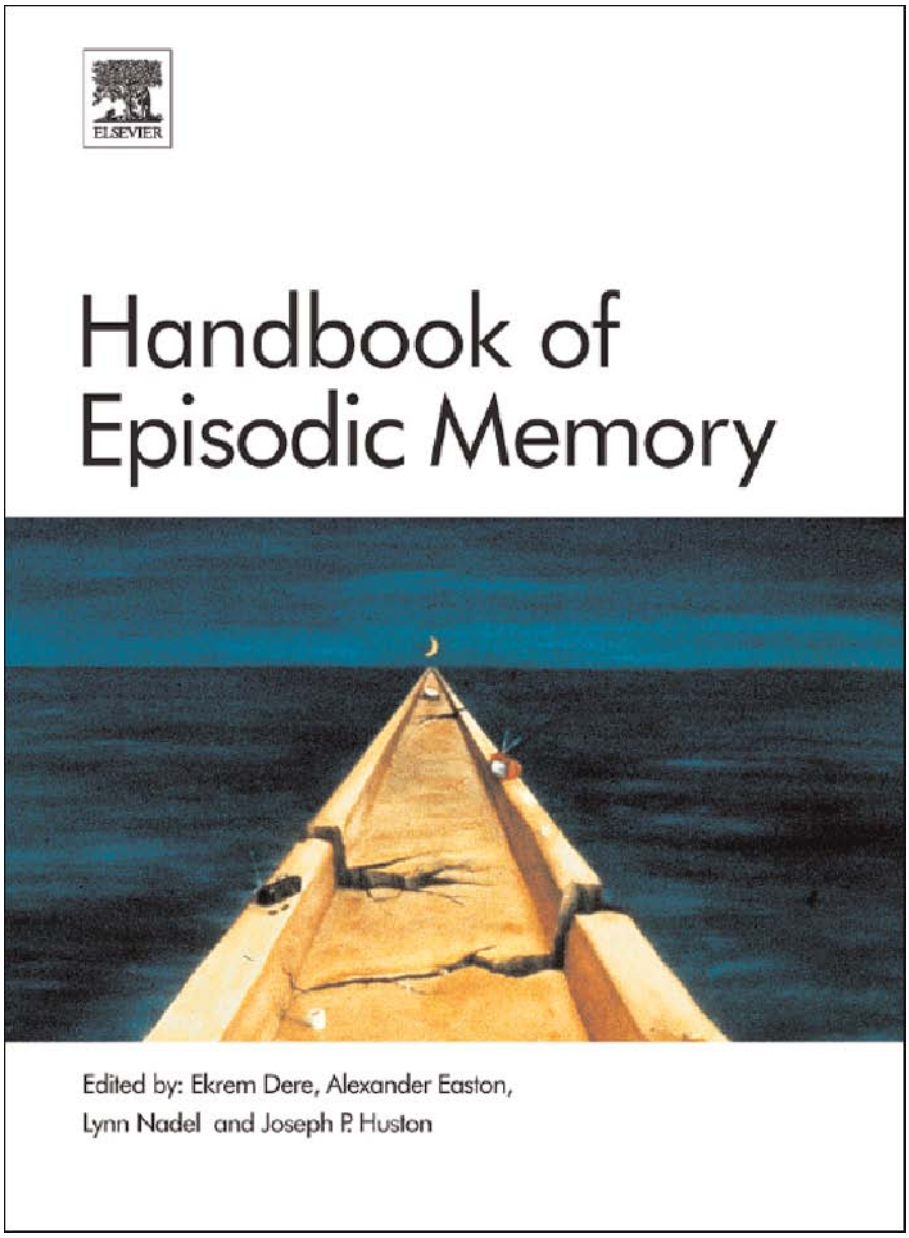




From Edmund T. Rolls, The primate hippocampus and episodic memory. In: J. P. Huston, editor: Handbook of Behavioral Neuroscience, Vol 18, Handbook of Episodic Memory, Ekrem Dere, Alexander Easton, Lynn Nadel and Joseph P. Huston. The Netherlands: Elsevier, 2008, pp. 417–438.
ISBN: 978-0-444-53174-2

CHAPTER 4.2

# The primate hippocampus and episodic memory

Edmund T. Rolls*

*Department of Experimental Psychology, University of Oxford, Oxford, UK*

**Abstract:** Humans can remember where in space they have seen an object. In this example of one-trial event or episodic memory, an association must be formed between an object and a place "out there," and the place can be recalled with the object as a recall cue, or the object can be recalled with the place as a recall cue. Hippocampal spatial view neurons in primates provide allocentric representations of a view of space "out there." The responses depend on where the monkey is looking; and can be updated by idiothetic (self-motion) inputs provided by eye movements when the view is hidden. In a room-based object–place memory task, some hippocampal neurons respond to the objects shown, some to the places viewed, and some to combinations of the places viewed and the objects present in those locations. In an object–place recall task when the location in space at which an object has been seen is recalled by the presentation of the object, some primate hippocampal neurons maintain their responding to the object recall cue in a delay period without the object being visible while the place is being recalled; and other neurons respond to the place being recalled. Other spatial view neurons form associations with the rewards present at particular locations in space. These findings and computational neuroscience neural network models of how the hippocampus could implement episodic memory help to show how the primate including human hippocampus is involved in episodic memory.

**Keywords:** spatial view neurons; object–place memory; macaque; recall; nonhuman primate; computational neuroscience; neural network models of memory; hippocampus

## I. Introduction

In this chapter the functions of the primate hippocampal system in episodic memory are considered. Episodic memory, the memory of a particular episode, requires the ability to remember particular events, and to distinguish them from other events. An event consists of a set of items that occur together, such as seeing a particular object or person's face in a particular place. An everyday example might be remembering where one was for dinner, who was present, what was eaten, what was discussed, and the time at which it occurred. The spatial context is almost always an important part of an episodic memory, and it may be partly for this reason that episodic memory is linked to the functions of the hippocampal system, which is involved in spatial processing and memory. In this chapter I will show that the primate hippocampus has a special representation of space that makes it particularly appropriate for episodic memory in primates including humans, for it is a representation of space "out there." This enables memories to be formed of what one has seen at a particular place, even if one has not been to the place. This is not possible with rodent place cells, which respond to the place where the rodent is located. I will show that the primate hippocampus has more than only a spatial representation, for it also represents objects that are seen at particular places, and

*Corresponding author.
E-mail: Edmund.Rolls@oxcns.org

*E. Dere, A. Easton, L. Nadel and J.P. Huston (Eds.) Handbook of Episodic Memory*, Vol. 18
ISBN 978-0-444-53174-2



DOI: 10.1016/S1569-7339(08)00223-3

rewards that are found at particular places in spatial scenes. I will also show that primate hippocampal neurons are activated when a memory must be recalled from a part of a memory, for example, when the place at which an object was shown must be recalled when the object is seen alone. The ability to recall a whole memory from a partial cue is an important property of episodic memory. I then consider how the network architecture of the hippocampal system may enable it to implement episodic memory, and to recall the whole of a memory back in the neocortex when a partial retrieval cue is present. I focus in this Chapter on a fundamental component of episodic memory, the ability to store and retrieve the memory of a particular single event involving an association between items such as the place, and the object or reward seen at that place. Episodic memory in the sense of a series of linked events requires this type of event memory, and could be implemented by linking together a series of events.

## II. Systems-level functions of the primate hippocampus

Any theory of the hippocampus must state at the systems level what is computed by the hippocampus. Some of the relevant evidence comes from the effects of damage to the hippocampus, the responses of neurons in the hippocampus during behavior, and the systems-level connections of the hippocampus, described in more detail elsewhere (Rolls and Kesner, 2006; Rolls, 2008). As described in the introduction, evidence of spatial representations in the hippocampus, and of whether these spatial representations can be combined with object representations in memory tasks such as object–place memory, are relevant to the understanding of the functions of the primate hippocampus in episodic memory. In object–place memory, the place where an object was seen must be remembered. A related task is whole scene memory, in which a monkey must learn which locations in each of a number of scenes if touched, lead to reward. The task can be thought of as one in which a place in a scene and reward must be associated together by learning.

### *II.A. 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 were 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 nontrial-unique stimuli and with many object–place pairs would be impaired by neurotoxic hippocampal lesions.) Further, 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). 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 the 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; Jarrard, 1993; Cassaday and Rawlins, 1997; Martin et al., 2000;

Kesner et al., 2004), to utilize spatial cues or bridge delays (Rawlins, 1985; Kesner, 1998; Kesner and Rolls, 2001; Kesner et al., 2004; Rolls and Kesner, 2006), or to perform relational operations on remembered material (Eichenbaum, 1997).

Many of these memory functions are important in event or episodic memory, in which the ability to remember what happened where, on a typically single occasion is important. It will be 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, remembering 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 on the basis of such accumulating information the semantic representation (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).

### *II.B. Systems-level neurophysiology of the primate hippocampus*

The systems-level neurophysiology of the hippocampus shows what information could be stored or processed by the hippocampus. 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 has been reviewed by Rolls and Xiang (2006), and a summary is provided here because it provides a perspective relevant to the understanding of 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).

#### *II.B.1. Spatial view neurons in the primate hippocampus*

The primate hippocampus contains spatial cells that respond when the monkey looks at a certain part of the 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 about 2.4% combine information about the position in space with information about the object that is in that position in the 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. If the view details are obscured by curtains and darkness, then some spatial view neurons (especially those in CA1 and less those in CA3) continue to respond when the monkey looks towards the spatial view field, showing that these neurons can be updated for at least short periods by idiothetic

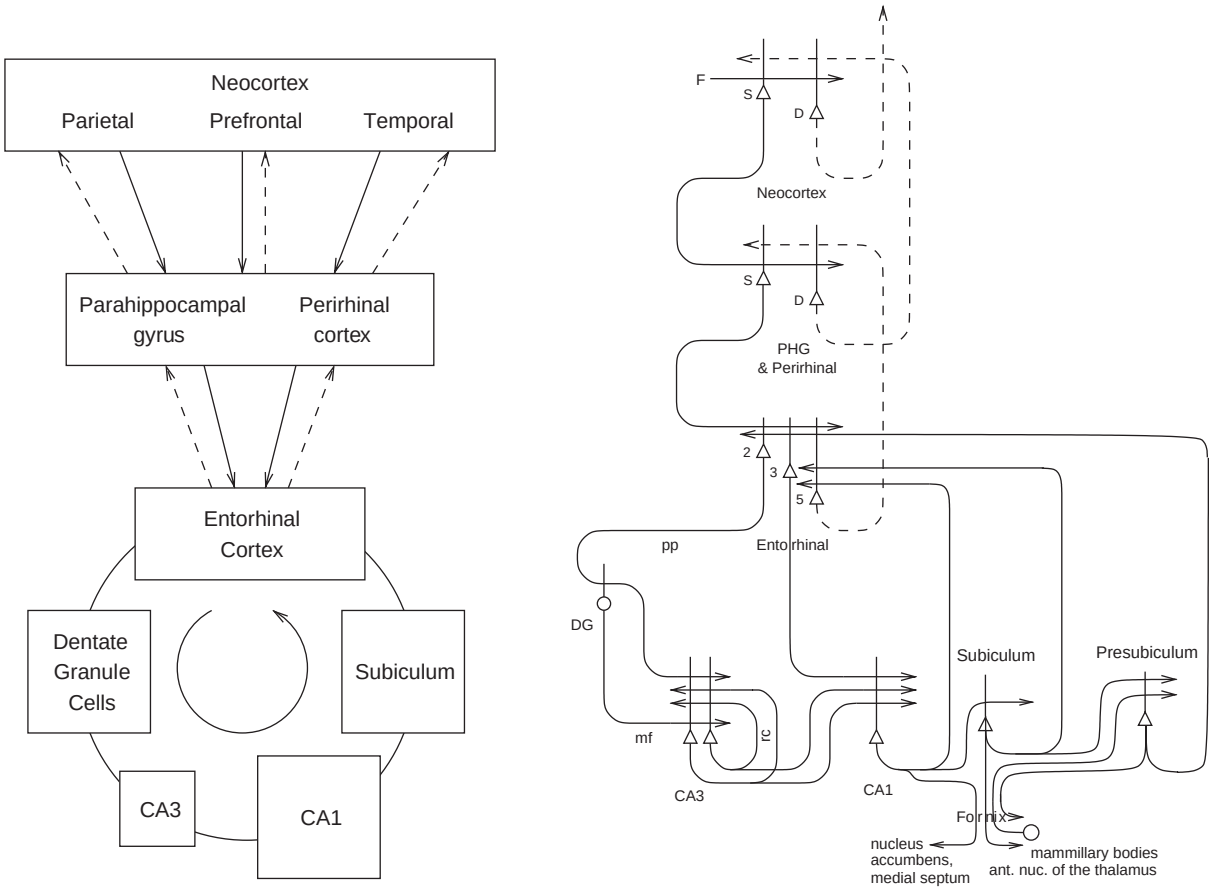


Fig. 1. Forward connections (solid lines) from areas of cerebral association neocortex via the parahippocampal gyrus and perirhinal cortex, and entorhinal cortex, to the hippocampus; and backprojections (dashed lines) via the hippocampal CA1 pyramidal cells, subiculum, and parahippocampal gyrus to the 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 backprojections. Left: Block diagram. Right: More detailed representation of some of the principal excitatory neurons in the pathways. Abbreviations: 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.

(self-motion) cues including eye position and head-direction signals (Rolls et al., 1997b; Robertson et al., 1998).

### *II.B.2. 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. 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 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, are 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 backprojections 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, it is argued below that the hippocampus may be able to make a special contribution to event or episodic memory, by enabling in the CA3 network with its very widespread recurrent collateral connections and association between any one item with any other item to form an arbitrary association to represent an event.

### *II.B.3. 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 Fig. 2. 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. A number of these neurons had activity that was related to the recall process. An example of one of these neurons is shown in Fig. 3. The neuron had activity that was greater to object one 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 was being recalled from the object, this type of neuron continued to respond as if the object was present, that it has kept the representation of the object active after the object was no longer visible, and the place to touch was being recalled. 16 of the neurons responded in this way, and an additional 6 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

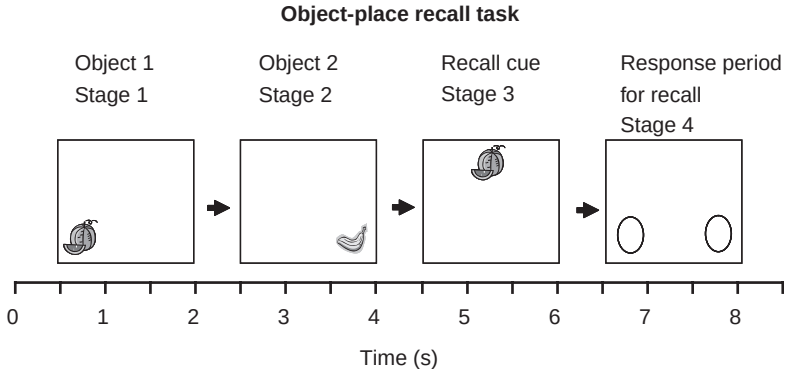


Fig. 2. The object–place recall task. One trial is shown. After a 0.5 s tone to indicate the start of a trial, in stage 1 one of two objects (O1) is shown at one of the places (P1). (The object and the place are chosen randomly on each trial.) To ensure that the monkey sees the stimulus, the monkey can touch the screen at the place to obtain one drop of juice reward by licking. After a 0.5 s 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 s delay, in stage 3, the recall cue, one of the objects chosen at random, is shown at the top of the screen in the middle. (One drop of fruit juice was available as in stage 1.) After a 0.5 s 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. (See Color Plate 4.2.2 in Color Plate Section.)

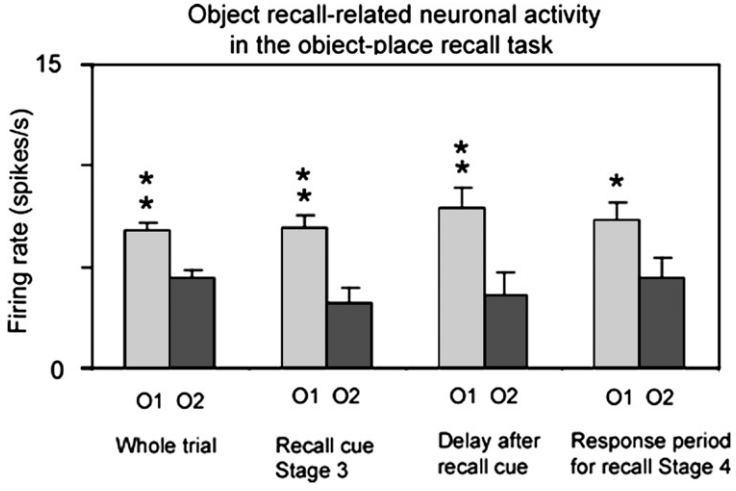


Fig. 3. Activity of a neuron with responses related to one of the objects used in the object–place recall task. The firing rate to object 1 (O1) and object 2 (O2) are shown (mean firing rate in spikes/s across trials ± SEM). The first histogram pair (on the left) shows the responses to the two objects measured throughout the trial whenever O1 or O2 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 s delay period after one of the objects had been shown in stage 3 as the recall cue. The neuron continued to respond more after O1 than after O2 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. $^{**} = p < 0.01$; $^{*} = p < 0.05$.

shown in stages 1 or 2 of each trial. An example of one of these neurons is shown in Fig. 4. 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 on 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 due to some response bias the monkey touched the incorrect side, the neuron could still respond according to the correct recalled location. 13 neurons had differential responses 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 was 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 arise by chance within the 347 tests performed and were able to reject this with

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

| | |
|---|---|
| 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 |

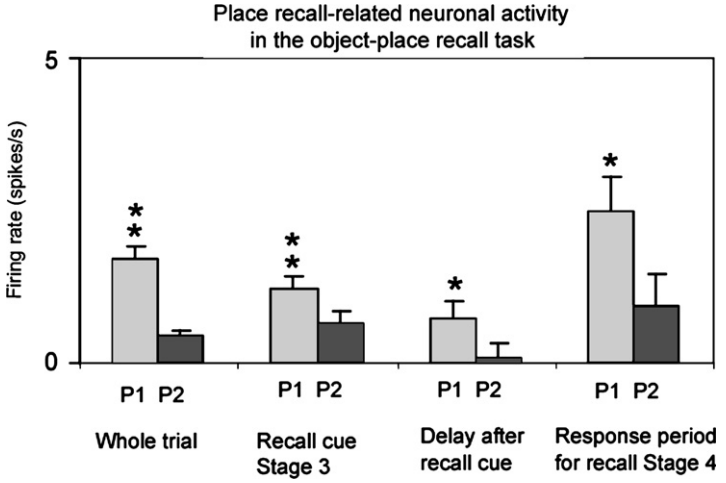


Fig. 4. Activity of a neuron with responses related to the left place (P1) in the object–place recall task. The firing rate to place 1 (P1) and place 2 (P2) are shown (mean firing rate in spikes/s across trials ± SEM). The first histogram pair (on the left) shows the responses to the two places measured 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 P1 or P2. The third histogram pair shows the neuronal responses in the 0.5 s delay period after one of the objects had been shown in stage 3 as the recall cue. The neuron responded more when P1 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. $^{**} = p < 0.01$; $^{*} = p < 0.05$.

$p < 0.05$. 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) to the most effective object or place was $7.2 \pm 0.6$ spikes/s ($\pm$SEM), and their mean spontaneous rate was $3.2 \pm 0.6$ spikes/s.

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 the object–place recall process that is hypothesized to take place in the CA3 cells. 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).

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. 2), 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.

#### *II.B.4. 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 video monitor had one location which if touched yielded a preferred fruit juice reward, and a second location which 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. 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).

### *II.C. Systems-level anatomy*

To understand the functions of the primate hippocampus in event or episodic memory, it is necessary to understand which other parts of the brain it receives information from. Does it for example, receive object as well as spatial

information in terms of its connectivity? 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; Rolls and Kesner, 2006; Rolls, 2008) (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 projects to 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), though there are other outputs (Rolls and Kesner, 2006). These are the pathways that are likely to be involved in the recall of information from the hippocampus.

## III. The operation of hippocampal circuitry as a memory system

In this section, I consider how event or episodic memories might be learned and retrieved by hippocampal circuitry, and in addition retrieved back into the neocortex.

### *III.A. Hippocampal circuitry*

See Fig. 1 and see also Amaral and Witter (1989); Storm-Mathiesen et al. (1990); Amaral (1993); Witter et al. (2000b); Naber et al. (2001); and 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 are projected to CA3 cells via the mossy fibers (mf), which provide a *sparse* but possibly powerful connection to the $3 \times 10^5$ CA3 pyramidal cells in the rat. Each CA3 cell receives approximately 50 mossy fibers inputs, so that the sparseness of this connectivity is thus 0.005%. By contrast, there are many more – possibly weaker – direct perforant path inputs also from layer 2 of the entorhinal cortex onto each CA3 cell, in the rat of the order of $4 \times 10^3$. The largest number of synapses (about $1.2 \times 10^4$ 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; Amaral et al., 1990; Ishizuka et al., 1990; Amaral and Witter, 1995), so that effectively the CA3 system provides a single network, with a connectivity of approximately 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).

### *III.B. CA3 as an autoassociation memory*

#### *III.B.1. Arbitrary associations, and pattern completion in recall*

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 which 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).) The architecture of an autoassociation network is shown in Fig. 5, and the learning rule for the change in the synaptic weight is as shown in Eq. (1) (Rolls and Treves, 1998; Rolls and Deco, 2002).

$$\delta w_{ij} = \mathrm{k} \cdot r_i \cdot r'_j \tag{1}$$

where k is a constant, $r_i$ the activation of the dendrite (the postsynaptic term), $r'_j$ the presynaptic firing rate, and $w_{ij}$ 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,

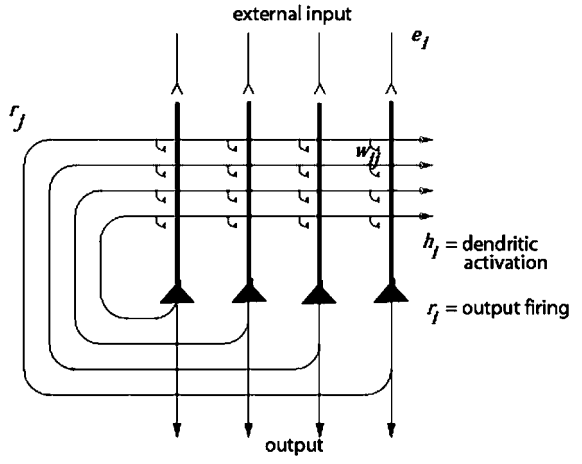


Fig. 5. The architecture of an autoassociation or attractor neural network (CANN).

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 which 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 in the remainder of Section III.B and in more detail elsewhere (Rolls and Kesner, 2006; Rolls, 2008).

### *III.B.2. Storage capacity*

We have performed quantitative analyzes 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 $C^{\mathrm{RC}}$ 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.[1] The neuronal population sparseness $a$ of the representation can

[1] Each 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.

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

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

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,

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

where $k$ is a factor that depends weakly on the detailed structure of the rate distribution, on the connectivity pattern, etc., but is roughly in the order of 0.2–0.3 (Treves and Rolls, 1991). For example, for $C^{\mathrm{RC}} = 12{,}000$ and $a = 0.02$, $p_{\max}$ is calculated to be approximately 36,000. This analysis emphasizes the utility of having a sparse representation in the hippocampus, for this enables many different memories to be stored. [The sparseness $a$ in this equation is strictly the population sparseness (Treves and Rolls, 1991; Franco et al., 2007). The population sparseness $a^{\mathrm{p}}$ 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 $a^{\mathrm{s}}$ would be measured by the distribution of firing rates to a set of stimuli, which would take a long time. The selectivity or sparseness $a^{\mathrm{s}}$ 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).] (I note that the sparseness estimates obtained by measuring early gene changes, which are effectively population sparsenesses, would be expected to depend greatly on the range of environments or stimuli in which this was measured. If the environment was restricted to one stimulus, this would reflect the population sparseness. If the environment was changing, the measure from early gene changes would be rather undefined, 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). Simulations that are fully consistent with the analytic theory are provided by Simmen et al. (1996) and Rolls et al. (1997b).

A number of 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, 2008; Rolls, 1996). 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.

*III.B.3. Recall and completion*

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 and completed from any part (Rolls and Treves, 1998; Rolls and Kesner, 2006; Rolls, 2008). For example, in an object–place autoassociation memory, an object could be recalled from a place retrieval cue and vice versa. In a test of this, Day, Langston, and Morris (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.

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 Roullet, 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; Gold and Kesner, 2005). 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.

*III.B.4. 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), Taylor (1999), Samsonovich and McNaughton (1997), Battaglia and Treves (1998), Stringer et al. (2002b), Stringer et al. (2002a), Stringer et al. (2004), Stringer and Rolls (2002), and Rolls and Stringer (2005) (see Rolls and Deco, 2002; Rolls, 2008) and are described briefly next.

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 view, spatial position, head direction, etc. 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 has started to represent the current state (head direction, position, 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 Fig. 5. The main difference is that the patterns stored in a

CANN are continuous patterns, with each neuron having broadly tuned firing which 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 Fig. 6. 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. 6). 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.

Space is continuous and object representations are discrete. If these representations are to be combined, for example, in 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. 6), 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).

If spatial representations are stored in the hippocampus, the important issue arises in terms

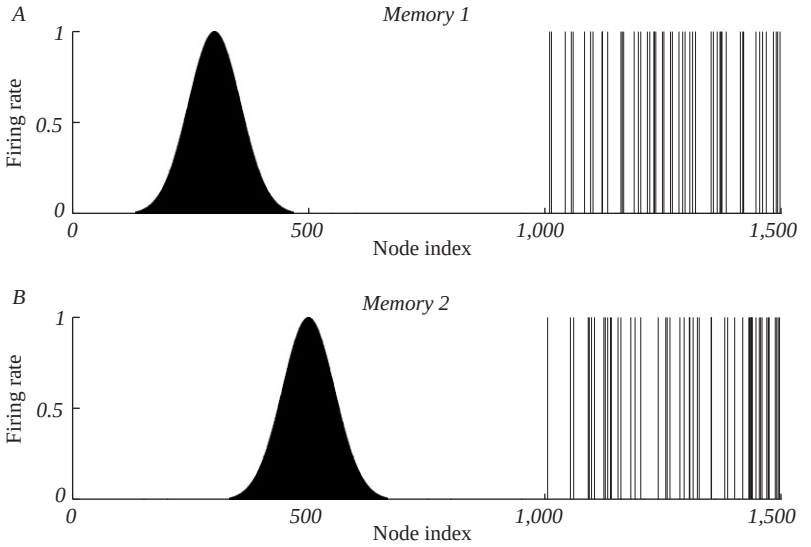


Fig. 6. The types of firing patterns stored in continuous attractor networks are illustrated for the patterns present on neurons 1–1,000 for Memory 1 (when the firing is that produced when the spatial state represented is that for location 300), and for Memory 2 (when the firing is that produced when 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 1,001–1,500. 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, etc., and each memory contains part that is continuously distributed to represent physical space, and part that represents a discrete event or object.

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), many different maps can be simultaneously stored in a continuous attractor network (Battaglia and Treves, 1998).

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 elsewhere (McNaughton et al., 2006; Samsonovich and McNaughton, 1997; Stringer et al., 2002a, 2002b) and has been applied to 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.

*III.B.5. Mossy fibers 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, 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. 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. The hypothesis is that the mossy fiber sparse connectivity solution performs the appropriate function for learning to operate correctly in CA3. 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 and other points, we predicted that the mossy fibers may be necessary for new learning in the hippocampus, but may not be necessary for the recall of existing memories from the hippocampus (Treves and Rolls, 1992; Rolls and Treves, 1998; Rolls, 2008). 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.

*III.B.6. 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. On the other hand, an autoassociative memory network needs afferent inputs to apply the retrieval cue to the network. We have shown (Treves and Rolls, 1992) that the perforant path system is likely to be the one involved in relaying the cues that initiate retrieval.

### *III.C. Dentate granule cells*

The theory is that the dentate granule cell stage of hippocampal processing which precedes the CA3 stage acts as a competitive network in a number of 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). An important property for episodic memory is that the dentate by acting in this way would perform pattern separation, enabling the hippocampus to store different memories of even similar events, and this prediction has been confirmed (Gilbert et al., 2001; Goodrich-Hunsaker et al., 2005; Rolls and Kesner, 2006; Rolls, 2008).

As just described, the dentate granule cells could be 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).

### *III.D. CA1 cells*

The CA3 cells connect to the CA1 cells by the Schaeffer collateral synapses. The associative modifiability in this connection helps the full information present in CA3 to be retrieved in the CA1 neurons (Treves and Rolls, 1994; Rolls, 1995; Treves, 1995; Schultz and Rolls, 1999).

### *III.E. Backprojections to the neocortex and memory recall*

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 backprojections 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. 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 backprojections (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, which 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 backprojections, 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, 1989c, 1990a, b), or in organizing actions.

The hypothesis of the architecture with which this would be achieved is shown in Fig. 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 backprojections allow for divergence back to neocortical areas. The way in which I suggest that the backprojection 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 towards 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 backprojecting synapses to the entorhinal cortex. Here the backprojecting synapses from active backprojection 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 with cortical areas, as shown in Fig. 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 backprojected activity (indirectly) from the hippocampus which would by pattern association from the backprojecting 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 backprojecting 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 backprojection pathway, as made evident in Fig. 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 chapter 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 hippocampo-directed recall operation, is that the signal generated from the hippocampally retrieved pattern of activity, and carried backwards towards 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 backprojecting 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 (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. ($k'$ 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 latter with that shown here, 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 $C^{HBP} = C^{RC} a_{nc}/a_{CA3}$, 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.

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 backprojections. 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 backprojecting 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 backprojections thus provides a quantitative account of why the cerebral cortex has as many backprojection as forward projection connections. Further aspects of the operation of the backprojecting systems are described elsewhere (Rolls, 2008).

## IV. Discussion

The present theory holds that the hippocampus is used for the formation of episodic memories using autoassociation. This function is often necessary for successful spatial computation, but is not itself spatial computation. Instead, I believe that spatial computation is more likely to be performed in the neocortex (utilizing information if necessary recalled from the hippocampus). Consistent with this view, hippocampal damage impairs the ability to learn new environments but not to perform spatial computations such as finding one's way to a place in a familiar environment, whereas damage to the parietal cortex and parahippocampal cortex can lead to problems such as topographical and other spatial agnosias, in humans (see Kolb and Whishaw, 2003). This is consistent with spatial computations normally being performed in the neocortex. (In monkeys, there is evidence for a role of the parietal cortex in allocentric spatial computation. For example, monkeys with parietal cortex lesions are impaired at performing a landmark task, in which the object to be chosen is signified by the proximity to it of a "landmark" (another object) (Ungerleider and Mishkin, 1982)).

A theory closely related to the present theory of how the hippocampus operates has been developed by McClelland et al (1995). It is very similar to the theory we have developed (Rolls, 1987, 1989a, b, c, 2008; Treves and Rolls, 1992, 1994) at the systems level, except that it takes a stronger position on the gradient of retrograde amnesia, emphasizes that recall from the hippocampus of episodic information is used to help build semantic representations in the neocortex, and holds that the last set of synapses that are modified rapidly during the learning of each episode are those between the CA3 and the CA1 pyramidal cells, as described above (see Fig. 1). It also emphasizes the important point that the hippocampal and neocortical memory systems may be quite different, with the hippocampus specialized for the rapid learning of single events or episodes, and the neocortex for the slower learning of semantic representations which may necessarily benefit from the many exemplars needed to shape the semantic representation.

In this chapter we have thus seen that there is neurophysiological evidence that different representations of the type important in episodic memory, including object and place, and reward

and place, are brought together in the primate hippocampus, and even that the whole representation can be completed from a partial retrieval cue in a one-trial object–place memory task. It appears to be a property of the hippocampus that it is involved in associations when one of the associates is place (Rolls and Kesner, 2006). We have moreover seen that the representation of space in the primate hippocampus provided by spatial view cells is appropriate for episodic memory in primates including humans, for it is a representation of space "out there," which is prototypical of the spatial representation that is involved in episodic memory in primates including humans. We have also provided an overview of a theory of how the hippocampal system could implement episodic memory (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-Pew Centre in Cognitive Neuroscience. The author has performed the experimental and theoretical work, which 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: Mountcastle, V.B. (Ed.), Handbook of Physiology Section 1: The Nervous System. American Physiological Society, Washington, pp. 211–294.

Amaral, D.G. (1993) Emerging principles of intrinsic hippocampal organisation. Curr. Opin. Neurobiol., 3: 225–229.

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: Aggleton, J.P. (Ed.), The Amygdala. Wiley-Liss, New York, pp. 1–66.

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: Paxinos, G. (Ed.), The Rat Nervous System. Academic Press, San Diego, pp. 443–493.

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. (1998) Attractor neural networks storing multiple space representations: a model for hippocampal place fields. Phys. Rev. E, 58: 7738–7753.

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: Bolhuis, J.J. (Ed.), Brain, Perception, and Memory: Advances in Cognitive Neuroscience. Oxford University Press, Oxford, pp. 279–298.

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.

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 spatial 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.

Eichenbaum, H. (1997) Declarative memory: insights from cognitive neurobiology. Annu. Rev. Psychol., 48: 547–572.

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.

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., 37B: 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. (2008) The interactions and dissociations of the dorsal hippocampus subregions: how the denate gyrus CA3, and CA1 process spatial information. Behav. Neurosci., 122: 16–26.

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.

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., 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.

Jeffery, K.J. and Hayman, R. (2004) Plasticity of the hippocampal place cell representation. Rev. Neurosci., 15: 309–331.

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., 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. and Rolls, E.T. (2001) Role of long term synaptic modification in short term memory. Hippocampus, 11: 240–250.

Kolb, B. and Whishaw, I.Q. (2003) Fundamentals of Human Neuropsychology. Worth, New York.

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: Hawkins, R.D. and Bower, G.H. (Eds.), Computational Models of Learning in Simple Neural Systems. Academic Press, San Diego, pp. 243–305.

Lynch, M.A. (2004) Long-term potentiation and memory. Physiol. Rev., 84: 87–136.

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.

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. Rev., 16: 193–220.
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.
McNaughton, B.L. and Morris, R.G.M. (1987) Hippocampal synaptic enhancement and information storage within a distributed memory system. Trends Neurosci., 10(10): 408–415.
McNaughton, B.L. and Nadel, L. (1990) Hebb-Marr networks and the neurobiological representation of action in space. In: Gluck, M.A. and Rumelhart, D.E. (Eds.), Neuroscience and Connectionist Theory. Erlbaum, Hillsdale.
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.
Morris, R.G.M. (1989) Does synaptic plasticity play a role in information storage in the vertebrate brain? In: Morris, R.G.M. (Ed.), Parallel Distributed Processing: Implications for Psychology and Neurobiology. Oxford University Press, Oxford, pp. 248–285.
Moscovitch, M., Rosenbaum, R.S., Gilboa, A., Addis, D.R., Westmacott, R., Grady, C., McAndrews, M.P., Levine, B., Black, S., Winocur, G. and Nadel, L. (2005) Functional neuroanatomy of remote episodic, semantic and spatial memory: a unified account based on multiple trace theory. J. Anat., 207: 35–66.
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: Paillard, J. (Ed.), Brain and Space. Oxford University Press, Oxford, pp. 296–333.
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. and Tonegawa, S. (2002) Requirement for hippocampal CA3 NMDA receptors in associative memory recall. Science, 297: 211–218.
O'Keefe, J. (1984) Spatial memory within and without the hippocampal system. In: Seifert, W. (Ed.), Neurobiology of the Hippocampus. Academic Press, London, pp. 375–403.
O'Keefe, J. and Nadel, L. (1978) The Hippocampus as a Cognitive Map. Clarendon Press, Oxford.
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: 4059–4167.
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.
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.
Rolls, E.T. (1987) Information representation, processing and storage in the brain: analysis at the single neuron level. In: Changeux, J.-P. and Konishi, M. (Eds.), The Neural and Molecular Bases of Learning. Wiley, Chichester, pp. 503–540.
Rolls, E.T. (1989a) Functions of neuronal networks in the hippocampus and cerebral cortex in memory. In: Cotterill, R.M.J. (Ed.), Models of Brain Function. Cambridge University Press, Cambridge, pp. 15–33.
Rolls, E.T. (1989b) Functions of neuronal networks in the hippocampus and neocortex in memory. In: Byrne, J.H. and Berry, W.O. (Eds.), Neural Models of Plasticity: Experimental and Theoretical Approaches. Academic Press, San Diego, pp. 240–265.
Rolls, E.T. (1989c) The representation and storage of information in neuronal networks in the primate cerebral cortex and hippocampus. In: Durbin, R. Miall, C. and Mitchison, G. (Eds.), The Computing Neuron. Addison-Wesley, Wokingham, England, pp. 125–159.
Rolls, E.T. (1990a) Functions of the primate hippocampus in spatial processing and memory. In: Olton, D.S. and Kesner, R.P. (Eds.), Neurobiology of Comparative Cognition. L. Erlbaum, Hillsdale, pp. 339–362.

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. (1995) A model of the operation of the hippocampus and entorhinal cortex in memory. Int. J. Neural Syst., 6: 51–70.

Rolls, E.T. (1996) Roles of long term potentiation and long term depression in neuronal network operations in the brain. In: Fazeli, M.S. and Collingridge, G.L. (Eds.), Cortical Plasticity. Bios, Oxford, pp. 223–250.

Rolls, E.T. (1999) Spatial view cells and the representation of place in the primate hippocampus. Hippocampus, 9: 467–480.

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., 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. 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., Stringer, S.M. and Elliot, T. (2006) Entorhinal cortex grid cells can map to hippocampal place cells by competitive learning. Netw.: Comput. Neural Syst., 17: 447–465.

Rolls, E.T., Stringer, S.M. and Trappenberg, T.P. (2002) A unified model of spatial and episodic memory. Proc. R. Soc. Lond. B, 269: 1087–1093.

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., 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. 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., Xiang, J.-Z. and Franco, L. (2005) Object, space and object-space representations in the primate hippocampus. J. Neurophysiol., 94: 833–844.

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.

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-Mathiesen, 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., 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.

Stringer, S.M., Rolls, E.T., Trappenberg, T.P. and Araujo, I.E.T. (2002a) Self-organizing continuous attractor networks and path integration: two-dimensional models of place cells. Netw.: Comput. Neural Syst., 13: 429–446.

Stringer, S.M., Trappenberg, T.P., Rolls, E.T. and Araujo, I.E.T. (2002b) Self-organizing continuous attractor networks and path integration: one-dimensional models of head direction cells. Netw.: Comput. Neural Syst., 13: 217–242.

Suzuki, W.A. and Amaral, D.G. (1994a) Perirhinal and parahippocampal cortices of the macaque monkey-cortical afferents. J. Comp. Neurol., 350(4): 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: 117–132.

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. (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.

Ungerleider, L.G. and Mishkin, M. (1982) Two cortical visual systems. In: Ingle, D.J. Goodale, M.A. and Mansfield, R.J.W. (Eds.), Analysis of Visual Behavior. MIT Press, Cambridge, pp. 549–586.

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.

Witter, M.P. (1993) Organization of the entorhinal-hippocampal system: a review of current anatomical data. Hippocampus, 3: 33–44.

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.

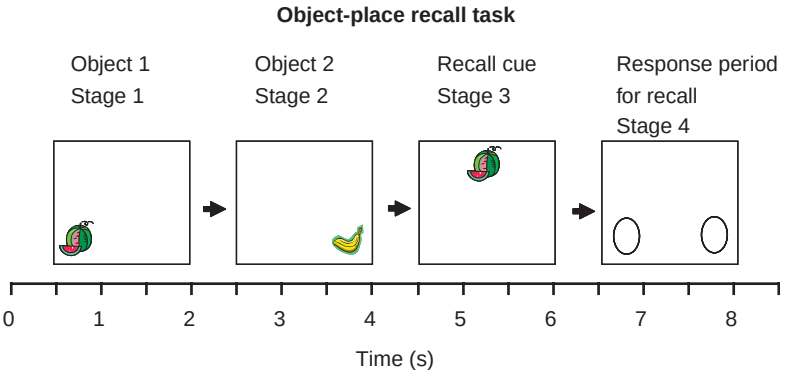


Plate 4.2.2. The object–place recall task. One trial is shown. After a 0.5 s tone to indicate the start of a trial, in stage 1 one of two objects (O1) is shown at one of the places (P1). (The object and the place are chosen randomly on each trial.) To ensure that the monkey sees the stimulus, the monkey can touch the screen at the place to obtain one drop of juice reward by licking. After a 0.5 s 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 s delay, in stage 3, the recall cue, one of the objects chosen at random, is shown at the top of the screen in the middle. (One drop of fruit juice was available as in stage 1.) After a 0.5 s 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.

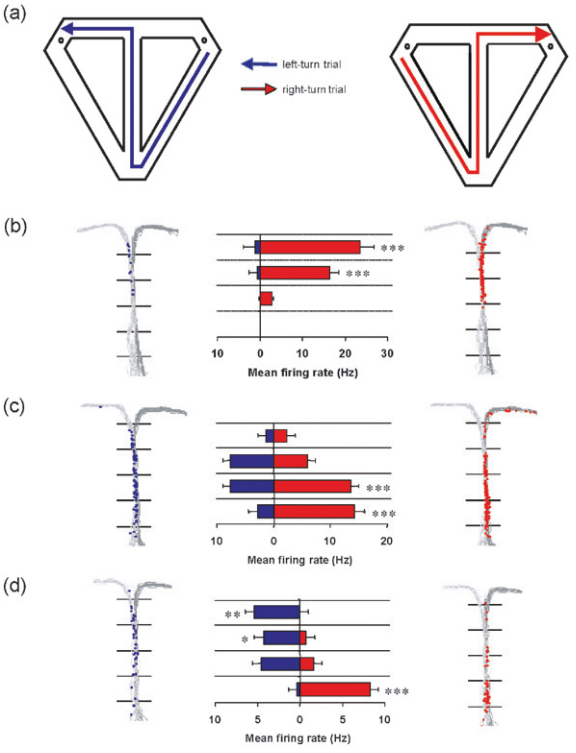


Plate 4.3.2. Hippocampal activity related to behavioral context. Context-dependent activity in CA1 on the continuous spatial alternation T-maze task (Wood et al., 2000). (a) Schematic representation of the modified T-maze. Left-turn trials are illustrated in blue and right-turn trials are illustrated in red. (b) Example of cell that fired almost exclusively on right-turn trials. Paths taken on left-turn trials (light gray) and right-turn trials (dark gray) are illustrated in the left and right plots with spikes fired on left-turn trials shown as blue dots and spikes fired in right-turn trials shown as red dots. The central stem is divided into four sectors and the mean firing rate of the cell in each sector for each type of trial is plotted. (c) Example of a cell that had a higher firing rate on right than left-turn trials. (d) Example of a cell that fired in different places on the central stem during left and right-turn trials. (Panels (b)–(d) Reprinted from Wood et al. (2000), with permission from Elsevier.)