Computational Models of Hippocampal Functions☆

Edmund T Rolls, Oxford Centre for Computational Neuroscience, Oxford, United Kingdom; and University of Warwick, Coventry, United Kingdom

Introduction

In this chapter a computational approach to the function of the hippocampus in memory is described and compared to other approaches. The theory is quantitative, and takes into account the internal- and systems-level connections of the hippocampus, the effects on memory of damage to different parts of the hippocampus, and the responses of hippocampal neurons recorded during memory tasks. The theory was developed by Rolls (1987, 1989b,c, 1990b, 1996b, 2008, 2016a), Treves and Rolls (1992, 1994), and with other colleagues (Kesner and Rolls, 2015; Rolls and Kesner, 2006; Rolls and Stringer, 2005; Rolls et al., 2002). The theory was preceded by work of Marr (1971) who developed a mathematical model, which although not applied to particular networks within the hippocampus and dealing with binary neurons and binary synapses, which utilized heavily the properties of the binomial distribution, was important in utilizing computational concepts and in considering how recall could occur in a network with recurrent collateral connections. Analyses of these autoassociation or attractor networks developed rapidly (Amit, 1989; Gardner-Medwin, 1976; Hopfield, 1982; Kohonen, 1977; Rolls and Treves, 1998; Treves and Rolls, 1991). Rolls (1987, 1989b, 1990b) produced a theory of the hippocampus in which the CA3 neurons operated as an autoassociation memory to store episodic memories including object and place memories, and the dentate granule cells operated as a preprocessing stage for this by performing pattern separation so that the mossy fibers could act to set up different representations for each memory to be stored in the CA3 cells. He suggested that the CA1 cells operate as a recoder for the information recalled from the CA3 cells to a partial memory cue, so that the recalled information would be represented more efficiently to enable recall, via the backprojection synapses, of activity in the neocortical areas similar to that which had been present during the original episode. At about the same time McNaughton and Morris (1987) suggested that the CA3 network might be an autoassociation network, and that the mossy fiber to CA3 connections might implement "detonator" synapses. The concepts that the diluted mossy fiber connectivity might implement selection of a new random set of CA3 cells for each new memory, and that a direct perforant path input to CA3 was needed to initiate retrieval were introduced by Treves and Rolls (1992). Since then, many investigators have contributed to our understanding of hippocampal computation, with some of these approaches described in section Comparison With Other Theories of Hippocampal Function and throughout the chapter.

☆*Change History*: March 2016. ET Rolls updated the text and references to this chapter, and deleted Figure 8.

This is an update of E.T. Rolls, 1.33 – Computational Models of Hippocampal Functions. In: Learning and Memory: A Comprehensive Reference, edited by John H. Byrne, Academic Press, Oxford, 2008, Pages 641–665.

A Theory of Hippocampal Function

Systems-Level Functions of the 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, which are described in more detail elsewhere (Rolls, 2008; Rolls and Kesner, 2006).

Evidence From the Effects of Damage to the Hippocampus

Damage to the hippocampus or to some of its connections such as the fornix in monkeys produces deficits in learning about the places of objects and about the places where responses should be made (Buckley and Gaffan, 2000). For example, macaques and humans with damage to the hippocampal system or fornix are impaired in object–place memory tasks in which not only the objects seen, but where they were seen must be remembered (Banta Lavenex and Lavenex, 2009; Burgess et al., 2002; Crane and Milner, 2005; Gaffan, 1994). 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). Furthermore, neurotoxic lesions that selectively damage the primate hippocampus impair spatial scene memory, tested by the ability to remember where in a scene to touch to obtain reward (Murray et al., 1998). Rats with hippocampal lesions are impaired in using environmental spatial cues to remember particular places, to perform object–place memory tasks, or to bridge delays (Cassaday and Rawlins, 1997; Jarrard, 1993; Kesner et al., 2004; Kesner and Rolls, 2015; Martin et al., 2000; O'Keefe and Nadel, 1978). These memory functions are important in event or episodic memory, in which the ability to remember what happened where on typically a single occasion is important. In humans, functional neuroimaging shows that the hippocampal system is activated by allocentric spatial including scene processing (Burgess, 2008; Chadwick et al., 2013; Hassabis et al., 2009; Maguire, 2014).

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.

The Necessity to Recall Information From the Hippocampus

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 (Kesner and Rolls, 2015; Rolls, 1989b,d, 2016a; Treves and Rolls, 1994). For example, remembering many particular journeys could help build a geographical 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 much accumulating information the semantic representation (McClelland et al., 1995; Moscovitch et al., 2005; Rolls, 1989b, 2016a; Treves and Rolls, 1994). The present theory 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 backprojections (see **Fig. 1**).

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 (Rolls, 2016a; Rolls and Xiang, 2006), and a brief summary is provided here because it provides a perspective relevant to understanding the function of the human hippocampus that is somewhat different from that provided by the properties of place cells in rodents, which have been reviewed elsewhere (Hartley et al., 2014; Jeffery, 2011; Jeffery et al., 2004; Jeffery and Hayman, 2004; McNaughton et al., 1983; Muller et al., 1991; Neunuebel and Knierim, 2012; O'Keefe, 1984; O'Keefe and Dostrovsky, 1971).

The primate hippocampus contains spatial view cells that respond when the monkey looks at a certain part of space, for example, at one quadrant of a video monitor while the monkey is performing an object–place memory task in which he must remember where on the monitor he has seen particular images (Rolls et al., 1989). Approximately 9% of the hippocampal neurons have such spatial view fields, and about 2.4% combine information about the position in space with information about the object that is in that position in space (Rolls et al., 1989). The representation of space is for the majority of hippocampal neurons in allocentric not egocentric coordinates (Feigenbaum and Rolls, 1991). These spatial view cells can be recorded while monkeys move themselves round the test environment by walking (or running) on all fours (Georges-François et al., 1999; Robertson et al., 1998; Rolls et al., 1997a, 1998). 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 toward the spatial view field, showing that these neurons can be updated for at least short periods by idiothetic (self-motion) cues including eye position and head direction signals (Robertson et al., 1998; Rolls et al., 1997b). There is some evidence consistent in humans (Ekstrom et al., 2003) consistent with these findings. In rodents, grid cells which have repeated peaks of "place" firing as the animal traverses a space are found in the medial entorhinal cortex (Giocomo et al., 2011; Hafting et al., 2005; Moser et al., 2015). In primates, there is now evidence that

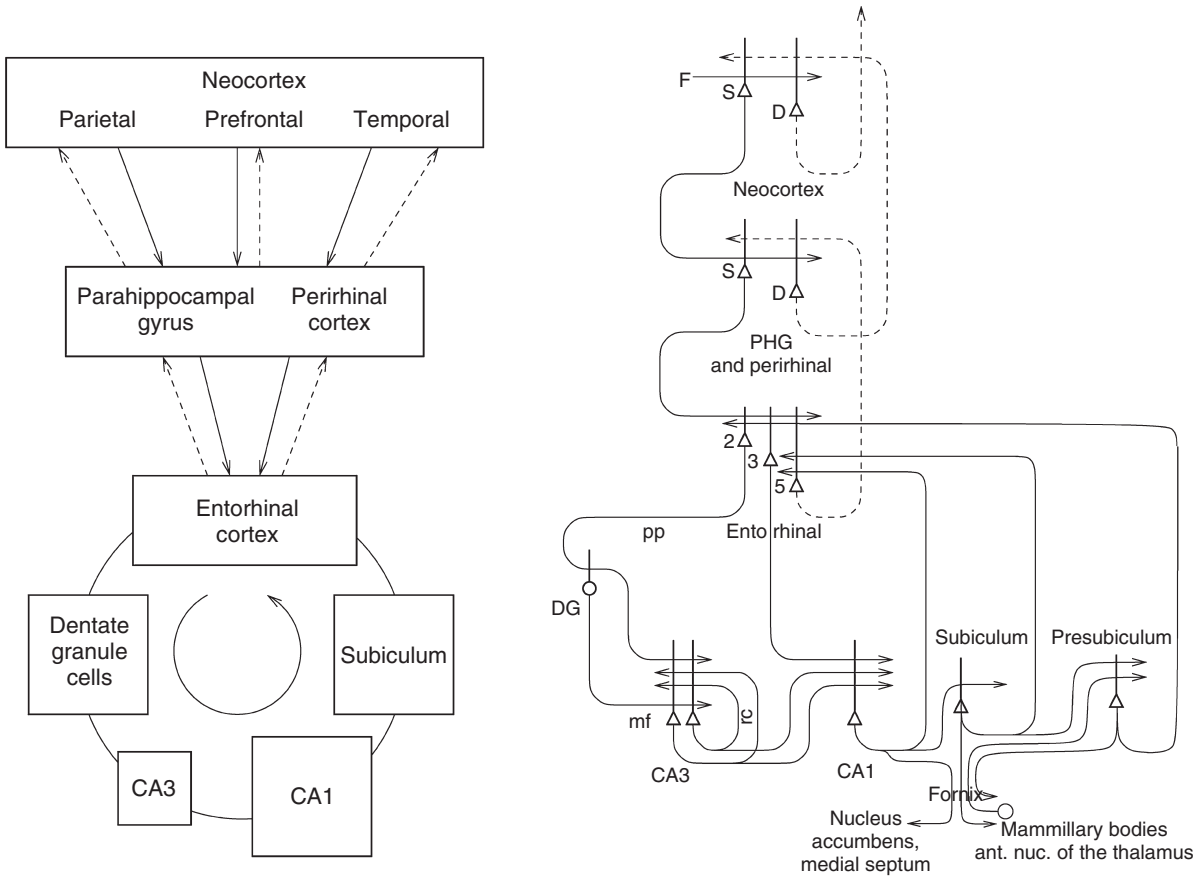

Figure 1 Forward connections (*solid lines*) from areas of cerebral association neocortex via the parahippocampal gyrus and perirhinal cortex, and entorhinal cortex, to the hippocampus; and backprojections (*dashed lines*) via the hippocampal CA1 pyramidal cells, subiculum, and parahippocampal gyrus 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. *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.

there is an analogous but spatial view–based grid cell–like representation in the entorhinal cortex, with neurons having gridlike firing as the monkey moves the eyes across a spatial scene (Killian et al., 2012). Further support for this type of representation of space being viewed "out there" rather than where one is located as for rat place cells is that cells in the human entorhinal cortex with spatial view gridlike properties have now been described (Jacobs et al., 2013).

A fundamental question about the function of the primate including human hippocampus 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 independent of location; other neurons (13%) responded to the spatial view independent 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).

Primate hippocampal neuronal activity has also been shown to be related to the recall of memories. In a one-trial object–place recall task, images of an object in one position on a screen, and of a second object in a different position on the screen, were shown successively. Then one of the objects was shown at the top of the screen, and the monkey had to recall the position in which it had been shown earlier in the trial, and to touch that location (Rolls and Xiang, 2006). In addition to neurons that responded to the

objects or places, a new type of neuronal response was found in which 5% of hippocampal neurons had place-related responses when a place was being recalled by an object cue.

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 (Kondo et al., 2005; Pitkanen et al., 2002; Rolls, 2015b). 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. Out 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; Redila et al., 2014; Tabuchi et al., 2003).

Systems-Level Anatomy

The primate hippocampus receives inputs via the entorhinal cortex (area 28) and the highly developed parahippocampal gyrus (areas TF and TH) as well as the perirhinal cortex from the ends of many processing streams of the cerebral association cortex, including the visual and auditory temporal lobe association cortical areas, the prefrontal cortex, and the parietal cortex (Amaral, 1987; Andersen et al., 2007; Kondo et al., 2005; Lavenex et al., 2004; Suzuki and Amaral, 1994b; Van Hoesen, 1982; van Strien et al., 2009) (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 (Carmichael and Price, 1995; Kondo et al., 2005; Pitkanen et al., 2002; Stefanacci et al., 1996; Suzuki and Amaral, 1994a).

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 (Delatour and Witter, 2002; van Haeften et al., 2003; Van Hoesen, 1982; Witter, 1993) (see **Fig. 1**), though there are other outputs (Kesner and Rolls, 2015; Rolls, 2016a).

The Operation of Hippocampal Circuitry as a Memory System

Hippocampal Circuitry (see Fig. 1; Amaral, 1993; Amaral and Witter, 1989; Andersen et al., 2007; Kondo et al., 2009; Lavenex et al., 2004; Naber et al., 2001; Storm-Mathiesen et al., 1990; Witter, 2007; Witter et al., 2000b)

Projections from the entorhinal cortex layer 2 reach the granule cells (of which there are 10^6 in the rat) in the dentate gyrus, via the perforant path (Witter, 1993). The granule cells project to CA3 cells via the mossy fibers (mf), which provide a *sparse* but possibly powerful connection to the 3×10^5 CA3 pyramidal cells in the rat. Each CA3 cell receives approximately 50 mossy fiber 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 4×10^3. The largest number of synapses (about 1.2×10^4 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 et al., 1990; Amaral and Witter, 1989, 1995; Ishizuka et al., 1990), 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. Of considerable interest, the CA3–CA3 recurrent collateral system is even more extensive in macaques than in rats (Kondo et al., 2009). In humans, there may be separate CA3 networks in the two hemispheres, for the left hippocampal system specializes in language-based memories, whereas the right hippocampal system specializes in spatial memory (Banks et al., 2012). The underlying computational reason for this is that language does not work by spatial locations, that is humans do not use and need a word-place episodic memory. 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**).

Dentate Granule Cells

The theory is that the dentate granule cell stage of hippocampal processing which precedes the CA3 stage acts in a number of ways to produce during learning by pattern separation the sparse yet efficient (i.e., nonredundant) representation in CA3 neurons that is required for the autoassociation implemented by CA3 to perform well (Kesner and Rolls, 2015; Rolls, 1989b, 1996b, 2008, 2013a,b, 2016a,b; Treves and Rolls, 1992). By pattern separation I mean that the correlations between different memory patterns represented by a population of neurons become reduced.

The first way is that the perforant path—dentate granule cell system with its Hebb-like modifiability is suggested to act as a competitive learning network to remove redundancy from the inputs producing a more orthogonal, sparse, and categorized set of outputs (Rolls, 1987, 1989b, 2016a; Rolls and Treves, 1998) (competitive networks are described elsewhere (Hertz et al.,

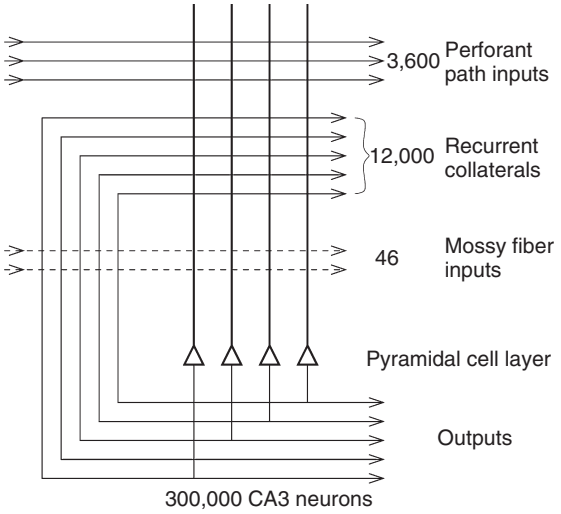

Figure 2 The numbers of connections from three different sources onto each CA3 cell from three different sources in the rat. After Rolls, E.T., Treves, A., 1998. Neural Networks and Brain Function. Oxford University Press, Oxford; Treves, A., Rolls, E.T., 1992. Computational constraints suggest the need for two distinct input systems to the hippocampal CA3 network. Hippocampus 2, 189–199.

1991; Rolls, 2016a,b; Rolls and Treves, 1998), together with Matlab code to simulate them (Rolls, 2016a)). The nonlinearity in the NMDA receptors may help the operation of such a competitive net, for it ensures that only the most active neurons left after the competitive feedback inhibition have synapses that become modified and thus learn to respond to that input (Rolls, 2016a). Because of the feedback inhibition, the competitive process may result in a relatively constant number of strongly active dentate neurons relatively independent of the number of active perforant path inputs to the dentate cells. The operation of the dentate granule cell system as a competitive network may also be facilitated by a Hebb rule of the form:

$$\delta w_{ij} = k \cdot r_i \left(r'_j - w_{ij} \right) \tag{1}$$

where k is a constant, r_i is the activation of the dendrite (the postsynaptic term), r'_j is the presynaptic firing rate, w_{ij} is the synaptic weight, and r'_j and w_{ij} are in appropriate units (Rolls, 2016a). Incorporation of a rule such as this which implies heterosynaptic long-term depression as well as long-term potentiation (see Levy et al., 1990; Levy and Desmond, 1985) makes the sum of the synaptic weights on each neuron remain roughly constant during learning (cf. Oja, 1982; Rolls, 2016a; see Rolls and Treves, 1998).

This functionality could be used to help build hippocampal place cells in rats from the grid cells present in the medial entorhinal cortex (Giocomo et al., 2011; Moser et al., 2015). Each grid cell responds to a set of places in a spatial environment, with the places to which a cell responds set out in a regular grid. Different grid cells have different phases (positional offsets) and grid spacings (or frequencies). We (Rolls et al., 2006) have simulated the dentate granule cells as a system that receives as inputs the activity of a population of entorhinal cortex grid cells as the animal traverses a spatial environment, and have shown that the competitive net builds dentate-like place cells from such entorhinal grid cell inputs (see **Fig. 3**). This occurs because the firing states of entorhinal cortex cells that are active at the same time when the animal is in one place become associated together by the learning in the competitive net, yet each dentate cell represents primarily one place because the dentate representation is kept sparse, thus helping to implement symmetry breaking (Rolls et al., 2006). The same competitive learning could, it is suggested (Rolls, 2013a), be involved in the conversion of primate entorhinal spatial view grid cells (Killian et al., 2012) into primate spatial view cells.

With respect to the medial entorhinal cortex grid cells, there are a number of theories of their computational bases (Giocomo et al., 2011). One attractive theory is that different temporal delays in the neural circuitry of the medial entorhinal cortex are related to different temporal adaptation time courses, and result in the different sizes of spatial grids that are found in the medial entorhinal cortex (Kropff and Treves, 2008). This is essentially a timing hypothesis; and it is an interesting new idea that the timing mechanism that may be inherent in the medial entorhinal cortex might also be a source of the temporal delay information needed to form sequence memories in CA1, using the direct projection from medial entorhinal cortex to CA1 to introduce timing information in the form of neurons that fire at different temporal times in a sequence (Eichenbaum, 2014; Howard and Eichenbaum, 2015; Howard et al., 2014; Macdonald et al., 2011), which is what is needed for the formation of odor, object, and spatial sequence memories in CA1, and would also be useful in the formation of spatial grid cells in the entorhinal cortex.

The second way is also a result of the competitive learning hypothesized to be implemented by the dentate granule cells (Rolls, 1987, 1989b, 2016a). It is proposed that this allows overlapping (or very similar) inputs to the hippocampus to be separated, in the following way (see also Rolls, 1996b). Consider three patterns B, W, and BW, where BW is a linear combination of B and W. (To make the example very concrete, we could consider binary patterns where B = 10, W = 01, and BW = 11.) Then the memory system is required to associate B with reward, W with reward, but BW with punishment. Without the hippocampus, rats might have more

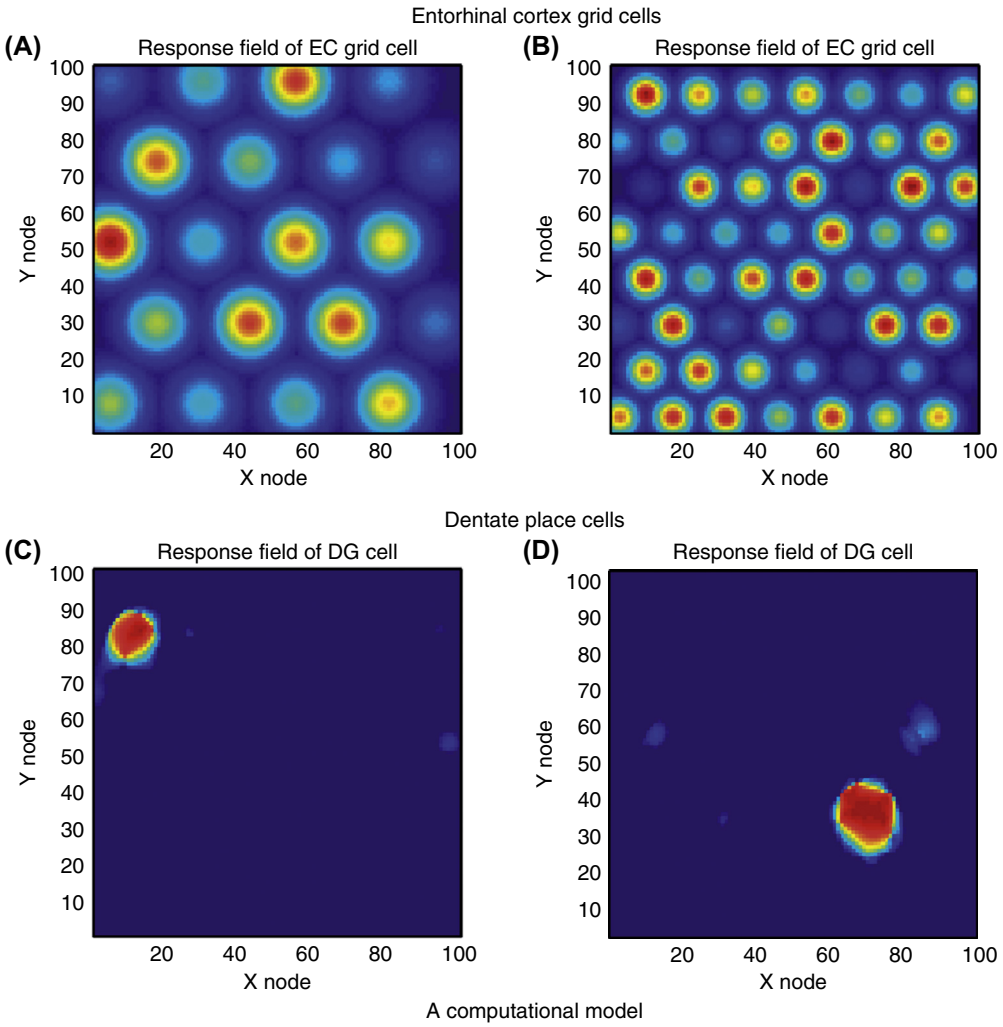

Figure 3 Simulation of competitive learning in the dentate gyrus to produce place cells from the entorhinal cortex (EC) grid cell inputs. (A, B) Firing rate profiles of two EC grid cells with frequencies of four and seven cycles. In the simulation, cells with frequencies of 4–7 cycles were used, and with 25 phases or spatial offsets. (A phase is defined as an offset in the X and Y directions, and five offset values were used in each direction.) The standard deviation of the peak heights was set to 0.6. (C, D) Firing rate profiles of two dentate gyrus (DG) cells after competitive network training with the Hebb's rule. After Rolls, E.T., Stringer, S.M., Elliot, T., 2006. Entorhinal cortex grid cells can map to hippocampal place cells by competitive learning. Netw. Comput. Neural Syst. 17, 447–465.

difficulty in solving such problems, particularly when they are spatial, for the dentate–CA3 system in rodents is characterized by being implicated in spatial memory. However, it is a property of competitive neuronal networks that they can separate such overlapping patterns, as has been shown elsewhere (Rolls, 2008, 2016a,b; Rolls and Treves, 1998); normalization of synaptic weight vectors is required for this property. It is thus an important part of hippocampal neuronal network architecture that there is a competitive network that precedes the CA3 autoassociation system. Without the dentate gyrus, if a conventional autoassociation network were presented with the mixture BW having learned B and W separately, then the autoassociation network would produce a mixed output state, and would therefore be incapable of storing separate memories for B, W, and BW. It was therefore suggested that competition in the dentate gyrus is one of the powerful computational features of the hippocampus, and that could enable it to help solve spatial pattern separation tasks (Rolls and Kesner, 2006).

This computational hypothesis and its predictions have been tested. Rats with dentate gyrus lesions are impaired at a metric spatial pattern separation task (Gilbert et al., 2001; Goodrich-Hunsaker et al., 2005; Kesner and Rolls, 2015) (see **Fig. 4**). The recoding of grid cells in the entorhinal cortex (Hafting et al., 2005) into small place field cells in the dentate granule cells that has been modeled (Rolls et al., 2006) can also be considered to be a case where overlapping inputs must be recoded so that different spatial components can be treated differently. I note that Sutherland and Rudy's configural learning hypothesis was similar, but was not tested with spatial pattern separation. Instead, when tested with, for example, tone and light combinations, it was not consistently found that the hippocampus was important (O'Reilly and Rudy, 2001; Sutherland and Rudy, 1991). I suggest that application of the

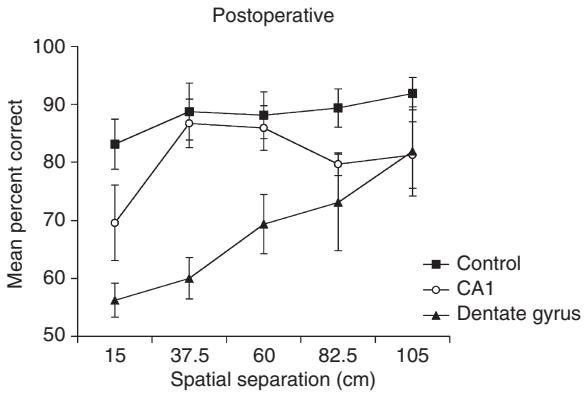

Figure 4 Pattern separation impairment produced by dentate gyrus lesions. Mean percent correct performance as a function of spatial separation of control group, CA1 lesion group, and dentate gyrus lesion group on postoperative trials. A graded impairment was found as a function of the distance between the places only following dentate gyrus lesions. After Gilbert, P.E., Kesner, R.P., Lee, I., 2001. Dissociating hippocampal subregions: double dissociation between dentate gyrus and CA1. Hippocampus 11, 626–636.

configural concept, but applied to spatial pattern separation, may capture part of what the dentate gyrus acting as a competitive network could perform, particularly when a large number of such overlapping spatial memories must be stored and retrieved.

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

A fourth way is that as suggested and explained in section Mossy Fiber Inputs to the CA3 Cells and Pattern Separation, the dentate granule cell—mossy fiber input to the CA3 cells may be powerful and its use particularly during learning would be efficient in forcing a new pattern of firing onto the CA3 cells during learning.

Fifth, expansion recoding can decorrelate input patterns, and this can be performed by a stage of competitive learning with a large number of neurons (Rolls, 2016a). A similar mechanism appears to be implemented by the dentate granule cells, which are numerous (1×10^6 in the rat, compared to 300,000 CA3 cells), have associatively modifiable synapses (required for a competitive network), and strong inhibition provided by the inhibitory interneurons. This may not represent expansion of numbers relative to the number of entorhinal cortex cells, but the principle of a large number of dentate granule cells, with competitive learning and strong inhibition through inhibitory interneurons, would produce a decorrelation of signals like that achieved by expansion recoding (Rolls, 2008, 2013a).

Sixth, adult neurogenesis in the dentate gyrus may perform the computational role of facilitating pattern separation for new patterns, by providing new dentate granule cells with new sets of random connections to CA3 neurons (Deng et al., 2010; Rolls, 2010). Consistent with the dentate spatial pattern separation hypothesis (Kesner and Rolls, 2015; Rolls, 1989b, 1996b, 2016a; Treves and Rolls, 1992, 1994), in mice with impaired dentate neurogenesis, spatial learning in a delayed nonmatching-to-place task in the radial arm maze was impaired for arms that were presented with little separation, but no deficit was observed when the arms were presented farther apart (Clelland et al., 2009). Consistently, impaired neurogenesis in the dentate also produced a deficit for small spatial separations in an associative object-in-place task (Clelland et al., 2009).

In the ways just described, the dentate granule cells could be particularly important in helping to build and prepare spatial representations for the CA3 network. The actual representation of space in the primate hippocampus includes a representation of spatial view, whereas in the rat hippocampus it is of the place where the rat is. The representation in the rat may be related to the fact that with a much less-developed visual system than the primate, the rat's representation of space may be defined more by the olfactory and tactile as well as distant visual cues present, and may thus tend to reflect the place where the rat is. However, the spatial representations in the rat and primate could arise from essentially the same computational process as follows (de Araujo et al., 2001; Rolls, 1999). 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 degrees, 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 degrees, a combination of visual features formed over such a wide visual angle would effectively define a position in space that is a place (de Araujo et al., 2001).

Although spatial view cells are present in the parahippocampal areas (Georges-François et al., 1999; Robertson et al., 1998; Rolls et al., 1997a, 1998), and neurons with placelike fields (often as grid cells (Moser et al., 2015)) are found in the medial entorhinal cortex, there are backprojections from the hippocampus to the entorhinal cortex and thus to parahippocampal areas, and these backprojections could enable the hippocampus to influence the spatial representations found in the entorhinal cortex and parahippocampal gyrus. On the other hand, as described earlier, the gridlike place cells in the medial entorhinal cortex could, if transformed

by the competitive net functionality of the dentate cells, result in the place cell activity (without a repeating grid) that is found in dentate and rat hippocampal neurons.

In primates, spatial view cells represent a scene view allocentrically, as described in section Systems-Level Neurophysiology of the Primate Hippocampus. How could such spatial view representations be formed, in which the relative spatial position of features in a scene is encoded? We have proposed that this involves competitive learning analogous to that used to form place cells in rats, but in primates operating on the representations of objects that reach the hippocampus from the inferior temporal visual cortex (Rolls, 2012b, 2016a). Rolls et al. (2008) have demonstrated in a unifying computational approach that competitive network processes operating in areas such as the parahippocampal cortex, the entorhinal cortex, and/or the dentate granule cells could form unique views of scenes by forming a sparse representation of these object- or feature-tuned inferior temporal cortex ventral visual stream representations that have some spatial asymmetry.

CA3 as an Autoassociation Memory

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 Andersen et al., 2007; Jackson, 2013; Lynch, 2004; Morris, 1989, 2003; Morris et al., 2003; Nakazawa et al., 2004; Nakazawa et al., 2003; Wang and Morris, 2010). On the basis of the evidence summarized earlier, Rolls (1987, 1989b, 1991, 2013b, 2016a) and others (Levy, 1989; McNaughton, 1991; McNaughton and Morris, 1987) 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 elsewhere (Hertz et al., 1991; Rolls, 2016a; Rolls and Treves, 1998). The architecture of an autoassociation network is shown in **Fig. 4**, and the learning rule is as shown in Eq. (1) except that the subtractive term could be the presynaptic firing rate (Rolls, 2016a; Rolls and Treves, 1998).

The hypothesis is that because the CA3 operates effectively as a single network, it can allow arbitrary associations between inputs originating from very different parts of the cerebral cortex to be formed. These might involve associations between information originating in the temporal visual cortex about the presence of an object, and information originating in the parietal cortex about where it is. I note that although there is some spatial gradient in the CA3 recurrent connections, so that the connectivity is not fully uniform (Ishizuka et al., 1990; Witter, 2007), 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. It is very interesting indeed that in primates (macaques), the associational projections from CA3 to CA3 travel extensively along the longitudinal axis, and overall the radial, transverse, and longitudinal gradients of CA3 fiber distribution, clear in the rat, are much more subtle in the nonhuman primate brain (Kondo et al., 2009). The implication is that in primates, the CA3 network operates even more as a single network than in rodents.

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 CA3 as an Autoassociation Memory and in more detail elsewhere (Rolls, 2008; Rolls and Kesner, 2006).

Storage Capacity

We have performed quantitative analyses of the storage and retrieval processes in the CA3 network (Treves and Rolls, 1991, 1992). We have extended previous formal models of autoassociative memory (see Amit, 1989; Hopfield, 1982) by analyzing a network with graded response units, so as to represent more realistically the continuously variable rates at which neurons fire, and with incomplete connectivity (Treves, 1990; Treves and Rolls, 1991). We have found that in general the maximum number $p_{\max}$ of firing patterns that can be (individually) retrieved is proportional to the number C^{RC} of (associatively) modifiable recurrent collateral synapses per cell, by a factor that increases roughly with the inverse of the sparseness a of the neuronal representation.[1] The sparseness of response (or selectivity) of a single cell to a set of stimuli (which in the brain has approximately the same value as the sparseness of the response of the population of neurons to any one stimulus, which can in turn be thought of as the proportion of neurons that is active to any one stimulus if the neurons had binary responses, see Franco et al., 2007) is defined as

$$a = \left(\sum_{i=1,n} r_i/n \right)^2 \Big/ \sum_{i=1,n} (r_i^2/n) \tag{2}$$

[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.

where r_i is the firing rate to the i'th stimulus in the set of n stimuli. The sparseness ranges from $1/n$, when the cell responds to only one stimulus, to a maximal value of 1.0, attained when the cell responds with the same rate to all stimuli. Approximately,

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

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). [The sparseness a in this equation is strictly the population sparseness (Franco et al., 2007; Treves and Rolls, 1991). The population sparseness a^p would be measured by measuring the distribution of firing rates of all neurons to a single stimulus at a single time. The single cell sparseness or selectivity a^s would be measured by the distribution of firing rates to a set of stimuli, which would take a long time. These concepts are elucidated by Franco et al. (2007). The sparseness estimates obtained by measuring early gene changes, which are effectively population sparsenesses, would thus 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.] 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. Third, for most associative networks to store information efficiently, heterosynaptic long-term depression (as well as LTP) is required (Collingridge et al., 2010; Fazeli and Collingridge, 1996; Rolls, 2016a; Rolls and Treves, 1990, 1998; Treves and Rolls, 1991). Simulations that are fully consistent with the analytical theory are provided by Rolls (1995, 2012a), Simmen et al. (1996), and Rolls et al. (1997b).

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

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

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

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

$$i_p > a \ln(1/a) \tag{5}$$

As the information content of each stored pattern i_p depends on the storage process, we see how the retrieval capacity analysis, coupled with the notion that the system is organized so as to be an efficient memory device in a quantitative sense, leads to a constraint on the storage process.

A number of points that arise are treated elsewhere (Kesner and Rolls, 2015; Rolls, 2016a). 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 (Kesner and Rolls, 2015; Rolls, 1996a, 2016a). The limited capacity of the CA3 system does also provide one of the arguments that some transfer of information from the hippocampus to neocortical memory stores may be useful (see Treves and Rolls, 1994). Given its limited capacity, the hippocampus might be a useful store for only a limited period, which might be in the order of days, weeks, or months. This period may well depend on the acquisition rate of new episodic memories. If the animal were in a constant and limited environment, then as new information is not being added to the hippocampus, the representations in the hippocampus would remain stable and persistent. These hypotheses have clear experimental implications, both for recordings from single neurons and for the gradient of retrograde amnesia, both of which might be expected to depend on whether the environment is stable or frequently changing. They show that the conditions under which a gradient of retrograde amnesia might be demonstrable would be when large numbers of new memories are being acquired, not when only a few memories (few in the case of the hippocampus being less than a few hundred) are being learned.

Recall

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

A prediction of the theory is thus that the CA3 recurrent collateral associative connections enable arbitrary associations to be formed between whatever is represented in the hippocampus, in that, for example, any place could be associated with any object, and in that the object could be recalled with a spatial recall cue, or the place with an object recall cue.

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

In a hippocampus subregion test of this, rats in a study phase are shown one object in one location, and then a second object in another location. (There are 50 possible objects, and 48 locations.) In the test phase, the rat is shown one object in the start box, and then after a 10 s delay it must go to the correct location (choosing between two marked locations). CA3 lesions made after training in the task produced chance performance on this one-trial object–place recall task (Kesner et al., 2008; Kesner and Rolls, 2015). A control-fixed visual conditional to place task with the same delay was not impaired, showing that it is a recall after one-trial (or rapid) learning that is impaired. In the context of arbitrary associations between whatever is represented in CA3, the theory also predicts that cued place–object recall tasks and cued place–odor recall tasks should be impaired by CA3 lesions.

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

In a neurophysiological investigation of one-trial object–place learning followed by recall of the spatial position in which to respond when shown the object, Rolls and Xiang (2006) showed that some primate hippocampal (including CA3) neurons respond to an object cue with the spatial position in which the object had been shown earlier in the trial. Thus, some hippocampal neurons appear to reflect spatial recall given an object recall cue.

Completion

Another fundamental property is that the recall can be complete even from a small fragment. Thus, it is a prediction that when an incomplete retrieval cue is given, CA3 may be especially important in the retrieval process. Tests of this prediction of a role for CA3 in pattern completion have been performed as follows (Kesner and Rolls, 2015). Rats were tested on a cheese board with a black curtain with four extramaze cues surrounding the apparatus. (The cheese board is like a dry land water maze with 177 holes on a 119 cm diameter board.) Rats were trained to move a sample phase object covering a food well that could appear in one of five possible spatial locations. During the test phase of the task, following a 30 s delay, the animal needs to find the same food well to receive reinforcement with the object now removed. After reaching stable performance in terms of accuracy to find the correct location, rats received lesions in CA3. During postsurgery testing, four extramaze cues were always available during the sample phase. However, during the test phase zero, one, two, or three cues were removed in different combinations. The results indicate that controls performed well on the task regardless of the availability of one, two, three, or all cues, suggesting intact spatial pattern completion. Following the CA3 lesion, however, there was an impairment in accuracy compared to the controls especially when only one or two cues were available, suggesting impairment in spatial pattern completion in CA3-lesioned rats (Gold and Kesner, 2005) (see **Fig. 6**). A useful aspect of this task is that the test for the ability to remember a spatial location learned in

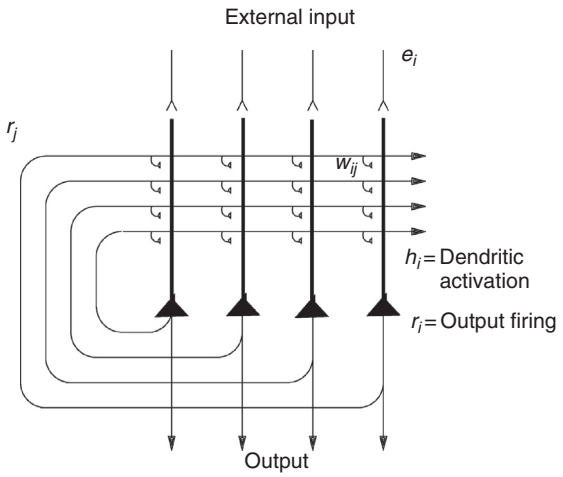

Figure 5 The architecture of a continuous attractor neural network. The architecture is the same as that of a discrete attractor neural network.

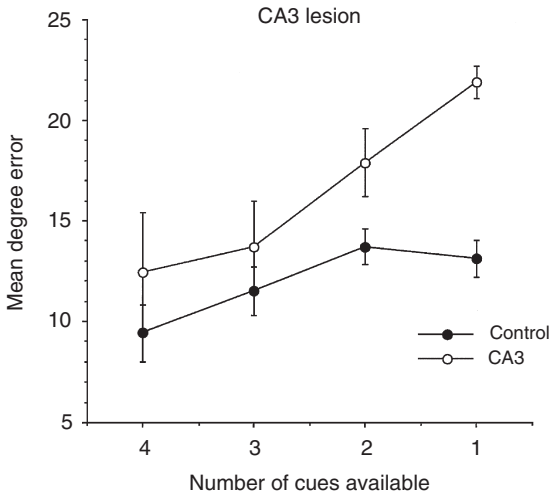

Figure 6 Pattern completion impairment produced by CA3 lesions. The mean (and sem) degree of error in finding the correct place in the cheese board task when rats were tested with one, two, three, or four of the cues available. A graded impairment in the CA3 lesion group as a function of the number of cues available was found. The task was learned in the study phase with the four cues present. The performance of the control group is also shown. After Gold, A.E., 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.

one presentation can be tested with varying number of available cues, and many times in which the locations vary, to allow for accurate measurement of pattern completion ability when the information stored on the single presentation must be recalled.

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

Continuous, Spatial, Patterns, and CA3 Representations

The fact that spatial patterns, which imply continuous representations of space, are represented in the hippocampus has led to the application of continuous attractor models to help understand hippocampal function. This has been necessary, because space is inherently continuous, because the firing of place and spatial view cells is approximately Gaussian as a function of the distance away from the preferred spatial location, because these cells have spatially overlapping fields, and because the theory is that these cells in CA3 are connected by Hebb-modifiable synapses. This specification would inherently lead the system to operate as a continuous attractor network. Continuous attractor network models have been studied by Amari (1977), Zhang (1996), Taylor (1999), Samsonovich and McNaughton (1997), Battaglia and Treves (1998a), Stringer et al. (2002a,b), Stringer et al. (2004), Stringer and Rolls (2002), and Rolls and Stringer (2005) (see Rolls, 2016a) and are described 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 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 is 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 ensure that there is only one activity packet. Continuous attractor networks can be thought of as very similar to autoassociation or discrete attractor networks (Rolls, 2016a), 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. 7. 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. 7). A consequent difference is that the CANN can maintain its firing at any location in the trained continuous space, whereas a discrete attractor or autoassociation network moves its population of active neurons toward one of the previously learned attractor states, and thus implements the recall of a particular previously learned pattern from an incomplete or noisy (distorted) version of one of the previously learned patterns.

The energy landscape of a discrete attractor network has separate energy minima, each one of which corresponds to a learned pattern, whereas the energy landscape of a continuous attractor network is flat, so that the activity packet remains stable with

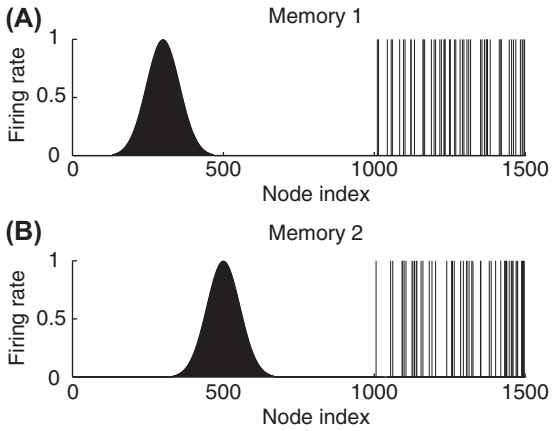

Figure 7 The types of firing patterns stored in continuous attractor networks are illustrated for the patterns present on neurons 1–1000 for Memory 1 (when the firing is that produced when the spatial state represented is that for location 300), and for Memory 2 (when the firing is that produced when the spatial state represented is that for location 500). The continuous nature of the spatial representation results from the fact that each neuron has a Gaussian firing rate that peaks at its optimal location. This particular mixed network also contains discrete representations that consist of discrete subsets of active binary firing rate neurons in the range 1001–1500. The firing of these latter neurons can be thought of as representing the discrete events that occur at the location. Continuous attractor networks by definition contain only continuous representations, but this particular network can store mixed continuous and discrete representations, and is illustrated to show the difference of the firing patterns normally stored in separate continuous attractor and discrete attractor networks. For this particular mixed network, during learning, Memory 1 is stored in the synaptic weights, then Memory 2, etc., and each memory contains part that is continuously distributed to represent physical space, and part that represents a discrete event or object.

continuous firing wherever it is started in the state space (Rolls, 2016a; Rolls and Deco, 2010). (The state space refers to the set of possible spatial states of the animal in its environment, e.g., the set of possible places in a room.)

So far we have said that the neurons in the continuous attractor network are connected to each other by synaptic weights w_{ij} that are a simple function, for example Gaussian, of the distance between the states of the agent in the physical world (e.g., head directions, spatial views etc.) represented by the neurons. In many simulations, the weights are set by formula to have weights with these appropriate Gaussian values. However, Stringer et al. (2002b) showed how the appropriate weights could be set up by learning. They started with the fact that since the neurons have broad tuning that may be Gaussian in shape, nearby neurons in the state space will have overlapping spatial fields, and will thus be coactive to a degree that depends on the distance between them. They postulated that therefore the synaptic weights could be set up by associative learning based on the coactivity of the neurons produced by external stimuli as the animal moved in the state space. For example, head direction cells are forced to fire during learning by visual cues in the environment that produce Gaussian firing as a function of head direction from an optimal head direction for each cell. The learning rule is simply that the weights w_{ij} from head direction cell j with firing rate r_j^{HD} to head direction cell i with firing rate r_i^{HD} are updated according to an associative (Hebb) rule

$$\delta w_{ij} = k r_i^{\mathrm{HD}} r_j^{\mathrm{HD}} \tag{6}$$

where δw_{ij} is the change of synaptic weight and k is the learning rate constant. During the learning phase, the firing rate r_i^{HD} of each head direction cell i might be the following Gaussian function of the displacement of the head from the optimal firing direction of the cell

$$r_i^{\mathrm{HD}} = e^{-s_{\mathrm{HD}}^2/2\sigma_{\mathrm{HD}}^2}, \tag{7}$$

where s_{HD} is the difference between the actual head direction x (in degrees) of the agent and the optimal head direction x_i for head direction cell i, and σ_{HD} is the standard deviation. Stringer et al. (2002b) showed that after training at all head directions, the synaptic connections develop strengths that are an almost Gaussian function of the distance between the cells in head direction space.

Combined Continuous and Discrete Memory Representations in the Same (e.g., CA3) Network and Episodic Memory

Space is continuous, and object representations are discrete. If these representations are to be combined in, for example, an object–place memory, then we need to understand the operation of networks that combine these representations. It has now been shown that attractor networks can store both continuous patterns and discrete patterns (as illustrated in Fig. 7), and can thus be used to store, for example, the location in (continuous, physical) space (e.g., the place "out there" in a room represented by spatial view cells) where an object (a discrete item) is present (Rolls et al., 2002).

The Capacity of a Continuous Attractor Network and Multiple Charts

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

Idiothetic Update by Path Integration

We have considered how spatial representations could be stored in continuous attractor networks, and how the activity can be maintained at any location in the state space in a form of short-term memory when the external (e.g., visual) input is removed. However, many networks with spatial representations in the brain can be updated by internal, self-motion (i.e., idiothetic), cues even when there is no external (e.g., visual) input. Path integration can be implemented in recurrent attractor networks as described elsewhere for the entorhinal cortex (Giocomo et al., 2011; Zilli, 2012) and hippocampal CA3 (Kesner and Rolls, 2015; Rolls, 2016a).

The Dynamics of the Recurrent Network

The analysis described earlier about the capacity of a recurrent network such as the CA3 considered steady-state conditions of the firing rates of the neurons. The question arises of how quickly the recurrent network would settle into its final state. With reference to the CA3 network, how long does it take before a pattern of activity, originally evoked in CA3 by afferent inputs, becomes influenced by the activation of recurrent collaterals? In a more general context, recurrent collaterals between the pyramidal cells are an important feature of the connectivity of the cerebral neocortex. How long would it take these collaterals to contribute fully to the activity of cortical cells? If these settling processes took in the order of hundreds of ms, they would be much too slow to contribute usefully to cortical activity, whether in the hippocampus or the neocortex (Panzeri et al., 2001; Rolls, 2016a).

It has been shown that if the neurons are treated not as McCulloch–Pitts neurons which are simply "updated" at each iteration, or cycle of time steps (and assume the active state if the threshold is exceeded), but instead are analyzed and modeled as "integrate-and-fire" neurons in real continuous time, then the network can effectively "relax" into its recall state very rapidly, in one or two time constants of the synapses (Battaglia and Treves, 1998b; Rolls, 2016a; Rolls and Treves, 1998; Treves, 1993). This corresponds to perhaps 20 ms in the brain. One factor in this rapid dynamics of autoassociative networks with brainlike "integrate-and-fire" membrane and synaptic properties is that with some spontaneous activity, some of the neurons in the network are close to threshold already before the recall cue is applied, and hence some of the neurons are very quickly pushed by the recall cue into firing, so that information starts to be exchanged very rapidly (within 1–2 ms of brain time) through the modified synapses by the neurons in the network. The progressive exchange of information starting early on within what would otherwise be thought of as an iteration period (of perhaps 20 ms, corresponding to a neuronal firing rate of 50 spikes/s), is the mechanism accounting for rapid recall in an autoassociative neuronal network made biologically realistic in this way. Further analysis of the fast dynamics of these networks, if they are implemented in a biologically plausible way with "integrate-and-fire" neurons, is provided elsewhere (Panzeri et al., 2001; Rolls, 2016a; Rolls and Treves, 1998; Treves, 1993).

Mossy Fiber Inputs to the CA3 Cells and Pattern Separation

For the CA3 to operate with high capacity as an autoassociation or attractor memory, the sets of CA3 neurons that represent each event to be stored and later recalled need to be as uncorrelated from each other as possible. Correlations between patterns reduce the memory capacity of an autoassociation network (Kohonen, 1977, 1984; Kohonen et al., 1981; Marr, 1971; Rolls, 2008; Rolls and Treves, 1998; Sompolinsky, 1987), and because storage capacity is at a premium in an episodic memory system, there are several mechanisms that reduce the correlations between the firing of the population vectors of CA3 neuron firing each one of which represents a different event to be stored in memory. In the theoretical physics approach to the capacity of attractor networks, it is indeed assumed that the different vectors of firing rates to be stored are well separated from each other, by drawing each vector of firing at random, and by assuming very large (infinite) numbers of neurons in each pattern (Hopfield, 1982; Rolls and Treves, 1998).

I have proposed that there are several mechanisms that help achieve this pattern separation, namely the mossy fiber pattern separation effect produced by the small number of connections received by a CA3 neuron from mossy fibers which dominate the CA3 cell firing; the expansion recoding, and the sparse representation provided by the dentate granule cells that form the mossy fiber synapses; and the sparseness of the CA3 cell representation (Kesner and Rolls, 2015; Rolls, 2016a,b). Neurogenesis of dentate granule cells is a fifth potential contributor to achieving pattern separation of CA3 cell firing (Aimone et al., 2010; Deng et al., 2010; Johnston et al., 2016).

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 (Cerasti and Treves, 2010; Rolls, 1987, 1989b; Treves and Rolls, 1992). (The strong effects likely to be mediated by the mossy fibers were also emphasized by McNaughton and Morris (1987) and McNaughton and Nadel (1990).) We hypothesize that the mossy fiber input appears to be particularly appropriate in several ways. First of all, the fact that mossy fiber synapses are large and located very close to the soma makes them relatively powerful in activating the postsynaptic cell. [This should not be taken to imply that a CA3 cell can be fired by a single mossy fiber excitatory postsynaptic potential

(EPSP).] Second, the firing activity of dentate granule cells appears to be very sparse (Jung and McNaughton, 1993; Leutgeb et al., 2007) 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.[2] Third, nonassociative plasticity of mossy fibers (see Brown et al., 1989; Brown et al., 1990) might have a useful effect in enhancing the signal-to-noise ratio, in that a consistently firing mossy fiber would produce nonlinearly amplified currents in the postsynaptic cell, which would not happen with an occasionally firing fiber (Treves and Rolls, 1992). This plasticity, and also learning in the dentate, would also have the effect that similar fragments of each episode (e.g., the same environmental location) recurring on subsequent occasions would be more likely to activate the same population of CA3 cells, which would have potential advantages in terms of economy of use of the CA3 cells in different memories, and in making some link between different episodic memories with a common feature, such as the same location in space. Fourth, with only a few, and powerful, active mossy fiber inputs to each CA3 cell, setting a given sparseness of the representation provided by CA3 cells would be simplified, for the EPSPs produced by the mossy fibers would be Poisson distributed with large membrane potential differences for each active mossy fiber. Setting the average firing rate of the dentate granule cells would effectively set the sparseness of the CA3 representation, without great precision being required in the threshold setting of the CA3 cells (Rolls et al., 1997b). Part of what is achieved by the mossy fiber input may be setting the sparseness of the CA3 cells correctly, which, as shown earlier, is very important in an autoassociative memory store. Fifth, the nonassociative and sparse connectivity properties of the mossy fiber connections to CA3 cells may be appropriate for an episodic memory system which can learn very fast, in one trial. The hypothesis is that the sparse connectivity would help arbitrary relatively uncorrelated sets of CA3 neurons to be activated for even somewhat similar input patterns without the need for any learning of how best to separate the patterns, which in a self-organizing competitive network would take several repetitions (at least) of the set of patterns. The mossy fiber solution may thus be adaptive in a system that must learn in one trial, and for which the CA3 recurrent collateral learning requires uncorrelated sets of CA3 cells to be allocated for each (one-trial) episodic memory. The hypothesis is that the mossy fiber sparse connectivity solution performs the appropriate function without the mossy fiber system having to learn by repeated presentations of how best to separate a set of training patterns. The perforant path input would, the quantitative analysis shows, not produce a pattern of firing in CA3 that contains sufficient information for learning (Treves and Rolls, 1992).

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

As acetylcholine turns down the efficacy of the recurrent collateral synapses between CA3 neurons (Giocomo and Hasselmo, 2007; 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 (Giocomo and Hasselmo, 2007; Hasselmo et al., 1995), when presumably it may be particularly relevant to allocate some of the limited memory capacity to new memories.

Perforant Path Inputs to CA3 Cells

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

An autoassociative memory network needs afferent inputs also in the other mode of operation, i.e., when it retrieves a previously stored pattern of activity. We have shown (Treves and Rolls, 1992) that the requirements on the organization of the afferents are in this case very different, implying the necessity of a second, separate input system, which we have identified with the perforant path to CA3. In brief, the argument is based on the notion that the cue available to initiate retrieval might be rather small, i.e., the distribution of activity on the afferent axons might carry a small correlation, $q << 1$, with the activity distribution present during learning. In order not to lose this small correlation altogether, but rather transform it into an input current in the CA3 cells that carries a sizable signal—which can then initiate the retrieval of the full pattern by the recurrent collaterals—one needs a large number

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

of associatively modifiable synapses. This is expressed by the formulas that give the specific signal S produced by sets of associatively modifiable synapses, or by nonassociatively modifiable synapses: if C^{AFF} is the number of afferents per cell,

$$S_{ASS} \sim \frac{\sqrt{C^{AFF}}}{\sqrt{p}} q \quad S_{NONASS} \sim \frac{1}{\sqrt{C^{AFF}}} q. \tag{8}$$

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

CA1 Cells

Associative Retrieval at the CA3 to CA1 (Schaffer Collateral) Synapses

The CA3 cells connect to the CA1 cells by the Schaffer collateral synapses. The following arguments outline the advantage of this connection being associatively modifiable, and apply independent of the relative extent to which the CA3 or the direct entorhinal cortex inputs to CA1 drive the CA1 cells during the learning phase.

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

Recoding in CA1 to Facilitate Retrieval to the Neocortex

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

Another argument on the operation of the CA1 cells is also considered to be related to the CA3 autoassociation effect. In this, several arbitrary patterns of firing occur together on the CA3 neurons, and become associated together to form an episodic or "whole scene" memory. It is essential for this CA3 operation that several different sparse representations are present conjunctively to form the association. Moreover, when completion operates in the CA3 autoassociation system, all the neurons firing in the original conjunction can be brought into activity by only a part of the original set of conjunctive events. For these reasons, a memory in the CA3 cells consists of several different simultaneously active ensembles of activity. To be explicit, the parts A, B, C, D, and E of a particular episode would each be represented, roughly speaking, by its own population of CA3 cells, and these five populations would be linked together by autoassociation. It is suggested that the CA1 cells, which receive these groups of simultaneously active ensembles, can detect the conjunctions of firing of the different ensembles that represent the episodic memory, and allocate by competitive learning neurons to represent at least larger parts of each episodic memory (Kesner and Rolls, 2015; Rolls, 1987, 1989b, 2016a). In relation to the simple example mentioned earlier, some CA1 neurons might code for ABC, and others for BDE, rather than having to maintain independent representations in CA1 of A, B, C, D, and E. This implies a more efficient representation, in the sense that when eventually after many further stages, neocortical neuronal activity is recalled (as discussed later), each neocortical cell need not be accessed by all the axons carrying each component A,B,C,D, and E, but instead by fewer axons carrying larger fragments, such as ABC, and BDE. This process is performed by competitive networks, which self-organize to find categories in the input space, where each category is represented by a set of simultaneously active inputs (Rolls, 2000, 2016a,b; Rolls and Treves, 1998).

CA1 Inputs From CA3 Versus Direct Entorhinal Inputs

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

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

CA1 and Sequence Memory for Objects and Odors: Time Cells

Selective lesions of the CA1 implicate this region in the memory for sequences of objects and odors (Kesner and Rolls, 2015). In humans, the hippocampus becomes activated when the temporal order of events is being processed (Lehn et al., 2009). How might this be implemented? As we have just seen, the entorhinal cortex does send direct inputs to CA1 that bypass CA3 (see **Fig. 1**), and with respect to the medial entorhinal cortex grid cells, there are a number of theories of their computational bases (Giocomo et al., 2011). One attractive theory is that different temporal delays in the neural circuitry of the medial entorhinal cortex related to a temporal adaptation process might result in the different sizes of spatial grids that are found in the medial entorhinal cortex (Kropff and Treves, 2008). This is essentially a timing hypothesis; and it is an interesting new idea that the timing mechanism that may be inherent in the medial entorhinal cortex might also be a source of the temporal delay/timing information needed to form sequence memories in CA1, using the direct projection from medial entorhinal cortex to CA1 to introduce timing information in the form of neurons that fire at different temporal times in a sequence, which is what is needed for the formation of odor, object, and spatial sequence memories in CA1 (see section Dentate Granule Cells). Consistent with this hypothesis, timing cells are found in the medial entorhinal cortex (Kraus et al., 2013a,b). Object (including odor) information could reach CA1 from the lateral entorhinal cortex. The time and object information might be combined in CA1 by the associative process inherent in a competitive network (Rolls, 2008) to form time and object combination neurons (with time originating from medial and object from lateral entorhinal cortex). Replay of the time information on another trial could allow the same object and time combination neurons to be retrieved by generalization, and thus effectively for the position of the object in the sequence to be retrieved.

This hypothesis is consistent with neurophysiological evidence of Macdonald et al. (2011) showing that neurons in the rat hippocampus have firing rates that reflect which temporal part of the task is current. In particular, a sequence of different neurons is activated at successive times during a time delay period. The tasks used included an object–odor-paired associate nonspatial task with a 10 s delay period between the visual stimulus and the odor. The new evidence also shows that a large proportion of hippocampal neurons fire in relation to individual events in a sequence being remembered (e.g., a visual object or odor), and some to combinations of the event and the time in the delay period (Macdonald et al., 2011).

These interesting neurophysiological findings indicate that rate encoding is being used to encode time, that is, the firing rates of different neurons are high at different times within a trial, delay period, etc (Kraus et al., 2013a,b; Macdonald et al., 2011; Rolls and Deco, 2010).

Backprojections to the Neocortex—A Hypothesis

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. This may be assisted as described earlier by the direct perforant path input to CA1. This might allow details of the input key for the recall process, as well as the possibly less information-rich memory of the whole episode recalled from the CA3 network, to contribute to the firing of CA1 neurons. 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, 1989b). 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; Rolls, 1990a,b, 2016a), 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 (Kesner and Rolls, 2015; Rolls, 1989b, 2016a). During the setting up of a new episodic memory, there would be strong feedforward activity progressing toward the hippocampus. During the episode, the CA3 synapses would be modified, and via the CA1 neurons and the subiculum, a pattern of activity would be produced on the 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 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.

Backprojections to the Neocortex—Quantitative Analysis

A plausible requirement for a successful hippocampo-directed recall operation, is that the signal generated from the hippocampally retrieved pattern of activity, and carried backward toward neocortex, remain undegraded when compared to the noise due, at each stage, to the interference effects caused by the concurrent storage of other patterns of activity on the same 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 earlier.) 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 $C^{\mathrm{HBP}} = C^{\mathrm{RC}} a_{\mathrm{nc}}/a_{\mathrm{CA3}}$, 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 aforementioned 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}, i.e., 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, 2016a), including the advantage of diluted connectivity in this pattern association system (Rolls, 2015a, 2016a).

Comparison With Other Theories of Hippocampal Function

The overall theory described here is close in different respects to those of a number of other investigators (Brown and Zador, 1990; Eichenbaum et al., 1992; Marr, 1971; McClelland et al., 1995; McNaughton and Nadel, 1990; Moscovitch et al., 2005; Preston and Eichenbaum, 2013; Squire, 1992; Wang and Morris, 2010; Winocur and Moscovitch, 2011) and of course priority is not claimed on all the propositions put forward here.

Some theories postulate that the hippocampus performs spatial computation. The theory of O'Keefe and Nadel (1978) that the hippocampus implements a cognitive map placed great emphasis on spatial function. It supposed that the hippocampus at least

holds information about allocentric space in a form that enables rats to find their way in an environment even when novel trajectories are necessary, that is it permits an animal to "go from one place to another independent of particular inputs (cues) or outputs (responses), and to link together conceptually parts of the environment which have never been experienced at the same time." O'Keefe (1990) extended this analysis and produced a computational theory of the hippocampus as a cognitive map, in which the hippocampus performs geometrical spatial computations. Key aspects of the theory are that the hippocampus stores the centroid and slope of the distribution of landmarks in an environment, and stores the relationships between the centroid and the individual landmarks. The hippocampus then receives information as inputs about where the rat currently is, and where the rat's target location is, and computes geometrically the body turns and movements necessary to reach the target location. In this sense, the hippocampus is taken to be a spatial computer, which produces an output which is very different from its inputs. This is in contrast to the present theory, in which the hippocampus is a memory device, which is able to recall what was stored in it, using as input a partial cue. A prototypical example in Rolls' theory is the learning of object–place association memory and the recall of the whole memory from a part, which can be used as a model of event or episodic memory. The theory of O'Keefe postulates that the hippocampus actually performs a spatial computation. A later theory (Barry and Burgess, 2014; Burgess, 2008; Burgess et al., 1994, 2000) also makes the same postulate, but now, for example, the firing of place cells is determined by the distance and approximate bearing to landmarks, and the navigation is performed by increasing the strength of connections from place cells to "goal cells," and then performing gradient-ascent style search for the goal using the network.

McNaughton et al. (1991) have also proposed that the hippocampus is involved in spatial computation. They propose a "compass" solution to the problem of spatial navigation along novel trajectories in known environments, postulating that distances and bearings (i.e., vector quantities) from landmarks are stored, and that computation of a new trajectory involves vector subtraction by the hippocampus. They postulate that a linear associative mapping is performed, using as inputs a "cross-feature" (combination) representation of (head) angular velocity and (its time integral) head direction, to produce as output the future value of the integral (head direction) after some specified time interval. The system can be reset by learned associations between local views of the environment and head direction, so that when later a local view is seen, it can lead to an output from the network which is a (corrected) head direction. They suggest that some of the key signals in the computational system can be identified with the firing of hippocampal cells (e.g., local view cells) and subicular cells (head direction cells). It should be noted that this theory requires a (linear) associative mapping with an output (head direction) different in form from the inputs (head angular velocity over a time period, or local view). This is pattern association (with the conditioned stimulus local view, and the unconditioned stimulus head direction), not autoassociation, and it has been postulated that this pattern association can be performed by the hippocampus (cf. McNaughton and Morris, 1987). This theory is again in contrast to the present theory, in which the hippocampus operates as a memory to store events that occur at the same time, and can recall the whole memory *from any part* of what was stored. A pattern associator uses a conditioned stimulus to map an input to a pattern of firing in an output set of neurons which is like that produced in the output neurons by the unconditioned stimulus. A description of pattern associations and autoassociators in a neurobiological context is provided by Rolls (2016a). The present theory is fully consistent with the presence of "spatial view" cells and whole body motion cells in the primate hippocampus (Rolls, 1999; Rolls and O'Mara, 1993; Rolls and Xiang, 2006) (or place or local view cells in the rat hippocampus, and head direction cells in the presubiculum), for it is often important to store and later recall where one has been (views of the environment, body turns made, etc.), and indeed such (episodic) memories are required for navigation by "dead reckoning" in small environments.

The present theory thus 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, 2009). 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 (Kesner and Rolls, 2015; Newman et al., 2012; Rolls, 1987, 1989a,b,d, 2008, 2016a; 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 earlier (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.

Lisman et al. (2005) have considered how the memory of sequences could be implemented in the hippocampus. This theory of sequential recall within the hippocampus is inextricably linked to the internal timing within the hippocampus imposed he believes by the theta and gamma oscillations, and this makes it difficult to recall each item in the sequence as it is needed. It is not specified how one would read out the sequence information, given that the items are only 12 ms apart. The Jensen and Lisman (1996) model

requires short-time constant NMDA channels, and is therefore unlikely to be implemented in the hippocampus. Hasselmo and Eichenbaum (2005) have taken up some of these sequence ideas and incorporated them into their model, which has its origins in the Rolls and Treves model (Rolls, 1989b; Treves and Rolls, 1992, 1994), but proposes, for example, that sequences are stored in entorhinal cortex layer 3. The proposal that acetylcholine could be important during encoding by facilitating CA3–CA3 LTP, and should be lower during retrieval (Hasselmo et al., 1995; Newman et al., 2012), is an important concept.

Another type of sequence memory uses synaptic adaptation to effectively encode the order of the items in a sequence (Deco and Rolls, 2005). This could be implemented in recurrent networks such as the CA3 or the prefrontal cortex.

In this chapter, we have thus seen that quantitative approaches to the functions of the hippocampus in memory are being developed by a number of investigators, and that these theories are consistent with the quantitative circuitry of the hippocampus as well as with neuronal recordings and the effects of lesions. Moreover, tests of the theory described here are in general consistent with the predictions of the theory, as described elsewhere (Kesner and Rolls, 2015).

Acknowledgments

Different parts of the research described here were supported by Program 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

Aimone, J.B., Deng, W., Gage, F.H., 2010. Adult neurogenesis: integrating theories and separating functions. Trends Cogn. Sci. 14, 325–337.
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 DC, pp. 211–294.
Amaral, D.G., 1993. Emerging principles of intrinsic hippocampal organisation. Curr. Opin. Neurobiol. 3, 225–229.
Amaral, D.G., Ishizuka, N., Claiborne, B., 1990. Neurons, numbers and the hippocampal network. Prog. Brain Res. 83, 1–11.
Amaral, D.G., Witter, M.P., 1989. The three-dimensional organization of the hippocampal formation: a review of anatomical data. Neuroscience 31, 571–591.
Amaral, D.G., 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.
Andersen, P., Morris, R.G.M., Amaral, D.G., Bliss, T.V.P., O'Keefe, J., 2007. The Hippocampus Book. Oxford University Press, London.
Banks, S.J., Sziklas, V., Sodums, D.J., Jones-Gotman, M., 2012. fMRI of verbal and nonverbal memory processes in healthy and epileptogenic medial temporal lobes. Epilepsy Behav. 25, 42–49.
Banta Lavenex, P., Lavenex, P., 2009. Spatial memory and the monkey hippocampus: not all space is created equal. Hippocampus 19, 8–19.
Barry, C., Burgess, N., 2014. Neural mechanisms of self-location. Curr. Biol. 24, R330–R339.
Battaglia, F.P., Treves, A., 1998a. Attractor neural networks storing multiple space representations: a model for hippocampal place fields. Phys. Rev. E 58, 7738–7753.
Battaglia, F.P., Treves, A., 1998b. Stable and rapid recurrent processing in realistic auto-associative memories. Neural Comput. 10, 431–450.
Brown, T.H., Ganong, A.H., Kairiss, E.W., Keenan, C.L., Kelso, S.R. (Eds.), 1989. Long-term Potentiation in Two Synaptic Systems of the Hippocampal Brain Slice. Academic Press, San Diego.
Brown, T.H., Kairiss, E.W., Keenan, C.L., 1990. Hebbian synapses: biophysical mechanisms and algorithms. Annu. Rev. Neurosci. 13, 475–511.
Brown, T.H., Zador, A., 1990. The hippocampus. In: Shepherd, G. (Ed.), The Synaptic Organisation of the Brain. Oxford University Press, New York, pp. 346–388.
Brun, V.H., Otnass, M.K., Molden, S., Steffenach, H.A., Witter, M.P., Moser, M.B., Moser, E.I., 2002. Place cells and place recognition maintained by direct entorhinal-hippocampal circuitry. Science 296, 2243–2246.
Buckley, M.J., 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., 2008. Spatial cognition and the brain. Ann. N.Y. Acad. Sci. 1124, 77–97.
Burgess, N., Jackson, A., Hartley, T., O'Keefe, J., 2000. Predictions derived from modelling the hippocampal role in navigation. Biol. Cybern. 83, 301–312.
Burgess, N., Maguire, E.A., O'Keefe, J., 2002. The human hippocampus and spatial and episodic memory. Neuron 35, 625–641.
Burgess, N., Recce, M., O'Keefe, J., 1994. A model of hippocampal function. Neural Netw. 7, 1065–1081.
Carmichael, S.T., 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., Rawlins, J.N., 1997. The hippocampus, objects, and their contexts. Behav. Neurosci. 111, 1228–1244.
Cerasti, E., Treves, A., 2010. How informative are spatial CA3 representations established by the dentate gyrus? PLoS Comput. Biol. 6, e1000759.
Chadwick, M.J., Mullally, S.L., Maguire, E.A., 2013. The hippocampus extrapolates beyond the view in scenes: an fMRI study of boundary extension. Cortex 49, 2067–2079.
Clelland, C.D., Choi, M., Romberg, C., Clemenson Jr., G.D., Fragniere, A., Tyers, P., Jessberger, S., Saksida, L.M., Barker, R.A., Gage, F.H., Bussey, T.J., 2009. A functional role for adult hippocampal neurogenesis in spatial pattern separation. Science 325, 210–213.
Collingridge, G.L., Peineau, S., Howland, J.G., Wang, Y.T., 2010. Long-term depression in the CNS. Nat. Rev. 11, 459–473.
Crane, J., Milner, B., 2005. What went where? Impaired object-location learning in patients with right hippocampal lesions. Hippocampus 15, 216–231.
Day, M., Langston, R., 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., 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.
Deco, G., Rolls, E.T., 2005. Sequential memory: a putative neural and synaptic dynamical mechanism. J. Cogn. Neurosci. 17, 294–307.
Delatour, B., 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.
Deng, W., Aimone, J.B., Gage, F.H., 2010. New neurons and new memories: how does adult hippocampal neurogenesis affect learning and memory? Nat. Rev. 11, 339–350.
Eichenbaum, H., 2014. Time cells in the hippocampus: a new dimension for mapping memories. Nat. Rev. 15, 732–744.
Eichenbaum, H., Otto, T., Cohen, N.J., 1992. The hippocampus - what does it do? Behav. Neural Biol. 57, 2–36.

Ekstrom, A.D., Kahana, M.J., Caplan, J.B., Fields, T.A., Isham, E.A., Newman, E.L., Fried, I., 2003. Cellular networks underlying human spatial navigation. Nature 425, 184–188.
Fazeli, M.S., Collingridge, G.L. (Eds.), 1996. Cortical Plasticity: LTP and LTD. Bios, Oxford.
Feigenbaum, J.D., Rolls, E.T., 1991. Allocentric and egocentric spatial information processing in the hippocampal formation of the behaving primate. Psychobiology 19, 21–40.
Florian, C., 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., 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.
Gardner-Medwin, A.R., 1976. The recall of events through the learning of associations between their parts. Proc. R. Soc. Lond. B 194, 375–402.
Georges-François, P., Rolls, E.T., 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., Lee, I., 2001. Dissociating hippocampal subregions: double dissociation between dentate gyrus and CA1. Hippocampus 11, 626–636.
Giocomo, L.M., Hasselmo, M.E., 2007. Neuromodulation by glutamate and acetylcholine can change circuit dynamics by regulating the relative influence of afferent input and excitatory feedback. Mol. Neurobiol. 36, 184–200.
Giocomo, L.M., Moser, M.B., Moser, E.I., 2011. Computational models of grid cells. Neuron 71, 589–603.
Gold, A.E., 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., Kesner, R.P., 2005. Effects of hippocampus sub-regional lesions for metric and topological spatial information processing. Soc. Neurosci. Abstr. 647.641.
Hafting, T., Fyhn, M., Molden, S., Moser, M.B., Moser, E.I., 2005. Microstructure of a spatial map in the entorhinal cortex. Nature 436, 801–806.
Hargreaves, E.L., Rao, G., Lee, I., Knierim, J.J., 2005. Major dissociation between medial and lateral entorhinal input to dorsal hippocampus. Science 308, 1792–1794.
Hartley, T., Lever, C., Burgess, N., O'Keefe, J., 2014. Space in the brain: how the hippocampal formation supports spatial cognition. Philos. Trans. R. Soc. Lond. 369, 20120510.
Hassabis, D., Chu, C., Rees, G., Weiskopf, N., Molyneux, P.D., Maguire, E.A., 2009. Decoding neuronal ensembles in the human hippocampus. Curr. Biol. 19, 546–554.
Hasselmo, M.E., Eichenbaum, H.B., 2005. Hippocampal mechanisms for the context-dependent retrieval of episodes. Neural Netw. 18, 1172–1190.
Hasselmo, M.E., Schnell, E., 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., Palmer, R.G., 1991. An Introduction to the Theory of Neural Computation. Addison-Wesley, Wokingham.
Hölscher, C., Jacob, W., Mallot, H.A., 2003. Reward modulates neuronal activity in the hippocampus of the rat. Behav. Brain Res. 142, 181–191.
Hopfield, J.J., 1982. Neural networks and physical systems with emergent collective computational abilities. Proc. Natl. Acad. Sci. U.S.A. 79, 2554–2558.
Howard, M.W., Eichenbaum, H., 2015. Time and space in the hippocampus. Brain Res. 1621, 345–354.
Howard, M.W., MacDonald, C.J., Tiganj, Z., Shankar, K.H., Du, Q., Hasselmo, M.E., Eichenbaum, H., 2014. A unified mathematical framework for coding time, space, and sequences in the hippocampal region. J. Neurosci. 34, 4692–4707.
Ishizuka, N., Weber, J., Amaral, D.G., 1990. Organization of intrahippocampal projections originating from CA3 pyramidal cells in the rat. J. Comp. Neurol. 295, 580–623.
Jackson, M.B., 2013. Recall of spatial patterns stored in a hippocampal slice by long-term potentiation. J. Neurophysiol. 110, 2511–2519.
Jacobs, J., Weidemann, C.T., Miller, J.F., Solway, A., Burke, J.F., Wei, X.X., Suthana, N., Sperling, M.R., Sharan, A.D., Fried, I., Kahana, M.J., 2013. Direct recordings of grid-like neuronal activity in human spatial navigation. Nat. Neurosci. 16, 1188–1190.
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., 2011. Place cells, grid cells, attractors, and remapping. Neural Plast. 2011, 182602.
Jeffery, K.J., Anderson, M.I., Hayman, R., Chakraborty, S., 2004. A proposed architecture for the neural representation of spatial context. Neurosci. Biobehav. Rev. 28, 201–218.
Jeffery, K.J., Hayman, R., 2004. Plasticity of the hippocampal place cell representation. Rev. Neurosci. 15, 309–331.
Jensen, O., Lisman, J.E., 1996. Theta/gamma networks with slow NMDA channels learn sequences and encode episodic memory: role of NMDA channels in recall. Learn. Mem. (Cold Spring Harbor, N.Y) 3, 264–278.
Johnston, S.T., Shtrahman, M., Parylak, S., Goncalves, J.T., Gage, F.H., 2016. Paradox of pattern separation and adult neurogenesis: a dual role for new neurons balancing memory resolution and robustness. Neurobiol. Learn. Mem. 129, 60–68.
Jung, M.W., McNaughton, B.L., 1993. Spatial selectivity of unit activity in the hippocampal granular layer. Hippocampus 3, 165–182.
Kesner, R.P., Hunsaker, M.R., Warthen, M.W., 2008. The CA3 subregion of the hippocampus is critical for episodic memory processing by means of relational encoding in rats. Behav. Neurosci. 122, 1217–1225.
Kesner, R.P., Lee, I., Gilbert, P., 2004. A behavioral assessment of hippocampal function based on a subregional analysis. Rev. Neurosci. 15, 333–351.
Kesner, R.P., Rolls, E.T., 2015. A computational theory of hippocampal function, and tests of the theory: new developments. Neurosci. Biobehav. Rev. 48, 92–147.
Killian, N.J., Jutras, M.J., Buffalo, E.A., 2012. A map of visual space in the primate entorhinal cortex. Nature 491, 761–764.
Kohonen, T., 1977. Associative Memory: A System Theoretical Approach. Springer, New York.
Kohonen, T., 1984. Self-Organization and Associative Memory. Springer-Verlag, Berlin.
Kohonen, T., Oja, E., Lehtio, P., 1981. Storage and processing of information in distributed memory systems. In: Hinton, G.E., Anderson, J.A. (Eds.), Parallel Models of Associative Memory. Lawrence Erlbaum, Hillsdale, pp. 129–167.
Kolb, B., Whishaw, I.Q., 2009. Fundamentals of Human Neuropsychology. Worth, New York.
Kondo, H., Lavenex, P., Amaral, D.G., 2009. Intrinsic connections of the macaque monkey hippocampal formation: II. CA3 connections. J. Comp. Neurol. 515, 349–377.
Kondo, H., Saleem, K.S., Price, J.L., 2005. Differential connections of the perirhinal and parahippocampal cortex with the orbital and medial prefrontal networks in macaque monkeys. J. Comp. Neurol. 493, 479–509.
Kraus, B.J., Brandon, M.P., Robinson, R.J., Connerney, M.A., Hasselmo, M.E., Eichenbaum, H., 2013a. Grid cells are time cells. Soc. Neurosci. Abstr. 769.719.
Kraus, B.J., Robinson 2nd, R.J., White, J.A., Eichenbaum, H., Hasselmo, M.E., 2013b. Hippocampal "time cells": time versus path integration. Neuron 78, 1090–1101.
Kropff, E., Treves, A., 2008. The emergence of grid cells: intelligent design or just adaptation? Hippocampus 18, 1256–1269.
Lassalle, J.M., Bataille, T., 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., Amaral, D.G., 2000. Hippocampal-neocortical interaction: a hierarchy of associativity. Hippocampus 10, 420–430.
Lavenex, P., Suzuki, W.A., Amaral, D.G., 2004. Perirhinal and parahippocampal cortices of the macaque monkey: intrinsic projections and interconnections. J. Comp. Neurol. 472, 371–394.
Lehn, H., Steffenach, H.A., van Strien, N.M., Veltman, D.J., Witter, M.P., Haberg, A.K., 2009. A specific role of the human hippocampus in recall of temporal sequences. J. Neurosci. 29, 3475–3484.
Leutgeb, J.K., Leutgeb, S., Moser, M.B., Moser, E.I., 2007. Pattern separation in the dentate gyrus and CA3 of the hippocampus. Science 315, 961–966.
Levy, W.B., 1989. A computational approach to hippocampal function. In: Hawkins, R.D., Bower, G.H. (Eds.), Computational Models of Learning in Simple Neural Systems. Academic Press, San Diego, pp. 243–305.
Levy, W.B., Colbert, C.M., Desmond, N.L., 1990. Elemental adaptive processes of neurons and synapses: a statistical/computational perspective. In: Gluck, M.A., Rumelhart, D.E. (Eds.), Neuroscience and Connectionist Theory. Erlbaum, Hillsdale, NJ, pp. 187–235.
Levy, W.B., Desmond, N.L., 1985. The rules of elemental synaptic plasticity. In: Levy, W.B., Anderson, J.A., Lehmkuhle, S. (Eds.), Synaptic Modification, Neuron Selectivity, and Nervous System Organization. Erlbaum, Hillsdale, NJ, pp. 105–121.

Lisman, J.E., Talamini, L.M., Raffone, A., 2005. Recall of memory sequences by interaction of the dentate and CA3: a revised model of the phase precession. Neural Netw. 18, 1191–1201.
Lynch, M.A., 2004. Long-term potentiation and memory. Physiol. Rev. 84, 87–136.
Macdonald, C.J., Lepage, K.Q., Eden, U.T., Eichenbaum, H., 2011. Hippocampal "time cells" bridge the gap in memory for discontiguous events. Neuron 71, 737–749.
Maguire, E.A., 2014. Memory consolidation in humans: new evidence and opportunities. Exp. Physiol. 99, 471–486.
Malkova, L., 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.
Marr, D., 1971. Simple memory: a theory for archicortex. Phil. Trans. Roy. Soc. Lond. B 262, 23–81.
Martin, S.J., Grimwood, P.D., 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., 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., 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., Chen, L.L., Markus, E.J., 1991. "Dead reckoning", landmark learning, and the sense of direction: a neurophysiological and computational hypothesis. J. Cogn. Neurosci. 3, 190–202.
McNaughton, B.L., Morris, R.G.M., 1987. Hippocampal synaptic enhancement and information storage within a distributed memory system. Trends Neurosci. 10, 408–415.
McNaughton, B.L., Nadel, L., 1990. Hebb-Marr networks and the neurobiological representation of action in space. In: Gluck, M.A., Rumelhart, D.E. (Eds.), Neuroscience and Connectionist Theory. Erlbaum, Hillsdale, NJ, pp. 1–64.
Morris, R.G., 2003. Long-term potentiation and memory. Philos. Trans. R. Soc. Lond. 358, 643–647.
Morris, R.G., Moser, E.I., Riedel, G., Martin, S.J., Sandin, J., Day, M., 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. 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., 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.
Moser, M.B., Rowland, D.C., Moser, E.I., 2015. Place cells, grid cells, and memory. Cold Spring Harb. Perspect. Biol. 7, a021808.
Muller, R.U., Kubie, J.L., Bostock, E.M., Taube, J.S., 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., 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., 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., McHugh, T.J., Wilson, M.A., Tonegawa, S., 2004. NMDA receptors, place cells and hippocampal spatial memory. Nat. Rev. Neurosci. 5, 361–372.
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., Tonegawa, S., 2002. Requirement for hippocampal CA3 NMDA receptors in associative memory recall. Science 297, 211–218.
Nakazawa, K., Sun, L.D., Quirk, M.C., Rondi-Reig, L., Wilson, M.A., Tonegawa, S., 2003. Hippocampal CA3 NMDA receptors are crucial for memory acquisition of one-time experience. Neuron 38, 305–315.
Neunuebel, J.P., Knierim, J.J., 2012. Spatial firing correlates of physiologically distinct cell types of the rat dentate gyrus. J. Neurosci. 32, 3848–3858.
Newman, E.L., Gupta, K., Climer, J.R., Monaghan, C.K., Hasselmo, M.E., 2012. Cholinergic modulation of cognitive processing: insights drawn from computational models. Front. Behav. Neurosci. 6, 24.
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., 1990. A computational theory of the hippocampal cognitive map. Prog. Brain Res. 83, 301–312.
O'Keefe, J., Dostrovsky, J., 1971. The hippocampus as a spatial map: preliminary evidence from unit activity in the freely moving rat. Brain Res. 34, 171–175.
O'Keefe, J., Nadel, L., 1978. The Hippocampus as a Cognitive Map. Clarendon Press, Oxford.
O'Reilly, R.C., Rudy, J.W., 2001. Conjunctive representations in learning and memory: principles of cortical and hippocampal function. Psychol. Rev. 108, 311–345.
Oja, E., 1982. A simplified neuron model as a principal component analyser. J. Math. Biol. 15, 267–273.
Panzeri, S., Rolls, E.T., Battaglia, F., Lavis, R., 2001. Speed of feedforward and recurrent processing in multilayer networks of integrate-and-fire neurons. Netw. Comput. Neural Syst. 12, 423–440.
Pitkanen, A., Kelly, J.L., 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.
Preston, A.R., Eichenbaum, H., 2013. Interplay of hippocampus and prefrontal cortex in memory. Curr. Biol. 23, R764–R773.
Redila, V., Kinzel, C., Jo, Y.S., Puryear, C.B., Mizumori, S.J., 2014. A role for the lateral dorsal tegmentum in memory and decision neural circuitry. Neurobiol. Learn. Mem.
Robertson, R.G., Rolls, E.T., 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., 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., Berry, W.O. (Eds.), Neural Models of Plasticity: Experimental and Theoretical Approaches. Academic Press, San Diego, pp. 240–265.
Rolls, E.T., 1989c. Parallel distributed processing in the brain: implications of the functional architecture of neuronal networks in the hippocampus. In: Morris, R.G.M. (Ed.), Parallel Distributed Processing: Implications for Psychology and Neurobiology. Oxford University Press, Oxford, pp. 286–308.
Rolls, E.T., 1989d. The representation and storage of information in neuronal networks in the primate cerebral cortex and hippocampus. In: Durbin, R., Miall, C., 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., Kesner, R.P. (Eds.), Neurobiology of Comparative Cognition. L.Erlbaum, Hillsdale, NJ, 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., 1996a. Roles of long term potentiation and long term depression in neuronal network operations in the brain. In: Fazeli, M.S., Collingridge, G.L. (Eds.), Cortical Plasticity. Bios, Oxford, pp. 223–250.
Rolls, E.T., 1996b. A theory of hippocampal function in memory. Hippocampus 6, 601–620.
Rolls, E.T., 1999. Spatial view cells and the representation of place in the primate hippocampus. Hippocampus 9, 467–480.

Rolls, E.T., 2000. Memory systems in the brain. Annu. Rev. Psychol. 51, 599–630.
Rolls, E.T., 2008. Memory, Attention, and Decision-Making: A Unifying Computational Neuroscience Approach. Oxford University Press, Oxford.
Rolls, E.T., 2010. A computational theory of episodic memory formation in the hippocampus. Behav. Brain Res. 215, 180–196.
Rolls, E.T., 2012a. Advantages of dilution in the connectivity of attractor networks in the brain. Biol. Inspired Cogn. Archit. 1, 44–54.
Rolls, E.T., 2012b. Invariant visual object and face recognition: neural and computational bases, and a model, VisNet. Front. Comput. Neurosci. 6 (35), 1–70.
Rolls, E.T., 2013a. The mechanisms for pattern completion and pattern separation in the hippocampus. Front. Syst. Neurosci. 7, 74.
Rolls, E.T., 2013b. A quantitative theory of the functions of the hippocampal CA3 network in memory. Front. Cell Neurosci. 7, 98.
Rolls, E.T., 2015a. Diluted connectivity in pattern association networks facilitates the recall of information from the hippocampus to the neocortex. Prog. Brain Res. 219, 21–43.
Rolls, E.T., 2015b. Limbic systems for emotion and for memory, but no single limbic system. Cortex 62, 119–157.
Rolls, E.T., 2016a. Cerebral Cortex: Principles of Operation. Oxford University Press, Oxford.
Rolls, E.T., 2016b. Pattern separation, completion, and categorisation in the hippocampus and neocortex. Neurobiol. Learn. Mem. 129, 4–28.
Rolls, E.T., Deco, G., 2010. The Noisy Brain: Stochastic Dynamics as a Principle of Brain Function. Oxford University Press, Oxford.
Rolls, E.T., 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., 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., O'Mara, S., 1993. Neurophysiological and theoretical analysis of how the primate hippocampus functions in memory. In: Ono, T., Squire, L.R., Raichle, M.E., Perrett, D.I., Fukuda, M. (Eds.), Brain Mechanisms of Perception and Memory: From Neuron to Behavior. Oxford University Press, New York, pp. 276–300.
Rolls, E.T., Robertson, R.G., Georges-François, P., 1997a. Spatial view cells in the primate hippocampus. Eur. J. Neurosci. 9, 1789–1794.
Rolls, E.T., Treves, A., Foster, D., 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., 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., 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., Trappenberg, T.P., 2002. A unified model of spatial and episodic memory. Proc. R. Soc. Lond. B 269, 1087–1093.
Rolls, E.T., Treves, A., 1990. The relative advantages of sparse versus distributed encoding for associative neuronal networks in the brain. Network (Bristol, England) 1, 407–421.
Rolls, E.T., Treves, A., 1998. Neural Networks and Brain Function. Oxford University Press, Oxford.
Rolls, E.T., Treves, A., Robertson, R.G., Georges-François, P., Panzeri, S., 1998. Information about spatial view in an ensemble of primate hippocampal cells. J. Neurophysiol. 79, 1797–1813.
Rolls, E.T., Tromans, J., Stringer, S.M., 2008. Spatial scene representations formed by self-organizing learning in a hippocampal extension of the ventral visual system. Eur. J. Neurosci. 28, 2116–2127.
Rolls, E.T., Xiang, J.-Z., 2005. Reward-spatial view representations and learning in the hippocampus. J. Neurosci. 25, 6167–6174.
Rolls, E.T., 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., Franco, L., 2005. Object, space and object-space representations in the primate hippocampus. J. Neurophysiol. 94, 833–844.
Samsonovich, A., McNaughton, B.L., 1997. Path integration and cognitive mapping in a continuous attractor neural network model. J. Neurosci. 17, 5900–5920.
Schultz, S., Rolls, E.T., 1999. Analysis of information transmission in the Schaffer collaterals. Hippocampus 9, 582–598.
Simmen, M.W., Treves, A., Rolls, E.T., 1996. Pattern retrieval in threshold-linear associative nets. Network (Bristol, England) 7, 109–122.
Sompolinsky, H., 1987. The theory of neural networks: the Hebb rule and beyond. In: Heidelberg colloquium on Glassy Dynamics. Springer, pp. 485–527.
Squire, L.R., 1992. Memory and the hippocampus: a synthesis from findings with rats, monkeys and humans. Psychol. Rev. 99, 195–231.
Stefanacci, L., Suzuki, W.A., 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., Ottersen, O.P. (Eds.), 1990. Understanding the Brain through the hippocampus. Elsevier, Oxford.
Stringer, S.M., 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., 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., Trappenberg, T.P., Araujo, I.E.T., 2002a. Self-organizing continuous attractor networks and path integration. Two-dimensional models of place cells. Network Comput. Neural Syst. 13, 429–446.
Stringer, S.M., Trappenberg, T.P., Rolls, E.T., Araujo, I.E.T., 2002b. Self-organizing continuous attractor networks and path integration: one-dimensional models of head direction cells. Network Comput. Neural Syst. 13, 217–242.
Sutherland, R.J., Rudy, J.W., 1991. Exceptions to the rule of space. Hippocampus 1, 250–252.
Suzuki, W.A., Amaral, D.G., 1994a. Perirhinal and parahippocampal cortices of the macaque monkey - cortical afferents. J. Comp. Neurol. 350, 497–533.
Suzuki, W.A., 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., 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., 1993. Mean-field analysis of neuronal spike dynamics. Network (Bristol, England) 4, 259–284.
Treves, A., 1995. Quantitative estimate of the information relayed by Schaffer collaterals. J. Comput. Neurosci. 2, 259–272.
Treves, A., Rolls, E.T., 1991. What determines the capacity of autoassociative memories in the brain? Network (Bristol, England) 2, 371–397.
Treves, A., 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., Rolls, E.T., 1994. A computational analysis of the role of the hippocampus in memory. Hippocampus 4, 374–391.
Ungerleider, L.G., Mishkin, M., 1982. Two cortical visual systems. In: Ingle, D.J., Goodale, M.A., Mansfield, R.J.W. (Eds.), Analysis of Visual Behavior. MIT Press, Cambridge, MA, pp. 549–586.
van Haeften, T., Baks-te-Bulte, L., Goede, P.H., Wouterlood, F.G., 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.
van Strien, N.M., Cappaert, N.L., Witter, M.P., 2009. The anatomy of memory: an interactive overview of the parahippocampal-hippocampal network. Nat. Rev. 10, 272–282.
Wang, S.H., Morris, R.G., 2010. Hippocampal-neocortical interactions in memory formation, consolidation, and reconsolidation. Annu. Rev. Psychol. 61, 49–79. C41–44.
Winocur, G., Moscovitch, M., 2011. Memory transformation and systems consolidation. J. Int. Neuropsychol. Soc. 17, 766–780.
Witter, M.P., 1993. Organization of the entorhinal-hippocampal system: a review of current anatomical data. Hippocampus 3, 33–44.
Witter, M.P., 2007. Intrinsic and extrinsic wiring of CA3: indications for connectional heterogeneity. Learn. Mem. 14, 705–713.

Witter, M.P., Naber, P.A., van Haeften, T., Machielsen, W.C., Rombouts, S.A., Barkhof, F., Scheltens, P., Lopes da Silva, F.H., 2000a. Cortico-hippocampal communication by way of parallel parahippocampal-subicular pathways. Hippocampus 10, 398–410.
Witter, M.P., Van Hoesen, G.W., Amaral, D.G., 1989. Topographical organisation of the entorhinal projection to the dentate gyrus of the monkey. J. Neurosci. 9, 216–228.
Witter, M.P., Wouterlood, F.G., Naber, P.A., 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., 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.
Zilli, E.A., 2012. Models of grid cell spatial firing published 2005-2011. Front. Neural Circuits 6, 16.