D. 1971. Simple memory: A theory for archicortex. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **262**: 23–81.
- Martin, S.J., Grimwood, P.D., and Morris, R.G. 2000. Synaptic plasticity and memory: An evaluation of the hypothesis. *Annu. Rev. Neurosci.* **23**: 649–711.
- McClelland, J.L., McNaughton, B.L., and O'Reilly, R.C. 1995. Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychol. Rev.* **102**: 419–457.
- McNaughton, B.L. 1991. Associative pattern completion in hippocampal circuits: New evidence and new questions. *Brain Res. Brain Res. Rev.* **16**: 193–220.
- McNaughton, B.L. and Morris, R.G.M. 1987. Hippocampal synaptic enhancement and information storage within a distributed memory system. *Trends Neurosci.* **10**: 408–415.
- McNaughton, B.L. and Nadel, L. 1990. Hebb-Marr networks and the neurobiological representation of action in space. In *Neuroscience and connectionist theory* (eds. M.A. Gluck and D.E. Rumelhart). Erlbaum, Hillsdale, NJ.
- McNaughton, B.L., Barnes, C.A., and O'Keefe, J. 1983. The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. *Exp. Brain Res.* **52**: 41–49.
- McNaughton, B.L., Battaglia, F.P., Jensen, O., Moser, E.I., and Moser, M.-B. 2006. Path integration and the neural basis of the 'cognitive map.' *Nat. Rev. Neurosci.* **7**: 663–678.
- Morris, R.G.M. 1989. Does synaptic plasticity play a role in information storage in the vertebrate brain? In *Parallel distributed processing: Implications for psychology and neurobiology* (ed. R.G.M. Morris), pp. 248–285. Oxford University Press, Oxford.
- Morris, R.G. 2003. Long-term potentiation and memory. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **358**: 643–647.
- Morris, R.G., Moser, E.I., Riedel, G., Martin, S.J., Sandin, J., Day, M., and O'Carroll, C. 2003. Elements of a neurobiological theory of the hippocampus: The role of activity-dependent synaptic plasticity in memory. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **358**: 773–786.
- Moscovitch, M., Rosenbaum, R.S., Gilboa, A., Addis, D.R., Westmacott, R., Grady, C., McAndrews, M.P., Levine, B., Black, S., Winocur, G., et al. 2005. Functional neuroanatomy of remote episodic, semantic and spatial memory: A unified account based on multiple trace theory. *J. Anat.* **207**: 35–66.
- Moser, E.I. 2004. Hippocampal place cells demand attention. *Neuron* **42**: 183–185.
- Moser, M.B. and Moser, E.I. 1998. Functional differentiation in the hippocampus. *Hippocampus* **8**: 608–619.
- Muller, R.U., Kubie, J.L., Bostock, E.M., Taube, J.S., and Quirk, G.J. 1991. Spatial firing correlates of neurons in the hippocampal formation of freely moving rats. In *Brain and space* (ed. J. Paillard), pp. 296–333. Oxford University Press, Oxford.
- Muller, R.U., Ranck Jr., J.B., and Taube, J.S. 1996. Head direction cells: Properties and functional significance. *Curr. Opin. Neurobiol.* **6**: 196–206.
- Murray, E.A., Baxter, M.G., and Gaffan, D. 1998. Monkeys with rhinal cortex damage or neurotoxic hippocampal lesions are impaired on spatial scene learning and object reversals. *Behav. Neurosci.* **112**: 1291–1303.
- Naber, P.A., Lopes da Silva, F.H., and Witter, M.P. 2001. Reciprocal connections between the entorhinal cortex and hippocampal fields CA1 and the subiculum are in register with the projections from CA1 to the subiculum. *Hippocampus* **11**: 99–104.
- Nakazawa, K., Quirk, M.C., Chitwood, R.A., Watanabe, M., Yeckel, M.F., Sun, L.D., Kato, A., Carr, C.A., Johnston, D., Wilson, M.A., et al. 2002. Requirement for hippocampal CA3 NMDA receptors in associative memory recall. *Science* **297**: 211–218.
- Oja, E. 1982. A simplified neuron model as a principal component analyzer. *J. Math. Biol.* **15**: 267–273.
- O'Keefe, J. 1984. Spatial memory within and without the hippocampal system. In *Neurobiology of the hippocampus* (ed. W. Seifert), pp. 375–403. Academic Press, London.
- O'Keefe, J. and Dostrovsky, J. 1971. The hippocampus as a spatial map: Preliminary evidence from unit activity in the freely moving rat. *Brain Res.* **34**: 171–175.
- O'Keefe, J. and Nadel, L. 1978. *The hippocampus as a cognitive map*. Clarendon Press, Oxford.
- O'Reilly, R.C. and Rudy, J.W. 2001. Conjunctive representations in learning and memory: Principles of cortical and hippocampal function. *Psychol. Rev.* **108**: 311–345.
- Panzeri, S., Rolls, E.T., Battaglia, F., and Lavis, R. 2001. Speed of information retrieval in multilayer networks of integrate-and-fire neurons. *Network Comput. Neural Syst.* **12**: 423–440.
- Parkinson, J.K., Murray, E.A., and Mishkin, M. 1988. A selective mnemonic role for the hippocampus in monkeys: Memory for the location of objects. *J. Neurosci.* **8**: 4159–4167.
- Petrides, M. 1985. Deficits on conditional associative-learning tasks after frontal- and temporal-lobe lesions in man. *Neuropsychologia* **23**: 601–614.
- Pitkanen, A., Kelly, J.L., and Amaral, D.G. 2002. Projections from the lateral, basal, and accessory basal nuclei of the amygdala to the entorhinal cortex in the macaque monkey. *Hippocampus* **12**: 186–205.
- Ranck, J.B.J. 1985. Head direction cells in the deep cell layer of dorsal presubiculum in freely moving rats. In *Electrical activity of the archicortex* (eds. G. Buzsáki and C.H. Vanderwolf), pp. 217–220. Akademiai Kiado, Budapest.
- Rawlins, J.N.P. 1985. Associations across time: The hippocampus as a temporary memory store. *Behav. Brain Sci.* **8**: 479–496.
- Robertson, R.G., Rolls, E.T., and Georges-François, P. 1998. Spatial view cells in the primate hippocampus: Effects of removal of view details. *J. Neurophysiol.* **79**: 1145–1156.
- Robertson, R.G., Rolls, E.T., Georges-François, P., and Panzeri, S. 1999. Head direction cells in the primate pre-subiculum. *Hippocampus* **9**: 206–219.
- Rolls, E.T. 1987. Information representation, processing and storage in the brain: Analysis at the single neuron level. In *The neural and molecular bases of learning* (eds. J.-P. Changeux and M. Konishi), pp. 503–540. Wiley, Chichester.
- Rolls, E.T. 1989a. Functions of neuronal networks in the hippocampus and cerebral cortex in memory. In *Models of brain function* (ed. R.M.J. Cotterill), pp. 15–33. Cambridge University Press, Cambridge.
- Rolls, E.T. 1989b. Functions of neuronal networks in the hippocampus and neocortex in memory. In *Neural models of plasticity: Experimental and theoretical approaches* (eds. J.H. Byrne and W.O. Berry), pp. 240–265. Academic Press, San Diego.
- Rolls, E.T. 1989c. The representation and storage of information in neuronal networks in the primate cerebral cortex and hippocampus. In *The computing neuron* (eds. R. Durbin et al.), pp. 125–159. Addison-Wesley, Wokingham, UK.
- Rolls, E.T. 1990a. Functions of the primate hippocampus in spatial processing and memory. In *Neurobiology of comparative cognition* (eds. D.S. Olton and R.P. Kesner), pp. 339–362. Erlbaum, Hillsdale, NJ.
- Rolls, E.T. 1990b. Theoretical and neurophysiological analysis of the functions of the primate hippocampus in memory. *Cold Spring Harb. Symp. Quant. Biol.* **55**: 995–1006.
- Rolls, E.T. 1991. Functions of the primate hippocampus in spatial and non-spatial memory. *Hippocampus* **1**: 258–261.
- Rolls, E.T. 1992. Neurophysiological mechanisms underlying face processing within and beyond the temporal cortical visual areas. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **335**: 11–21.
- Rolls, E.T. 1994. Neurophysiological and neuronal network analysis of how the primate hippocampus functions in memory. In *The memory system of the brain* (ed. J. Delacour), pp. 713–744. World Scientific,

- London.
- Rolls, E.T. 1995. A model of the operation of the hippocampus and entorhinal cortex in memory. *Int. J. Neural Syst.* **6**: 51–70.
- Rolls, E.T. 1996a. Roles of long term potentiation and long term depression in neuronal network operations in the brain. In *Cortical plasticity* (eds. M.S. Fazeli and G.L. Collingridge), pp. 223–250. Bios, Oxford.
- Rolls, E.T. 1996b. A theory of hippocampal function in memory. *Hippocampus* **6**: 601–620.
- Rolls, E.T. 1999. Spatial view cells and the representation of place in the primate hippocampus. *Hippocampus* **9**: 467–480.
- Rolls, E.T. 2000. Memory systems in the brain. *Annu. Rev. Psychol.* **51**: 599–630.
- Rolls, E.T. 2003. Consciousness absent and present: A neurophysiological exploration. *Prog. Brain Res.* **144**: 95–106.
- Rolls, E.T. 2008. *Memory, attention, and decision-making: A unifying computational neuroscience approach*. Oxford University Press, Oxford.
- Rolls, E.T. and Deco, G. 2002. *Computational neuroscience of vision*. Oxford University Press, Oxford.
- Rolls, E.T. and Kesner, R.P. 2006. A computational theory of hippocampal function, and empirical tests of the theory. *Prog. Neurobiol.* **79**: 1–48.
- Rolls, E.T. and Stringer, S.M. 2005. Spatial view cells in the hippocampus, and their idiothetic update based on place and head direction. *Neural Netw.* **18**: 1229–1241.
- Rolls, E.T. and Treves, A. 1990. The relative advantages of sparse versus distributed encoding for associative neuronal networks in the brain. *Network* **1**: 407–421.
- Rolls, E.T. and Treves, A. 1998. *Neural networks and brain function*. Oxford University Press, Oxford.
- Rolls, E.T. and Xiang, J.-Z. 2005. Reward-spatial view representations and learning in the hippocampus. *J. Neurosci.* **25**: 6167–6174.
- Rolls, E.T. and Xiang, J.-Z. 2006. Spatial view cells in the primate hippocampus, and memory recall. *Rev. Neurosci.* **17**: 175–200.
- Rolls, E.T., Miyashita, Y., Cahusac, P.M.B., Kesner, R.P., Niki, H., Feigenbaum, J., and Bach, L. 1989. Hippocampal neurons in the monkey with activity related to the place in which a stimulus is shown. *J. Neurosci.* **9**: 1835–1845.
- Rolls, E.T., Robertson, R.G., and Georges-François, P. 1997a. Spatial view cells in the primate hippocampus. *Eur. J. Neurosci.* **9**: 1789–1794.
- Rolls, E.T., Treves, A., Foster, D., and Perez-Vicente, C. 1997b. Simulation studies of the CA3 hippocampal subfield modelled as an attractor neural network. *Neural Netw.* **10**: 1559–1569.
- Rolls, E.T., Treves, A., Robertson, R.G., Georges-François, P., and Panzeri, S. 1998. Information about spatial view in an ensemble of primate hippocampal cells. *J. Neurophysiol.* **79**: 1797–1813.
- Rolls, E.T., Stringer, S.M., and Trappenberg, T.P. 2002. A unified model of spatial and episodic memory. *Proc. R. Soc. Lond. B. Biol. Sci.* **269**: 1087–1093.
- Rolls, E.T., Xiang, J.-Z., and Franco, L. 2005. Object, space and object-space representations in the primate hippocampus. *J. Neurophysiol.* **94**: 833–844.
- Rolls, E.T., Stringer, S.M., and Elliot, T. 2006. Entorhinal cortex grid cells can map to hippocampal place cells by competitive learning. *Network Comput. Neural Syst.* **17**: 447–465.
- Rupniak, N.M.J. and Gaffan, D. 1987. Monkey hippocampus and learning about spatially directed movements. *J. Neurosci.* **7**: 2331–2337.
- Samsonovich, A. and McNaughton, B.L. 1997. Path integration and cognitive mapping in a continuous attractor neural network model. *J. Neurosci.* **17**: 5900–5920.
- Schultz, S. and Rolls, E.T. 1999. Analysis of information transmission in the Schaffer collaterals. *Hippocampus* **9**: 582–598.
- Simmen, M.W., Treves, A., and Rolls, E.T. 1996. Pattern retrieval in threshold-linear associative nets. *Network* **7**: 109–122.
- Skaggs, W.E., Knierim, J.J., Kudrimoti, H.S., and McNaughton, B.L. 1995. A model of the neural basis of the rat's sense of direction. In *Advances in neural information processing systems* (eds. G. Tesoro et al.), pp. 173–180. MIT Press, Cambridge, MA.
- Smith, M.L. and Milner, B. 1981. The role of the right hippocampus in the recall of spatial location. *Neuropsychologia* **19**: 781–793.
- Stefanacci, L., Suzuki, W.A., and Amaral, D.G. 1996. Organization of connections between the amygdaloid complex and the perirhinal and parahippocampal cortices in macaque monkeys. *J. Comp. Neurol.* **375**: 552–582.
- Storm-Mathieson, J., Zimmer, J., and Ottersen, O.P., eds. 1990. *Understanding the brain through the hippocampus*. Elsevier, Oxford.
- Stringer, S.M. and Rolls, E.T. 2002. Invariant object recognition in the visual system with novel views of 3D objects. *Neural Comput.* **14**: 2585–2596.
- Stringer, S.M. and Rolls, E.T. 2006. Self-organizing path integration using a linked continuous attractor and competitive network: Path integration of head direction. *Network Comput. Neural Syst.* **17**: 419–445.
- Stringer, S.M., Rolls, E.T., Trappenberg, T.P., and de Araujo, I.E.T. 2002a. Self-organizing continuous attractor networks and path integration. Two-dimensional models of place cells. *Network Comput. Neural Syst.* **13**: 429–446.
- Stringer, S.M., Trappenberg, T.P., Rolls, E.T., and de Araujo, I.E.T. 2002b. Self-organizing continuous attractor networks and path integration: One-dimensional models of head direction cells. *Network Comput. Neural Syst.* **13**: 217–242.
- Stringer, S.M., Rolls, E.T., and Trappenberg, T.P. 2004. Self-organising continuous attractor networks with multiple activity packets, and the representation of space. *Neural Netw.* **17**: 5–27.
- Stringer, S.M., Rolls, E.T., and Trappenberg, T.P. 2005. Self-organizing continuous attractor network models of hippocampal spatial view cells. *Neurobiol. Learn. Mem.* **83**: 79–92.
- Sutherland, R.J. and Rudy, J.W. 1991. Exceptions to the rule of space. *Hippocampus* **1**: 250–252.
- Suzuki, W.A. and Amaral, D.G. 1994a. Perirhinal and parahippocampal cortices of the macaque monkey—Cortical afferents. *J. Comp. Neurol.* **350**: 497–533.
- Suzuki, W.A. and Amaral, D.G. 1994b. Topographic organization of the reciprocal connections between the monkey entorhinal cortex and the perirhinal and parahippocampal cortices. *J. Neurosci.* **14**: 1856–1877.
- Tabuchi, E., Mulder, A.B., and Wiener, S.I. 2003. Reward value invariant place responses and reward site associated activity in hippocampal neurons of behaving rats. *Hippocampus* **13**: 1117–1132.
- Taube, J.S., Muller, R.U., and Ranck, J.B.J. 1990. Head-direction cells recorded from the postsubiculum in freely moving rats. 1: Description and quantitative analysis. *J. Neurosci.* **10**: 420–435.
- Taube, J.S., Goodridge, J.P., Golob, E.J., Dudchenko, P.A., and Stackman, R.W. 1996. Processing the head direction signal: A review and commentary. *Brain Res. Bull.* **40**: 477–486.
- Taylor, J.G. 1999. Neural “bubble” dynamics in two dimensions: Foundations. *Biol. Cybern.* **80**: 393–409.
- Treves, A. 1990. Graded-response neurons and information encodings in autoassociative memories. *Phys. Rev. A* **42**: 2418–2430.
- Treves, A. 1993. Mean-field analysis of neuronal spike dynamics. *Network* **4**: 259–284.
- Treves, A. 1995. Quantitative estimate of the information relayed by Schaffer collaterals. *J. Comput. Neurosci.* **2**: 259–272.
- Treves, A. and Rolls, E.T. 1991. What determines the capacity of autoassociative memories in the brain? *Network* **2**: 371–397.
- Treves, A. and Rolls, E.T. 1992. Computational constraints suggest the need for two distinct input systems to the hippocampal CA3 network. *Hippocampus* **2**: 189–199.
- Treves, A. and Rolls, E.T. 1994. A computational analysis of the role of the hippocampus in memory. *Hippocampus* **4**: 374–391.
- van Haeften, T., Baks-te-Bulte, L., Goede, P.H., Wouterlood, F.G., and Witter, M.P. 2003. Morphological and numerical analysis of synaptic interactions between neurons in deep and superficial layers of the entorhinal cortex of the rat. *Hippocampus* **13**: 943–952.
- Van Hoesen, G.W. 1982. The parahippocampal gyrus. New observations regarding its cortical connections in the monkey. *Trends Neurosci.* **5**: 345–350.
- Wallace, D.G. and Whishaw, I.Q. 2003. NMDA lesions of Ammon's horn and the dentate gyrus disrupt the direct and temporally paced homing displayed by rats exploring a novel environment: Evidence for a role of the hippocampus in dead reckoning. *Eur. J. Neurosci.* **18**: 513–523.
- Whishaw, I.Q., Hines, D.J., and Wallace, D.G. 2001. Dead reckoning (path integration) requires the hippocampal formation: Evidence from spontaneous exploration and spatial learning tasks in light (allothetic) and dark (idiothetic) tests. *Behav. Brain Res.* **127**: 49–69.
- Witter, M.P. 1993. Organization of the entorhinal-hippocampal system: A review of current anatomical data. *Hippocampus* **3**: 33–44.
- Witter, M.P., Van Hoesen, G.W., and Amaral, D.G. 1989. Topographical organisation of the entorhinal projection to the dentate gyrus of the monkey. *J. Neurosci.* **9**: 216–228.
- Witter, M.P., Naber, P.A., van Haeften, T., Machielsen, W.C., Rombouts, S.A., Barkhof, F., Scheltens, P., and Lopes da Silva, F.H. 2000a. Cortico-hippocampal communication by way of parallel parahippocampal-subicular pathways. *Hippocampus* **10**: 398–410.
- Witter, M.P., Wouterlood, F.G., Naber, P.A., and Van Haeften, T. 2000b. Anatomical organization of the parahippocampal-hippocampal network. *Ann. N. Y. Acad. Sci.* **911**: 1–24.
- Wood, E.R., Dudchenko, P.A., and Eichenbaum, H. 1999. The global record of memory in hippocampal neuronal activity. *Nature* **397**: 613–616.
- Zhang, K. 1996. Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: A theory. *J. Neurosci.* **16**: 2112–2126.

Received May 17, 2007; accepted in revised form July 18, 2007.