Neuroscience and Biobehavioral Reviews 48 (2015) 92–147

Contents lists available at ScienceDirect

Neuroscience and Biobehavioral Reviews

journal homepage: www.elsevier.com/locate/neubiorev

Review

A computational theory of hippocampal function, and tests of the theory: New developments

Raymond P. Kesner[a], Edmund T. Rolls[b,c,*]

[a] University of Utah, Department of Psychology, Salt Lake City, UT 84112, USA
[b] Oxford Centre for Computational Neuroscience, Oxford, UK
[c] University of Warwick, Department of Computer Science, Coventry CV4 7AL, UK

ARTICLE INFO

Article history:
Received 5 July 2014
Received in revised form 24 October 2014
Accepted 12 November 2014
Available online 20 November 2014

Keywords:
Hippocampus
Episodic memory
Recall
Autoassociation
Cortical Backprojections
CA3
CA1
Dentate granule cells

ABSTRACT

The aims of the paper are to update Rolls' quantitative computational theory of hippocampal function and the predictions it makes about the different subregions (dentate gyrus, CA3 and CA1), and to examine behavioral and electrophysiological data that address the functions of the hippocampus and particularly its subregions. Based on the computational proposal that the dentate gyrus produces sparse representations by competitive learning and via the mossy fiber pathway forces new representations on the CA3 during learning (encoding), it has been shown behaviorally that the dentate gyrus supports spatial pattern separation during learning. Based on the computational proposal that CA3–CA3 autoassociative networks are important for episodic memory, it has been shown behaviorally that the CA3 supports spatial rapid one-trial learning, learning of arbitrary associations where space is a component, pattern completion, spatial short-term memory, and spatial sequence learning by associations formed between successive items. The concept that the CA1 recodes information from CA3 and sets up associatively learned backprojections to neocortex to allow subsequent retrieval of information to neocortex, is consistent with findings on consolidation. Behaviorally, the CA1 is implicated in processing temporal information as shown by investigations requiring temporal order pattern separation and associations across time; and computationally this could involve associations in CA1 between object and timing information that have their origins in the lateral and medial entorhinal cortex respectively. The perforant path input from the entorhinal cortex to DG is implicated in learning, to CA3 in retrieval from CA3, and to CA1 in retrieval after longer time intervals ("intermediate-term memory") and in the temporal sequence memory for objects.

Contents

* Corresponding author at: Oxford Centre for Computational Neuroscience, Oxford, UK.
E-mail address: Edmund.Rolls@oxcns.org (E.T. Rolls).
URL: www.oxcns.org (E.T. Rolls).

http://dx.doi.org/10.1016/j.neubiorev.2014.11.009
0149-7634/

1. Introduction

In this paper a computational theory of hippocampal function developed by Rolls (1987, 1989a,b,c, 1996b, 2010b, 2013b,c), Treves and Rolls (1992, 1994) and with other colleagues (Rolls and Stringer, 2005; Rolls et al., 2002) is refined and further developed, and tests of the theory based especially on subregion analyses of the hippocampal system are described. The relation of this theory to other computational approaches to hippocampal function is described. This paper follows the general format of the earlier joint paper by Rolls and Kesner (2006), but updates the theory and the tests of the theory. The aims of this paper are thus to update a particular computational theory of hippocampal function (which remains the only quantitative theory of hippocampal function in memory and its recall to the neocortex), and the predictions it makes about the different hippocampal subregions (dentate gyrus, CA3 and CA1); and to update the empirical tests of these predictions by especially subregion analysis of hippocampal function. The paper is intended to provide an updated position or landmark description of the theory, and how it has been tested, and this combination makes this paper a unique contribution.

The theory was originally developed as described next, and 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.[1] The model was assessed by Willshaw and Buckingham (1990). Early work of Gardner-Medwin (1976) showed how progressive recall could operate in a network of binary neurons with binary synapses. Rolls (1987) described a theory of the hippocampus presented to the Dahlem conference in 1985 on the Neural and Molecular Bases of Learning 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.[2] 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. This

[1] Marr (1971) showed how a network with recurrent collaterals could complete a memory using a partial retrieval cue, and how sparse representations could increase the number of memories stored (see also Willshaw and Buckingham, 1990). The analysis of these autoassociation or attractor networks was developed by Kohonen (1997, 1984) and Hopfield (1982), and the value of sparse representations was quantified by Treves and Rolls (1991). Marr (1971) did not specify the functions of the dentate granule cells vs the CA3 cells vs the CA1 cells (which were addressed in the Rolls (1989a,b) papers and by Treves and Rolls (1992, 1994)), nor how retrieval to the neocortex of hippocampal memories could be produced, for which a quantitative theory was developed by Treves and Rolls (1994). In addition, Treves and Rolls (1994) and Rolls and Treves (1998) have argued that approaches to neurocomputation which base their calculations on what would happen in the tail of an exponential, Poisson, or binomial distribution are very fragile.

[2] McNaughton and Morris (1987) at about the same time suggested that the CA3 network might be an autoassociation network, and that the mossy fiber to CA3 connections might implement 'detonator' synapses. However, 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). Contributions by Levy (e.g. 1989); McNaughton (1991); Hasselmo; Lisman; McClelland et al. (1995), and many others, are described below.

theory was developed further (Rolls, 1989a,b,c,d, 1990a,b), including further details about how the backprojections could operate (Rolls, 1989b,d), and how the dentate granule cells could operate as a competitive network (Rolls, 1989a). Quantitative aspects of the theory were then developed with A. Treves, who brought the expertise of theoretical physics applied previously mainly to understand the properties of fully connected attractor networks with binary neurons (Amit, 1989; Hopfield, 1982), to bear on the much more diluted connectivity of the recurrent collaterals found in real biological networks (e.g. 2% between CA3 pyramidal cells in the rat), in networks of neurons with graded (continuously variable) firing rates, graded synaptic strengths, and sparse representations in which only a small proportion of the neurons is active at any one time, as is found in the hippocampus (Rolls, 2008a; Treves, 1990; Treves and Rolls, 1991). These developments in understanding quantitatively the operation of more biologically relevant recurrent networks with modifiable synapses were applied quantitatively to the CA3 region (Treves and Rolls, 1991), and to the issue of why there are separate mossy fiber and perforant path inputs to the CA3 cells of the hippocampus (Treves and Rolls, 1992). The whole model of the hippocampus was described in more detail, and a quantitative treatment of the theory of recall by backprojection pathways in the brain was provided by Treves and Rolls (1994). The speed of operation of the CA3 system, and of the cortico-cortical recall backprojections, was addressed in a number of new developments (Battaglia and Treves, 1998b; Panzeri et al., 2001; Simmen et al., 1996a; Treves, 1993). Rolls (1995) produced a simulation of the operation of the major parts of the hippocampus from the entorhinal cortex through the dentate, CA3, and CA1 cells back to the hippocampus, which established the quantitative feasibility of the whole theory, and raised a number of important issues considered in Section 2.2.8. Further developments of the theory (Rolls, 2010b, 2013b,c), and new developments introduced here, are described below.

Tests of the theory including analysis of the functions of different subregions of the hippocampus are described in Section 3.

2. The theory

2.1. 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.

2.1.1. Evidence from the effects of damage to the hippocampus in primates

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 (Burgess et al., 2002; Crane and Milner, 2005; Gaffan, 1994; Gaffan and Saunders, 1985; Parkinson et al., 1988; Smith and Milner, 1981). 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). (Rolls further predicts that a more difficult object–place learning task with non-trial-unique stimuli and with many object–place pairs would be impaired by neurotoxic hippocampal lesions.) Further, neurotoxic lesions that selectively damage the primate hippocampus impair spatial scene memory, tested by the ability to remember where in a scene to touch to obtain reward (Murray et al., 1998). Also, fornix lesions impair conditional left–right discrimination learning, in which the visual appearance of an object specifies whether a response is to be made to the left or the right (Rupniak and Gaffan, 1987). A comparable deficit is found in humans (Petrides, 1985). Fornix sectioned monkeys are also impaired in learning on the basis of a spatial cue which object to choose (e.g. if two objects are on the left, choose object A, but if the two objects are on the right, choose object B) (Gaffan and Harrison, 1989a). Monkeys with fornix damage are also impaired in using information about their place in an environment. For example, there are learning impairments when which of two or more objects the monkey had to choose depended on the position of the monkey in the room (Gaffan and Harrison, 1989b). More recently, Lavenex and colleagues have described deficits produced by hippocampal damage in monkeys performing allocentric spatial memory tasks (Banta Lavenex and Lavenex, 2009). One such task involved freely moving in an environment using allocentric spatial room cues to remember the locations of inverted cups that contained food. This is a food reward-allocentric place association task. Rats with hippocampal lesions are also impaired in using environmental spatial cues to remember particular places (Cassaday and Rawlins, 1997; Jarrard, 1993; Kesner et al., 2004; Martin et al., 2000; O'Keefe and Nadel, 1978), to utilize spatial cues or to bridge delays (Kesner, 1998; Kesner et al., 2004; Kesner and Rolls, 2001; Rawlins, 1985; Rolls and Kesner, 2006), or to perform relational operations on remembered material (Eichenbaum, 1997). 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).

Whereas the hippocampus is involved in spatial, object–place, and object-temporal sequence memory, the perirhinal cortex is involved in recognition memory, as shown by the effects of lesions (Baxter and Murray, 2001; Buckley and Gaffan, 2006; Malkova et al., 2001). Moreover, and consistently, the primate perirhinal cortex contains neurons related to the degree of long-term familiarity of objects, with neuronal responses increasing over hundreds of presentations (Hölscher et al., 2003b; Rolls, 2008a; Rolls et al., 2005a) (see further Section 2.1.5). It has also been suggested that the perirhinal cortex puts together combinations of features to represent complex objects by building conjunctive/configural representations (Bussey et al., 2005). However, when we recorded in the primate perirhinal cortex, we did not find much evidence for very specifically tuned object neurons, but did find that neuronal activity is related to short-term perceptual memory (Hölscher and Rolls, 2002). In the light of this, and the fact that inferior temporal visual cortex neurons do provide very good invariant representations of objects (Rolls, 2012b; Rolls and Treves, 2011), the possibility is suggested that one function of the perirhinal cortex may be to provide a visual short-term memory, which is usually required for the tasks such as oddity with multiple distractors that led to the conjunctive/configural hypothesis (Bussey et al., 2005). Indeed, the visual short-term memory capability of the perirhinal cortex is probably important also for its role in delayed paired-associate learning (Fujimichi et al., 2010).

The extensive literature on the effects of hippocampal damage in rats is considered in Section 3, in particular the parts of it that involve selective lesions to different subregions of the hippocampus and allow the theory described here to be tested. The theory described here is intended to be as relevant to rodents as to primates, the main difference being that the spatial representation used in the hippocampus of rats is about the place where the rat is, whereas the spatial representation in primates includes a representation of space "out there", as represented by spatial view cells (see Section 2.1.3). Although it is argued that the hippocampus can form

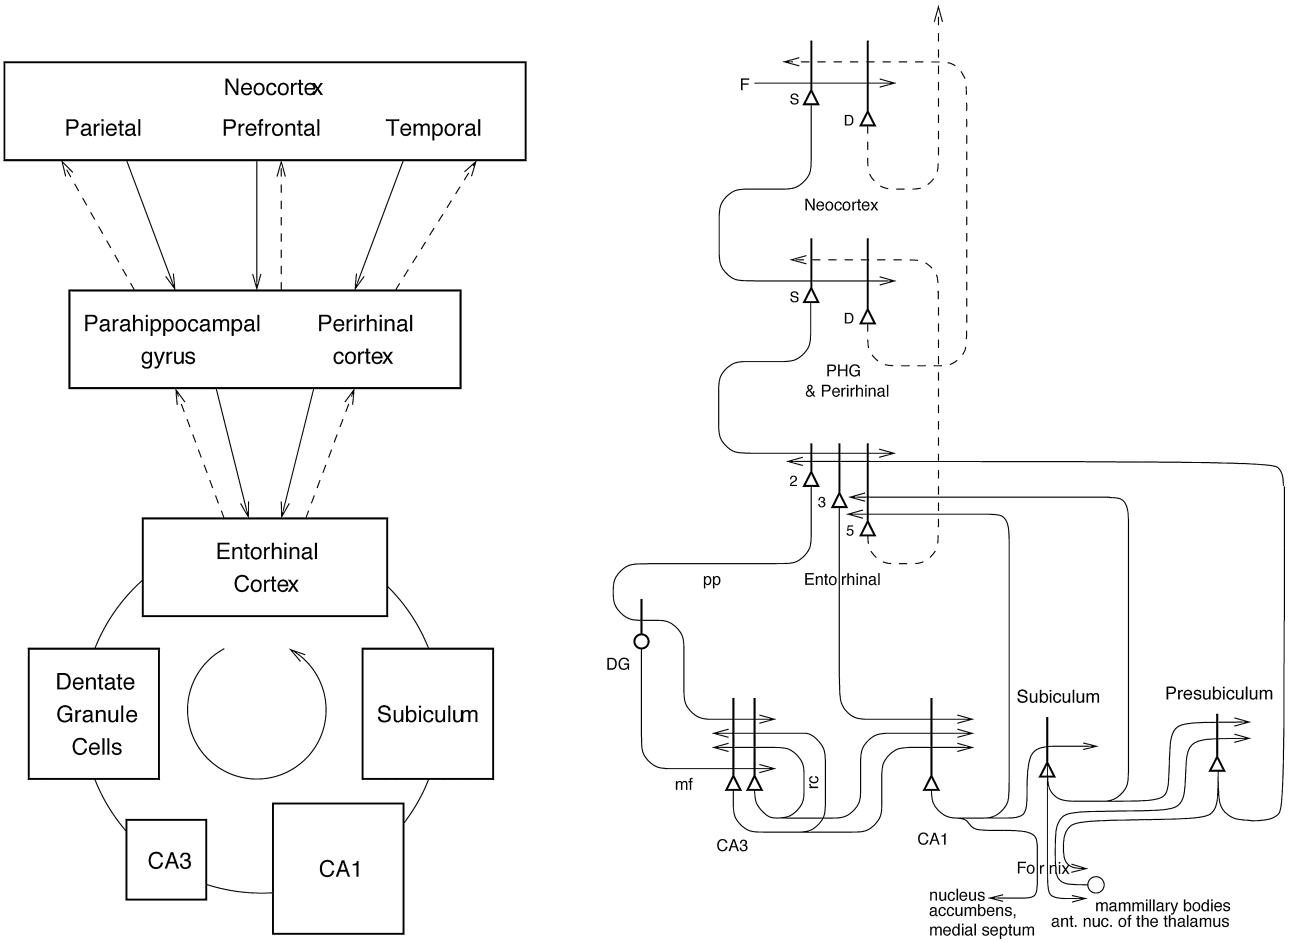

Fig. 1. Forward connections (solid lines) from areas of cerebral association neocortex via the parahippocampal gyrus and perirhinal cortex, and entorhinal cortex, to the hippocampus; and backprojections (dashed lines) via the hippocampal CA1 pyramidal cells, subiculum, and parahippocampal gyrus to the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

associations between places and other inputs rapidly to help form an episodic memory, it is shown in Section 3 that in addition some types of hippocampus-dependent learning may occur more slowly, including trace conditioning; paired associate learning; and the formation of spatial representations of a new environment (Wilson and McNaughton, 1993).

2.1.2. *The necessity to recall information from the hippocampus, episodic memory, and consolidation*

The information about episodic events recalled from the hippocampus could be used to help form semantic memories (Rolls, 1989b,d, 1990a; Treves and Rolls, 1994) or neocortical 'schemas' (Preston and Eichenbaum, 2013). For example, remembering many particular journeys could help to build a geographic cognitive map in the neocortex. The hippocampus and neocortex would thus be complementary memory systems, with the hippocampus being used for rapid, 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; Rolls, 1989b; Treves and Rolls, 1994). A view close to this in which a hippocampal episodic memory trace helps to build a neocortical semantic memory has also been developed by Moscovitch et al. (2005). The quantitative evidence considered below (Section 2.2.3.1) and by Treves and Rolls (1994) on the storage capacity of the CA3 system indicates that episodic memories could be stored in the hippocampus for at least a considerable period. The theory described below shows how information could be retrieved within the hippocampus, and how this retrieved information could enable the activity in neocortical areas that was present during the original storage of the episodic event to be reinstated, thus implementing recall, by using hippocampo-neocortical backprojections (see Fig. 1). The recalled information could then be stored in the neocortex, implementing a systems-level type of consolidation. The general concept that backprojections from the entorhinal cortex to the neocortex are involved in systems-level consolidation has been discussed by others, though when the recall is to the anterior cingulate cortex (Takehara-Nishiuchi, 2014), this may be particularly for the retrieval of emotional components of episodic memory (Rolls, 2014b).

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

What we mean by an episodic memory of the type in which the hippocampus could be involved is described further next. A hippocampus-dependent episodic memory is unstructured in the sense that the information is stored as an association between representations such as spatial location, objects, and position in a temporal sequence. An episodic memory could be a single event capturing what happens in a time period in the order of 1–10 s (for example a memory of who was present at one's fifth birthday, what they were wearing, where they sat, where the party was, etc.), or a sequence of events with some information about the temporal order of the events. A hippocampus-dependent episodic memory includes spatial and/or temporal e.g. sequence information. All this is in contrast to a neocortical semantic memory which might utilize many such episodic memories, for example of separate train journeys, to form a topological map. Such a semantic memory might need to benefit from information about exceptions, which contributes to the type of structure needed for a semantic memory. For example we may be told that a bat, although it flies, is not a bird because it is a mammal which nurses its young, and this information allows restructuring of a semantic representation of what constitutes a bird beyond the simple association and generalization from limited experience that birds fly. Because this active restructuring is an important part of semantic memory that may frequently benefit from language, we propose that semantic memories are best formed during waking using thought and reasoning, and not during sleep. Indeed, we propose that episodic events that are recalled actively in this way from the hippocampus can become incorporated into neocortical representations and thus consolidated into semantic including autobiographical memories. If episodic material (such as where and with whom one had dinner three days ago) is not recalled in this way, then it will be gradually overwritten in hippocampal memory as new sparse distributed representations are added to the store, in the way that has been analyzed quantitatively (Parisi, 1986; Rolls, 2008a; Rolls and Treves, 1998). In contrast, some have suggested that the transfer of information from the hippocampus to the neocortex occurs especially during sleep (Marr, 1971; Schwindel and McNaughton, 2011). Indeed, it has been shown that after learning in hippocampal-dependent tasks, neocortical representations may change (Schwindel and McNaughton, 2011). Although this has been interpreted as the transfer of memories from the hippocampus to the neocortex (Schwindel and McNaughton, 2011), it should be noted that if the hippocampal representation changes as a result of learning, then the altered representation in CA1 will, even with fixed synaptic connections back to neocortex, alter neocortical firing, with no learning or actual 'transfer' involved. (This occurs whenever one vector comprised by a population of neurons firing changes and influences another vector of neuronal firing through fixed connections.) Indeed, for the reasons given above, active transfer of episodic memories from the hippocampus to the neocortex to form semantic memories is more likely Rolls proposes during the guided construction and correction of semantic memories in the cortex during waking. The actual addition of new information to the neocortical semantic representation could of course itself be fast, for example when one recalls a particular train journey, and incorporates information from that into a geographical representation. An additional property of hippocampal episodic memories is that when recalled to the neocortex they can be used flexibly for many types of behavioral output (because the information is back in the neocortex). We further note that rapid learning is not a unique feature of episodic memory, with for example the stimulus-reward learning and its reversal implemented in the primate orbitofrontal cortex possible in one trial (Rolls, 2014a,b).

This raises the issue of the possible gradient of retrograde amnesia following hippocampal damage, and of whether information originally hippocampus-dependent for acquisition gradually becomes "consolidated" in the neocortex and thereby become hippocampus-independent over time (Debiec et al., 2002; McGaugh, 2000). The issue of whether memories stored some time before hippocampal damage are less impaired than more recent memories, and whether the time course is minutes, hours, days, weeks or years is a debated issue (Gaffan, 1993; Squire, 1992). (In humans, there is evidence for a gradient of retrograde amnesia; in rats and monkeys, hippocampal damage in many studies appears to impair previously learned hippocampal-type memories, suggesting that in these animals, at least with the rather limited numbers of different memories that need to be stored in the tasks used, the information remains in the hippocampus for long periods). If there is a gradient of retrograde amnesia related to hippocampal damage, then this suggests that information may be retrieved from the hippocampus if it is needed, allowing the possibility of incorporating the retrieved information into neocortical memory stores. If on the other hand there is no gradient of retrograde amnesia related to hippocampal damage, but old as well as recent memories of the hippocampal type are stored in the hippocampus, and lost if it is damaged, then again this implies the necessity of a mechanism to retrieve information stored in the hippocampus, and to use this retrieved information to affect neural circuits elsewhere (for if this were not the case, information stored in the hippocampus could never be used for anything). The current perspective is thus that whichever view of the gradient of retrograde amnesia is correct, information stored in the hippocampus will need to be retrieved and affect other parts of the brain in order to be used. 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. The backprojections from the hippocampus to the neocortex are one of the two major outputs of the hippocampus (see Fig. 1). The backprojections are most likely to be involved in what is described by humans as recall, and in enabling information about an episode captured on the fly to be incorporated into long-term, possibly semantic, neocortical stores with a rich associative structure (cf. McClelland et al., 1995). As a result of such neocortical recall, action may be initiated.

The other major set of outputs from the hippocampus projects via the fimbria/fornix system to the anterior nucleus of the thalamus (both directly and via the mammillary bodies), which in turn projects to the cingulate cortex. This may provide an output for more action-directed use of information stored in the hippocampus, for example in the initiation of conditional spatial responses in a visual conditional spatial response task (Miyashita et al., 1989; Rupniak and Gaffan, 1987). In such a task, a rapid mapping must be learned between a visual stimulus and a spatial response, and a new mapping must be learned each day. The hippocampus is involved in this rapid visual to spatial response mapping (Rupniak and Gaffan, 1987), and the way in which hippocampal circuitry may be appropriate for this is that the CA3 region enables signals originating from very different parts of the cerebral cortex to be associated rapidly together (see below).

2.1.3. *Systems-level neurophysiology of the primate hippocampus*

The systems level neurophysiology of the hippocampus shows what information could be stored or processed by the hippocampus. To understand how the hippocampus works it is not sufficient to state just that it can store information – one needs to know what information. The systems-level neurophysiology of the primate hippocampus has been reviewed by Rolls and Xiang (2006), and a 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

(see Jeffery et al., 2004; Jeffery and Hayman, 2004; McNaughton et al., 1983; Muller et al., 1991; O'Keefe, 1984).

2.1.3.1. Spatial view neurons. The primate hippocampus contains spatial cells that respond when the monkey looks at a certain part of space, for example at one quadrant of a video monitor while the monkey is performing an object–place memory task in which he must remember where on the monitor he has seen particular images (Rolls et al., 1989). Approximately 9% of the hippocampal neurons have such spatial view fields, and 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 latter point shows that information from very different parts of the cerebral cortex (parietal for spatial information, and inferior temporal for visual information about objects) is brought together onto single neurons in the primate hippocampus. The representation of space is for the majority of hippocampal neurons in allocentric not egocentric coordinates (Feigenbaum and Rolls, 1991; Georges-François et al., 1999).

In rodents, place cells, which respond when a rat is at a particular place, are found in the hippocampus (Hartley et al., 2014; Jeffery, 2011; Jeffery et al., 2004; Jeffery and Hayman, 2004; McNaughton et al., 1983; Muller et al., 1991; O'Keefe, 1984; O'Keefe and Dostrovsky, 1971). Place cells are found in regions CA3 and CA1 (with smaller place fields in the dentate granule cells) (Neunuebel and Knierim, 2012). The representation is allocentric (as contrasted with egocentric), in that the neurons fire whenever the rat is at the place, typically independently of the head direction of the rat. To analyze whether such cells might be present in the primate hippocampus, Rolls and O'Mara (1993, 1995) recorded the responses of hippocampal cells when macaques were moved in a small chair or robot on wheels in a cue-controlled testing environment. The most common type of cell responded to the part of space at which the monkeys were looking, independently of the place where the monkey was. Ono et al. (1993) performed studies on the representation of space in the primate hippocampus while the monkey was moved in a cab to different places in a room. They found that 13.4% of hippocampal formation neurons fired more when the monkey was at some but not other places in the test area, but it was not clear whether the responses of these neurons occurred to the place where the monkey was independently of spatial view, or whether the responses of place-like cells were view dependent.

In rats, place cells fire best during active locomotion by the rat (Foster et al., 1989). To investigate whether place cells might be present in monkeys if active locomotion was being performed, we (Georges-François et al., 1999; Robertson et al., 1998; Rolls et al., 1997a, 1998) recorded from single hippocampal neurons while monkeys move themselves round the test environment by walking (or running) on all fours. We found that these hippocampal "spatial view neurons" responded significantly differently for different allocentric spatial views and had information about spatial view in their firing rate, but did 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. This experiment by Robertson et al. (1998) (see also Rolls et al., 1997b) shows that primate hippocampal spatial view neurons can be updated for at least short periods by idiothetic (self-motion) cues including eye position and head direction signals. Consistent with this, a small drift is sometimes evident after a delay when the view details are obscured, consistent with inaccuracies in the temporal path integration of these signals, which is then corrected by showing the visual scene again.

It is essential to measure the firing rate of a primate hippocampal cell with different head directions so that different spatial views can be compared, as testing with just one head direction (Matsumura et al., 1999) cannot provide evidence that will distinguish a place cell from a spatial view cell. These points will need to be borne in mind in future studies of hippocampal neuronal activity in primates including humans (cf. Ekstrom et al., 2003; Fried et al., 1997; Kreiman et al., 2000), and simultaneous recording of head position, head direction, and eye position, as described by Rolls and colleagues (Georges-François et al., 1999; Robertson et al., 1998; Rolls et al., 1997a, 1998) will be needed. To distinguish spatial view from place cells it will be important to test neurons while the primate or human is in one place with all the different spatial views visible from that place; and to also test the same neuron when the organism is in a different place, but at least some of the same spatial views are visible, as has been done in our primate recording.

In rodents, grid cells which have repeated peaks of 'place' firing as the animal traverses a space are found in the medial entorhinal cortex (Hafting et al., 2005). In primates, there is now evidence that there is an analogous but spatial view-based grid-cell like representation in the entorhinal cortex, with neurons having grid-like firing as the monkey moves the eyes across a spatial scene (Killian et al., 2012). Similar competitive learning processes may transform these entorhinal cortex 'spatial view grid cells' into hippocampal spatial view cells, and may help with the idiothetic (produced in this case by movements of the eyes) update of spatial view cells (Robertson et al., 1998). The presence of spatial view grid cells in the entorhinal cortex of primates (Killian et al., 2012) is of course predicted from the presence of spatial view cells in the primate CA3 and CA1 regions (Georges-François et al., 1999; Robertson et al., 1998; Rolls, 2008a; Rolls et al., 1997a, 1998; Rolls and Xiang, 2006). 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 grid-like properties have now been described (Jacobs et al., 2013).

2.1.3.2. Object–place neurons. 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. (2005b) made recordings from single hippocampal formation neurons while macaques performed an object–place memory task which required the monkeys to learn associations between objects, and where they were shown in a room. Some neurons (10%) responded differently to different objects independently of location; other neurons (13%) responded to the spatial view independently of which object was present at the location; and some neurons (12%) responded to a combination of a particular object and the place where it was shown in the room. These results show that there are separate as well as combined representations of objects and their locations in space in the primate hippocampus. This is a property required in an episodic memory system, for which associations between objects and the places where they are seen, is 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., 2005b).

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). Results consistent with our object–place neurons in primates are that Diamond and colleagues have now shown using the vibrissa somatosensory input for the 'object' system, that rat hippocampal neurons respond to object–place combinations, objects, or places, and there is even a reward–place association system in rats (Itskov et al., 2011) similar to that in primates described below. This brings the evidence from rats closely into line with the evidence from primates of hippocampal neurons useful for object–place episodic associative memory.

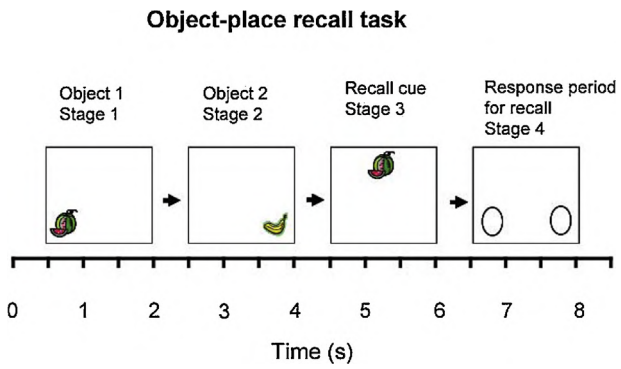

Fig. 2. The object–place recall task. One trial is shown. After a 0.5 s tone to indicate the start of a trial, in Stage 1 one of two objects (O1) is shown at one of the places (P1). (The object and the place are chosen randomly on each trial.) To ensure that the monkey sees the stimulus, the monkey can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

2.1.3.3. Recall-related neurons in the primate hippocampus. It has now been possible to investigate directly, neurophysiologically, the hippocampal recall process in primates (Rolls and Xiang, 2006). We used a visual object–place memory task because this is prototypical of episodic memory. It has been shown that a one-trial odor–place recall memory task is hippocampal-dependent in rodents (Day et al., 2003). We designed a one-trial object–place recall task, in which the whole memory was recalled from a part of it. The task is illustrated in Fig. 2. Images of new objects were used each day, and within a day the same objects were used, so that with non-trial unique objects within a day, the recall task is quite difficult. Recordings were made from 347 neurons in the hippocampus of a macaque performing the object–place recall task. The following types of neurons were found in the task (Rolls and Xiang, 2006).

One type of neuron had responses that occurred to one of the objects used in the task. A number of these neurons had activity that was related to the recall process. For example, one of these neurons had activity that was greater to object one not only when it was shown in stages 1, 2 and 3 of the task, but also in the delay period following stage 3 when the object was no longer visible, and in stage 4, when also the object was no longer visible and the macaque was touching the remembered location of that object. Thus while the location was being recalled from the object, this type of neuron continued to respond as if the object was present, that is it kept the representation of the object active after the object was no longer visible, and the place to touch was being recalled. 16 of the neurons responded in this way, and an additional 6 had object-related firing that did not continue following stage 3 of the task in the recall period. The difference of the firing rates of these 22 neurons to the different objects were in many cases highly statistically significant (e.g. $p < 10^{-6}$) (Rolls and Xiang, 2006). None of these neurons had differential responses for the different places used in the object–place recall task.

A second type of neuron had responses related to the place (left or right) in which an object was shown in stages 1 or 2 of each trial. For example, one of these neurons responded more when an object was shown in the left position (P1) than in the right position (P2) on the screen. Interestingly, when the recall object was shown in stage 3 of the trial in the top center of the screen, the neuron also responded as if the left position (P1) was being processed on trials on which the left position had to be recalled. This firing continued in the delay period after the recall cue had been removed at the end of stage 3, and into stage 4. Thus this type of neuron appeared to reflect the recall of the position on the screen at which the object had been represented. Analysis of trials on which errors were made indicated that the responses were not just motor response related, for if due to some response bias the monkey touched the incorrect side, the neuron could still respond according to the correct recalled location. 13 neurons had differential responses to the different places P1 and P2, and continued to show place-related activity in the recall part of the task, stage 3. Five other neurons had left–right place-related responses without a memory recall component, in that they did not respond in stage 3 of the task, when a non-spatial recall stimulus was being shown, and a place should be being recalled. The population of 18 neurons as a population had statistically significant place-related responses (Rolls and Xiang, 2006). The new finding is that 13 of the neurons had place-related responses when a place was being recalled by an object cue.

In addition to the neurons described above, 3 further neurons responded to particular combinations of objects and places, e.g. to object 1 when it was shown in place 1, but not to other combinations. The recording sites of the object and of the place neurons were within the hippocampus proper (Rolls and Xiang, 2006). The mean firing rate of the population of responsive neurons to the most effective object or place was 7.2 ± 0.6 spikes/s ($\pm$sem), and their mean spontaneous rate was 3.2 ± 0.6 spikes/s.

These findings (Rolls and Xiang, 2006) are the first we know in the primate hippocampus of neuronal activity that is related to recall. It is particularly interesting that the neurons with continuing activity to the object after it had disappeared in the recall phase of the task could reflect the operation of the object–place recall process that is hypothesized to take place in the CA3 cells. By continuing to respond to the object while the place is being recalled in the task, the object-related neurons could be part of the completion of the whole object–place combination memory from an autoassociation or attractor process in CA3 (Rolls and Kesner, 2006). Consistent with these findings, and with the computational theory, it has now been reported that human hippocampal neurons are activated during recall (Gelbard-Sagiv et al., 2008).

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

2.1.3.4. Reward–place neurons in the primate hippocampus. The primate anterior hippocampus (which corresponds to the rodent ventral hippocampus) receives inputs from brain regions involved in reward processing such as the amygdala and orbitofrontal cortex (Carmichael and Price, 1995; Pitkanen et al., 2002; Stefanacci et al., 1996; Suzuki and Amaral, 1994a). To investigate how this affective input may be incorporated into primate hippocampal function, Rolls and Xiang (2005) recorded neuronal activity while macaques performed a reward–place association task in which each spatial scene shown on a video monitor had one location which if touched yielded a preferred fruit juice reward, and a second location which yielded a less preferred juice reward. Each scene had

different locations for the different rewards. Of 312 hippocampal neurons analyzed, 18% responded more to the location of the preferred reward in different scenes, and 5% to the location of the less preferred reward (Rolls and Xiang, 2005). When the locations of the preferred rewards in the scenes were reversed, 60% of 44 neurons tested reversed the location to which they responded, showing that the reward–place associations could be altered by new learning in a few trials. The majority (82%) of these 44 hippocampal reward–place neurons tested did not respond to object–reward associations in a visual discrimination object–reward association task. Thus the primate hippocampus contains a representation of the reward associations of places "out there" being viewed, and this is a way in which affective information can be stored as part of an episodic memory, and how the current mood state may influence the retrieval of episodic memories. There is consistent recent evidence that rewards available in a spatial environment can influence the responsiveness of rodent place neurons (Hölscher et al., 2003a; Redila et al., 2014; Tabuchi et al., 2003).

2.1.3.5. Primate hippocampal neurons with activity related to conditional spatial responses. In another type of task for which the primate hippocampus is needed, conditional spatial response learning, in which the monkeys had to learn which spatial response to make to different stimuli, that is, to acquire associations between visual stimuli and spatial responses, 14% of hippocampal neurons responded to particular combinations of visual stimuli and spatial responses (Miyashita et al., 1989). The firing of these neurons could not be accounted for by the motor requirements of the task, nor wholly by the stimulus aspects of the task, as demonstrated by testing their firing in related visual discrimination tasks. These results showed that single hippocampal neurons respond to combinations of the visual stimuli and the spatial responses with which they must become associated in conditional response tasks, and are consistent with the computational theory described above according to which part of the mechanism of this learning involves associations between visual stimuli and spatial responses learned by single hippocampal neurons. In a following study by Cahusac et al. (1993), it was found that during such conditional spatial response learning, 22% of this type of neuron analyzed in the hippocampus and parahippocampal gyrus altered their responses so that their activity, which was initially equal to the two new stimuli, became progressively differential to the two stimuli when the monkey learned to make different responses to the two stimuli (Cahusac et al., 1993). These changes occurred for different neurons just before, at, or just after the time when the monkey learned the correct response to make to the stimuli, and are consistent with the hypothesis that when new associations between objects and places (in this case the places for responses) are learned, some hippocampal neurons learn to respond to the new associations that are required to solve the task. Similar findings have been described by Wirth et al. (2003).

2.1.4. 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; Amaral et al., 1992; Andersen et al., 2007; Kondo et al., 2009; Lavenex et al., 2004; Rolls, 2008a; Rolls and Kesner, 2006; Suzuki and Amaral, 1994b; Van Hoesen, 1982; van Strien et al., 2009; Witter, 2007; Witter et al., 1989b, 2000b) (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; Pitkanen et al., 2002; Stefanacci et al., 1996; Suzuki and Amaral, 1994a). (The connections are analogous in the rat, although areas such as the temporal lobe visual cortex areas are not well developed in rodents, and the parahippocampal gyrus may be represented by the dorsal perirhinal cortex, Burwell et al., 1995). In rodents, the medial entorhinal cortex contains spatial representations in the form of grid cells (Fyhn et al., 2004; Giocomo et al., 2011; Hafting et al., 2005) and the projections to the hippocampus involve NMDA receptors (Martinez et al., 2002). The lateral entorhinal cortex contains non-spatial (e.g. object, and odor) representations (Takehara-Nishiuchi, 2014) and the projections to the hippocampus involve opioid receptors (Martinez et al., 2002). Layer 2 of the entorhinal cortex projects to dentate gyrus, and CA3a,b,c. Layer 3 of the entorhinal cortex projects to CA1, with the medial entorhinal cortex projecting more strongly to distal CA1, and the lateral entorhinal cortex projecting more strongly to proximal CA1 (Amaral and Witter, 1995; Nakamura et al., 2013; Witter, 1993, 2007).

Given that some topographic segregation is maintained in the afferents to the hippocampus through the perirhinal and parahippocampal cortices (Amaral and Witter, 1989; van Strien et al., 2009), it may be that these areas are able to subserve memory within one of these topographically separated areas, of for example visual object, *or* spatial, *or* olfactory information. In contrast, the final convergence afforded by the hippocampus into one network in CA3 (see Fig. 1) may be especially appropriate for an episodic memory typically involving arbitrary associations between any of the inputs to the hippocampus, e.g. spatial, vestibular related to self-motion, visual object, olfactory, and auditory (see below). There are also direct subcortical inputs from for example the amygdala and septum (Amaral, 1986).

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). CA1 and subiculum also project to subcortical areas via the fimbria/fornix such as the mammillary bodies and anterior thalamic nuclei (see Fig. 1). In addition to the projections originating in CA1, projections out of Ammon's horn also originate in CA3. Many researchers have reported that CA3 projects to the lateral and medial septal nuclei (Amaral and Witter, 1995; Gaykema et al., 1991; Risold and Swanson, 1997). The lateral septum also has projections to the medial septum (Jakab and Leranth, 1995), which in turn projects to subiculum and eventually entorhinal cortex (Amaral and Witter, 1995; Jakab and Leranth, 1995).

2.1.5. Perirhinal cortex, recognition memory, and long-term familiarity memory

The functions of the perirhinal cortex are different to those of the hippocampus, and are reviewed briefly here partly because it provides afferents to the hippocampus via the entorhinal cortex, and partly to make it clear that there are memory functions at one time attributed to the hippocampus that a hippocampal theory need not account for, given the memory functions now known to be performed by the perirhinal cortex.

The perirhinal cortex is involved in object recognition memory in that damage to the perirhinal cortex produces impairments in object recognition memory tasks in which several items intervene between the sample presentation of a stimulus and its presentation again as a match stimulus (Malkova et al., 2001; Zola-Morgan et al., 1989, 1994). Indeed, damage to the perirhinal cortex rather than to the hippocampus is believed to underlie the impairment in recognition memory found in amnesia in humans associated with medial temporal lobe damage (Brown and Aggleton, 2001; Buckley

and Gaffan, 2000; Winters and Bussey, 2005). Neurophysiologically, it has been shown that many inferior temporal cortex (area TE) neurons (Rolls, 2000a; Rolls and Deco, 2002), which provide visual inputs to the perirhinal cortex (Suzuki and Amaral, 1994a,b), respond more to the first, than to the second, presentation of a stimulus in a running recognition task with trial-unique stimuli (Baylis and Rolls, 1987). In this task, there is typically a presentation of a novel stimulus, and after a delay which may be in the order of minutes or more and in which other stimuli may be shown, the stimulus is presented again as "familiar", and the monkey can respond to obtain food reward. Most neurons responded more to the "novel" than to the "familiar" presentation of a stimulus, where "familiar" in this task reflects a change produced by seeing the stimulus typically once (or a few times) before. (A small proportion of neurons respond more to the familiar (second) than to the novel (first) presentation of each visual stimulus.) In the inferior temporal cortex this memory spanned up to 1–2 intervening stimuli between the first (novel) and second (familiar) presentations of a given stimulus (Baylis and Rolls, 1987), and as recordings are made more ventrally, toward and within the perirhinal cortex, the memory span increases to several or more intervening stimuli (Brown and Xiang, 1998; Wilson et al., 1990; Xiang and Brown, 1998), allowing the perirhinal cortex to contribute to recognition memory over larger memory spans.

In a similar task though typically performed with non-trial-unique stimuli, a delayed matching-to-sample task with up to several intervening stimuli, some perirhinal cortex neurons respond more to the match stimulus than to the sample stimulus (Miller et al., 1998). This has been modeled by interaction with a prefrontal cortex attractor network that holds the first item in a short term memory, and then shows enhanced responding when that item is repeated (Renart et al., 2001; Rolls and Deco, 2002). Many other neurons in this task respond more to the sample ("novel") than to the match ("familiar") presentations of the stimuli, and this short term memory is reset at the start of the next trial (Hölscher and Rolls, 2002). The resetting at the start of each trial shows that the perirhinal cortex in actively involved in the task demands.

A third type of memory in which the perirhinal cortex is implicated is paired associate learning (a model of semantic long term memory), which is represented by a population of neurons in a restricted part of area 36 where the neuronal responses may occur after learning the paired association to both members of a pair of pictures separated by a delay used in the paired association task (Hirabayashi et al., 2013; Miyashita et al., 1998, 1996). The learning of this task may be facilitated by the short-term memory-related firing that follows the first member of the pair during the delay period (Hirabayashi et al., 2013).

Evidence that the perirhinal cortex is involved in a fourth type of memory, long-term familiarity memory, comes from a neuronal recording study in which it was shown that perirhinal cortex neuronal responses in the rhesus macaque gradually increase in magnitude to a set of stimuli as that set is repeated for 400 presentations each 1.3 s long (Hölscher et al., 2003b). The single neurons were recorded in the perirhinal cortex in monkeys performing a delayed matching-to-sample task with up to 3 intervening stimuli, using a set of very familiar visual stimuli used for several weeks. When a novel set of stimuli was introduced, the neuronal responses were on average only 47% of the magnitude of the responses to the familiar set of stimuli. It was shown that the responses of the perirhinal cortex neurons gradually increased over hundreds of presentations of the new set of (initially novel) stimuli to become as large as to the already familiar stimuli. The mean number of 1.3 s presentations to induce this effect was 400 occurring over 7–13 days. These results show that perirhinal cortex neurons represent the very long-term familiarity of visual stimuli. A representation of the long-term familiarity of visual stimuli may be important for many aspects of social and other behavior, and part of the impairment in temporal lobe amnesia may be related to the difficulty of building representations of the degree of familiarity of stimuli. It has been shown that long-term familiarity memory can be modeled by neurons with synapses with a small amount of long-term potentiation which occurs at every presentation of a given stimulus, but which also incorporate hetero-synaptic long-term depression to make the learning self-limiting (Rolls et al., 2005a).

The perirhinal cortex thus has visual representations of objects, and some at least of the perirhinal cortex neurons have view-invariant representations. The perirhinal cortex, via it projections to the lateral entorhinal cortex and lateral perforant path, could thus introduce object information into the hippocampus for object–place associations (see Fig. 1). The parahippocampal gyrus, in which cells with spatial information are present in macaques (Rolls and Xiang, 2006; Rolls et al., 2005b) could via connections to the medial entorhinal cortex introduce spatial information into the hippocampus (see Fig. 1), and there is evidence consistent with this in rodents (Hargreaves et al., 2005). In contrast to the perirhinal cortex, the final convergence afforded by the hippocampus into one network in CA3 (see Fig. 1) may enable the hippocampus proper to implement an event or episodic memory typically involving arbitrary associations between any of the inputs to the hippocampus, e.g. spatial, visual object, olfactory, and auditory (see below). In contrast to the hippocampus, the functions of the perirhinal cortex do not involve associations to spatial representations, and indeed many of the memory-related functions of the perirhinal cortex involve unimodal visual representations.

2.2. The operation of hippocampal circuitry as a memory system

Given the systems level hypothesis about what the hippocampus performs, and the neurophysiological evidence about what is represented in the primate hippocampus, the next step is to consider how using its internal connectivity and synaptic modifiability the hippocampus could store and retrieve many memories, and how retrieval within the hippocampus could lead to retrieval of the activity in the neocortex that was present during the original learning of the episode. To develop understanding of how this is achieved, we have developed a computational theory of the operation of the hippocampus (Rolls, 1987, 1989a,b,d, 1990a,b, 2010b, 2013b,c; Rolls and Kesner, 2006; Rolls et al., 2006; Treves and Rolls, 1991, 1992, 1994). This theory, and new developments of it, are outlined next. Some background is that many of the synapses in the hippocampus show associative modification as shown by long-term potentiation, and there is evidence that this synaptic modification is involved in learning (Andersen et al., 2007; Takeuchi et al., 2014; Wang and Morris, 2010).

2.2.1. Hippocampal circuitry (see Fig. 1 and 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 (DG), via the perforant path (pp) (Witter, 1993). In the dentate gyrus there is a set of excitatory interneurons in the hilus that interconnect granule cells and a layer of inhibitory interneurons that provide recurrent inhibition (Lisman et al., 2005; Myers and Scharfman, 2011; Scharfman and Myers, 2012), but we do not at this stage incorporate this into our computational model, as more quantitative information on the connectivity of these neurons, including the numbers of such neurons, and the numbers of synaptic connections onto each type of neuron from different sources, would be helpful. The granule cells project to CA3 cells via the mossy fibers

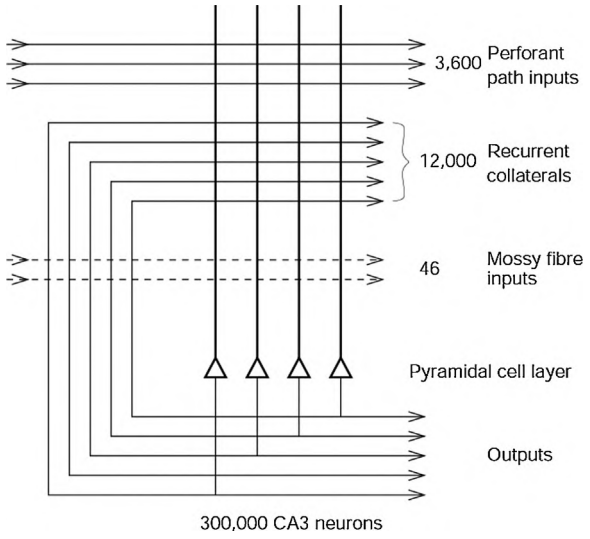

Fig. 3. The numbers of connections from three different sources onto each CA3 cell from three different sources in the rat.
After Rolls and Treves (1998), Treves and Rolls (1992).

(mf), which provide a *sparse* (i.e. diluted) 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. 3). 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 rat 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 rodents and probably in primates, the CA3 systems in the two hemispheres are connected by recurrent collaterals to form essential a single CA3 network. 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).

2.2.2. 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 the sparse yet efficient (i.e. non-redundant) representation in CA3 neurons that is required for the autoassociation performed in CA3 to perform well (Rolls, 1989a,b,d, 1996b, 2008a, 2013b,c; Rolls and Kesner, 2006; Treves and Rolls, 1992). An important property for episodic memory is that the dentate by acting in this way would perform pattern separation (or orthogonalization) (Rolls, 1989b, 2008a, 2013b,c; Rolls and Kesner, 2006; Rolls et al., 2006; Treves and Rolls, 1992), enabling the hippocampus to store different memories of even similar events, and this prediction has been confirmed as described in Section 3. By pattern separation we 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 (Rolls, 2008a) to remove redundancy from the inputs producing a more orthogonal, sparse, and categorized set of outputs (Rolls, 1987, 1989a,b,d, 1990a,b). The non-linearity in the NMDA receptors may help the operation of such a competitive net, for it ensures that only the most active neurons left after the competitive feedback inhibition have synapses that become modified and thus learn to respond to that input (Rolls, 1989a). We note that if the synaptic modification produced in the dentate granule cells lasts for a period of more than the duration of learning the episodic memory, then it could reflect the formation of codes for regularly occurring combinations of active inputs that might need to participate in different episodic memories. We note on the other hand that even with a single exposure to a new input, as might occur in episodic memory, a competitive network can usefully separate representations (cf. Hasselmo and Wyble, 1997). Because of the non-linearity in the NMDA receptors, the non-linearity of the competitive interactions between the neurons (produced by feedback inhibition and non-linearity in the activation function of the neurons), need not be so great (Rolls, 1989a, 2008a). Because of the feedback inhibition, the competitive process may result in a relatively constant number of strongly active dentate neurons relatively independently of the number of active perforant path inputs to the dentate cells. The operation of the dentate granule cell system as a competitive network may also be facilitated by a Hebb rule of the form:

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

where k is a constant, r_i is the activation of the dendrite (the post-synaptic term), r'_j is the presynaptic firing rate, w_{ij} is the synaptic weight, and r'_j and w_{ij} are in appropriate units (Rolls, 1989a, 2008a). 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; see Rolls, 1989a; Rolls and Deco, 2002; 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 (Hafting et al., 2005). Each grid cell responds to a set of places in a spatial environment, with the places to which a cell responds set out in a regular grid. Different grid cells have different phases (positional offsets) and grid spacings (or frequencies) (Hafting et al., 2005). We (Rolls et al., 2006) have simulated the dentate granule cells as a system that receives as inputs the activity of a population of entorhinal cortex grid cells as the animal traverses a spatial environment, and have shown that the competitive net builds dentate-like place cells from such entorhinal grid cell inputs (see Fig. 4). 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 action of 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). As noted above, the same competitive learning could it is suggested (Rolls, 2013b) be involved in the conversion of primate entorhinal spatial view grid cells (Killian et al., 2012) into primate spatial view cells. The presence of spatial view grid cells in the entorhinal cortex of primates (Killian et al., 2012) is of course predicted from the presence of spatial view cells in the primate CA3 and CA1 regions

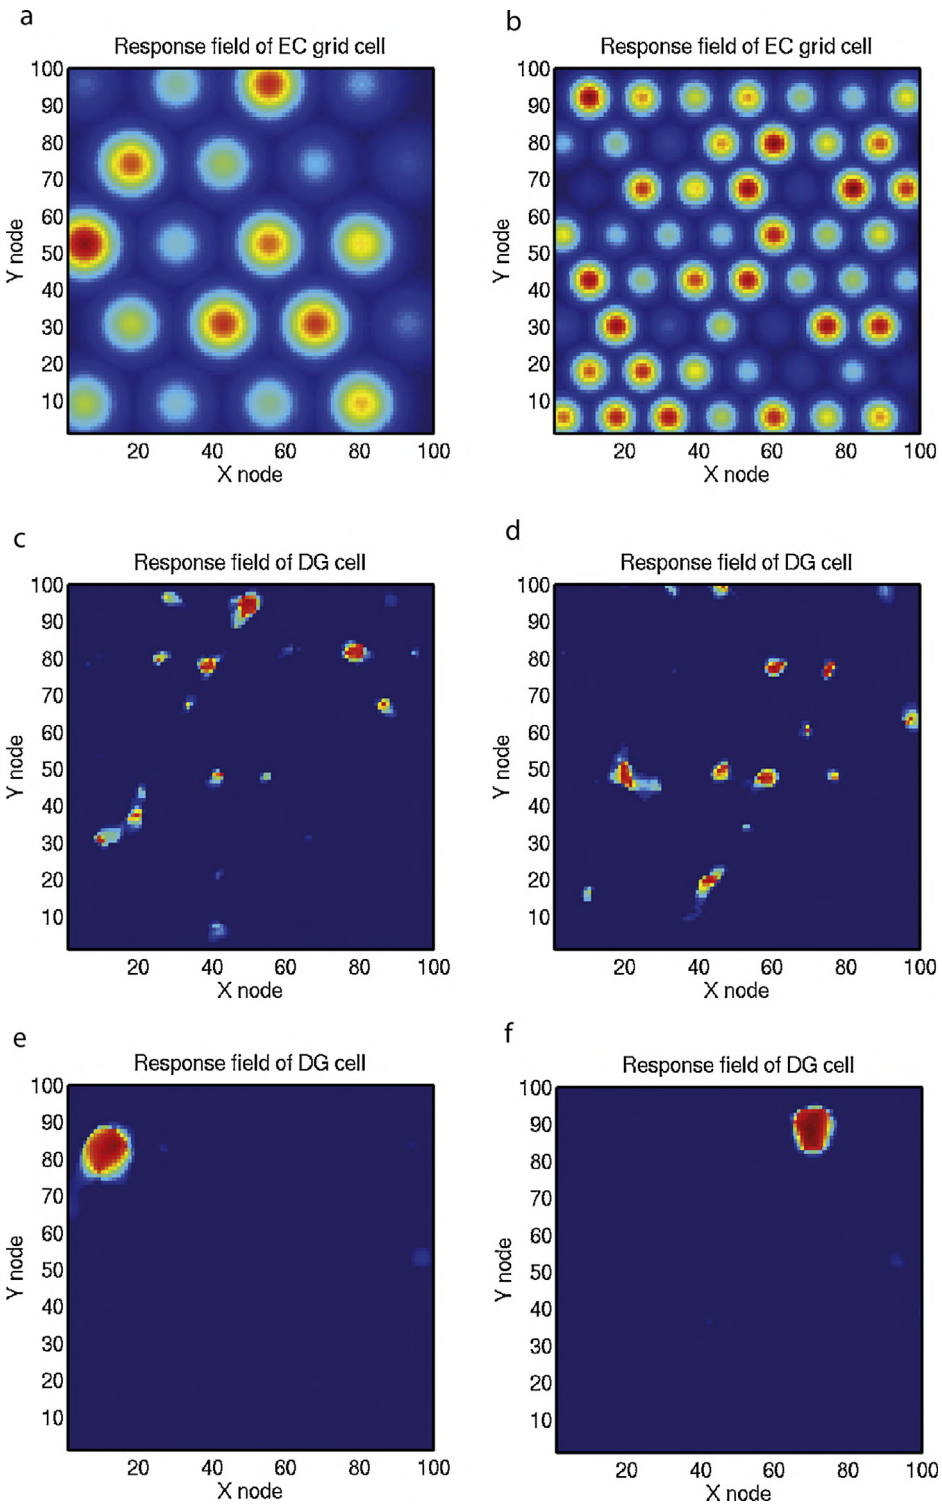

Fig. 4. Simulation of competitive learning in the dentate gyrus to produce place cells from the entorhinal cortex grid cell inputs. (a and b) Firing rate profiles of two entorhinal cortex (EC) grid cells with frequencies of 4 and 7 cycles. (c and d) Firing rate profiles of two dentate gyrus (DG) cells with no training using competitive learning. (e and f) Firing rate profiles of two dentate gyrus (DG) cells trained using competitive learning.
After Rolls et al. (2006).

(Georges-François et al., 1999; Robertson et al., 1998; Rolls, 2008a; Rolls et al., 1997a, 1998; Rolls and Xiang, 2006). 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 grid-like properties have now been described (Jacobs et al., 2013).

In primates, spatial view cells represent a scene view allocentrically, as described in Section 2.1.3. How could such spatial view representations be formed, in which the relative spatial position of features in a scene is encoded? Rolls and colleagues 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 et al., 2008). Rolls and colleagues have shown that in complex natural scenes the receptive fields of inferior temporal cortex neurons become reduced in size and asymmetric with respect to the fovea so that they represent information not only about the object being viewed but also about its position relative to the fovea (Aggelopoulos et al., 2005; Rolls, 2008a, 2012b). Building on this discovery, 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 which have some spatial asymmetry into spatial view cells in the hippocampus.

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, which is what is needed for the formation of odor, object and spatial sequence memories in CA1 (see Section 3), 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, 1989a,b,d, 1990a,b, 1994). It is proposed that this allows overlapping (or very similar) inputs to the hippocampus to be separated, in the following way (see also Rolls, 1996b). Consider three patterns B, W and BW where BW is a linear combination of B and W. (To make the example very concrete, we could consider binary patterns where B = 10, W = 01 and BW = 11.) Then the memory system is required to associate B with reward, W with reward, but BW with punishment. Without the hippocampus, rats might have more difficulty in solving such problems, particularly when they are spatial, for the dentate/CA3 system in rodents is characterized by being implicated in spatial memory. However, it is a property of competitive neuronal networks that they can separate such overlapping patterns, as has been shown elsewhere (Rolls, 1989a; Rolls and Treves, 1998; Rolls, 2008a); normalization of synaptic weight vectors is required for this property. It is thus an important part of hippocampal neuronal network architecture that there is a competitive network that precedes the CA3 autoassociation system. Without the dentate gyrus, if a conventional autoassociation network were presented with the mixture BW having learned B and W separately, then the autoassociation network would produce a mixed output state, and would therefore be incapable of storing separate memories for B, W and BW. It is suggested therefore that 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. Consistent with this, rats with dentate gyrus lesions are impaired at a metric spatial pattern separation task (Gilbert et al., 2001). The recoding of grid cells in the entorhinal cortex (Hafting et al., 2005) into small place field cells in the dentate granule cells that was modeled by 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. We 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). Rolls suggests that application of the configural concept, but applied to spatial pattern separation, may capture part of what the dentate gyrus acting as a competitive network could perform, particularly when a large number of such overlapping spatial memories must be stored and retrieved.

The third way in which the dentate gyrus is hypothesized to contribute to the sparse and relatively orthogonal representations in CA3 arises because of the very low contact probability in the mossy fiber–CA3 connections, and is described in Section 2.2.3.8 and by Treves and Rolls (1992).

A fourth way is that as suggested and explained in Section 2.2.3.11, 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, 2008a). A mechanism like this 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, 2008a, 2013b).

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 (Rolls, 2010b). Consistent with the dentate spatial pattern separation hypothesis (Rolls, 1989b,d, 1996b, 2008a; Treves and Rolls, 1992, 1994), in mice with impaired dentate neurogenesis, spatial learning in a delayed non-matching-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).

The dentate granule cell/mossy fiber system may be one of the few places in the mammalian brain where neurogenesis may occur, because here the proposed computational role of the dentate granule cells is just to select new subsets of CA3 neurons for new learning. That is, the dentate granule cells, according to the theory, do not store the information, which instead is stored by the modification that then occurs in the CA3–CA3 recurrent collateral synapses between the selected CA3 neurons. In contrast, in most other brain areas where memories are stored and learning occurs, the learning is implemented by synaptic modification of the neurons in the circuit, and any loss of existing neurons and replacement by new neurons would result in loss of the learning or memory (Rolls, 2008a).

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°, result in the formation of a spatial view cell, the effective trigger for which will thus be a combination of visual features within a relatively small part of space. In contrast, in the rat, given the very extensive visual field subtended by the rodent retina, which may extend over 180–270°, a combination of visual features formed over such a wide visual angle would effectively define a position in space that is a place. The actual processes by which the hippocampal formation cells would come to respond to feature combinations could be similar in rats and monkeys, involving for example competitive learning in the dentate granule cells. This competitive learning has the required properties, of associating together co-occurrent features, and of ensuring by the competition that each feature combination (representing a place or spatial view) is very different from that of other places or spatial views (pattern separation, an orthogonalizing effect), and is also a sparse representation.

Consistent with this, dentate place fields in rats are small (Jung and McNaughton, 1993; Leutgeb et al., 2007). Although this computation could be performed to some extent also by autoassociation learning in CA3 pyramidal cells, and competitive learning in CA1 pyramidal cells (Rolls and Treves, 1998; Treves and Rolls, 1994), the combined effect of the dentate competitive learning and the mossy fiber low probability but strong synapses to CA3 cells would enable this part of the hippocampal circuitry to be especially important in separating spatial representations, which as noted above are inherently continuous. The prediction thus is that during the learning of spatial tasks in which the spatial discrimination is difficult (for example because the places are close together), then the dentate system should be especially important. Tests of this prediction are described below.

Another interesting hypothesis to account for spatial view cells in primates but place cells in rodents is based on slow learning, a principle used in vision to account for representations that become invariant for objects as compared to views or positions of the objects because learning with a short memory trace emphasizes slowly varying properties (in this case the object while it transforms over view and position), given natural world statistics (Földiák, 1991; Rolls, 2012b; Wallis et al., 1993; Wiskott and Sejnowski, 2002). The hypothesis in the case of place cells is that in rodents the spatial view would change more rapidly than the place of the rat due to head movements by the rat, so neurons would learn place fields by slow learning (Franzius et al., 2007). The hypothesis is that in primates due to the ability to visually fixate and foveate a location in space such as a target even while moving or running toward it, the view may change less rapidly than the place, so spatial view cells would be formed by slow learning (Franzius et al., 2007).

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 place-like fields (though in some cases as a grid (Hafting et al., 2005)) are found in the medial entorhinal cortex (Brun et al., 2002; Fyhn et al., 2004; Moser, 2004; Moser and Moser, 1998), 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 above, the grid-like place cells in the medial entorhinal cortex could if transformed by the competitive net functionality of the dentate cells result in the place cell activity (without a repeating grid) that is found in dentate and rat hippocampal neurons.

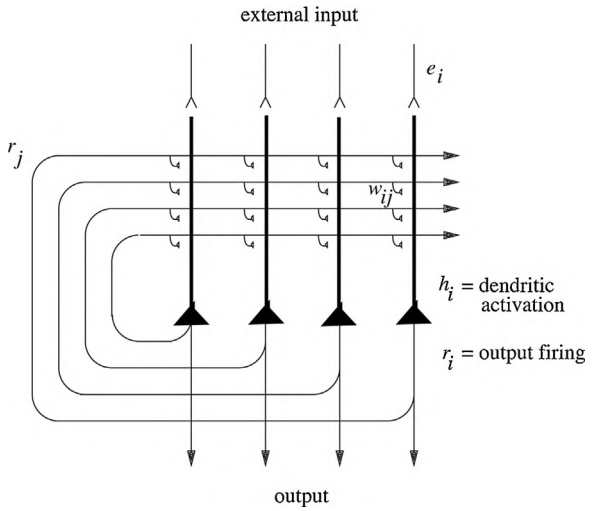

Fig. 5. The architecture of a continuous attractor neural network (CANN). The architecture is the same as that of a discrete attractor neural network.

Computations such as those in the dentate granule cells could also be useful in the formation of spatial view hippocampal representations. In primates, spatial view cells represent a scene view allocentrically, as described in Section 2.1.3. 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 et al., 2008). As noted above, 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.

2.2.3. CA3 as an autoassociation memory

2.2.3.1. Arbitrary associations, and pattern completion in recall. Many of the synapses in the hippocampus show associative modification as shown by long-term potentiation, and this synaptic modification appears to be involved in learning (see Andersen et al., 2007; Jackson, 2013; Lynch, 2004; Morris, 2003, 1989; Morris et al., 2003; Nakazawa et al., 2004, 2003; Wang and Morris, 2010). On the basis of the evidence summarized above, Rolls (1987, 1989a,b,d, 1990a,b, 1991, 2008a) 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 in detail elsewhere (Amit, 1989; Hertz et al., 1991; Rolls, 2010a; Rolls and Deco, 2002, 2010; Rolls and Treves, 1998), including *Memory, Attention, and Decision-Making* (Rolls, 2008a) (with its Appendices available at www.oxcns.org)). The architecture of an autoassociation network is shown in Fig. 5, and the learning rule is as shown in Eq. (1) except that the subtractive

term could be the presynaptic firing rate (Rolls, 2008a; 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. We 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.

The autoassociation or attractor memory is different from pattern association memory, in which a visual stimulus might become associated with a taste by associative synaptic modification. Later presentation of the visual stimulus would retrieve the taste representation. However, presentation of the taste would not retrieve the visual representation, and this is an important and fundamental difference between autoassociation and pattern association, as described in detail elsewhere (Rolls, 2008a, 2014a; Rolls and Treves, 1998). Pattern association memories generalize, and autoassociation memories complete incomplete representations (Rolls, 2008a). Further, competitive networks can generalize, but may not complete perfectly depending on what input is being clamped into the network (Rolls, 2008a).

Tests of the attractor hypothesis of CA3 operation based on the effects of CA3 lesions or CA3 NMDA knockouts on pattern completion are described in Section 3.2. Another type of test of the autoassociation (or attractor) hypothesis for CA3 has been to train rats in different environments, e.g. a square and a circular environment, and then test the prediction of the hypothesis that when presented with an environment ambiguous between these, hippocampal neurons will fall into an attractor state that represents one of the two previously learned environments, but not a mixture of the two environments. Evidence consistent with the hypothesis has been found (Wills et al., 2005). In a particularly dramatic example, it has been found that within each theta cycle, hippocampal pyramidal neurons may represent one or other of the learned environments (Jezek et al., 2011). This was measured by recording in an ambiguous environment, in which it was found that the CA3 representation 'flickered' between two possible environments. This is an indication, predicted by Rolls and Treves (1998) (page 118), that autoassociative memory recall can take place sufficiently rapidly to be complete within one theta cycle (120 ms), and that theta cycles could provide a mechanism for a fresh retrieval process to occur after a reset caused by the inhibitory part of each theta cycle, so that the memory can be updated rapidly to reflect a continuously changing environment, and not remain too long in an attractor state.

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 2.2.3.

2.2.3.1.1. 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 Hopfield, 1982; Amit, 1989) by analyzing a network with graded response units, so as to represent more realistically the continuously variable rates at which neurons fire, and with incomplete connectivity (Treves, 1990; Treves and Rolls, 1991). We have found that in general the maximum number p_{max} of firing patterns that can be (individually) retrieved is proportional to the number C^{RC} of (associatively) modifiable recurrent collateral synapses per cell, by a factor that increases roughly with the inverse of the sparseness a of the neuronal representation.[3] 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

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

where r_i is the firing rate to the ith 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,

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

where k is a factor that depends weakly on the detailed structure of the rate distribution, on the connectivity pattern, etc., but is roughly in the order of 0.2–0.3 (Treves and Rolls, 1991). [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 $C^{RC} = 12{,}000$ 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, in order for most associative networks to store information efficiently, heterosynaptic Long Term Depression (as well as LTP) is required (Fazeli and Collingridge, 1996; Rolls and Deco, 2002; Rolls and Treves, 1990, 1998; Treves and Rolls, 1991). Simulations that are fully consistent with the analytic theory are provided by Rolls (1995, 2012a), Simmen et al. (1996b) and Rolls et al. (1997b).

[3] 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.

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}:

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

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:

$$i_p > a \ \ln\left(\frac{1}{a}\right) \quad (5)$$

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

A number of points deserve comment. First, if it is stated that a certain number of memories is the upper limit of what could be stored in a given network, then the question is sometimes asked, what constitutes a memory? The answer is precise. Any one memory is represented by the firing rates of the population of neurons that are stored by the associative synaptic modification, and can be correctly recalled later. The firing rates in the primate hippocampus might be constant for a period of for example 1 s in which the monkey was looking at an object in one position in space; synaptic modification would occur in this time period (cf. the time course of LTP, which is sufficiently rapid for this); and the memory of the event would have been stored. The quantitative analysis shows how many such random patterns of rates of the neuronal population can be stored and later recalled correctly. If the rates were constant for 5 s while the rat was at one place or monkey was looking at an object at one position in space, then the memory would be for the pattern of firing of the neurons in this 5 s period. (The pattern of firing of the population refers to the rate at which each neuron in the population of CA3 neurons is firing.)

Second, the question sometimes arises of whether the CA3 neurons operate as an attractor network. An attractor network is one in which a stable pattern of firing is maintained once it has been started. Autoassociation networks trained with modified Hebb rules can store the number of different memories, each one expressed as a stable attractor, indicated in Eq. (3). However, the hippocampal CA3 cells do not necessarily have to operate as a stable attractor: instead, it would be sufficient for the present theory if they can retrieve stored information in response to a partial cue initiating retrieval. The partial cue would remain on during recall, so that the attractor network would be operating in the clamped condition (see Rolls and Treves, 1998). The completion of the partial pattern would then provide more information than entered the hippocampus, and the extra information retrieved would help the next stage to operate. Demonstrations of this by simulations that are fully consistent with the analytic theory are provided by Rolls (1995), Simmen et al. (1996b), and Rolls et al. (1997b).

Third, in order for most associative networks to store information efficiently, heterosynaptic Long Term Depression (as well as LTP) is required (Rolls, 1996a; Rolls and Treves, 1990; Treves and Rolls, 1991). Without heterosynaptic LTD, there would otherwise always be a correlation between any set of positively firing inputs acting as the input pattern vector to a neuron. LTD effectively enables the average firing of each input axon to be subtracted from its input at any one time, reducing the average correlation between different pattern vectors to be stored to a low value (Rolls, 1996a).

Fourth, the firing of CA3 cells is relatively sparse, and this helps to decorrelate different population vectors of CA3 cell firing for different memories (Rolls, 2013b,c). (Sparse representations are more likely to be decorrelated with each other (Rolls, 2008a).) Evidence on the sparseness of the CA3 cell representation in rats includes evidence that CA3 cell ensembles may support the fast acquisition of detailed memories by providing a locally continuous, but globally orthogonal spatial representation, onto which new sensory inputs can rapidly be associated (Leutgeb and Leutgeb, 2007). In the macaque hippocampus, in which spatial view cells are found (Georges-François et al., 1999; Robertson et al., 1998; Rolls et al., 1997a, 1998), for the representation of 64 locations around the walls of the room, the mean single cell sparseness a^s was 0.34, and the mean population sparseness a^p was 0.33 (Rolls, 2008a; Rolls and Treves, 2011; Rolls et al., 1998). For comparison, the corresponding values for inferior temporal cortex neurons tuned to objects and faces were 0.77 (Franco et al., 2007; Rolls, 2008a; Rolls and Treves, 2011); for taste and oral texture neurons in the insular cortex the population sparseness was 0.71; for taste and oral texture neurons in the orbitofrontal cortex was 0.61; and for taste and oral texture neurons in the amygdala was 0.81 (Rolls, 2008a; Rolls and Treves, 2011). Thus the evidence is that the hippocampal CA3/pyramidal cell representation is more sparse in macaques than in neocortical areas and the amygdala, and this is consistent with the importance in hippocampal CA3 of using a sparse representation to produce a large memory capacity. Although the value of a for the sparseness of the representation does not seem especially low, it must be remembered that these are graded firing rate representations, and the graded nature appears to increase the value of a (Rolls, 2008a; Rolls and Treves, 2011).

Fifth, given that the memory capacity of the hippocampal CA3 system is limited, it is necessary to have some form of forgetting in this store, or other mechanism to ensure that its capacity is not exceeded. (Exceeding the capacity can lead to a loss of much of the information retrievable from the network.) Heterosynaptic LTD could help this forgetting, by enabling new memories to overwrite old memories (Rolls, 1996a, 2008a). 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.

The potential link to the gradient of retrograde amnesia is that the retrograde memories lost in amnesia are those not yet consolidated in longer term storage (in the neocortex). As they are still held in the hippocampus, their number has to be less than the storage capacity of the (presumed) CA3 autoassociative memory. Therefore the time gradient of the amnesia provides not only a measure of a characteristic time for consolidation, but also an upper bound on the rate of storage of new memories in CA3. For example, if one

were to take as a measure of the time gradient in the monkey, say, 5 weeks (about 50,000 mins, Squire, 1992) and as a reasonable estimate of the capacity of CA3 in the monkey e.g. $p = 50{,}000$, then one would conclude that there is an upper bound on the rate of storage in CA3 of not more than one new memory per minute, on average. (This might be an average over many weeks; the fastest rate might be closer to 1 per sec, see Treves and Rolls, 1994.) These quantitative considerations are consistent with the concept described in Section 2.1 that retrieval of episodic memories from the hippocampus for a considerable time after they have been stored would be useful in helping the neocortex to build a semantic memory (McClelland et al., 1995; Moscovitch et al., 2005; Rolls, 1989b,d; Rolls, 1990a; Treves and Rolls, 1994)

2.2.3.1.2. Recall and completion. A fundamental property of the autoassociation model of the CA3 recurrent collateral network is that the recall can be symmetric, that is, the whole of the memory can be retrieved 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.

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 are described in Section 3.2. Further evidence for pattern completion has been observed using imaging with voltage-sensitive dye in the CA3 region of a rat hippocampal slice. Following the induction of long-term potentiation from two stimulation sites activated simultaneously, stimulation at either of the two sites produced the whole pattern of activation that could be produced from both stimulation sites before LTP, thus demonstrating pattern completion in CA3 (Jackson, 2013).

Especially important in assessing the implications of tests of completion and the other properties of CA3 is that the theory sets out how the system operates when large numbers of memories, in the order of thousands, are to be stored and retrieved, and this is difficult to test adequately in behavioral experiments. Effects found when the storage and retrieval of just a few memories are tested may not reflect well the operation of the system when it is heavily loaded, as it is expected to be when operating in the natural environment.

In a rather different approach, it has been suggested that the CA3–CA3 synapses do not modify during the learning of an event/episodic memory but instead CA3 operates as a fixed sequence memory (Cheng, 2013), but this loses the desirable property of completion by autoassociation in CA3, and makes the system very sensitive to small changes in the firing of dentate neurons and hence which CA3 neurons are activated by the mossy fibers.

2.2.3.1.3. CA3 as a short term attractor memory. Another fundamental property of the autoassociation model of the CA3 recurrent collateral network is that it can implement a short term memory by maintaining the firing of neurons using the excitatory recurrent collateral connections. A stable attractor can maintain one memory active in this way for a considerable period, until a new input pushes the attractor to represent a new location or memory. For example, if one place were being held in a CA3 place short-term memory, then if the rat moved to a new place, the CA3 representation would move to represent the new place, and the short term memory held in the CA3 network would be lost. It is thus predicted that when the hippocampus is used as a short term memory with ongoing neuronal activity used to represent the short term memory, than maintenance of the memory will be very dependent on what happens in the delay period. If the animal is relatively isolated from its environment and does not move around in a well-defined spatial environment (as can be achieved by placing a vertical cylinder/bucket around a rat in the delay period), then the memory may be maintained for a considerable period in the CA3 (and even updated by idiothetic, self-motion, inputs as described below). However, the CA3 short-term memory will be very sensitive to disruption/interference, so that if the rat is allowed to move in the spatial environment during the delay period, then it is predicted that CA3 will not be able to maintain correctly the spatial short term memory. In the circumstances where a representation must be kept active while the hippocampus (or inferior temporal visual cortex or parietal cortex) must be updating its representation to reflect the new changing perceptual inputs, which must be represented in these brain areas for the ongoing events to have high-level representations, the prefrontal cortex is thought computationally to provide an off-line buffer store (Rolls, 2008a). The predictions are thus that if there is no task in a delay period (and no need for an off-line buffer store), the hippocampus may be sufficient to perform the short term memory function (such as remembering a previous location), provided that there is little distraction in the delay period. On the other hand, short-term memory deficits are predicted to be produced by prefrontal cortex lesions if there are intervening stimuli in a delay period in a task.

2.2.3.2. 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 developed by Samsonovich and McNaughton (1997), Battaglia and Treves (1998a), Stringer et al. (2002a, 2002b, 2004), Stringer and Rolls (2002) and Rolls and Stringer (2005) (see Rolls, 2008a), and are described next.

A class of network that can maintain the firing of its neurons to represent any location along a continuous physical dimension such as spatial position, head direction, etc. is a "Continuous Attractor" neural network (CANN). 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 to ensure that there is only one activity packet. Continuous attractor networks can be thought of as very similar to autoassociation or discrete attractor networks (Rolls, 2008a), and have the same architecture, as illustrated in Fig. 5. The main difference is that the patterns stored in a CANN are continuous patterns, with each neuron having broadly tuned firing which decreases with for example a Gaussian function as the distance from the optimal firing location of the cell is varied, and with different neurons having tuning that overlaps throughout the space. Such tuning is illustrated in Fig. 6. For comparison, autoassociation networks normally have discrete

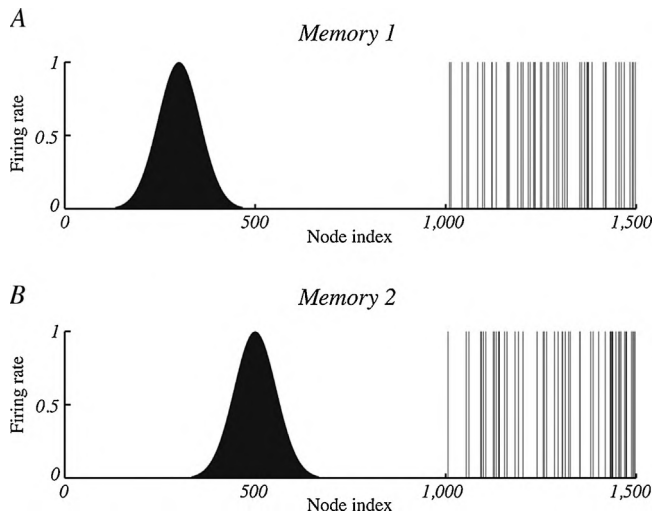

Fig. 6. The types of firing patterns stored in continuous attractor networks are illustrated for the patterns present on neurons 1–1000 for Memory 1 (when the firing is that produced when the spatial state represented is that for location 300), and for Memory 2 (when the firing is that produced when 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.

(separate) patterns (each pattern implemented by the firing of a particular subset of the neurons), with no continuous distribution of the patterns throughout the space (see Fig. 6). A consequent difference is that the CANN can maintain its firing at any location in the trained continuous space, whereas a discrete attractor or autoassociation network moves its population of active neurons toward one of the previously learned attractor states, and thus implements the recall of a particular previously learned pattern from an incomplete or noisy (distorted) version of one of the previously learned patterns. The energy landscape of a discrete attractor network (Rolls, 2008a) has separate energy minima, each one of which corresponds to a learned pattern, whereas the energy landscape of a continuous attractor network is flat, so that the activity packet remains stable with continuous firing wherever it is started in the state space. (The state space refers to the set of possible spatial states of the animal in its environment, e.g. the set of possible places in a room.) Continuous attractor networks are described next, as they are very likely to apply to the operation of systems with spatial representations and recurrent connections, such as the CA3 neurons. Continuous attractor networks have been studied by for example Amari (1977), Zhang (1996), Taylor (1999), and Stringer et al. (2002b).

The generic model of a continuous attractor is as follows. (The model is described in the context of head direction cells, which represent the head direction of rats (Muller et al., 1996; Taube et al., 1996) and macaques (Robertson et al., 1999), and can be reset by visual inputs after gradual drift in darkness.) The model is a recurrent attractor network with global inhibition. It is different from a Hopfield attractor network (Hopfield, 1982) primarily in that there are no discrete attractors formed by associative learning of discrete patterns. Instead there is a set of neurons that 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) represented by the neurons. Neurons that represent similar states (locations in the state space) of the agent in the physical world have strong synaptic connections, which can be set up by an associative learning rule, as described in Section 2.2.3.2. The network updates its firing rates by the following "leaky-integrator" dynamical equations (where head direction is used as a prototypical example, but could be replaced by place, or spatial view). The continuously changing activation h_i^{HD} of each head direction cell i is governed by the equation

$$\frac{dh_i^{\mathrm{HD}}(t)}{dt} = -h_i^{\mathrm{HD}}(t) + \frac{\phi_0}{C^{\mathrm{HD}}} \sum_j (w_{ij} - w^{inh}) r_j^{\mathrm{HD}}(t) + I_i^V, \quad (6)$$

where r_j^{HD} is the firing rate of head direction cell j, w_{ij} is the excitatory (positive) synaptic weight from head direction cell j to cell i, w^{inh} is a global constant describing the effect of inhibitory interneurons, and τ is the time constant of the system.[4] The term $-h_i^{\mathrm{HD}}(t)$ indicates the amount by which the activation decays (in the leaky integrator neuron) at time t. (The network is updated in a typical simulation at much smaller timesteps than the time constant of the system, τ.) The next term in Eq. (6) is the input from other neurons in the network r_j^{HD} weighted by the recurrent collateral synaptic connections w_{ij} (scaled by a constant ϕ_0 and C^{HD} which is the number of synaptic connections received by each head direction cell from other head direction cells in the continuous attractor). The term I_i^V represents a visual input to head direction cell i. Each term I_i^V is set to have a Gaussian response profile in most continuous attractor networks, and this sets the firing of the cells in the continuous attractor to have Gaussian response profiles as a function of where the agent is located in the state space (see e.g. Fig. 6), but the Gaussian assumption is not crucial. (It is known that the firing rates of head direction cells in both rats (Muller et al., 1996; Taube et al., 1996) and macaques (Robertson et al., 1999) is approximately Gaussian.) When the agent is operating without visual input, in memory mode, then the term I_i^V is set to zero. The firing rate r_i^{HD} of cell i is determined from the activation h_i^{HD} and the sigmoid function

$$r_i^{\mathrm{HD}}(t) = \frac{1}{1 + e^{-2\beta(h_i^{\mathrm{HD}}(t) - \alpha)}}, \quad (7)$$

where α and β are the sigmoid threshold and slope, respectively.

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 co-active 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 co-activity 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

[4] Note that here we use r rather than y to refer to the firing rates of the neurons in the network, remembering that, because this is a recurrently connected network (see Fig. 3), the output from a neuron y_i might be the input x_j to another neuron.

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

$$\delta w_{ij} = k r_i^{HD} r_j^{HD} \quad (8)$$

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

$$r_i^{HD} = e^{-s_{HD}^2/2\sigma_{HD}^2}, \quad (9)$$

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. Interestingly if a non-linearity is introduced into the learning rule that mimics the properties of NMDA receptors by allowing the synapses to modify only after strong postsynaptic firing is present, then the synaptic strengths are still close to a Gaussian function of the distance between the connected cells in head direction space. They showed that after training, the continuous attractor network can support stable activity packets in the absence of visual inputs provided that global inhibition is used to prevent all the neurons becoming activated. (The exact stability conditions for such networks have been analyzed by Amari, 1977). Thus Stringer et al. (2002b) demonstrated biologically plausible mechanisms for training the synaptic weights in a continuous attractor using a biologically plausible local learning rule.

2.2.3.3. 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 and discrete patterns, 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). In this network, when events are stored that have both discrete (object) and continuous (spatial) aspects, then the whole place can be retrieved later by the object, and the object can be retrieved by using the place as a retrieval cue (see Fig. 6). The combined continuous and discrete attractor network described by Rolls et al. (2002) shows that in brain regions where the spatial and object processing streams are brought together, then a single network can represent and learn associations between both types of input. Indeed, in brain regions such as the hippocampal system, it is essential that the spatial and object processing streams are brought together in a single network, for it is only when both types of information are in the same network that spatial information can be retrieved from object information, and vice versa, which is a fundamental property of episodic memory (Rolls and Deco, 2002; Rolls and Treves, 1998).

2.2.3.4. 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 capacity of a continuous attractor network can be approached on the following bases. First, as there are no discrete attractor states, but instead a continuous physical space is being represented, some concept of spatial resolution must be brought to bear, that is the number of different positions in the space that can be represented. Second, the number of connections per neuron in the continuous attractor will directly influence the number of different spatial positions (locations in the state space) that can be represented. Third, the sparseness of the representation can be thought of as influencing the number of different spatial locations (in the continuous state space) that can be represented, in a way analogous to that described for discrete attractor networks (Battaglia and Treves, 1998a). That is, if the tuning of the neurons is very broad, then fewer locations in the state space may be represented. Fourth, and very interestingly, if representations of different continuous state spaces, for example maps or charts of different environments, are stored in the same network, there may be little cost of adding extra maps or charts. The reason for this is that a large part of the interference between the different memories stored in such a network arises from the correlations between the different positions in any one map, which are typically relatively high because quite broad tuning of individual cells is common. In contrast, there are in general low correlations between the representations of places in different maps or charts, and thus many different maps can be simultaneously stored in a continuous attractor network (Battaglia and Treves, 1998a).

2.2.3.5. Idiothetic update by path integration. We have considered how spatial representations could be stored in continuous attractor networks, and how the activity can be maintained at any location in the state space in a form of short term memory when the external (e.g. visual) input is removed. However, many networks with spatial representations in the brain can be updated by internal, self-motion (i.e. idiothetic), cues even when there is no external (e.g. visual) input. The way in which path integration could be implemented in recurrent networks such as the CA3 system in the hippocampus or in related systems is described next. Single-cell recording studies have shown that some neurons represent the current position along a continuous physical dimension or space even when no inputs are available, for example in darkness. Examples include neurons that represent the positions of the eyes (i.e. eye direction with respect to the head), the place where the animal is looking in space, head direction, and the place where the animal is located. In particular, examples of such classes of cells include head direction cells in rats (Muller et al., 1996; Ranck, 1985; Taube et al., 1996, 1990) and primates (Robertson et al., 1999), which respond maximally when the animal's head is facing in a particular preferred direction; place cells in rats (Markus et al., 1995; McNaughton et al., 1983; Muller et al., 1991; O'Keefe, 1984; O'Keefe and Dostrovsky, 1971) that fire maximally when the animal is in a particular location; and spatial view cells in primates that respond when the monkey is looking toward a particular location in space (Georges-François et al., 1999; Robertson et al., 1998; Rolls et al., 1997a).

One approach to simulating the movement of an activity packet produced by idiothetic cues (which is a form of path integration whereby the current location is calculated from recent movements) is to employ a look-up table that stores (taking head direction cells as an example), for every possible head direction and head rotational velocity input generated by the vestibular system, the corresponding new head direction (Samsonovich and McNaughton, 1997). Another approach involves modulating the strengths of the recurrent synaptic weights in the continuous attractor on one but not the other side of a currently represented position, so that the stable position of the packet of activity, which requires symmetric connections in different directions from each node, is lost, and the packet moves in the direction of the temporarily increased weights, although no possible biological implementation was proposed

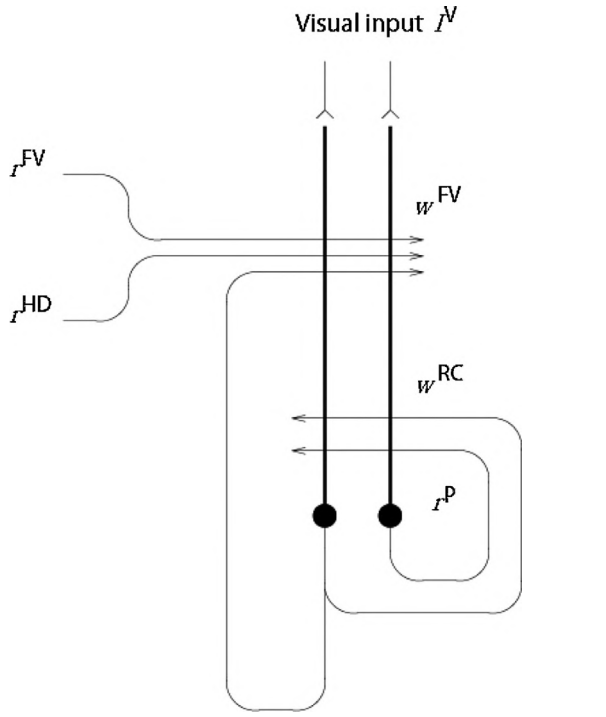

Fig. 7. Neural network architecture for two-dimensional continuous attractor models of place cells. There is a recurrent network of place cells with firing rates r^P,which receives external inputs from three sources: (i) the visual system, (ii) a population of head direction cells with firing rates r^{HD}, and (iii) a population of forward velocity cells with firing rates r^{FV}. The recurrent weights between the place cells are denoted by w^{RC}, and the idiothetic weights to the place cells from the forward velocity cells and head direction cells are denoted by w^{FV}.

of how the appropriate dynamic synaptic weight changes might be achieved (Zhang, 1996). Another mechanism (for head direction cells) (Skaggs et al., 1995) relies on a set of cells, termed (head) rotation cells, which are co-activated by head direction cells and vestibular cells and drive the activity of the attractor network by anatomically distinct connections for clockwise and counter-clockwise rotation cells, in what is effectively a look-up table. However, these proposals did not show how the synaptic weights for this path integration could be achieved by a biologically plausible learning process.

Stringer et al. (2002b) introduced a proposal with more biological plausibility about how the synaptic connections from idiothetic inputs to a continuous attractor network can be learned by a self-organizing learning process. The mechanism associates a short term memory trace of the firing of the neurons in the attractor network reflecting recent movements in the state space (e.g. of places), with an idiothetic velocity of movement input (see Fig. 7). This has been applied to head direction cells (Stringer et al., 2002b), rat place cells (Stringer et al., 2002a, 2002b), and primate spatial view cells (Rolls and Stringer, 2005; Stringer et al., 2004, 2005). These attractor networks provide a basis for understanding cognitive maps, and how they are updated by learning and by self-motion. The implication is that to the extent that path integration of place or spatial view representations is performed within the hippocampus itself, then the CA3 system is the most likely part of the hippocampus to be involved in this, because it has the appropriate recurrent collateral connections. Consistent with this, Whishaw and colleagues (Maaswinkel et al., 1999; Wallace and Whishaw, 2003; Whishaw et al., 2001) have shown that path integration is impaired by hippocampal lesions. Path integration of head direction is reflected in the firing of neurons in the presubiculum, and mechanisms outside the hippocampus probably implement path integration for head direction.

Cognitive maps (O'Keefe and Nadel, 1978) can be understood by the operations of attractor networks, and how they are updated by learning and by self-motion (Rolls, 2008a). It has been argued that the bumpiness of the CA3 representation of space is more consistent with episodic memory storage, as argued in this paper, than with spatial path integration using the CA3 system as a continuous attractor network implementing path integration (Cerasti and Treves, 2013; Stella et al., 2013).

Indeed, the update of spatial representations by idiothetic inputs need not occur within the hippocampus proper, and consistent with this, idiothetic update occurs in the entorhinal cortex grid cell system (Giocomo et al., 2011; Kropff and Treves, 2008; Zilli, 2012).

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

It has been shown that if the neurons are treated not as McCulloch-Pitts neurons which are simply "updated" at each iteration, or 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 and Deco, 2002; 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 brain-like 'integrate-and-fire' membrane and synaptic properties is that with some spontaneous activity, some of the neurons in the network are close to threshold already before the recall cue is applied, and hence some of the neurons are very quickly pushed by the recall cue into firing, so that information starts to be exchanged very rapidly (within 1–2 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 in Section 7.7 of Rolls and Deco (2002), in Appendix A5 of Rolls and Treves (1998), by Treves (1993), by Panzeri et al. (2001), and by Rolls (2008a).

2.2.3.7. Memory for sequences. Evidence from rodents that the hippocampus is involved in the memory for spatial sequences is described in Section 3.3.1. Ways in which the recurrent collateral connections in CA3 might contribute to this are described in this section. However, we note that CA1 may be involved in the memory for sequences of objects and odors, and consider in Section 2.2.4.4 how this may be implemented.

One of the first extensions of the standard autoassociator paradigm that has been explored in the literature is the capability to store and retrieve not just individual patterns, but whole sequences of patterns (Kohonen, 1977; Kohonen et al., 1981; Willshaw, 1981). Hopfield (1982) suggested that this could be achieved by adding to the standard connection weights, which associate a pattern with itself, a new, asymmetric component, that associates a pattern with the next one in the sequence. In practice this scheme does not work very well, unless the new component is made to operate on a slower time scale than the purely autoassociative component (Kleinfeld, 1986; Sompolinsky and Kanter, 1986). With two different time scales, the autoassociative component can stabilize a pattern for a while, before the heteroassociative component moves the network, as it were, into the next pattern. The heteroassociative retrieval cue for the next pattern in the sequence is just the previous pattern in the sequence. A particular type of "slower" operation occurs if the asymmetric component acts after a delay τ. In this case, the network sweeps through the sequence, staying for a time of order τ in each pattern.

One can see how the necessary ingredient for the storage of sequences is only a minor departure from purely Hebbian learning: in fact, the (symmetric) autoassociative component of the weights can be taken to reflect the Hebbian learning of strictly simultaneous conjunctions of pre- and post-synaptic activity, whereas the (asymmetric) heteroassociative component can be implemented by Hebbian learning of each conjunction of postsynaptic activity with presynaptic activity shifted a time τ in the past. Both components can then be seen as resulting from a generalized Hebbian rule, which increases the weight whenever postsynaptic activity is paired with presynaptic activity occurring within a given time range, which may extend from a few hundred milliseconds in the past up to include strictly simultaneous activity. This is similar to a temporal trace rule (Rolls, 2008a; Rolls and Deco, 2002), which itself matches very well the observed conditions for induction of Long-Term Potentiation, and appears entirely plausible. The learning rule necessary for learning sequences, though, is more complex than a simple trace rule in that the time-shifted conjunctions of activity that are encoded in the weights must in retrieval produce activations that are time-shifted as well (otherwise one falls back into the Hopfield (1982) proposal, which does not quite work). The synaptic weights should therefore keep separate "traces" of what was simultaneous and what was time-shifted during the original experience, and this is not very plausible. Levy and colleagues (Levy et al., 1995; Wu et al., 1996) have investigated these issues further, and the temporal asymmetry that may be present in LTP has been suggested as a mechanism that might provide some of the temporal properties that are necessary for the brain to store and recall sequences (Abbott and Blum, 1996; Abbott and Nelson, 2000; Markram et al., 1998; Minai and Levy, 1993). A problem with this suggestion is that, given that the temporal dynamics of attractor networks are inherently very fast when the networks have continuous dynamics, and that the temporal asymmetry in LTP may be in the order of only milliseconds to a few tens of milliseconds, the recall of the sequences would be very fast, perhaps 10–20 ms per step of the sequence, with every step of a 10-step sequence effectively retrieved and gone in a quick-fire session of 100–200 ms.

Rolls and Stringer have suggested that the over-rapid replay of a sequence of memories stored in an autoassociation network such as CA3 if it included asymmetric synaptic weights to encode a sequence could be controlled by the physical inputs from the environment (Rolls, 2008a; Rolls and Kesner, 2006). If a sequence of places 1, 2, and 3 had been learned by the use of an asymmetric trace learning rule implemented in the CA3 network, then the firing initiated by place 1 would reflect (by the learned association) a small component of place 2 (which might be used to guide navigation to place 2, and which might be separated out from place 1 better by the competitive network action in CA1), but would not move fully away from representing place 1 until the animal moved away from place 1, because of the clamping effect on the CA3 firing of the external input representing place 1 to the recurrent network. In this way, the physical constraints of movements between the different places in the environment would control the speed of readout from the sequence memory. A simulation to illustrate this is shown in Rolls and Kesner (2006). The CA1 neurons might thus make a useful contribution to this sequential recall by acting as a competitive network that would separate out the patterns represented in CA3 for the current and the next place (Rolls and Kesner, 2006). We note that if an episodic memory with a spatial sequence is thought to consist of a temporally linked sequence of event memories, then the mechanism for sequence memory just described could implement this type of multiple sequential item episodic memory.

Lisman and colleagues (Jensen et al., 1996; Jensen and Lisman, 1996, 2005; Lisman and Redish, 2009; Lisman and Idiart, 1995; Lisman et al., 2005) have suggested that sequences might be encoded by utilizing the neuronal after-depolarization which follows a spike from a neuron by approximately 200 ms to tend to make neurons fire again after they are activated. Interaction with theta oscillations (5–8 Hz) and gamma oscillations (30–80 Hz) in the model would make a set of neurons coding for one item tend to fire first in a new theta cycle. Neurons encoding a second item in a sequence would tend to fire in the next gamma cycle 20 ms later. Neurons encoding a third item in a sequence would tend to fire in the next gamma cycle 20 ms later. To keep the neurons representing one item synchronized, LTP associated with NMDA receptors with a short time constant of e.g. 12 ms is needed. Every theta cycle, the whole sequence is replayed by the model with 20 ms between items. In one version of the model (Jensen and Lisman, 2005) this is suggested to be implemented in the neocortex. To form associations between successive items in a sequence, LTP associated with NMDA receptors with a long time constant of e.g. 150 ms is needed, and this it is suggested might be implemented in the hippocampus (Jensen and Lisman, 2005; Lisman et al., 2005). It is a difficulty of the model that every theta cycle, the whole sequence is replayed by the model with 20 ms between items, for if the next item in a sequence is needed (e.g. the third), then the model just replays at this high speed all the following items in a sequence, so a special decoding system would be needed to know which the next item is. The model is also very susceptible to noise. The model also only works if there is great regularity in the spike trains of the neurons, which should repeat at very regular intervals locked to theta, and this is inconsistent with the almost Poisson nature of the spike trains of single neurons found in most brain areas. Further, no experimental evidence has been found that successive items in a remembered sequence are replayed at gamma frequency every theta cycle.

Another way in which a delay could be inserted in a recurrent collateral path in the brain is by inserting another cortical area in the recurrent path. This could fit in with the cortico-cortical backprojection connections described in Fig. 1 with return projections to the entorhinal cortex, which might then send information back into the hippocampus, which would introduce some conduction delay (see Panzeri et al., 2001). Another possibility is that the connections between the deep and the superficial layers of the entorhinal cortex help to close a loop from entorhinal cortex through hippocampal circuitry and back to the entorhinal cortex (van Haeften et al., 2003). Another suggestion is that the CA3 system holds by continuing firing a memory for event 1, which can then be associated with event 2 at the CA3 to CA1 synapses if event 2 activates CA1 neurons by the direct entorhinal input (Levy, 1989).

Another type of sequence memory uses synaptic adaptation to effectively encode the order of the items in a sequence (Deco and Rolls, 2005c). Whenever the system is quenched into inactivity, the next member of the sequence emerges out of the spontaneous

activity, because the least recently activated member of the sequence has the least synaptic or neuronal adaptation. This could be implemented in recurrent networks such as the CA3 or the prefrontal cortex.

For recency short term memory, the sequence type of memory would be possible, but difficult because once a sequence has been recalled, a system has to "manipulate" the items, picking out which is later (or earlier) in the sequence. A simpler possibility is to use a short term form of long term potentiation (LTP) which decays, and then the latest of two items can be picked out because when it is seen again, the neuronal response is larger than to the earlier items (Renart et al., 2001; Rolls and Deco, 2002). This type of memory does not explicitly encode sequences, but instead reflects just how recently an item occurred. This could be implemented at any synapses in the system (e.g. in CA3 or CA1), and does not require recurrent collateral connectivity.

2.2.3.8. The dilution of the CA3 recurrent collateral connectivity enhances memory storage capacity and pattern completion. Fig. 3 shows that in the rat, there are approximately 300,000 CA3 neurons, but only 12,000 recurrent collateral synapses per neuron. The dilution of the connectivity is thus 12,000/300,000 = 0.04. The connectivity is thus not complete, and complete connectivity in an autoassociation network would make it simple, for the connectivity between the neurons would then be symmetric (i.e. the connection strength from any one neuron to another is matched by a connection of the same strength in the opposite direction), and this guarantees energy minima for the basins of attraction that will be stable, and a memory capacity than can be calculated (Hopfield, 1982). We have shown how this attractor type of network can be extended to have similar properties with diluted connectivity, and also with sparse representations with graded firing rates (Rolls and Treves, 1990; Rolls and Webb, 2012; Treves, 1990, 1991; Treves and Rolls, 1991; Webb et al., 2011).

However, the question has recently been asked about whether there are any advantages to diluted autoassociation or attractor networks compared to fully connected attractor networks (Rolls, 2012a). One biological property that may be a limiting factor is the number of synaptic connections per neuron, which is 12,000 in the CA3–CA3 network just for the recurrent collaterals (see Fig. 3). The number may be higher in humans, allowing more memories to be stored in the hippocampus than order 12,000. I note that the storage of large number of memories may be facilitated in humans because the left and right hippocampus appear to be much less connected between the two hemispheres than in the rat, which effectively has a single hippocampus (Rolls, 2008a). In humans, with effectively two separate CA3 networks, one on each side of the brain, the memory storage capacity may be doubled, as the capacity is set by the number of recurrent collaterals per neuron in each attractor network (Eq. (3)). In humans, the right hippocampus may be devoted to episodic memories with spatial and visual components, whereas the left hippocampus may be devoted to memories with verbal/linguistic components, i.e. in which words may be part of the episode (e.g. who said what to whom and when) (Barkas et al., 2010; Bonelli et al., 2010; Sidhu et al., 2013).

The answer that has been suggested to why the connectivity of the CA3 autoassociation network is diluted (and why neocortical recurrent networks are also diluted), is that this may help to reduce the probability of having two or more synapses between any pair of randomly connected neurons within the network, which it has been shown greatly impairs the number of memories that can be stored in an attractor network, because of the distortion that this produces in the energy landscape (Rolls, 2012a). In more detail, the hypothesis proposed is that the diluted connectivity allows biological processes that set up synaptic connections between neurons to arrange for there to be only very rarely more than one synaptic connection between any pair of neurons. If probabilistically there were more than one connection between any two neurons, it was shown by simulation of an autoassociation attractor network that such connections would dominate the attractor states into which the network could enter and be stable, thus strongly reducing the memory capacity of the network (the number of memories that can be stored and correctly retrieved), below the normal large capacity for diluted connectivity. Diluted connectivity between neurons in the cortex thus has an important role in allowing high capacity of memory networks in the cortex, and helping to ensure that the critical capacity is not reached at which overloading occurs leading to an impairment in the ability to retrieve any memories from the network (Rolls, 2012a). The diluted connectivity is thus seen as an adaptation that simplifies the genetic specification of the wiring of the brain, by enabling just two attributes of the connectivity to be specified (e.g. from a CA3 to another CA3 neuron chosen at random to specify the CA3 to CA3 recurrent collateral connectivity), rather than which particular neuron should connect to which other particular neuron (Rolls, 2012a; Rolls and Stringer, 2000). Consistent with this hypothesis, there are NMDA receptors with the genetic specification that they are NMDA receptors on neurons of a particular type, CA3 neurons (as shown by the evidence from CA3-specific vs CA1-specific NMDA receptor knockouts) (Nakazawa et al., 2004, 2002, 2003; Rondi-Reig et al., 2001). A consequence is that the vector of output neuronal firing in the CA3 region, i.e. the number of CA3 neurons, is quite large (300,000 neurons in the rat). The large number of elements in this vector may have consequences for the noise in the system, as we will see below.

The role of dilution in the connectivity of the CA3 recurrent collateral connectivity includes enabling this large number of separate memories to be recalled from any part of each memory, that is, in pattern completion (Rolls, 2012a).

The dilution of the CA3–CA3 recurrent collateral connectivity at 0.04 may be greater dilution than that in a local neocortical area, which is in the order of 0.1 (Rolls, 2008a, 2012a). This is consistent with the hypothesis that the storage capacity of the CA3 system is at a premium, and so the dilution is kept to a low value (i.e. great dilution), as then there is lower distortion of the basins of attraction and hence the memory capacity is maximized (Rolls, 2012a).

2.2.3.9. Noise and stability produced by the diluted connectivity and the graded firing rates in the CA3–CA3 attractor network. Many processes in the brain are influenced by the noise or variability of neuronal spike firing (Deco et al., 2013; Faisal et al., 2008; Rolls and Deco, 2010). The action potentials are generated in a way that frequently approximates a Poisson process, in which the spikes for a given mean firing rate occur at times that are essentially random (apart from a small effect of the refractory period), with a coefficient of variation of the interspike interval distribution (CV) near 1.0 (Rolls and Deco, 2010). The sources of the noise include quantal transmitter release, and noise in ion channel openings (Faisal et al., 2008). The membrane potential is often held close to the firing threshold, and then small changes in the inputs and the noise in the neuronal operations cause spikes to be emitted at almost random times for a given mean firing rate. Spiking neuronal networks with balanced inhibition and excitation currents and associatively modified recurrent synaptic connections can be shown to possess a stable attractor state where neuron spiking is approximately Poisson too (Amit and Brunel, 1997; Miller and Wang, 2006). The noise caused by the variability of individual neuron spiking which then affects other neurons in the network can play an important role in the function of such recurrent attractor networks, by causing for example an otherwise stable network to jump into a decision state (Deco and Rolls, 2006; Rolls and Deco, 2010; Wang, 2002). Attractor networks with this type of spiking-related noise are used in the brain for memory recall, and for decision-making, which in

terms of the neural mechanism are effectively the same process (Rolls, 2008a). Noise in attractor networks is useful for memory and decision-making, for it makes them non-deterministic, and this contributes to new solutions to problems, and indeed to creativity (Rolls, 2014a; Rolls and Deco, 2010).

To investigate the extent to which the diluted connectivity affects the dynamics of attractor networks in the cerebral cortex (which includes the hippocampus), we simulated an integrate-and-fire attractor network taking decisions between competing inputs with diluted connectivity of 0.25 or 0.1 but the same number of synaptic connections per neuron for the recurrent collateral synapses within an attractor population as for full connectivity (Rolls and Webb, 2012). The results indicated that there was less spiking-related noise with the diluted connectivity in that the stability of the network when in the spontaneous state of firing increased, and the accuracy of the correct decisions increased. The decision times were a little slower with diluted than with complete connectivity. Given that the capacity of the network is set by the number of recurrent collateral synaptic connections per neuron, on which there is a biological limit, the findings indicate that the stability of cortical networks, and the accuracy of their correct decisions or memory recall operations, can be increased by utilizing diluted connectivity and correspondingly increasing the number of neurons in the network (which may help to smooth the noise), with little impact on the speed of processing of the cortex. Thus diluted connectivity can decrease cortical spiking-related noise, and thus enhance the reliability of memory recall, which includes completion from a partial recall cue (Rolls and Webb, 2012).

Representations in the neocortex and in the hippocampus are often distributed with graded firing rates in the neuronal populations (Rolls and Treves, 2011). The firing rate probability distribution of each neuron to a set of stimuli is often exponential or gamma (Rolls and Treves, 2011). These graded firing rate distributed representations are present in the hippocampus, both for place cells in rodents and for spatial view cells in the primate (Georges-François et al., 1999; McNaughton et al., 1983; O'Keefe and Speakman, 1987; O'Keefe, 1979; Robertson et al., 1998; Rolls, 2008a; Rolls et al., 1997a, 1998; Rolls and Treves, 2011). In processes in the brain such as memory recall in the hippocampus or decision-making in the cortex that are influenced by the noise produced by the close to random spike timings of each neuron for a given mean rate, the noise with this graded type of representation may be larger than with the binary firing rate distribution that is usually investigated. In integrate-and-fire simulations of an attractor decision-making network, we showed that the noise is indeed greater for a given sparseness of the representation for graded, exponential, than for binary firing rate distributions (Webb et al., 2011). The greater noise was measured by faster escaping times from the spontaneous firing rate state when the decision cues are applied, and this corresponds to faster decision or reaction times. The greater noise was also evident as less stability of the spontaneous firing state before the decision cues are applied. The implication is that spiking-related noise will continue to be a factor that influences processes such as decision-making, signal detection, short-term memory, and memory recall and completion (including in the CA3 network) even with the quite large networks found in the cerebral cortex. In these networks there are several thousand recurrent collateral synapses onto each neuron. The greater noise with graded firing rate distributions has the advantage that it can increase the speed of operation of cortical circuitry (Webb et al., 2011). The graded firing rates also by operating in a non-linear network effectively increase the sparseness of the representation, and this itself is a pattern separation effect (Webb et al., 2011).

2.2.3.10. Pattern separation of CA3 cell populations encoding different memories. 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, 2008a; Rolls and Treves, 1998), 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).

We have proposed that there are several mechanisms that help to 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. Neurogenesis of dentate granule cells is a fifth potential contributor to achieving pattern separation of CA3 cell firing. The five factors are described in this paper, and elsewhere (Rolls, 2008a, 2013b). It is remarked that some of this architecture may be special to the hippocampus, and not found in the neocortex, because of the importance of storing and retrieving large numbers of (episodic) memories in the hippocampus. The neocortex in contrast is more concerned with building new representations for which competitive learning is more important, and thus neocortical circuitry does not use a mossy fiber system to produce new random sets of neurons activated (Rolls, 2008a).

2.2.3.11. Mossy fiber inputs to the CA3 cells. We hypothesize that the mossy fiber inputs force efficient information storage by virtue of their strong and sparse influence on the CA3 cell firing rates (Rolls, 1987, 1989b,d; Treves and Rolls, 1992). (The strong effects likely to be mediated by the mossy fibers were also emphasized by McNaughton and Morris (1987) and McNaughton and Nadel (1990)). We hypothesize that the mossy fiber input appears to be particularly appropriate in several ways. First of all, the fact that mossy fiber synapses are large and located very close to the soma makes them relatively powerful in activating the postsynaptic cell (this should not be taken to imply that a CA3 cell can be fired by a single mossy fiber EPSP). Second, the firing activity of dentate granule cells appears to be very sparse (Jung and McNaughton, 1993; 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.[5] The hypothesis is that the mossy fiber sparse connectivity solution performs the appropriate function to enable learning to operate correctly in CA3 (Cerasti and Treves,

[5] 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.

2010; Treves and Rolls, 1992). 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). Third, nonassociative plasticity of mossy fibers (see Brown et al., 1989, 1990) might have a useful effect in enhancing the signal-to-noise ratio, in that a consistently firing mossy fiber would produce nonlinearly amplified currents in the postsynaptic cell, which would not happen with an occasionally firing fiber (Treves and Rolls, 1992). This plasticity, and also learning in the dentate, would also have the effect that similar fragments of each episode (e.g. the same environmental location) recurring on subsequent occasions would be more likely to activate the same population of CA3 cells, which would have potential advantages in terms of economy of use of the CA3 cells in different memories, and in making some link between different episodic memories with a common feature, such as the same location in space. Fourth, with only a few, and powerful, active mossy fiber inputs to each CA3 cell, setting a given sparseness of the representation provided by CA3 cells would be simplified, for the EPSPs produced by the mossy fibers would be Poisson distributed with large membrane potential differences for each active mossy fiber. Setting the average firing rate of the dentate granule cells would effectively set the sparseness of the CA3 representation, without great precision being required in the threshold setting of the CA3 cells (Rolls et al., 1997b). Part of what is achieved by the mossy fiber input may be setting the sparseness of the CA3 cells correctly, which, as shown above, is very important in an autoassociative memory store. Fifth, the non-associative and sparse connectivity properties of the mossy fiber connections to CA3 cells may be appropriate for an episodic memory system 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 argument based on information suggests, then, that an input system with the characteristics of the mossy fibers is essential during learning, in that it may act as a sort of (unsupervised) teacher that effectively strongly influences which CA3 neurons fire based on the pattern of granule cell activity. This establishes an information-rich neuronal representation of the episode in the CA3 network (see further Treves and Rolls, 1992). 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).

The particular property of the small number of mossy fiber connections onto a CA3 cell, approximately 46 (see Fig. 3), is that this has a randomizing effect on the representations set up in CA3, so that they are as different as possible from each other (Rolls, 1989b,d, 2008a; Rolls and Kesner, 2006; Rolls and Treves, 1998; Treves and Rolls, 1992). (This means for example that place cells in a given environment are well separated to cover the whole space.) The result is that any one event or episode will set up a representation that is very different from other events or episodes, because the set of CA3 neurons activated for each event is random. This is then the optimal situation for the CA3 recurrent collateral effect to operate, for it can then associate together the random set of neurons that are active for a particular event (for example an object in a particular place), and later recall the whole set from any part. It is because the representations in CA3 are unstructured, or random, in this way that large numbers of memories can be stored in the CA3 autoassociation system, and that interference between the different memories is kept as low as possible, in that they are maximally different from each other (Hopfield, 1982; Rolls, 2008a; Rolls and Treves, 1998; Treves and Rolls, 1991).

The requirement for a small number of mossy fiber connections onto each CA3 neuron applies not only to discrete (Treves and Rolls, 1992) but also to spatial representations, and some learning in these connections, whether associative or not, can help to select out the small number of mossy fibers that may be active at any one time to select a set of random neurons in the CA3 (Cerasti and Treves, 2010). Any learning may help by reducing the accuracy required for a particular number of mossy fiber connections to be specified genetically onto each CA3 neuron. The optimal number of mossy fibers for the best information transfer from dentate granule cells to CA3 cells is in the order of 35–50 (Cerasti and Treves, 2010; Treves and Rolls, 1992). The mossy fibers also make connections useful for feedforward inhibition in CA3 (Acsady et al., 1998), which is likely to be useful to help in the sparse representations being formed in CA3.

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) (see below).

We (Rolls and Kesner, 2006) have hypothesized that non-associative plasticity of mossy fibers (see Brown et al., 1989, 1990) might have a useful effect in enhancing the signal-to-noise ratio, in that a consistently firing mossy fiber would produce nonlinearly amplified currents in the postsynaptic cell, which would not happen with an occasionally firing fiber (Treves and Rolls, 1992). This plasticity, and also learning in the dentate, would also have the effect that similar fragments of each episode (e.g. the same environmental location) recurring on subsequent occasions would be more likely to activate the same population of CA3 cells, which would have potential advantages in terms of economy of use of the CA3 cells in different memories, and in making some link between different episodic memories with a common feature, such as the same location in space. Consistent with this, dentate neurons that fire repeatedly are more effective in activating CA3 neurons (Henze et al., 2002).

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.

2.2.3.12. Perforant path inputs to CA3 cells initiate recall in CA3 and contribute to generalization. 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, $q \ll 1$, with the activity distribution present during learning. In order not to lose this small correlation altogether, but rather transform it into an input current in the CA3 cells that carries a sizable signal – which can then initiate the retrieval of the full pattern by the recurrent collaterals – one needs a large number of associatively modifiable synapses. This is expressed by the formulas that give the specific signal S produced by sets of associatively modifiable synapses, or by nonassociatively modifiable synapses: if C^{AFF} is the number of afferents per cell,

$$S_{ASS} \sim \frac{\sqrt{C^{AFF}}}{\sqrt{p}} q; \quad S_{NONASS} \sim \frac{1}{\sqrt{C^{AFF}}} q. \tag{10}$$

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.

The concept is that to initiate retrieval, a numerically large input (the perforant path system, see Fig. 3) is useful so that even a partial cue is sufficient (see Eq. (17) of Treves and Rolls, 1992); and that the retrieval cue need not be very strong, as the recurrent collaterals (in CA3) then take over in the retrieval process to produce good recall (Rolls, 2008a; Treves and Rolls, 1992). In this scenario, the perforant path to CA3 synapses operate as a pattern associator, the quantitative properties of which are described elsewhere (Rolls, 2008a; Rolls and Treves, 1990, 1998). If an incomplete recall cue is provided to a pattern association network using distributed input representations, then most of the output pattern will be retrieved, and in this sense *pattern association networks generalize between similar retrieval patterns to produce the correct output firing* (Rolls, 2008a, 2014a), and this generalization performed at the perforant path synapses to CA3 cells helps in the further completion produced by the recurrent collateral CA3–CA3 autoassociation process.

In contrast, during storage, strong signals, in the order of mV for each synaptic connection, are provided by the mossy fiber inputs to dominate the recurrent collateral activations, so that the new pattern of CA3 cell firing can be stored in the CA3 recurrent collateral connections (Rolls, 2008a; Treves and Rolls, 1992).

Before leaving the CA3 cells, it is suggested that separate scaling of the three major classes of excitatory input to the CA3 cells (recurrent collateral, mossy fiber, and perforant path, see Fig. 1) could be independently performed, by virtue of the different classes of inhibitory interneuron which receive their own set of inputs, and end on different parts of the dendrite of the CA3 cells (cf. for CA1 Buhl et al., 1994; Gulyas et al., 1993). This possibility is made simpler by having these major classes of input terminate on different segments of the dendrites. Each of these inputs, and the negative feedback produced through inhibitory interneurons when the CA3 cells fire, should for optimal functioning be separately regulated (Rolls, 1995), and the anatomical arrangement of the different types of inhibitory interneuron might be appropriate for achieving this.

2.2.4. CA1 cells

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

The amount of information about each episode retrievable from CA3 has to be balanced 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 non-modifiable, 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 quantitatively analyzed by Schultz and Rolls (1999).

2.2.4.2. 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 in order to form the association. Moreover, when completion operates in the CA3 autoassociation system, all the neurons firing in the original conjunction can be brought into activity by only a part of the original set of conjunctive events. For these reasons, a memory in the CA3 cells consists of several different simultaneously active ensembles of activity. To be explicit, the parts A, B, C, D and E of a particular episode would each be represented, roughly speaking, by its own population of CA3 cells, and these five populations would be linked together by autoassociation. It is suggested that the CA1 cells, which receive these groups of simultaneously active ensembles, can detect the conjunctions of firing of the different ensembles that represent the episodic memory, and allocate by competitive learning neurons to represent at least larger parts of each episodic memory (Rolls, 1987, 1989a,b,d, 1990a, 1990b). In relation to the simple example above, some CA1 neurons might code for ABC, and others for BDE, rather than having to maintain independent representations in CA1 of A, B, C, D, and E. This implies a more efficient representation, in the sense that when eventually after many further stages, neocortical neuronal activity is recalled (as discussed below), each neocortical cell need not be accessed by all the axons carrying each component A,B,C,D and E, but instead by fewer axons carrying larger fragments, such as ABC, and BDE. This process is performed by competitive networks, which self-organize to find categories in the input space, where each category is represented by a set of simultaneously active inputs (Rolls, 2000b; Rolls and Deco, 2002; Rolls and Treves, 1998).

Concerning the details of operation of the CA1 system, we note that although competitive learning may capture part of how it is able to recode, the competition is probably not global, but instead would operate relatively locally within the domain of the connections of inhibitory neurons. This simple example is intended to show how the coding may become less componential and more conjunctive in CA1 than in CA3, but should not be taken to imply that the representation produced becomes more sparse. An example of what might be predicted follows. Given that objects and the places in which they are seen may be associated by learning in the hippocampus, we might predict that relatively more CA3 neurons might be encoding either objects or places (the components to be associated), and relatively more CA1 neurons might be encoding conjunctive combinations of objects and places. Although object, spatial view, and object-spatial view combination neurons have now been shown to be present in the primate hippocampus during a room-based object-spatial view task (Rolls et al., 2005b), the relative proportions of the types of neuronal response in CA3 vs CA1 have not yet been analyzed.

The conjunctive recoding just described applies to any one single item in memory. Successive memory items would be encoded by a temporal sequence of such conjunctively encoded representations. To the extent that CA1 helps to build conjunctively encoded separate memory items, this could help a sequence of such memories (i.e. memory items) to be distinct (i.e. orthogonal to each other). This may be contrasted with successive memory items as encoded in CA3, in which there might be more overlap because each memory would be represented by its components, and some of the components might be common to some of the successive memory items. Episodic memory has been used here to refer to the memory of a single event or item. It can also be used to describe a succession of items that occurred close together in time, usually at one place. CA1 could help the encoding of these successive items that are parts of a multi-item episodic memory by making each item distinct from the next, due to the conjunctive encoding of each item (Rolls and Stringer, in preparation). These arguments suggest that the CA1 could be useful in what is described as temporal pattern separation in Section 3.3.3.

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

Another possibility is that the perforant path input provides the strong forcing input to the CA1 neurons during learning, and that the output of the CA3 system is associated with this forced CA1 firing during learning (McClelland et al., 1995). During recall, an incomplete cue could then be completed in CA3, and the CA3 output would then produce firing in CA1 that would correspond to that present during the learning. This suggestion is essentially identical to that of Treves and Rolls (1994) about the backprojection system and recall, except that McClelland et al. (1995) suggest that the output of CA3 is associated at the CA3 to CA1 (Schaeffer collateral) synapses with the signal present during training in CA1, whereas in the theory of Treves and Rolls (1994), the output of the hippocampus consists of CA1 firing 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 below).

2.2.4.4. CA1 and sequence memory for objects and odors. Evidence described in Section 3 indicates that the CA1, but not (dorsal) CA3, subregion is involved in the memory for sequences of objects and odors.[6] 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). As we have seen above, 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 3.3). Consistent with this hypothesis, timing cells are found in the medial entorhinal cortex (Kraus et al., 2013a, 2013b). 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, 2008a) to form time and object combination neurons (with time originating

[6] The CA3 lesions described in Section 3 were in the dorsal part of CA3. It remains possible that the ventral part of CA3 contains neurons that respond to objects, odors, rewards, and perhaps time, and is involved in processing these by autoassociation. This could provide a source of time/object information to CA1. Tests on the effects of ventral CA3 lesions, or CA3 NMDA knockouts, on these tasks would be very interesting.

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, Eichenbaum and colleagues (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 non-spatial 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, 2013b; Macdonald et al., 2011; Rolls and Deco, 2010).

An alternative hypothesis would implicate the (perhaps ventral part of) CA3 in object-temporal sequence encoding. This hypothesis utilizes the slow transitions from one attractor to another which are a feature that arises at least in some networks in the brain due to the noise-influenced transitions from one state to another, and due to slow dynamics such as those produced by neuronal adaptation (Deco and Rolls, 2005c). The process would involve autoassociation between neurons that fire at different times in a trial and neurons that represent the object. These associations would provide the basis for the recall of the object from the time in a trial, or vice versa. The retrieval of object or temporal information from each other would occur in a way that is analogous to that shown for recalling the object from the place, or the place from the object (Rolls et al., 2002), but substituting the details of the properties of the 'time encoding' neurons (Macdonald et al., 2011) for what was previously the spatial (place) component. In addition, if the time encoding neurons simply cycled through their normal sequence during recall, this would enable the sequence of objects or events associated with each subset of time encoding neurons to be recalled correctly in the order in which they were presented.

2.2.4.5. Subcortical, vs neocortical backprojection, outputs from CA1. The CA1 network is on anatomical grounds thought to be important for the hippocampus to access the neocortex, and the route by which retrieval to the neocortex would occur (see Fig. 1). However, as also shown in Fig. 1, there are direct subcortical outputs of the hippocampus (Rolls, 2014b). One route is that outputs that originate from CA3 project via the fimbria directly to the medial septum (Gaykema et al., 1991) or indirectly via the lateral septum to the medial septum (Saper et al., 1979; Swanson and Cowan, 1977). The medial septum in turn provides cholinergic and GABAergic inputs back to the hippocampus, especially the CA3 subregion. In addition to projecting back to the hippocampus, the medial septum also projects to subiculum, entorhinal cortex and mammillary bodies. In turn, the mammillary bodies project to the anterior thalamus which projects to the cingulate cortex, and also to the entorhinal cortex via the pre and parasubiculum. The outputs from CA3 in the fimbria can thus potentially send information to the entorhinal cortex via direct projections from the medial septum, and via the anterior thalamus. A second subcortical output route is provided by the subiculum and CA1 projections via the fimbria/fornix to the lateral septum, and to the mammillary bodies and anterior thalamic nuclei, which in turn project to the cingulate cortex as well as to the prelimbic cortex (rat), which could provide a potential output to behavior (Vann and Aggleton, 2004), and to the entorhinal cortex. A third subcortical output route is that the subiculum and CA1 also project to the nucleus accumbens (Swanson and Cowan, 1977). How might the CA1 to entorhinal cortico-cortical backprojection system, and the different subcortical output routes from the hippocampus, implement different functions?

First, we note that the CA1 to entorhinal cortico-cortical backprojection system is numerically in terms of the number of connections involved large, and this as developed below is likely to be crucial in allowing large numbers of memories to be retrieved from the hippocampus back to the neocortex. This recall operation to the neocortex is described below, but we note here that once a memory has been retrieved to the neocortex, the neocortex then provides a way for hippocampal function to influence behavior. For example, in a place–object recall task, a place cue used to initiate recall of the whole place–object memory in the CA3 would then lead via the backprojections to the object being retrieved (via entorhinal cortex layer 5 and perirhinal cortex) back to the inferior temporal visual cortex (see Fig. 1 left), and this object representation could then guide behavior via the other outputs of the inferior temporal visual cortex (see Rolls and Deco, 2002; Rolls, 2008a). As objects are represented in neocortex, it is a strong prediction of the theory that place–object recall will involve this CA1 backprojection pathway back to the neocortex. Backprojections from CA1 via entorhinal cortex to parahippocampal gyrus and thus to parietal cortex or other parts of the neocortex with spatial representations could allow spatial representations to be retrieved in for example an object–place recall task, and this provides a possible route to spatial action (though an alternative subcortical route for spatial action is described below). In addition, there are more direct CA1, subicular, and entorhinal connections to the prefrontal cortex (Jay and Witter, 1991; Witter et al., 1989a), which could allow the hippocampus to provide inputs to prefrontal cortex networks involved in short term memory and in planning actions (Deco and Rolls, 2003, 2005a).

Second, we note that the part of the second subcortical pathway that projects from CA1 and subiculum to the mammillary bodies and anterior thalamus is also numerically large, with in the order of 10^6 axons in the fornix and mammillo-thalamic tract. This could therefore provide an information-rich output pathway to regions such as the cingulate cortex via which some types of (e.g. spatial) behavioral response could be specified by the hippocampus. In contrast, the CA3 projections via the fimbria to the medial septum described as part of the first subcortical route are much less numerous, with indeed for example only approximately 2000 cholinergic neurons in the rat medial septum (Moore et al., 1998). Given further that this first route returns mainly to the hippocampus and connected structures, this route may be more involved in regulating hippocampal function (via cholinergic and GABAergic influences), than in directly specifying behavior (cf. Hasselmo et al., 1995).

Differential predictions of the functions mediated via the cortical vs subcortical outputs of the hippocampal system are thus as follows. Place to object recall should involve the backprojections via CA1 to the neocortex, for that is where the objects are represented. But object to place (or spatial response) tasks *might not* require CA1 and backprojections to parts of the cortex with spatial representations, in so far as the hippocampal subcortical output needs then only to affect which behavioral (spatial) response is made, which it is conceivable could be produced by fornix/fimbria efferents which indirectly reach structures such as the cingulate cortex. To test the hypothesis of potentially different output systems for different types of behavior, a new rat paradigm could be run comparing object–place vs place–object cued recall tasks. The prediction is that object–place cued recall might depend on hippocampal outputs to subcortical structures such as the mammillary bodies etc., though cortical outputs might be an alternative route.

It would at least be of interest to know whether cued object–place recall is impaired by damaging hippocampal subcortical outputs. The stronger prediction is that place–object cued recall depends on CA1 and recall to the neocortex. It is noted that in either case the hippocampus itself may be useful, because object–place and place–object associations involve places, and can not be solved by object–reward association memory, which instead depends on the orbitofrontal cortex and amygdala (Rolls, 2005, 2014a).

2.2.5. Backprojections to the neocortex, episodic memory recall, and consolidation

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 (Rolls, 1995, 1996b, 2008a, 2010b; Treves and Rolls, 1994) is considered here in the context of recalling a complete memory representation in the complete set of cortical areas that provide inputs to the hippocampus (see Fig. 1).

It is suggested that the modifiable connections from the CA3 neurons to the CA1 neurons allow the whole episode in CA3 to be produced in CA1. This may be assisted as described above by the direct perforant path input to CA1. 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 (after Gold and Kesner, 2005) 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 (Markov et al., 2014; Rolls, 1989a,b,d). 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 the orbitofrontal and anterior cingulate cortex to retrieve the reward/affective aspects of an episodic memory (Rolls, 2014a,b)), and also areas of unimodal association cortex (e.g. inferior temporal visual cortex) (Lavenex and Amaral, 2000). 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 (sometimes called schemas (Preston and Eichenbaum, 2013)) can be considered as a form of memory consolidation (Rolls, 1989a,b,d, 1990a, 1990b, 2008a), 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 we suggest that the backprojection synapses are set up to have the appropriate strengths for recall is as follows (Rolls, 1989a,b,d, 2008a). 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 paper and by McClelland et al. (1995), that recall might be useful to the neocortex to help it build new semantic memories, which might inherently be a slow process and is not part of the theory of recall. It is an interesting possibility that recall might involve several cycles through the recall process. After the information fed back from the first pass with a recall cue from perhaps only one cortical area, information might gradually be retrieved to other cortical areas involved in the original memory, and this would then act as a better retrieval cue for the next pass.

The timing of the backprojecting activity would be sufficiently rapid, in that for example inferior temporal cortex (ITC) neurons become activated by visual stimuli with latencies of 90–110 ms and may continue firing for several hundred ms (Rolls, 1992); and hippocampal pyramidal cells are activated in visual object-and-place and conditional spatial response tasks with latencies of 120–180 ms (Miyashita et al., 1989; Rolls et al., 1989; Rolls and Xiang, 2006). Thus, backprojected activity from the hippocampus might be expected to reach association cortical areas such as the inferior temporal visual cortex within 60–100 ms of the onset of their firing, and there would be a several hundred ms period in which there would be conjunctive feedforward activation present with simultaneous backprojected signals in the association cortex.

During recall, the backprojection connections onto the distal synapses of cortical pyramidal cells would be helped in their efficiency in activating the pyramidal cells by virtue of two factors. The first is that with no forward input to the neocortical pyramidal cells, there would be little shunting of the effects received at the distal dendrites by the more proximal effects on the dendrite normally produced by the forward synapses. Further, without strong forward activation of the pyramidal cells, there would not be very strong feedback and feedforward inhibition via GABA cells, so that there would not be a further major loss of signal due to (shunting) inhibition on the cell body and (subtractive) inhibition on the dendrite. (The converse of this is that when forward inputs are present, as during normal processing of the environment rather than during recall, the forward inputs would, appropriately, dominate the activity of the pyramidal cells, which would be only influenced, not determined, by the backprojecting inputs (see Deco and Rolls, 2005b; Rolls, 1989b,d, 2008a).

The synapses receiving the backprojections would have to be Hebb-modifiable, as suggested by Rolls (1989b,d). This would solve

the de-addressing problem, which is the problem of how the hippocampus is able to bring into activity during recall just those cortical pyramidal cells that were active when the memory was originally being stored. The solution hypothesized (Rolls, 1989b,d) arises because modification occurs during learning of the synapses from active backprojecting neurons from the hippocampal system onto the dendrites of only those neocortical pyramidal cells active at the time of learning. Without this modifiability of cortical backprojections during learning at some cortical stages at least, it is difficult to see how exactly the correct cortical pyramidal cells active during the original learning experience would be activated during recall. Consistent with this hypothesis (Rolls, 1989b,d), there are NMDA receptors present especially in superficial layers of the cerebral cortex (Monaghan and Cotman, 1985), implying Hebb-like learning just where the backprojecting axons make synapses with the apical dendrites of cortical pyramidal cells. The quantitative argument that the backprojecting synapses in at least one stage have to be associatively modifiable parallels that applied to the pattern retrieval performed at the entorhinal to CA3 synapses (Treves and Rolls, 1992) and at the CA3 to CA1 synapses (Schultz and Rolls, 1999), and is that the information retrieved would otherwise be very low. The performance of pattern association networks is considered in detail by Rolls and Treves (Rolls, 2008a; Rolls and Treves, 1990, 1998) and other authors (Hertz et al., 1991). An interesting recent development is that dilution in the connectivity of pattern association networks by reducing the probability of two synapses onto a neuron from the same backprojecting input helps to maintain high capacity in the backprojecting system (Rolls, 2015). If there are some instances of two or more synapses onto a neuron from the same input axon in a pattern association network, the capacity of the system degrades because of distortion and interferences between the stored patterns (Rolls, 2015). It is also noted that the somewhat greater anatomical spread of the backprojection than the forward connections between two different stages in the hierarchy shown in Fig. 1 would not be a problem, for it would provide every chance for the backprojecting axons to find co-active neurons in an earlier cortical stage that are part of the representation that is relevant to the current memory being formed.

If the backprojection synapses are associatively modifiable, we may consider the duration of the period for which their synaptic modification should persist. What follows from the operation of the system described above is that there would be no point, indeed it would be disadvantageous, if the synaptic modifications lasted for longer than the memory remained in the hippocampal buffer store. What would be optimal would be to arrange for the associative modification of the backprojecting synapses to remain for as long as the memory persists in the hippocampus. This suggests that a similar mechanism for the associative modification within the hippocampus and for that of at least one stage of the backprojecting synapses would be appropriate. It is suggested that the presence of high concentrations of NMDA synapses in the distal parts of the dendrites of neocortical pyramidal cells and within the hippocampus may reflect the similarity of the synaptic modification processes in these two regions (cf. Kirkwood et al., 1993). It is noted that it would be appropriate to have this similarity of time course (i.e. rapid learning within 1–2 s, and slow decay over perhaps weeks) for at least one stage in the series of backprojecting stages from the CA3 region to the neocortex. Such stages might include the CA1 region, subiculum, entorhinal cortex, and perhaps the parahippocampal gyrus/perirhinal cortex. However from multimodal cortex (e.g. the parahippocampal gyrus) back to earlier cortical stages, it might be desirable for the backprojecting synapses to persist for a long period, so that some types of recall and top-down processing (Rolls, 1989b,d, 2008a; Rolls and Deco, 2002) mediated by the operation of neocortico-neocortical backprojecting synapses could be stable, and might not require modification during the learning of a new episodic memory.

An alternative hypothesis to that above is that rapid modifiability of backprojection synapses would be required only at the beginning of the backprojecting stream. Relatively fixed associations from higher to earlier neocortical stages would serve to activate the correct neurons at earlier cortical stages during recall. For example, there might be rapid modifiability from CA3 to CA1 neurons, but relatively fixed connections from there back (McClelland et al., 1995). For such a scheme to work, one would need to produce a theory not only of the formation of semantic memories in the neocortex, but also of how the operations performed according to that theory would lead to recall by setting up appropriately the backprojecting synapses.

We have noted elsewhere that backprojections, which included cortico-cortical backprojections, and backprojections originating from structures such as the hippocampus and amygdala, may have a number of different functions (Rolls, 1989a,b,d, 1990a,b, 1992, 2005, 2008a; Rolls and Deco, 2002) including implementing top-down attention by biased competition (Deco and Rolls, 2003, 2004, 2005a; Deco et al., 2005; Grabenhorst and Rolls, 2010; Rolls, 2008a, 2008b, 2013a; Rolls and Deco, 2002, 2006). The particular function with which we have been concerned here is how memories stored in the hippocampus might be recalled in regions of the cerebral neocortex, and this is not at all incompatible with such theories of top-down attentional control.

2.2.6. Backprojections to the neocortex – quantitative analysis

How many backprojecting fibers does one need to synapse on any given neocortical pyramidal cell, in order to implement the mechanism outlined above? Consider a polysynaptic sequence of backprojecting stages, from hippocampus to neocortex, as a string of simple (hetero-)associative memories in which, at each stage the input lines are those coming from the previous stage (closer to the hippocampus) (Rolls, 2008a; Treves and Rolls, 1994) (Fig. 1). (The interesting concept here is that one can treat for a capacity analysis the series of backprojection stages to the cerebral cortex which each involve a pattern association as an “unrolled” version of an autoassociator. Each backprojection pattern association stage would correspond to one iteration round the autoassociation system.) Implicit in this framework is the assumption that the synapses at each stage are modifiable and have been indeed modified at the time of first experiencing each episode, according to some Hebbian associative plasticity rule. A plausible requirement for a successful hippocampo-directed recall operation, is that the signal generated from the hippocampally retrieved pattern of activity, and carried backwards 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 above.) If p is equal to the number of memories held in the hippocampal memory, it is limited by the retrieval capacity of the CA3 network, $p_{\max}$. Putting together the formula for the latter with that shown here, one concludes that, roughly, the requirement implies that the number of afferents of (indirect) hippocampal origin to a given neocortical stage (C^{HBP}), must be $C^{\mathrm{HBP}} = C^{\mathrm{RC}} a_{\mathrm{nc}}/a_{\mathrm{CA3}}$, 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 in CA3.

The above requirement is very strong: even if representations were to remain as sparse as they are in CA3, which is unlikely, to avoid degrading the signal, C^{HBP} should be as large as C^{RC}, i.e. 12,000 in the rat. Moreover, other sources of noise not considered in the present calculation would add to the severity of the constraint, and partially compensate for the relaxation in the constraint that would result from requiring that only a fraction of the p episodes would involve any given cortical area. 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).

Although backprojections between any two adjacent areas in the cerebral cortex are approximately as numerous as forward projections, and much of the distal parts of the dendrites of cortical pyramidal cells are devoted to backprojections, the actual number of such connections onto each pyramidal cell may be on average only in the order of thousands. Further, not all might reflect backprojection signals originating from the hippocampus, for there are backprojections which might be considered to originate in the amygdala (see Amaral et al., 1992) or in multimodal cortical areas (allowing for example for recall of a visual image by an auditory stimulus with which it has been regularly associated). In this situation, one may consider whether the backprojections from any one of these systems would be sufficiently numerous to produce recall. One factor that may help here is that when recall is being produced by the backprojections, it may be assisted by the local recurrent collaterals between nearby (~1 mm) pyramidal cells which are a feature of neocortical connectivity. These would tend to complete a partial neocortical representation being recalled by the backprojections into a complete recalled pattern. (Note that this completion would be only over the local information present within a cortical area about e.g. visual input *or* spatial input; it provides a local "clean-up" mechanism, and could not replace the global autoassociation performed effectively over the activity of very many cortical areas which the CA3 could perform by virtue of its widespread recurrent collateral connectivity.) There are two alternative possibilities about how this would operate. First, if the recurrent collaterals showed slow and long-lasting synaptic modification, then they would be useful in completing the whole of long-term (e.g. semantic) memories. Second, if the neocortical recurrent collaterals showed rapid changes in synaptic modifiability with the same time course as that of hippocampal synaptic modification, then they would be useful in filling in parts of the information forming episodic memories which could be made available locally within an area of the cerebral neocortex.

This theory of recall by the backprojections thus provides a quantitative account of why the cerebral cortex has as many backprojection as forward projection connections (Rolls, 2008a).

A recent development has been on understanding the diluted connectivity in the back-projection pathways. The backprojection pathways have diluted connectivity in that each neocortical neuron may receive up to in the order of 10,000 backprojection synapses, yet the total number of neurons at each neocortical stage in the backprojection hierarchy shown in Fig. 1 has very many more neurons than this, at least 10 times as many neurons as there are synapses per neuron. This diluted connectivity in the backprojection pathways is likely to reduce the probability of more than one connection onto a receiving neuron in the backprojecting pathways. It has now been shown that if there were more than one connection onto each neuron in a pattern associator, this would reduce the capacity of the system, which in this case is the number of memories that can be recalled from the hippocampus to the neocortex (Rolls, 2015).

These concepts show how the backprojection system to neocortex can be conceptualized in terms of pattern completion, as follows. The information that is present when a memory is formed may be present in different areas of the cerebral cortex, for example of a face in a temporal cortex face area (Rolls, 2012b), of a spatial location in a neocortical location area, and of a reward received in the orbitofrontal cortex (Rolls, 2014a). To achieve detailed retrieval of the memory, reinstatement of the activity during recall of the neuronal activity during the original memory formation may be needed. This is what the backprojection system described could achieve, and is a form of completion of the information that was represented in the different cortical areas when the memory was formed. In particular, the concept of completion here is that if a recall cue from a visual object area is provided, then the emotional parts of the episodic memory can be recalled in the orbitofrontal cortex, and the spatial parts in parietal cortical areas, with the result that a complete memory is retrieved, with activity recalled into several higher-order cortical areas. Because such a wide set of different neocortical areas must be content-addressed, a multistage feedback system is required, to keep the number of synapses per neuron in the backprojection pathways down to reasonable numbers. (Having CA1 directly address neocortical areas would require each CA1 neuron to have tens of millions of synapses with cortical neurons. That is part of the computational problem solved by the multistage backprojection system shown in Fig. 1.) Thus, the backprojection system with its series of pattern associators can each be thought of as retrieving the complete pattern of cortical activity in many different higher-order cortical areas that was present during the original formation of the episodic memory.

2.2.7. Entorhinal cortex

In the medial entorhinal cortex, there are grid cells, which have regularly spaced peaks of firing in an environment, so that as a rat runs through an environment, a single neuron increases then decreases its firing a number of times as the rat traverses the environment (Fyhn et al., 2004; Giocomo et al., 2011; Hafting et al., 2005; McNaughton et al., 2006) (see Fig. 4). The grid cell system appears to provide ring continuous attractors that would be useful not only for spatial path integration (computing position based on self-motion) (Giocomo et al., 2011; McNaughton et al., 2006), but also for the timing information useful in sequence encoding for non-spatial as well as spatial information, as described in Section 3.3. 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 times in a sequence, which is what is needed for the formation of odor, object and spatial sequence memories in CA1 (see Section 3.3), and would also be useful in the formation of spatial grid cells in the entorhinal cortex. Consistent with this hypothesis, timing cells are found in the medial entorhinal cortex (Kraus et al., 2013a, 2013b). The lateral entorhinal

cortex may be more involved in representing object information (Hargreaves et al., 2005; Knierim et al., 2014; Witter et al., 1989a).

2.2.8. Simulations of hippocampal operation

In order to test the operation of the whole system for individual parts of which an analytic theory is available (Rolls, 2008a; Rolls and Treves, 1998; Schultz and Rolls, 1999; Treves and Rolls, 1992, 1994), Rolls (1995) simulated a scaled down version of the part of the architecture shown in Fig. 1 from the entorhinal cortex to the hippocampus and back to the entorhinal cortex. The network included competitive learning in the dentate gyrus to form a sparse representation, using this to drive the CA3 cells by diluted non-associative mossy fiber connectivity during learning, competitive learning in CA1, and pattern association in the retrieval back to entorhinal cortex. The analytic approaches to the storage capacity of the CA3 network, the role of the mossy fiber inputs to CA3 during learning and of the perforant path inputs to CA3 during retrieval, the functions of CA1, and the operation of the backprojections in recall, were all shown to be computationally plausible in the computer simulations. In the simulation, during recall, partial keys are presented to the entorhinal cortex, completion is produced by the CA3 autoassociation network, and recall is produced in the entorhinal cortex of the original learned vector. The network, which has 1000 neurons at each stage, can recall large numbers, which approach the calculated storage capacity, of different sparse random vectors presented to the entorhinal cortex.

One of the points highlighted by the simulation is that the network operated much better if the CA3 cells operated in binary mode (either firing or not), rather than having continuously graded firing rates (Rolls, 1995). The reason for this is that given that the total amount of information that can be stored in a recurrent network such as the CA3 network is approximately constant independently of how graded the firing rates are in each pattern (Treves, 1990), then if much information is used to store the graded firing rates in the firing of CA3 cells, fewer patterns can be stored. The implication of this is that in order to store many memories in the hippocampus, and to be able to recall them at later stages of the system in for example the entorhinal cortex and beyond, it may be advantageous to utilize relatively binary firing rates in the CA3 part at least of the hippocampus. This finding has been confirmed and clarified by simulation of the CA3 autoassociative system alone, and it has been suggested that the advantage of operation with binary firing rates may be related to the low firing rates characteristic of hippocampal neurons (Rolls et al., 1997b).

Another aspect of the theory emphasized by the results of the simulation was the importance of having effectively a single network provided in the hippocampus by the CA3 recurrent collateral network, for only if this operated as a single network (given the constraint of some topography present at earlier stages), could the whole of a memory be completed from any of its parts (Rolls, 1995).

Another aspect of the theory illustrated by these simulations was the information retrieval that can occur at the CA3–CA1 (Schaeffer collateral) synapses if they are associatively modifiable, and this led to a quantitative investigation of this (Schultz and Rolls, 1999).

3. Tests of the theory

Empirical evidence that tests the theory comes from a wide range of investigation methods, including the effects of selective lesions to different subregions of the hippocampal system, the activity of single neurons in different subregions of the hippocampal system, pharmacological and genetic manipulations of different parts of the system, the detailed neuroanatomy of the system, and functional neuroimaging studies and clinical neuropsychological studies in humans. Because in this analysis we are particularly interested in the contribution of different subregions of the hippocampal system, we focus here on the effects of selective lesions to different subregions of the hippocampal system, and the activity of single neurons in different subregions of the hippocampal system, together with pharmacological and genetic manipulations of different parts of the system. Having now described in Section 2 the overall theory of how different subregions of the hippocampal system contribute to its function, we now consider the effects of manipulations of different subregions of the hippocampal system, starting from the dentate gyrus, and working through to CA1.

3.1. Dentate gyrus (DG) subregion of the hippocampus

Based on the anatomy of the DG, its input and output pathways, and the development of a computational model, Rolls (1989b, 1996b, 2008a) (and the present paper Section 2.2.2) has suggested that the DG can act as a competitive learning network with Hebb-like modifiability to remove redundancy from the inputs producing a more orthogonal, sparse, and categorized set of outputs. One example of this in the theory is that to the extent that some entorhinal cortex neurons represent space as a grid (Hafting et al., 2005), the correlations in this type of encoding are removed to produce a representation of place, with each place encoded differently from other places (Jung and McNaughton, 1993). To the extent then that DG acts to produce separate representations of different places, it is predicted that the DG will be especially important when memories must be formed about similar places. We note that to form spatial representations, learned conjunctions of sensory inputs including vestibular, olfactory, visual, auditory, and somatosensory may be involved. DG may help to form orthogonal representations based on all these inputs, and could thus help to form orthogonal non-spatial as well as orthogonal spatial representations for use in CA3. In any case, the model predicts that for spatial information, the DG should play an important role in hippocampal memory functions when the spatial information is very similar, for example, when the places are close together.

3.1.1. Conjunctive encoding

The dDG has been shown to receive multiple sensory inputs, including vestibular, olfactory, visual, auditory, and somatosensory, from the perirhinal and lateral entorhinal cortex in conjunction with spatially organized grid cells from the medial entorhinal cortex (Hafting et al., 2005) to represent metric spatial representations. The perforant path input of the dDG can be divided into medial and lateral components. The medial component processes spatial information and the lateral component processes nonspatial (e.g. objects, odors) information (Hargreaves et al., 2005; Witter et al., 1989a). Based on the idea that the medial perforant path input into the dDG mediates spatial information via activation of NMDA receptors and the lateral perforant path input into the DG mediates visual object information via activation of opioid receptors, the following experiment was conducted. Using a paradigm developed by Poucet (1989), rats were tested for detection of a novel spatial change and detection of a novel visual object change while under the influence of direct infusions of AP5 (an NMDA antagonist) or naloxone (a μ opioid antagonist) into the dDG. Naloxone infusions into the dDG disrupted both novelty detection of a spatial location and a visual object, whereas AP5 infusions into the dDG disrupted only detection of a novel spatial location, but had no effect on detection of a novel object (Hunsaker et al., 2007b). These data suggest that the dDG uses conjunctive encoding of visual object and spatial information to provide for a spatial representation that may be based on metric information. Additional support for the idea that the dDG contributes to the organization of spatial information based on conjunctive encoding was provided in the

following experiment that emphasized acquisition of an association between a context and an odor cue (Morris et al., 2013b). In this experiment, dDG lesions as well as controls were trained in two Plexiglas boxes that represented two different contexts: Context 1 and Context 2. A context was defined as the total of all of the environmental cues in the apparatus, including floor texture, color of walls, and visual cues on the walls. Each animal was assigned two pseudo-randomly selected odors (Odor A and Odor B), which were used throughout all testing procedures. The task required that the rat learn to associate Odor A with Context 1 and Odor B with Context 2 in order to receive a food reward. Animals with dDG lesions demonstrated impairments in the ability to acquire the cue–context associations compared with the controls. Results from follow-up odor- and context-discrimination tests suggest that all rats acquired the discrimination tasks at similar rates. Therefore, it is unlikely that deficits in performance on the contextual associative learning task were due to an inability to discriminate between odors or contexts. The present findings suggest that the dDG may play an important role in conjunctive encoding of spatial paired with olfactory information.

3.1.2. Spatial pattern separation

It can clearly be demonstrated that single cells within the hippocampus are activated by most sensory inputs, including vestibular, olfactory, visual, auditory, and somatosensory, as well as higher-order integration of sensory stimuli (Hartley et al., 2014; Itskov et al., 2011; Komorowski et al., 2009; O'Mara et al., 1994; Rolls, 2008a; Rolls and Xiang, 2005, 2006; Rolls et al., 2005b). An important question is whether these sensory inputs via conjunctive encoding have a memory representation within the hippocampus. One possible role for the hippocampus in processing all sensory information might be to provide sensory markers to demarcate a spatial location, so that the hippocampus can more efficiently mediate spatial information. Thus, it is possible that one of the main processing functions of the hippocampus is to encode and separate spatial events from each other; this ensures that new highly processed sensory information is organized within the hippocampus and enhances the possibility of encoding and temporarily remembering one place as separate from another place, while reducing spatial interference. There have been a number of review articles describing functions of the dDG in supporting pattern separation (Carr et al., 2010; Gilbert and Brushfield, 2009; Kesner, 2013; Rolls, 2013b; Schmidt et al., 2012; Yassa and Stark, 2011), and this paper includes new topics including the determination of a role for the DG in context and odor pattern separation.

Rolls' model (Section 2) (Rolls, 1989b, 2008a, 2010b, 2013b) proposes that pattern separation is facilitated by sparse connections in the mossy fiber system, which connects dDG granular cells to dorsal dCA3 pyramidal neurons. Separation of patterns is accomplished based on the low probability that any two dCA3 neurons will receive mossy fiber input synapses from a similar subset of dDG cells. Mossy fiber inputs to dCA3 from dDG are suggested to be essential during learning and may influence which CA3 neurons fire based on the distributed activity within the dDG. Cells of the dDG are suggested to act as a competitive learning network with Hebb-like modifiability to reduce redundancy and produce sparse, orthogonal outputs. O'Reilly and McClelland (1994) and Shapiro and Olton (1994) also suggested that the mossy fiber connections between the dDG and dCA3 may support pattern separation.

If disruption of dDG function results in inefficient pattern separation, then deficits involving spatial tasks may occur when there is increased overlap or similarity among distal cues and presumably increased similarity among representations within the dDG. Remembering a specific location in, for example, an eight-arm maze, a water maze, or a spatial context in fear conditioning may be influenced by the degree of overlap among critical distal spatial cues. dDG-lesioned rats demonstrate deficits similar to complete hippocampal lesions when tested on a working memory version of the radial eight-arm maze (Emerich and Walsh, 1989; McLamb et al., 1988; Tilson et al., 1987; Walsh et al., 1986). In addition, rats with dDG lesions also showed deficits comparable to rats with complete hippocampal lesions when tested on the Morris water maze task if the start location was varied on each trial (Jeltsch et al., 2001; Nanry et al., 1989; Sutherland et al., 1983; Xavier et al., 1999).

Lee and Kesner (2004a) tested rats with dDG lesions on acquisition and retrieval of contextual fear conditioning. Rats with dDG lesions showed initial impairments in freezing behavior during acquisition of the task, but they eventually reached the freezing level of controls with subsequent testing. When retrieval of contextual fear was examined in rats with dDG lesions 24 h after acquisition, the animals showed a significant deficit in freezing compared to controls. Based on these studies it is clear that dDG lesions impair spatial memory. dDG-lesioned animals may be impaired as a result of impaired pattern separation; however, it is difficult to determine whether a particular memory process is impaired in these animals using these three paradigms.

To examine the contribution of the dDG to spatial pattern separation, Gilbert et al. (2001) tested rats with dDG lesions using a paradigm that measured short-term memory for spatial location information as a function of spatial similarity between spatial locations. Specifically, the study was designed to examine the role of the dDG subregion in discriminating spatial locations when rats were required to remember a spatial location based on distal environmental cues and to differentiate between the to-be-remembered location and a distractor location with different degrees of similarity or overlap among the distal cues.

Animals were tested using a cheeseboard maze apparatus (the cheeseboard is similar to a dry land water maze with 177 circular, recessed holes on a 119-cm-diameter board) on a delayed match-to-sample for a spatial location task. Animals were trained to displace an object that was randomly positioned to cover a baited food well in 1 of 15 locations along a row of food wells. Following a short delay, the animals were required to choose between objects that were identical to the sample-phase object: one object was in the same location as the sample-phase object and the second object was in a different location along the row of food wells. Rats were rewarded for displacing the object in the same spatial location as the sample-phase object (correct choice), but they received no reward for displacing the foil object (incorrect choice). Five spatial separations, from 15 cm to 105 cm, were used to separate the correct object and the foil object during the choice phase. The results indicated that rats with dDG lesions were significantly impaired at short spatial separations; however, during the choice phase performance of dDG-lesioned animals increased as a function of a greater spatial separation between the correct and foil objects. The performance of rats with dDG lesions matched control rats at the largest spatial separation. The graded nature of the impairment and the significant linear improvement in performance as a function of increased separation illustrate a deficit in pattern separation (Gilbert et al., 2001). Based on these results, it was concluded that lesions of the dDG decrease the efficiency of spatial pattern separation, which results in impairments on trials with increased spatial proximity and increased spatial similarity among working memory representations. Thus, the dDG may function to encode and to separate events in space producing spatial pattern separation. Such spatial pattern separation ensures that new highly processed sensory information is organized within the hippocampus, which in turn enhances the possibility of encoding and temporarily remembering one spatial location as separate from another.

There are other studies in the literature that have demonstrated that hippocampal lesions, including the dDG, can result in deficits

for spatial tasks that can be interpreted as a function of increased interference and impairment in the utilization of a spatial pattern separation process. One example will suffice. McDonald and White (1995) used a place preference procedure in an eight-arm maze. Two of the arms were selected and food was placed at the end of one arm and not at the end of the other arm. Fornix-lesioned rats acquired the place preference task as quickly as controls if the arm locations were opposite each other, but the fornix-lesioned rats were markedly impaired if the locations were adjacent to each other. Clearly, there would be greater spatial interference when the locations are adjacent to each other rather than far apart. In a more recent study (Morris et al., 2012), the role of the dDG was tested in mediating spatial pattern separation using a similar place-learning paradigm described by McDonald and White (1995). The results indicate that rats with dorsal dDG lesions and control lesions acquired the spatial discrimination for separate locations at similar rates. However, for the adjacent condition, dDG-lesioned rats required significantly more trials to reach the learning criterion than did controls.

3.1.3. Pattern separation and neurogenesis

Based on the observation that neurogenesis occurs in the DG and that new DG granule cells can be formed across time, it has been proposed that the dDG mediates a spatial pattern separation mechanism as well as generates patterns of episodic memories within remote memory (Aimone and Gage, 2011; Aimone et al., 2006). Thus far, it has been shown in mice that disruption of neurogenesis using low-dose X-irradiation was sufficient to produce a loss of newly born dDG cells. Further testing indicated impairments in spatial learning in a delayed non-matching-to-place task in the radial-arm maze. Specifically, impairment occurred for arms that were presented with little separation, but no deficit was observed when the arms were presented farther apart (Clelland et al., 2009), suggesting a spatial pattern separation deficit. These results are similar to those described for rats in the Morris et al. (2012) study Another study in mice provided evidence that the disruption of neurogenesis using lentivirus expression of a dominant Wnt protein produced a loss of newly born dDG cells; and these mice were tested in an associative object-in-place task with different spatial separations and observed to be impaired as a function of the degree of separation, again suggesting a spatial pattern separation deficit (Clelland et al., 2009). In a more recent study (Kesner et al., 2014) it was shown that DNA methyltransferase 1-c knockout mice are impaired relative to controls in the spatial pattern separation task (Goodrich-Hunsaker et al., 2008). These data suggest that neurogenesis in the dDG may contribute to the operation of spatial pattern separation. Thus, spatial pattern separation may play an important role in the acquisition of new spatial information, and there is a good possibility that the dDG may have been the subregion responsible for the impairments in the various tasks described above.

3.1.4. Temporal integration and remote memory neurogenesis

A novel role for the dDG in temporal processing for spatial information has begun to emerge based on the development of a computational model of neurogenesis (Aimone and Gage, 2011; Aimone et al., 2006). According to this model, adult born granule cells in the dDG contribute to a temporal associative integration process for events presented within the same time frame. Neurogenesis, or the proliferation of new neurons, has been shown to occur in two regions in the adult brain including the subventricular zone (SVZ) of the lateral ventricles and the subgranular zone (SGZ) of the DG hippocampal subregion (Kempermann et al., 2004). Thousands of granule cells (GC) are born daily; however, only a portion of cells survives and develops into fully mature GCs (Kempermann et al., 2004). Interestingly, immature GCs exhibit unique characteristics that differ from mature GCs. For example, immature GCs appear to be hyper-excitable and have a lower threshold for the induction of long-term potentiation (LTP) than adult GCs (Aimone et al., 2010; Ge et al., 2007; Schmidt-Hieber et al., 2004). However, during the maturation process, immature GCs begin to take on characteristics associated with fully developed GCs (Deng et al., 2010). Essentially, the continual influx of newly formed GCs into existing hippocampal circuitry gives rise to a continually changing population of GCs (Kempermann et al., 2004).

Based on the unique characteristics of young GCs during different stages of maturation, Aimone et al. (2006) suggested that young GCs might make a distinct contribution to memory formation. Specifically, they proposed that young GCs mediate a temporal integration process that operates to form associations among temporally contiguous events. In other words, events that occur close in time may be encoded by a similar set of young GCs, while events that occur farther apart in time may be encoded and represented by different cell populations allowing for the formation and separation of distinct memory representations (Aimone et al., 2010, 2009; Deng et al., 2010). Evidence provided by these models indicates that adult born GCs may provide a temporal tag to sparse spatial representations formed in the DG and, according to Aimone et al., stored in dCA3.

Although there is computational evidence to suggest that adult neurogenesis in the DG contributes to a temporal associative process for proximal spatial events, there is a paucity of behavioral evidence to support the temporal integration theory. In addition, the role of the dDG in the formation of temporal associations for spatial events has not been directly tested. Aimone and colleagues (2009) stressed the importance of developing new behavioral paradigms to test computational models of temporal pattern integration and the formation of temporal associations. Currently, there is a paucity of behavioral evidence to support the temporal integration theory. Therefore, we developed a novel behavioral paradigm to investigate the role of the dDG in temporal integration for proximal and distal spatial events (Morris et al., 2013a). Male Long–Evans rats were randomly assigned as control animals or to receive bilateral intracranial infusions of colchicine into the dDG. Following recovery from surgery, each rat was tested on a novel cued-recall of sequence paradigm for different spatial locations. In this task, the study phase was conducted across 2 consecutive days (Day 1 and Day 2) and consisted of two 5-min exploration sessions separated by a 3-min intersession interval (ISI) per day. The rats were allowed to explore the same object placed in designated spatial locations on a cheeseboard maze across 2 days (e.g. Day 1: spatial location A and 3 min later spatial location B; Day 2: spatial location C and 3 min later spatial location D). One week later, on the first day of the test phase (Day 8), animals were placed on the maze and allowed to explore an object positioned at spatial location A or C (A and C were used as recall cues) for 1 min followed by a 3-min ISI. After this interval, animals were given a 5-min preference test between spatial location B and D (positioned 106 cm apart). On the second day, rats that received A as the recall cue on Day 8 were tested on Day 9 with C as the recall cue. Similarly, if C was uses as a recall cue on Day 8, then A would be used as a recall cue on Day 9. The presentation order of spatial location recall cue for spatial location A and spatial location C was counterbalanced across subjects and across days. Control animals showed a significant preference for the spatial location previously paired with the cue (the temporal associate), but dDG lesioned animals did not show a preference for either spatial location during the preference test. These findings suggest that selective colchicine-induced dDG lesions are capable of disrupting the formation of temporal associations between spatial events presented within a 3 min time frame. Similar results were obtained for DNA methyl transferase 1c knockout mice but with a shorter time frame (30 s), suggesting that neurogenesis may

play an important role in the formation of temporal associations (Kesner et al., 2014).

The interpretation of the result (Morris et al., 2013a) is that the A–B and C–D preferences are a function of a pattern completion process in dCA3 rather than a disruption in detecting novelty. To answer this question in another test the preference test involved a new location B–F. Interestingly, data from the present investigation revealed that control animals continued to show a significant preference for the temporal associate (B when cued with A; D when cued with C) rather than the novel spatial location (E or F), suggesting that the formation of temporal associations for proximal spatial events may outweigh the effect of novelty preference. Although this finding is somewhat unexpected based on results from standard preference tasks (e.g. Dix and Aggleton, 1999; Ennaceur and Delacour, 1988), the task used in the present study is not a typical preference task; rather, it is a cued-recall task with a preference test portion. Based on this distinction, it is possible that a cued-recall based pattern completion process can account for the observed preference for the cued location in control animals. It has been suggested that temporal remote memories are initially processed in the dDG (via dDG based neurogenesis) (Aimone et al., 2010). This information is then stored in dCA3, a region implicated in pattern completion (Gold and Kesner, 2005; Hunsaker and Kesner, 2013; Nakazawa et al., 2002; Rolls, 1987, 1989b, 1989c, 2008a, 2010b, 2013c; Rolls and Treves, 1998; Treves and Rolls, 1994). In the present study spatial cued-recall paradigm, the animal can use a pattern completion process to generate the observed preference for the temporally paired spatial location. For example, based on pattern completion, a cue of A (following the initial temporal pairing of location A with location B) results in a preference for B, as indicated by more time spent exploring the paired temporal associate. Similarly, a cue of C (following the initial temporal pairing with location D) results in a preference for C. In contrast, dDG lesioned animals failed to demonstrate a significant preference for the novel spatial location, suggesting that prior results obtained for the cued recall for spatial location task were not driven by a novelty preference, but instead point to an encoding deficit. Control rats still prefer B and dDG lesioned rats are still impaired (Morris et al., 2013a).

3.1.5. Context-pattern separation for geometry of the environment

Do other subregions of the hippocampus engage in spatial pattern separation? Early research suggested that the dCA3 region may also engage in spatial pattern separation processes for a spatial representation of the geometry of the environment. This idea is supported by the finding of Tanila (1999), who showed that dCA3c place cells were able to maintain distinct representations of two visually identical environments and to selectively reactivate either one of the representation patterns depending on the experience of the rat. Also, Leutgeb et al. (2007) showed that when rats experienced a completely different environment, dCA3 place cells developed orthogonal representations of the different environments by changing their firing rates between the environments, whereas dCA1 place cells maintained similar responses.

To further test the role of dCA3 in mediating pattern separation, an experiment was conducted to determine whether the dDG or dCA3 regions cooperate to perform spatial pattern separation operations for specific spatial locations as well as the spatial geometry of the environment or whether the dDG performs spatial pattern separation on the basis of specific locations in space and the dCA3 performs spatial pattern separation on the basis of the geometry of the environment (Hunsaker et al., 2008b). Rats with lesions of dDG and dCA3a,b were given the opportunity to explore a white or black circular or square box of the same size as reported by Leutgeb et al. (2007), which contained two objects spaced 68 cm apart. After habituation to the box and objects, the rats received one of two transfer tests. In the first test, objects were changed to a 38-cm distance, but the box shape (geometry of the environment) remained the same; in the second test, the box shape (geometrical environment) was changed, but the distance between objects remained the same. Efficacy of the transfer test in terms of re-exploration of the metric change was based on a comparison between the degree of object exploration during the transfer session versus the degree of object exploration during the last session of habituation. Similarly, the efficacy of the transfer test in terms of re-exploration of the geometry of the environment was based on the number of grid crossings (activity level) and rearings during the transfer session versus the number of grid crossings and rearings during the last session of habituation. The results indicated that lesions of the dDG, but not dCA3a,b, disrupted both detection of metric changes in the spatial location of objects and changes in a geometrical environment. Thus far, these data are consistent with the prediction of Rolls' computational model (Rolls, 1989b, 2008a, 2010b, 2013b) that the dDG is the critical substrate for spatial pattern separation. These data are consistent with other findings (Leutgeb et al., 2007; Tanila, 1999) of a pattern separation function for geometrical environments within CA3c, which is anatomically and functionally a part of the dDG circuit.

It has been shown that the CA3 region can be divided into a CA3a, b, and c sub-areas (Li et al., 1994; Lorente de No, 1934). The CA3 a,b region have recurrent collaterals, but the CA3c region does not have recurrent collaterals. Also, there is a direct connection from CA3c to CA1 which means that the Ca3c region can bypass the Ca3a,b region. Finally, based on previous research (Buckmaster and Schwartzkroin, 1994; Li et al., 1994), it has been proposed that mossy cells receive excitatory inputs from granule cells and dCA3c pyramidal cells and integrate their inputs, which, in turn via excitatory recurrent axonal projections, activate many dDG cells. Such a circuit could integrate spatial location information and form representations of geometrical environments. Most of the recording of cells that respond to different environments as reported by (Leutgeb et al., 2007; Tanila, 1999) were based on electrode placements in the dCA3c area. The lesion data were based on lesions within dCA3a/b, but not dCA3c. Additional experiments with dCA3c lesions in contrast to dCA3a/b lesions were carried out (Hunsaker et al., 2008b). The results indicated that dorsal dCA3c lesions disrupted pattern separation processes only when the animal was required to detect a metric change in object location, but there was no apparent effect during the environmental change task. It must be noted, however, that dorsal dCA3c lesions never caused effects as dramatic as those caused by dDG lesions. One interpretation may be that the dDG selectively recruits dCA3c to assist in metric detection and not detection of the overall environmental change. This experiment (Hunsaker et al., 2008b) provides behavioral evidence that the dCA3c and the dDG may interact for spatial information processing. This effect was seen only in the condition under which the animal was required to detect a discrete metric change in object location, a task that has been shown to be particularly sensitive to dDG. Interestingly, this deficit was observed in object exploration but not rearing behaviors. Although the present behavioral data do not show any effects of dCA3a,b lesions, the dDG effect is clear. Additionally, the dorsal dCA3c lesion data suggest that there is a circuit involving the dCA3c and the dDG that is perhaps important for preprocessing spatial information prior to dorsal dCA3a, b processing stages.

3.1.6. Context-pattern separation for color of the environment

In addition to a role for dDG in supporting context based on the geometry of the environment, it has been suggested that the hippocampus also processes context based on background cues, such as colors (Anderson and Jeffery, 2003; Eacott and Norman, 2004). In order to test whether the dDG plays an important role in processing

of context based on colors and whether it is possible to generate a pattern separation effect based on different shades of black and white. Four shades of color boxes were constructed – white, light gray, dark gray and black and four shaded inserts were constructed to fit inside each box – white, light gray, dark gray and black. The inserts were constructed to cover half of each box, such that the two sides of each box had different shades of color. On the test day during the study phase dDG lesioned rats and controls were placed in one of the colored boxes, with a colored insert that was different than the color of the box for a 5 min period. Then the rat was placed into its home cage for ten minutes. During the 10 min retention interval either the insert or box color was changed to generate the context of the test phase. The study phase had a duration of 5 min. Context exploration was measured throughout the study and the test phase by the number of observed rearings each subject exhibited on the insert side and the box side (the side without the insert). Each rear constitutes removing both front paws from the ground. Each subject was tested for 12 days, using the number of rearings to measure novelty detection. Each experimental day was followed by a rest day or a 48 h break period, taking a total of 24 days for each subject to complete the task. Three context gradient levels were measured, taking the difference of the novel color (introduced during the test phase) and the familiar from the study phase. Level 1 context gradient was a slight color change–white to light gray, light gray to dark gray, dark gray to black, or the reverse of any of the former. Level 2 context was a medium color change – white to dark gray, light gray to black or the reverse of any of the former. Level 3 context was the maximum color change–white to black or the reverse. Each subject completed four trials at level 1, level 2 and level 3 context gradient transitions. These levels were used to establish the degree to which subjects noticed the context change. The results of the experiment indicates that control rats displayed a color pattern separation effect across levels of separation of context (shades of color), but the dDG lesioned rats did not display a color pattern separation across levels 1, 2, and 3 (Kesner and Musso, unpublished observations). To support the idea that rats in this task primarily process recognition for the color context, it can be shown that during the study phase there is a preference for the darker color for both groups, but during the test or recognition memory phase there is no preference for the darker color for both groups. Thus, the dDG plays an important role in context recognition for colors by reducing interference between shades of color/gray.

In additional studies it has been shown that dDG lesioned rats are impaired in recognizing changes in a visual scene which determined a correct response in a runway task, suggesting a problem with contextual information (Ahn and Lee, 2014). Furthermore, Neunuebel and Knierim (2014) showed that neural activity in the dDG reflected the degree of changes associated with angular rotation of cues, suggesting that dDG activity reflects a pattern separation process for visual cues.

3.1.7. A role for context in object recognition

It has been suggested that the hippocampus may provide contextual information by combining object and place information (Kesner, 2007; Rolls and Kesner, 2006). Previous research has shown that the dorsal dentate gyrus (dDG) subregion of hippocampus may play an important role in combining object and place information via conjunctive encoding (Kesner et al., 2012). In one of the most common paradigms to measure the importance of context associated with object recognition, two groups of animals are allowed to explore two similar objects during a study phase and then after a short delay (s to min) the animals are tested either in the same context but with of one of the two objects changed (labeled object recognition) or with one of the two objects changed and a context (a different color box or a different room) labeled (object-context recognition). Control animals explore the novel object more than the familiar in both the object and object-context situation. Hippocampal lesioned rats have no problems exploring the novel object more than the familiar object in the object recognition situation, but fail to explore the novel object more than the familiar object in the object-context recognition situation (Dellu et al., 1997; Mumby, 2001; O'Brien et al., 2006; Piterkin et al., 2008). Spanswick and Sutherland (2010) reported the same results following a loss of granule cells in the dDG following adrenalectomy. A different approach is to examine the effects of context in object recognition memory by using a black box (object recognition)) and using a clear box with available cues that define a spatial context (object-context recognition) (Dees and Kesner, 2013). Based on a 10 min retention interval between a study phase and a test phase, the results indicated that dDG lesioned rats are impaired when compared to controls in the object–context recognition test in the clear box. However, there were no reliable differences between the dDG lesioned rats and the control group for the object recognition test in the black box. Even though the dDG lesioned rats were more active in object exploration and activity within the boxes based on rearing responses, the familiarization gradients did not differ. These results suggest that the dDG lesioned rats are clearly impaired when there is an important contribution of context. In a more recent study Smith et al. (2014) showed that Ts65Dn/DnJ mice as a model of Down syndrome with extensive inhibition of dDG display in the object-context recognition paradigm the same results as reported by Dees and Kesner (2013).

3.1.8. Odor pattern separation and ventral dentate gyrus

It can be shown that the hippocampus is critical for pattern separation for olfactory cues when a mnemonic component is included. A recent report from Kesner and colleagues (Kesner et al., 2011) demonstrated that for a series of straight carbon chain alcohols that form an aliphatic series (i.e. the differences among chemicals is a linear series of carbons) (Cleland et al., 2002), the ventral, but not dorsal, hippocampus is critical for solving a task requiring the rat choose which of two odors was previously encountered during a working memory paradigm. Importantly, rodents performed more poorly on this task as the number of carbons separating the odors was reduced, suggesting increased olfactory interference. To verify there was not a deficit at the perceptual level, rats trained to discriminate odors, even as similar as one carbon apart, were able to do so long as the mnemonic contribution to task performance was minimized (i.e. no delay or memory demand required flexible use of the olfactory information).

To specifically assay pattern separation during this task, and more specifically the role of the ventral dentate gyrus (vDG) for odor processing, working memory and pattern separation for odor information was assessed in rats using a matching-to-sample for odors paradigm (Weeden et al., 2014). Odor separations of 1, 2, 3 or 4 were selected for each choice phase and represented the carbon chain difference between the study phase odor and the test phase odor. Once an animal reached a criterion of 80–90% correct across all odor separations based on the last 16 trials, rats received a control or vDG lesion and were retested on the task. On post-operative trials, there were no deficits at 15 s delay for either the controls or the vDG lesioned rats. However, when the delay was increased to 60 s rats with vDG lesions were significantly impaired at short odor separations, but performance of rats with ventral dentate gyrus lesions matched control rats at the largest odor based separation. The graded nature of the impairment and the significant linear improvement in performance as a function of increased separation illustrate a deficit in odor pattern separation. Based on these results, it was concluded that lesions of the vDG decrease the efficiency of odor based pattern separation, which results in impairments on trials with increased odor proximity and increased odor

similarity among working memory representations (Weeden et al., 2014). The data suggest that the ventral hippocampus, especially the vDG, but not the dorsal hippocampus, support pattern separation for odor information along the domain of carbon length in aliphatic series.

It is important to note that other forms of pattern separation do not involve the dDG subregion of the hippocampus. For example, temporal pattern separation is mediated by the dCA1 but not the dDG (Gilbert et al., 2001). Furthermore, hippocampal lesions including dDG do not produce a deficit for pattern separation of visual objects, reward values, or motor responses (Bussey et al., 2002; Gilbert and Kesner, 2002; Kesner and Gilbert, 2006). Instead, the ventral visual stream including the inferior temporal visual cortex performs pattern separation for visual objects (Rolls, 2008a, 2012b), the perirhinal cortex further subserves pattern separation for visual objects (Burke et al., 2011, 2010; Bussey et al., 2002; Gilbert and Kesner, 2003b), the orbitofrontal cortex (Rolls, 2014a) and amygdala (Gilbert and Kesner, 2002) subserve pattern separation for reward value, the caudate nucleus subserves pattern separation for motor responses (Kesner and Gilbert, 2006), and the pyriform cortex as well as ventral DG subserve pattern separation for odors (Chapuis and Wilson, 2012; Weeden et al., 2014).

3.1.9. Encoding vs retrieval

The theory postulates that the dentate/mossy fiber system is necessary for setting up the appropriate conditions for optimal storage of new information in the CA3 system (which could be called encoding); and that (especially when an incomplete cue is provided) the retrieval of information from CA3 is optimally cued by the direct entorhinal to CA3 connections.

In addition to the study of Lassalle et al. (2000) mentioned previously that the DG input into the CA3 is essential for encoding, but not retrieval of information, there is evidence from a study in which rats had 10 learning trials per day in a Hebb–Williams maze. Lee and Kesner (2004b) found that based on a within-day analysis DG lesions (but not lesions of the perforant path input to CA3) impaired the acquisition of this task, consistent with an encoding or learning impairment. However, when tested using a between days analysis, retrieval of what had been learned previously was not impaired by DG lesions, but was impaired by lesions of the perforant path input to CA3. This double dissociation is consistent with the hypothesis that the DG is important for optimal encoding, and the entorhinal to CA3 connections for optimal retrieval.

Further evidence for a DG mediation of spatial encoding comes from the observation that rats with DG lesions are impaired in learning the Morris water maze task when the start location varied on each trial (Nanry et al., 1989; Sutherland et al., 1983; Xavier et al., 1999). Under these conditions, spatial pattern separation may be at a premium, for different, partially overlapping, subsets of spatial cues are likely to be visible from the different starting locations. Further evidence consistent with the hypothesis that the DG is important in spatial learning (acquisition) is that rats with DG lesions are impaired on acquisition of spatial contextual fear conditioning (Lee and Kesner, 2004a).

However, we note that in the studies described in this section (encoding vs. retrieval), during acquisition of these tasks over a number of trials both the storage and the retrieval processes are likely to influence the rate of behavioral learning. Thus, even though in behavioral tasks the deficit can be described as an encoding deficit because it is apparent during initial learning, it is not easy to separate the actual underlying processes by which the DG may be important for setting up the representations for new learning in the CA3 system, and the perforant path to CA3 connections for initiating retrieval especially with a partial cue from CA3.

3.1.10. Conclusion on the dentate gyrus

In conclusion, based on the development of computational models of dDG and behavioral evidence based on dysfunction of dDG, it appears that the dDG mediates mnemonic processing of spatial information. The processes subserved by dDG include (a) the operation of conjunctive encoding of multiple sensory inputs, implying an integration of sensory inputs to determine a spatial representation, and (b) pattern separation of spatial (especially metric) information, involving the reduction of interference between similar spatial locations (c) pattern separation of context (d) importance of context in object recognition, and (e) temporal integration and remote memory and spatial pattern separation based in part on neurogenesis. In addition the vDG mediates odor pattern separation. We discuss evidence below that the CA1 but not the DG region may be involved in temporal pattern separation (Gilbert et al., 2001).

3.2. CA3 subregion of the hippocampus

In the model the CA3 system acts as an autoassociation system. This enables arbitrary (especially spatial in animals and likely language for humans as well) 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. The same system must be capable of fast (one-trial) learning if it is to contribute to forming new episodic memories, but it could also play an important role in encoding of new information requiring multiple trials. The same system is also predicted to be important in retrieval of hippocampus-dependent information when there is an incomplete retrieval cue. The same system has the property that it can maintain firing activity because of the associative recurrent collateral connections, and for this reason is likely to be important in hippocampus-dependent delay/short-term memory tasks. The CA3 recurrent collateral system could also make a contribution to the learning of sequences if it has some temporal asymmetry in its associative synapses, as described in Section 2.2.3.7. Tests of these proposed functions are described next.

3.2.1. Rapid encoding of new information

The theory proposes that associative learning in the CA3 recurrent collateral connections can occur rapidly, in as little as one trial, so that it can contribute to episodic memory. Evidence for one-trial learning in the hippocampus is provided by the studies to be discussed in more detail in Section 3.2.5 (on recall and arbitrary associations) (Day et al., 2003; Kesner et al., 2007; Rolls and Xiang, 2006), with the latter two studies implicating in particular the CA3 hippocampal subregion. In these studies, the one-trial learning was between objects or odors and places.

The situation is a little different in a task in which there is one-trial learning of just a spatial location. In one such task, delayed non-matching-to-place (DNMP) for a single spatial location in an 8-arm maze, Lee and Kesner (2002, 2003) showed that blockade of CA3 NMDA receptors with AP5 or CA3 (neurotoxic) lesions did not impair one-trial learning in a familiar environment, but did impair it in an initial period in a novel environment. The implication is that a place can be remembered over a 10 s delay even without the CA3, but that learning of new spatial environments does require the CA3 subregion. Consistent with this observation is the finding that hippocampal CA3 (dorsal plus ventral) lesions disrupt learning of the standard 8-arm maze where the rats have to learn not to go back to arms previously visited (Handelmann and Olton, 1981). Rats tend to use different sequences on every trial, thus requiring new learning on every trial. When learning a new spatial environment, associations may be made between the different visual cues and their spatial locations, and this type of learning is computationally

similar to that involved in object–place or odor–place association learning, providing a consistent interpretation of CA3 function.

Nakazawa et al. (2003) also reported similar findings with mice in which the function of CA3 NMDA receptors was disrupted. These mice were impaired in learning a novel platform location in a working memory water maze task, whereas they were normal in finding familiar platform locations.

Evidence that rapid learning is reflected in the responses of CA3 neurons comes from a study in which it was found that CA3 neurons alter their responses rapidly when rats encounter novel configurations of familiar cues for the first time (Lee et al., 2004a). Specifically, rats were trained to run clockwise on a ring track whose surface was composed of four different texture cues (local cues). The ring track was positioned in the center of a curtained area in which various visual landmarks were also available along the curtained walls. To produce a novel cue configuration in the environment, distal landmarks and local cues on the track were rotated in opposite directions (distal landmarks were rotated clockwise and local cues were rotated counterclockwise by equal amounts).

In sum, these results strongly suggest that rapid plastic changes in the CA3 network are essential in encoding novel information involving associations between objects and places, odors and places, or between landmark visual cues and spatial locations, and NMDA receptor-mediated plasticity mechanisms appear to play a significant role in the process.

3.2.2. The types of information associated in CA3

A prediction is 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. Arbitrary in this context refers to the possibility of what is provided for by non-localized connectivity within the CA3–CA3 system, in that any one representation in CA3 can potentially become associated with any other representation in CA3. The prediction is neutral with respect to what is actually represented in the CA3 system, which we note from the effects of lesions almost always includes space, at least in animals, but could include language in humans. The object representations might originate in the temporal cortical visual areas, and the spatial representations might utilize some information from the parietal cortex such as vestibular inputs related to self-motion (Rolls, 1996b) (see Section 2.1.4). Two important issues need to be clarified. First, is the hippocampus the only neural system that supports all types of stimulus–stimulus associations (Eichenbaum and Cohen, 2001)? Based on the observations that the primate orbitofrontal cortex is involved in some types of stimulus–stimulus association, including visual-to-taste and olfactory-to-taste association learning and more generally in rapid and reversible stimulus–reinforcer association learning (Grabenhorst and Rolls, 2011; Rolls, 2004, 2014a,b; Rolls and Grabenhorst, 2008), there appear to be other neural systems that can learn stimulus–stimulus associations. Second, what is the nature of the information that can be associated in this arbitrary way in the CA3 recurrent collateral system? The studies described next indicate that at least the non-human hippocampus is involved in associations when one of the components is spatial.

In order to directly test the involvement of the CA3 subregion of the hippocampus in spatial paired-associate learning, rats were trained on a successive discrimination go/no-go task to examine object–place and odor–place paired associate learning. Normal rats learn these tasks in approximately 240 trials. CA3 lesions impaired the learning of these object–place and odor–place paired associations (Gilbert and Kesner, 2003a). Furthermore, lesions of DG or CA1 did not produce a deficit, and the reasons for this are considered later.

In contrast, object–odor paired associate learning was not impaired by CA3 lesions (even with a delay between the object and the odor) (Kesner et al., 2005). Thus, not all types of paired associate learning require the rat hippocampus. This conclusion is supported by the fact that (large or total) hippocampal lesions do not impair the following types of association including odor–odor (Bunsey and Eichenbaum, 1996; Li et al., 1999), odor–reward (Wood et al., 2004), auditory–visual (Jarrard and Davidson, 1990), and object–object associations (Cho and Kesner, 1995; Murray et al., 1993). Thus, in rats paired associate learning appears to be impaired by CA3 lesions only when one of the associates is place. It is important to note that lesions of the parietal cortex and pre-limbic and infralimbic also disrupt the acquisition of object–place paired associate learning, suggesting that cortical regions can also associate objects with places.

3.2.3. Acquisition/encoding associated with multiple trials

The CA3 subregion also appears to be necessary in some tasks that require multiple trials to acquire the task, and a common feature of these tasks is that a new environment must be learned. For example, lesions of the CA3 (but not the CA1) subregion impair the acquisition of object–place and odor–place paired associate learning, a task that requires multiple trials to learn (Gilbert and Kesner, 2003a). Another multiple trial task in which a learning set is needed that is affected by CA3, but not by CA1, lesions is acquisition of a non-match to sample 1-trial spatial location task on an 8-arm maze with 10 s delays (DNMP). In this case the output from CA3 is via the fimbria, in that fimbria lesions that interrupt the CA3 hippocampal output without affecting the input also impair acquisition of this task. In another task that requires multiple trials, namely the Hebb–Williams maze, CA3, but not CA1, lesions impair within-day learning (encoding) and CA1, but not CA3, lesions impair retention retrieval across days. Finally, CA3 lesioned rats are impaired in the standard water maze task which requires multiple trials (Brun et al., 2002; Florian and Roullet, 2004) (although mice that lacked NMDA receptors in CA3 do not appear to be impaired in learning the water maze (Nakazawa et al., 2002)). These findings are consistent with the hypothesis that when a new environment is learned this may involve multiple associations between different environmental cues, both with each other and with idiothetic (self-motion cues), and that this type of learning depends on new associations formed in the CA3–CA3 connections (see Section 2.2.3.2 on learning in continuous attractor networks, and Section 2.2.3.3 on learning associations between continuous and discrete representations).

It appears that in tasks that require a large number of trials to learn such as the acquisition of the delayed non-matching to place task and probably the acquisition of the spatial tasks mentioned above, the object–place learning task, and the Hebb–Williams maze task, the CA3 and fimbria output are essential, but the CA1 does not play an important role. The fimbria output would enable CA3 to influence the lateral septal area, which could in turn influence the medial septum. Any alteration of medial septum functionality would alter cholinergic input to the hippocampus (most of which comes from the medial septum), and this would be expected to influence, probably detrimentally, LTP in the hippocampus (Hasselmo et al., 1995), and thus acquisition. Indeed, blockade of CA3 cholinergic receptors with scopolamine disrupts acquisition of the Hebb–Williams maze (Rogers and Kesner, 2003) and fear conditioning that is specific to the spatial context (Rogers and Kesner, 2004). However, it is possible that septal dysfunction disrupts the responsiveness of hippocampal neurons to environmental inputs and thereby altering attention processes which in turn could alter memory function (Vinogradova et al., 1996).

Fig. 8.

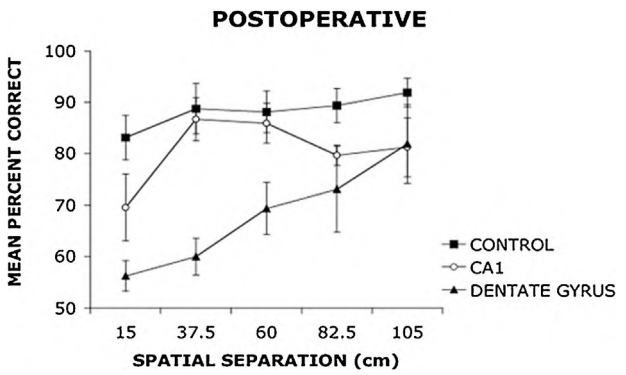

Fig. 8. Pattern separation impairment produced by dentate gyrus lesions. Mean percent correct performance as a function of spatial separation of control group, CA1 lesion group, and dentate gyrus lesion group on postoperative trials. A graded impairment was found as a function of the distance between the places only following dentate gyrus lesions.
After Gilbert et al. (2001).

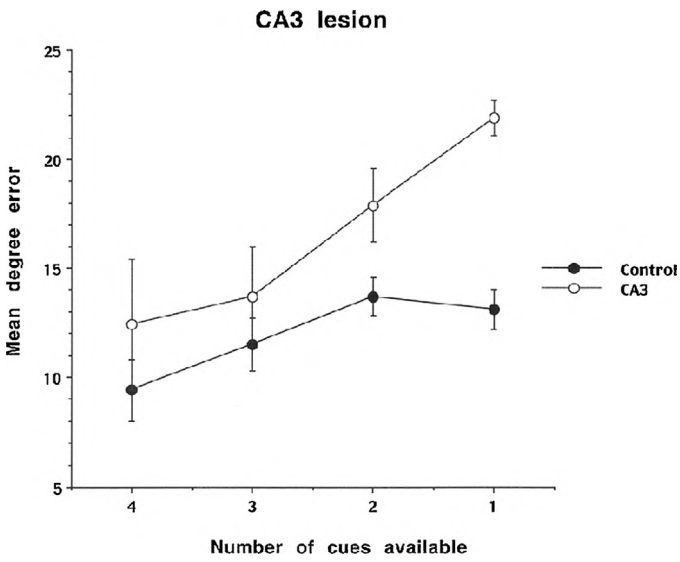

Fig. 9. Pattern completion impairment produced by CA3 lesions. The mean (and sem) degree of error in finding the correct place in the cheeseboard task when rats were tested with 1, 2, 3 or 4 of the cues available. A graded impairment in the CA3 lesion group as a function of the number of cues available was found. The task was learned in the study phase with the 4 cues present. The performance of the control group is also shown.
After Gold and Kesner (2005).

3.2.4. *Pattern completion*

The CA3 system is predicted to be important in the retrieval of hippocampus-dependent information when there is an incomplete retrieval cue. Support for the pattern completion process in CA3 can be found in lesion studies. 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 second delay, the animal needs to find the same food well in order to receive reinforcement with the object now removed. After reaching stable performance in terms of accuracy to find the correct location, rats received lesions in CA3. During post-surgery testing, four extramaze cues were always available during the sample phase. However, during the test phase zero, one, two, or three cues were removed in different combinations. The results indicate that controls performed well on the task regardless of the availability of one, two, three, or all cues, suggesting intact spatial pattern completion. Following the CA3 lesion, however, there was an impairment in accuracy compared to the controls especially when only one or two cues were available, suggesting impairment in spatial pattern completion in CA3-lesioned rats (Gold and Kesner, 2005) (see Fig. 9). A useful aspect of this task is that the test for the ability to remember a spatial location learned in one presentation can be tested with 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. Based on the observation that the lateral perforant path input into the dentate gyrus mediates visual object information via activation of opioid receptors, it was shown that naloxone (a μ opioid receptor antagonist) injected into the CA3 produced the same pattern completion problem (Kesner and Warthen, 2010).

In a different study Vazdarjanova and Guzowski (2004) placed rats in two different environments separated by ~30 min. The two environments differed greatly in that different objects were located in each room. The authors were able to monitor the time course of activations of ensembles of neurons in both CA3 and CA1, using a new immediate-early gene-based brain-imaging method (*Arc/H1a* catFISH). When the two environments were only modestly different, CA3 neurons exhibited higher overlap in their activity between the two environments compared to CA1 neurons. In another study (Lee et al., 2004b) recorded from ensembles of neurons using multiple electrodes in both CA1 and CA3 in freely behaving animals. By altering cue configuration in a ring track, Lee and colleagues demonstrated that the population spatial code in CA3 was less disrupted by modestly altered cue configurations. However, the population representation of space in CA1 was more easily disrupted by such moderately altered cue configurations. This can be considered as a pattern completion process in CA3, since the CA3 network maintained similar spatial representation of the environment (compared to the original spatial representation of the familiar cue configurations) even though the relationships among cues in the environment were altered.

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.

3.2.5. *Recall and arbitrary associations*

A prediction of the theory is 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, Langston and Morris (2003) trained rats in a study phase to learn in one trial an association between two flavors of food and two spatial locations. During a recall test phase they were presented with a flavor 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.

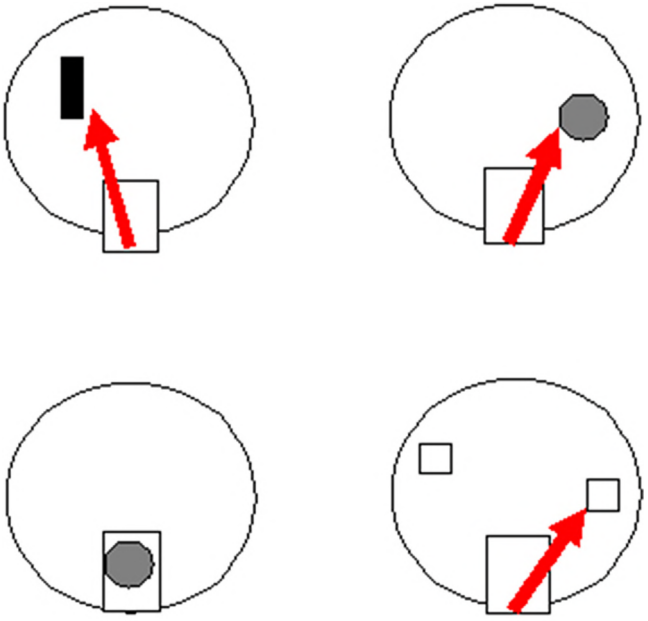

Fig. 10. Object-cued spatial location recall. Each shape represents a different object and the open squares represent spatial locations covered by neutral blocks. Each trial consisted of two spatial locations based study phases followed 15 s later by an object cue and 10 s later by a test between two previously experienced locations.
After Kesner et al. (2008).

In a development related to the Day et al. (2003) experiment, Kesner et al. (2008) designed a task in which a visual object could be recalled from a spatial location, and in an important extension, a spatial location could be recalled from an object. In this task, after training to displace objects, rats in the study phase are placed in the start box and when the door in front of the start box is opened the rats are allowed to displace one object in one location, and then after returning to the start box, the door is opened again and the rats are allowed to displace a second object in another location. There are 50 possible objects and 48 locations. In the test phase, the rat is shown one object (first or second randomized) in the start box as a cue, and then, after a 10-s delay, the door is opened and the rats must go to the correct location (choosing and displacing one of two identical neutral objects). The rats receive a reward for selecting the correct location that was associated with the object cue. (The task is illustrated in Fig. 10.) A spatial location-recall for a visual object task has also been developed. For the spatial recall for a visual object task, the study phase will be the same, but in this case in the test phase when the door is opened the rat is allowed to displace a neutral object in one location (first or second randomized) on the maze as a location cue, return to the start box, and then, after a 10-s delay, the door is opened and the rats must select the correct object (choosing and displacing one of two visual objects). The rats received a reward for selecting the correct visual object that was associated with the location cue. (The task is illustrated in Fig. 12.) Rats learn each task with 75% or better accuracy. Results indicate that CA3a,b lesions produce chance performance on both the one-trial object–place recall task, and on the one-trial place–object recall task (see Figs. 11 and 13) (Kesner et al., 2008) The potential implications of such results are that indeed the CA3a,b supports arbitrary associations as well as episodic memory based on one-trial learning. A control fixed visual conditional to place task with the same delay was not impaired, showing that it is recall after one-trial (or rapid) learning that is impaired.

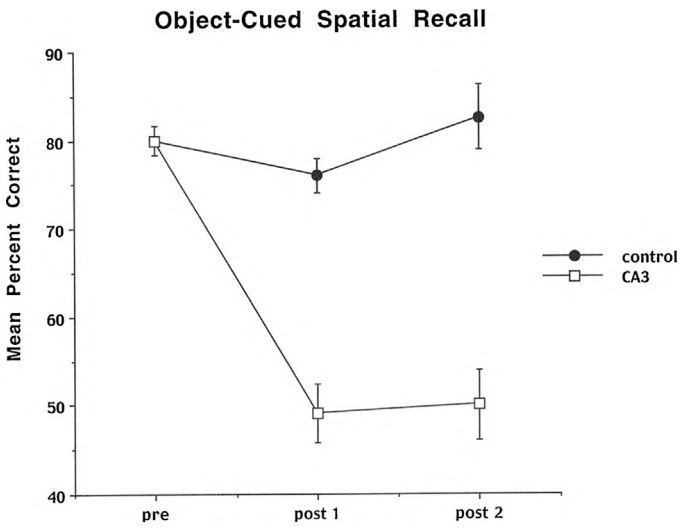

Fig. 11. Mean (and sem) percent correct performance for the control and CA3-lesioned rats on the object-cued spatial location recall task before (pre) and after (post 1 and post 2) surgery. The CA3 lesioned rats performed at chance levels after surgery.
After Kesner et al. (2008).

Additional support comes from the finding that in a similar one-trial object–place learning followed by recall of the spatial position in which to respond when shown the object, Rolls and colleagues (Rolls and Xiang, 2006; Rolls et al., 2005b) showed that some primate hippocampal (including CA3a,b) neurons respond with increased activity in the correct spatial location during and after the recall object cue. Thus, some hippocampal neurons appear to reflect spatial recall given an object recall cue. These data are consistent with the prediction of the standard computational model which

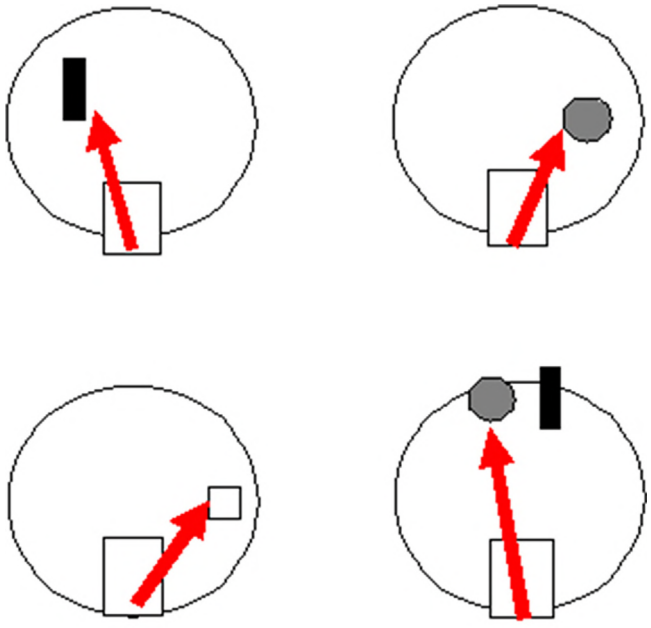

Fig. 12. Spatial location-cued object recall. Each shape represents a different object. The open square represents a neutral block placed on the correct location as a cue. Each trial consisted of two object-based study phases followed 15 s later by a spatial location cue and 10 s later a test between two previously experienced objects.
After Kesner et al. (2008).

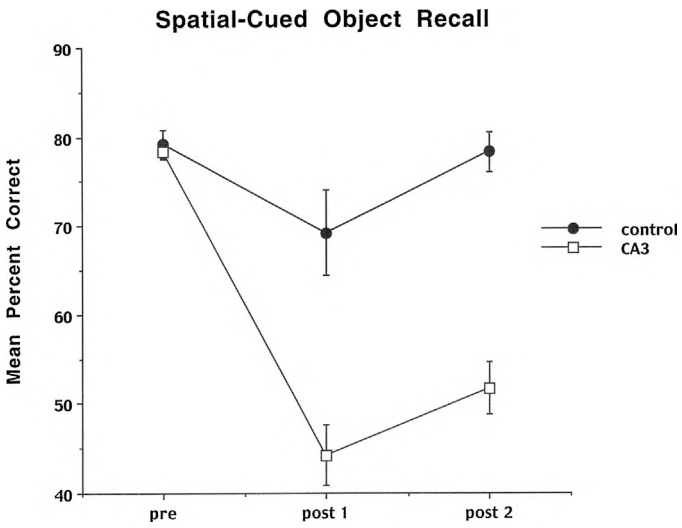

Fig. 13. Mean (and sem) percent correct performance for the control and CA3-lesioned rats on the spatial location-cued object recall task before (pre) and after (post 1 and post 2) surgery. The CA3 lesioned rats performed at chance levels after surgery.

After Kesner et al. (2008).

emphasizes the importance of CA3a,b in mediating the development of arbitrary associations.

There is anatomical support for CA3a,b involvement in support of the mediation of associative processes including arbitrary associations. The perforant path from the entorhinal cortex can be divided into a medial and lateral component. It has been suggested that the medial component processes spatial information and that the lateral component processes non-spatial (e.g. objects, odors) information (Hargreaves et al., 2005; Witter et al., 1989a). In one study, Ferbinteanu et al. (1999) showed that lesions of the medial perforant path disrupted water maze learning, whereas lateral perforant path lesions had no effect. In a more recent study based on the idea that the medial perforant path input into the CA3 mediates spatial information via activation of NMDA receptors and the lateral perforant path input into the CA3 mediates visual object information via activation of opioid receptors, the following experiment was conducted. Using a paradigm developed by Poucet (1989), rats were tested for the detection of a novel spatial change and the detection of a novel visual object change under the influence of direct infusions of AP5 (an NMDA antagonist) or naloxone (a μ opioid antagonist) into the CA3a,b. The results indicated that naloxone or AP5 infusions into the CA3a,b disrupted both novelty detection of a spatial location and a visual object (Hunsaker et al., 2007b). Based on electrophysiological data, there is associative long-term potentiation (LTP) between the medial or lateral perforant path and the intrinsic commissural/associational–CA3 synapses, demonstrated by the finding of an associative (cooperative) LTP between the medial and lateral perforant path inputs to the CA3 neurons (Martinez et al., 2002). This could provide a mechanism for object (via lateral perforant path)–place (via medial perforant path) associative learning, with either the object or the place during recall activating a CA3 neuron. Either place or object recall cues could thus be introduced by the associative medial and lateral perforant path connections to CA3 cells.

3.2.6. *Functional analysis of mossy fiber vs. perforant path inputs into CA3a,b*

The computational model (see Sections 2.2.2 and 2.2.3) suggests that the dentate granule cell/mossy fiber pathway to CA3 may be important during the learning of new associations in the CA3 network, and that part of the way in which it is important is that it helps by pattern separation to produce relatively sparse and orthogonal representations in CA3. In contrast, the theory predicts that the direct perforant path input to CA3 is important in initiating retrieval from the CA3 autoassociation network, especially with an incomplete retrieval cue.

In order to test the above mentioned ideas, a Hebb–Williams maze was used. In this task, rats were required to traverse a maze from the start to goal box in the most direct possible path. Learning across the final five trials of the first day compared to the first five trials of the first day of training was used as an index of encoding, and performance on the first five trials of the second day compared to the last five trials of the first day was used as an index of retrieval. Using this task, it was found that lesions of the DG or CA3a,b (or a crossed lesion) disrupt within-day learning on the Hebb–Williams maze, but that retrieval of information at the start of the following day is not impaired (Jerman et al., 2006; Lee and Kesner, 2004b). In contrast, lesions of the perforant path input to CA3a,b from entorhinal cortex disrupt retrieval (i.e. initial performance on the following day), but not learning within a day (Lee and Kesner, 2004b). These findings support the hypothesis that the CA3 processes both inputs from the DG to support pattern separation and inputs from the entorhinal cortex via the perforant path to support pattern completion. However, we note that in the studies described in this section (encoding vs. retrieval), during acquisition of these tasks over a number of trials both the storage and the retrieval processes are likely to influence the rate of behavioral learning. Thus, even though in behavioral tasks the deficit can be described as an encoding deficit because it is apparent during initial learning, it is not easy to separate the actual underlying processes by which the DG may be important for setting up the representations for new learning in the CA3 system, and the perforant path to CA3 connections for initiating retrieval especially with a partial cue from CA3.

The perforant path can be divided into a medial and lateral component. It has been suggested that the medial component processes spatial information and that the lateral component processes nonspatial (e.g. objects, odors) information (Hargreaves et al., 2005; Witter et al., 1989a). In one study Ferbinteanu et al. (1999) showed that lesions of the medial perforant path disrupted water maze learning, whereas lateral perforant path lesions had no effect. As predicted by the theory, there is associative LTP between the medial or lateral perforant path and the intrinsic commissural/associational–CA3 synapses) (Martinez et al., 2002). Either place or object recall cues could thus be introduced by the associative MPP and LPP connections to CA3 cells. This is consistent with the finding that disruption of the perforant path input impairs the multiple trial acquisition of an object–place paired associate task (perhaps because retrieval of what has been previously learned is impaired).

In addition, Martinez et al. (2002) demonstrated associative (cooperative) LTP between the medial and lateral perforant path inputs to the CA3 neurons. This could provide a mechanism for object (LPP)–place (MPP) associative learning, with either the object or the place during recall activating a CA3 neuron. However, this LPP to MPP cooperative LTP onto a CA3 neuron would be a low capacity memory system, in that there are only 3300 PP inputs to a CA3 cell, compared to 12,000 recurrent collaterals in the CA3–CA3 connections, and these numbers are the leading term in the memory capacity of pattern association and autoassociation memory systems (Rolls, 2008a; Rolls and Deco, 2002; Rolls and Treves, 1998).

We note that disruption of DG and mossy fiber input into CA3 do not produce a disruption in the acquisition of the object–place paired associate task (Gilbert and Kesner, 2003a) unless the stimuli are close together, implying that the DG contribution is important particularly when pattern separation is needed. The implication is

that sufficient input for object–place learning can be introduced into the CA3 system (which is required for this object–place learning) by the perforant path inputs provided that spatial pattern separation is not at a premium (or perhaps the storage of large numbers of different object–place associations is not at a premium, as this is the other condition for which the computational theory indicates that the DG is needed to help produce sparse and orthogonal representations in the CA3 system).

Furthermore, with respect to short-term memory, the DG and mossy fiber input into CA3 are not directly involved. This is illustrated by the observation that on a spatial pattern separation task, the performance of DG lesioned rats increased as a function of increased spatial separation between the target location and the foil location, whereas CA3 lesions were impaired across all spatial separations, indicating a potential impairment in short-term memory for spatial location information (Gilbert and Kesner, 2006; Gilbert et al., 2001).

Thus, from an anatomical perspective possible dissociations between DG and mossy fiber input into CA3 are primarily due to the observation that the CA3 subregion of the hippocampus has two major inputs with a direct connection from the DG via the mossy fibers and a direct input from the perforant path which bypasses the DG. In addition, short-term memory may be generated intrinsically in the CA3 region, but not in the DG.

3.2.7. *Cholinergic modulation of CA3 toward storage vs recall including pattern completion*

A circuit between the hippocampus and medial septum/diagonal of Broca has been characterized as a pathway from the hippocampus to the lateral septum (Gaykema et al., 1991; Raisman et al., 1966; Swanson and Cowan, 1977, 1979; Wyss et al., 1980). Briefly, CA3 subcortical efferents in the fimbria terminate in the lateral septum, medial septum, and diagonal band of Broca that result in net medial septum/diagonal band of Broca inhibition. CA1 subcortical efferents in the dorsal fornix also terminate in the lateral septum, medial septum, and diagonal band of Broca that result in net medial septum/diagonal band of Broca excitation. These differential effects occur because the CA3 and CA1 projections synapse onto distinct neuron populations. The medial septum and diagonal band of Broca send cholinergic efferents via the fimbria into the hippocampus that have been implicated in the hippocampal theta rhythm and modulation of learning and memory (Hasselmo, 1999, 2005; Hasselmo et al., 2002a; Hasselmo and Bower, 1993; Hasselmo and Fehlau, 2001; Hasselmo and Giocomo, 2006; Hasselmo et al., 2002b; Hasselmo and McGaughy, 2004; Hasselmo and Schnell, 1994; Hasselmo et al., 1995, 1996; McLennan and Miller, 1974; McNaughton and Miller, 1986; Rawlins et al., 1979).

Critically, these studies also demonstrated that acetylcholine in the hippocampus acts presynaptically by inhibiting glutamate release, presumably through M4 muscarinic acetylcholine receptor activation primarily at synapses in the stratum radiatum. GABAergic inputs are involved in hippocampal function as well, but act too rapidly to modulate encoding and consolidation/retrieval as operationally defined in most studies (i.e. behavioral encoding and retrieval are measured in seconds or minutes as opposed to the millisecond time course of the GABAergic modulation within theta phase precession and similar processes (cf. Wallenstein and Hasselmo, 1997)). It is suggested that these GABAergic projections are potentially involved in dynamics at millisecond timescales such as during theta cycles, and acetylcholine regulates these processes at longer time scales, such as the second to minute timescales encountered during cognitive testing. In CA1 and CA3, acetylcholine has a more robust inhibitory effect in stratum radiatum than stratum oriens, stratum lacunosum-moleculare or stratum lucidum in slice preparations. In the DG, there are dense cholinergic projections into the inner molecular layer as well as the hilus, and acetylcholine disinhibits granule neurons in the DG by reducing tonic inhibition.

The CA3 recurrent collaterals terminate primarily within the stratum radiatum, whereas the perforant path inputs terminate in the stratum lacunosum-moleculare, and the mossy fibers terminate largely in the stratum lucidum with additional synapses onto thorny excrescences in the stratum oriens and stratum radiatum. Despite a partial overlap in connectivity among these input pathways, this pattern of connectivity suggests that the mossy fiber pathway (at least synapses in the stratum oriens and stratum lucidum) and perforant pathway inputs (in the stratum lacunosum-moleculare) are not as dramatically affected by acetylcholine influx as the recurrent collateral and Schaffer collateral inputs (which are presynaptically inhibited by acetylcholine via M4 receptors). These data suggest that acetylcholine modulates the hippocampus primarily within the recurrent collaterals and Schaffer collaterals in the stratum radiatum by reducing synaptic transmission (Hasselmo and Schnell, 1994; Hasselmo et al., 1995, 1996).

An important additional effect of elevated acetylcholine levels in the hippocampus is the net result on DG granule neurons. Bilkey and Goddard (1985) demonstrated that activations of septal projections into the hippocampus resulted in disinhibition of granule neurons through inhibition of inhibitory interneurons that provide tonic inhibition to the granule neurons (i.e. disinhibition). Medial septal stimulation at magnitudes insufficient to result in evoked responses facilitated population spikes in the DG to medial perforant path stimulation in vivo. In other words, the tonic inhibition on the granule neurons from the inhibitory neurons in the hilus typically attenuating or shunting the activity levels of granule neurons is reduced by cholinergic influx – thereby increasing the levels of responsiveness to stimuli of the DG granule neurons under the same conditions that the influence of the recurrent collaterals in CA3 are minimized, supporting cholinergic models of encoding/retrieval dynamics in the hippocampus (cf. Hasselmo et al., 1996). Alternately, any reduction from the baseline levels of cholinergic influence in the DG would result in a net increase in inhibitory tone on granule neurons. This mechanism is important for models of hippocampal encoding and retrieval that require a relatively quiescent mossy fiber input to allow retrieval processes to be initiated via the perforant path projections from the entorhinal cortex (Rolls, 1996b, 2008a, 2010b, 2013b,c; Treves and Rolls, 1992, 1994).

It is important to note that the proposed mechanisms whereby the dentate gyrus engages in pattern separation benefits from a sparse representation, and from sufficient activation to produce some firing. The reduction of tonic inhibition in the DG produced by an increase in acetylcholine may help. The recruitment of CA3c via the mossy cells in the hilus which influence dentate granule cells (Myers and Scharfman, 2009, 2011; Scharfman and Myers, 2012) may be relatively unaffected by acetylcholine as there are few recurrent collaterals in CA3c in the stratum radiatum). Overall, the effect of acetylcholine is to increase the responsiveness to stimuli of the dentate granule cells, which in concert with the reduced efficacy of the CA3 recurrent collateral synapses would facilitate information transfer from the mossy fibers to the CA3 pyramidal cells and help to produce a pattern separation effect beneficial for storage in the CA3–CA3 recurrent collateral system (Giocomo and Hasselmo, 2007).

It has been demonstrated that disrupting CA3a,b sub-cortical efferents in the fimbria disrupts encoding, but not retrieval, of spatial information during learning of the Hebb–Williams maze task (Hunsaker et al., 2008c). Rats with CA3a,b outputs to the septal nuclei via the fimbria disrupted were impaired for encoding or within-day learning, but showed intact abilities to perform between day retrieval. Rats with CA1 outputs to the septal nuclei via the dorsal fornix/alveus disrupted showed intact abilities to perform within-day encoding but were impaired for between day

retrieval. Importantly, for these tasks the cholinergic inputs to the hippocampus via the fimbria were intact.

It is suggested that this effect is due to a critical role for cholinergic inputs to the hippocampus in biasing CA3a,b to pattern separation via reducing the recurrent collateral influence relative to that of the mossy fibers, since a disruption of these cholinergic projections would disrupt the ability of CA3a,b to encode spatial information (Hunsaker et al., 2007a, 2007b, 2008c, 2009; Hunsaker and Kesner, 2013). The consequence of increasing the cholinergic projections to CA3a,b would be a decrease in the synaptic transmission of the recurrent collaterals and thus less recurrent activation of CA3 relative to the mossy fiber inputs. Also, when acetylcholine levels increase in the DG, the inhibitory interneurons in the hilus are inhibited – thus reducing the amount of tonic inhibition on the granule neurons. The effect of this disinhibition is to increase mossy fiber neurotransmission (assuming increased transmission follows from increased responsiveness to stimulation), further biasing CA3a,b toward encoding information from the DG over recurrent activity from the CA3a,b recurrent collaterals.

Reducing acetylcholine by attenuating excitation of the medial septum/diagonal band of Broca cholinergic neurons would favor pattern completion and recall over pattern separation and encoding and push CA3a,b into the role of a working memory buffer by reducing the signal to noise level, and thus providing the best environment for the detection and completion of patterns from partial cues (Kesner and Rolls, 2001). As mentioned above, concurrent with the effects on CA3, reducing acetylcholine would reduce inhibitory tone on inhibitory interneurons in the hilus–thus increasing inhibition on the DG granule neurons, and subsequently reducing mossy fiber inputs to CA3a,b. Furthermore models also support this view by demonstrating that decreased acetylcholine levels in a septo-hippocampal classical conditioning learning model result in overall reductions in the learning rate (Myers et al., 1996).

The pattern of results for scopolamine (a cholinergic antagonist) and physostigmine (a cholinergic agonist) infusions into CA3a,b is intriguing. Scopolamine, but not physostigmine, infusions, into CA3a,b disrupt encoding. In contrast, physostigmine, but not scopolamine infusions, into CA3 a,b, disrupt recall. This was observed during a spatial exploration paradigm (Hunsaker et al., 2007b), during Hebb–Williams maze learning (Rogers and Kesner, 2003), and during delay fear-conditioning (Rogers and Kesner, 2004). In another study, Marques Pereira et al. (2005) showed that injections of 192IgG-saporin in the hippocampus, which produced a severe depletion of cholinergic cells in the medial septum, was sufficient to disrupt acquisition of the Hebb–Williams maze. These data suggest that cholinergic levels directly influence information processing in the hippocampus, supporting the assertion that acetylcholine levels are perfectly located in time and place to bias CA3a,b toward either pattern separation or pattern completion.

3.3. CA1 subregion of the hippocampus

The CA1 subregion of the hippocampus receives inputs from two major sources: the Schaffer collateral inputs from CA3, and the perforant path inputs from the entorhinal cortex (see Fig. 1). CA1 outputs are directed toward the subiculum, entorhinal cortex, prefrontal cortex and many other neural regions including the lateral septum, anterior thalamus and the mammillary bodies (Amaral and Witter, 1995). The anatomical and physiological characteristics suggest that the CA1 can operate as a competitive network and have triggered the development of computational models of CA1 in which for example the CA1 system is involved in the recall process by which backprojections to the cerebral neocortex allow neuronal activity during recall to reflect that present during the original learning (see Section 2.2.5). In this terminology, cortico-cortical backprojections originate from deep (layer 5) pyramidal cells, and terminate on the apical dendrites of pyramidal cells in layer one of the preceding cortical area (see Fig. 1). In contrast, cortico-cortical forward projections originate from superficial (layers 2 and 3) pyramidal cells, and terminate in the deep layers of the hierarchically next cortical area (see Fig. 1 and Rolls, 2008a). Within this context, the projections from the entorhinal cortex to DG, CA3 and CA1 can be thought of as forward projections; and from CA1 to entorhinal cortex, subiculum etc. as the start of the backprojection system (see Fig. 1).

The evidence reviewed next of the effects of selective damage to CA1 indicates that it makes a special contribution to temporal aspects of memory including associations over a delay period, and sequence memory. There is also evidence relating it to intermediate memory as well as consolidation. We note that the effects of damage to the CA1 may not be identical to the effects of damage to CA3, because CA1 has some inputs that bypass CA3 (e.g. the direct entorhinal/perforant path input); and CA3 has some outputs that bypass CA1. In particular, the CA3 output through the fimbria projects directly to the lateral septum and then to the medial septum or directly to the medial septum (Amaral and Witter, 1995; Gaykema et al., 1991; Risold and Swanson, 1997). The medial septum, in turn, provides cholinergic and GABAergic inputs to the hippocampus (Amaral and Witter, 1995).

3.3.1. Sequence memory

Sequence memory could be implemented in a recurrent collateral network such as CA3 with some temporal asymmetry in the associative synaptic modification learning rule (Section 2.2.3.7 and Rolls and Kesner, 2006). One behavioral test in which this type of memory may be needed is in complex maze learning, which is impaired by hippocampal including dorsal hippocampal and CA3 lesions. Maze learning could be implemented by pairwise associations between spatial views, or between responses if the task has become habit-based. However, it could also be implemented by associations between views and body responses (e.g. turns), and there might not be an explicitly sequential association involved in this case, as each body action would lead to the next view. Further, in the cases where complex maze learning is impaired by hippocampal lesions, it would remain to be shown that the deficit is due to sequence learning impairments rather than only to spatial learning deficits.

In another paradigm in which a fixed sequence is learned over multiple trials, it appears that pairwise sequential learning and not just temporal decay is required to account for the findings. In a sequential spatial location list learning paradigm, rats are trained by correction to visit the arms of an 8-arm maze in a particular sequence, with food obtained in each arm if a correct choice (made by orienting) is made. Even though lesions of the hippocampus impair the acquisition of this task (learned by control rats in 80 trials), there are no deficits when hippocampus lesions are made after reaching 90–100% correct performance (DeCoteau and Kesner, 2000). Thus the hippocampus may be necessary for learning new spatial sequences, but once learned, it appears that other brain systems can implement the behavior, perhaps using other egocentric or habit strategies.

A second experiment was designed to examine the effects of dorsal hippocampus, dorsal CA3, dorsal CA1, and control lesions on retention of a temporal sequence task (Hoang and Kesner, 2008). The rats were trained on the same sequential task for six spatial locations on a radial 8-arm maze. After initial training followed by surgery, it was found that all lesioned animals (dorsal hippocampus, dorsal CA3, dorsal CA1 or sham control) performed at very high levels for the original learned sequence when started from the beginning of the sequence. To test for temporal sequence completion, the animals were started at different positions in the previously learned sequence and expected to complete the

remainder of the sequence. The results indicate that control rats had no difficulty completing the sequence, regardless of starting point. In contrast, the rats with dorsal hippocampus, CA3a,b, or CA1 lesions were impaired and made many errors. These results suggest that the dorsal hippocampus and specifically the dorsal CA3a,b in conjunction with CA1 may mediate retrieval of temporal sequences via perhaps a temporal pattern completion process.

In a different task (Hunsaker et al., 2008a), rats learned trial unique sequences of spatial locations along a runway box. Each trial consisted of a study phase made up of the presentation of a linear sequence of four spatial locations marked by neutral blocks. After a 30 s interval the animal was given the test phase. The test phase consisted of the same sequence presented during the study phase, but one of the spatial locations was not marked by a block, but still contained a reward. The unmarked spatial location was pseudo-randomly distributed equally between the first, second, third, and fourth item in the sequence. To receive a reward, the rat had to visit the correct, unmarked spatial location. Once animals were able to reliably perform this short-term episodic memory task, they received lesions to either CA3 or CA1. Animals with lesions to either CA3 or CA1 had difficulty with short-term memory processing, although CA1 lesioned animals had a much greater deficit. However, when animals were trained on a non-episodic version of the same task, hippocampal lesions had no effect. The results suggest that the CA3 and CA1contribute to episodic based sequence pattern completion (Hunsaker et al., 2008a). The above data are consistent with the interpretation of episodic memory described in Section 2.1.2.

3.3.2. Associations across time

There is evidence implicating the hippocampus in mediating associations across time (Kesner, 1998; Rawlins, 1985). The subregion analyses described next show that CA1 lesions impair this. Computationally, a strong hypothesis would be that the CA3 system could provide the working memory necessary for hippocampus-dependent associations across time, and that the CA3 then influences the CA1 for this function to be implemented. The actual learning could involve holding one item active in CA3 but continuing firing in an attractor state until the next item in the sequence arrives, when it could be associated with the preceding item by temporally asymmetric synaptic associativity, as described in the model of Rolls and Stringer (see Rolls and Kesner, 2006). The computational suggestion thus is that associations across time could be implemented in the hippocampus by using the same functionality that may be used for sequence memory.

The CA1 subregion of the hippocampus may play a role in influencing the formation of associations whenever a time component (requiring a memory trace) is introduced between any two stimuli that need to be associated. Support for this idea comes from studies based on a classical conditioning paradigm. Lesions of the hippocampus in rabbits or rats disrupt the acquisition of eye-blink trace conditioning. In trace conditioning a short delay intervenes between the conditioned stimulus (CS) and the unconditioned stimulus (UCS). When, however, a UCS and CS overlap in time (delay conditioning), rabbits with hippocampal damage typically perform as well as controls (Moyer et al., 1990; Weiss et al., 1999). Based on a subregional analysis, it has now been shown that ventral, but not dorsal, CA1 lesions impair at least the retention (tested after 48 h) of trace fear conditioning (Rogers et al., 2006). Similar learning deficits in trace fear conditioning (but not in conditioning without a delay between the CS and UCS), were observed for mice that lacked NMDA receptors in both dorsal and ventral CA1 subregions of the hippocampus (Huerta et al., 2000). In an interesting study, McEchron et al. (2003) recorded from single cells in the CA1 region of the hippocampus during and after trace heart rate (fear) conditioning using either a 10 s or 20 s trace interval. They reported that a significant number of cells showed maximal firing on CS-alone retention trials timed to the 10 s or 20 s trace after CS offset. This could reflect the transition from an attractor state implemented in CA3 with its recurrent collaterals that represents the CS to an attractor state that represents the conditioned response (CR) (Deco and Rolls, 2003, 2005a; Rolls, 2008a; Rolls and Deco, 2002).

Based on a previous finding that the acquisition of an object–odor association is not dependent on the hippocampus, would adding a temporal component to an object–odor association task recruit the hippocampus and more specifically the CA1 region? To test this idea rats were given dorsal CA1, CA3, or control lesions prior to learning an object–trace–odor task. The task was run in a 115 cm linear box in which the rat was presented with an object for 10 s after which it was removed, followed by a 10 second trace period, and by the presentation of an odor 50 cm away. If the odor and the object were paired, the rat was to dig in the odor cup for a reward. If unpaired, the rat was to refrain from digging. Animals that had dorsal CA1 lesions were unable to make the association and never performed above chance, however CA3 lesioned rats displayed no deficit (Kesner et al., 2005). The observation of time cells during the trace in the object–trace–odor task supports the idea that the CA1 supports temporal processing of object and odor information. These results support the idea that the CA1 region is at least on the route for forming arbitrary associations across time even when there is no spatial component.

In a subsequent experiment rats with dorsal CA1, CA3 or control lesions were tested on an object–trace–place task (Hunsaker et al., 2006). In this case both CA1 and CA3 were impaired suggesting that CA3 contributes to temporal processing of information only when place is one of the components. Thus, it is likely that the dorsal CA3 is a source for timing information which contributes to spatial temporal processing. However, the dorsal CA3 does not appear to contribute timing information for object and odor information. A possible suggestion is that the ventral CA3 supports the processing of odor information which then could contribute timing information to CA1. Another possibility is that object and odor information can reach CA1 from the lateral entorhinal cortex and since timing cells are found in the medial entorhinal time and object as well as time and odor could form time and object or odor and object combination neurons.

3.3.3. Order memory

There are also data on order memory, which does not necessarily imply that a particular sequence has been learned, and could be recalled. Estes (1986) summarized data demonstrating that in human memory there are fewer errors for distinguishing items (by specifying the order in which they occurred) that are far apart in a sequence than those that are temporally adjacent. Other studies have also shown that order judgments improve as the number of items in a sequence between the test items increases (Banks, 1978; Chiba et al., 1994; Madsen and Kesner, 1995). This phenomenon is referred to as a temporal distance effect (sometimes referred to as a temporal pattern separation effect). It is assumed to occur because there is more interference for temporally proximal events than temporally distant events. Based on these findings, Gilbert et al. (2001) tested memory for the temporal order of items in a one-trial sequence learning paradigm. In the task, each rat was given one daily trial consisting of a sample phase followed by a choice phase. During the sample phase, the animal visited each arm of an 8-arm radial maze once in a randomly predetermined order and was given a reward at the end of each arm. The choice phase began immediately following the presentation of the final arm in the sequence. In the choice phase, two arms were opened simultaneously and the animal was allowed to choose between the two arms. To obtain a food reward, the animal had to enter the arm that occurred earliest in the sequence that it had just followed. Temporal separations of

0, 2, 4, and 6 were randomly selected for each choice phase. These values represented the number of arms in the sample phase that intervened between the two arms that were to be used in the test phase. After reaching criterion rats received DG or CA1 lesions or served as sham operated controls. Following surgery, control and DG lesioned rats matched their preoperative performance across all temporal separations. In contrast, rats with CA1 lesions performed at chance across 0, 2, or 4 temporal separations and a little better than chance in the case of a 6 separation. The results suggest that the CA1 subregion is involved in memory for spatial location as a function of temporal separation of spatial locations and that lesions of the CA1 decrease efficiency in temporal pattern separation. CA1 lesioned rats cannot separate events across time, perhaps due to an inability to inhibit interference that may be associated with sequentially occurring events. The increase in temporal interference impairs the rat's ability to remember the order of specific events.

Even though CA1 lesions produced a deficit in temporal pattern separation, some computational models (Levy, 1996; Lisman, 1999; Wallenstein and Hasselmo, 1997) and the model described above in Section 2.2.3.7 have suggested that the CA3 region is an appropriate part of the hippocampus to form a sequence memory, for example by utilizing synaptic associativity in the CA3–CA3 recurrent collaterals that has a temporally asymmetric component (Section 2.2.3.7). It was thus of interest to use the temporal order task to determine whether the CA3 region plays an important role in sequence memory. The CA3 lesions impaired the same sequence learning task described above that was also impaired by lesions of the CA1 region. However, this is likely to be the case only for spatial information.

The hippocampus is known to process spatial and temporal information independently (Kesner, 1998; O'Keefe and Nadel, 1978). In the previous experiments sequence learning and temporal pattern separation was assessed using spatial cues. Therefore, one possibility is that the CA1 and CA3 deficits were due to the processing of spatial rather than non-spatial temporal information. To determine if the hippocampus is critical for processing all domains of temporal information, it is necessary to use a task that does not depend on spatial information. Since the hippocampus does not mediate short-term (delayed match to sample) memory for odors (Dudchenko et al., 2000; Otto and Eichenbaum, 1992), it is possible to test whether the hippocampus plays a role in memory for the temporal sequence of odors (i.e. temporal separation effect). Memory for the temporal order for a sequence of odors was assessed in rats based on a varied sequence of five odors, using a similar paradigm described for sequences of spatial locations. Rats with hippocampal lesions were impaired relative to control animals for memory for all temporal distances for the odors, yet the rats were able to discriminate between the odors (Kesner et al., 2002). Fortin et al. (2002) reported similar results. In a further subregional analysis, rats with dorsal CA1 lesions show a mild impairment, but rats with ventral CA1 lesions show a severe impairment in memory for the temporal distance for odors. Thus, the CA1 appears to be involved in separating events in time for spatial and nonspatial information, so that one event can be remembered distinct from another event, but the dorsal CA1 might play a more important role than the ventral CA1 for spatial information, and conversely ventral CA1 might play a more important role than the dorsal CA1 for odor information. In a more recent experiment using a paradigm described by Hannesson et al. (2004), it can be shown that temporal order information for visual objects is impaired only for CA1, but not CA3 lesions (Hoge and Kesner, 2007). Thus, with respect to sequence learning or memory for order information one hypothesis that would be consistent with the data is that the CA1 is the critical substrate for sequence learning and temporal order or temporal pattern separation. This would be consistent with CA1 deficits in sequence completion of a spatial task and CA1 deficits in temporal order for spatial locations, odor and visual objects. Dorsal CA3 contributes to this temporal order or sequence process whenever spatial location is also important. When visual objects are used the CA3 region does not play a role. It is not know yet whether for odors the ventral CA3 plays an important role. It is predicted that it will not be involved.

We note that in these order tasks the animals are not required to recall the sequence (for example by retracing their steps). Instead, the animals are asked which of two items occurred earlier in the list. To implement this type of memory, some temporally decaying memory trace or temporally increasing memory trace via a consolidation process might provide a model (Marshuetz, 2005), and in such a model, temporally adjacent items would have memory traces of more similar strength and would be harder to discriminate than the strengths of the memory traces of more temporally distant items. (We assume that the animal could learn the rule of responding to the stronger or weaker trace.)

3.3.4. Intermediate memory

The following evidence suggests that the CA1 region is especially involved in an intermediate term memory. Rats with CA3, but not CA1 lesions were impaired in the acquisition of a delayed non-match to place task (in an 8 arm maze) with a 10 s delay. Also, when transferred to a new maze in a different room and using a 10 s delay, there was a deficit for CA3, but not CA1 lesioned rats. Deficits for CA1 emerged when rats were transferred to a 5 min delay. Comparable deficits at a 5 min delay were also found for CA3 lesioned rats. It should be noted that similar results as described for the lesion data were obtained following AP5 injections, which impaired transfer to the new environment when made into CA3 but not CA1. At 5 min delays AP5 injections into CA1 produced a sustained deficit in performance, whereas AP5 injections to CA3 did not produce a sustained deficit (Lee and Kesner, 2002, 2003).

Although the CA3 system is involved even at short delays in acquiring this spatial short-term memory task, it is not apparently involved in the task once it has been acquired (Lee and Kesner, 2003). The implication is that (after acquisition) the CA3 system is not required to operate as a short term attractor working memory to hold on line in the 10 s delay the place that has just been shown, although it could contribute to acquisition by operating in this way. The CA3 could also help in acquisition by providing a good spatial representation for the task, and in particular would enable the spatial cues in the environment to be learned. The suggestion therefore is that at the 5 min delay, the system must operate by associative connections from entorhinal cortex to CA1, which enable the sample place in the 8 arm maze to be remembered until the non-match part of the trial 5 min later, and the AP5 impairment after injection into CA1 is consistent with this.

Another implication is that the CA3 contribution to this task can be implemented partly independently of CA1. Consistent with this, lesions of the fimbria impair the acquisition of the delayed match to place task at a 10 s delay. With longer delays, the task may no longer be implemented by an attractor process in CA3, but may involve LTP to set up an episodic-like memory which appears in this case to depend especially on CA1. Thus, this evidence suggests two types of short term memory implementations in the hippocampus. One is a very short term (10 s) memory (typically involving space) that may be implemented in CA3. The other is an intermediate term (5 min) memory that appears to be a more episodic-like memory involving LTP in CA1. Also, it is likely that the direct entorhinal to CA1 connections may be of importance in that the CA1 lesion effect on the delayed non-match to place task at 5 min delays can also be produced by injections of apomorphine into CA1 which influence the entorhinal to CA1 connections.

Further evidence for some dissociation of an intermediate-term from a short-term memory in the hippocampus is that in the Hebb–Williams maze, learning/encoding measured within a day is impaired by CA3 and DG, but not by CA1 lesions. In contrast, retention/retrieval measured between days is impaired by CA1 but not by CA3 lesions. Thus, there may be an acquisition vs retention dissociation with CA3 vs CA1 lesions. Interestingly, the CA1 lesion effect on this task can also be produced by injections of apomorphine into CA1 which influence the entorhinal to CA1 connections, suggesting that on the following day when retrieval is at premium, the direct entorhinal to CA1 connections may play an important role.

In addition to the processing of temporal information, the CA1 subregion appears to have cellular processes that match it to longer term types of memory than CA3. To test the idea that the CA1 region might be involved in retrieval after long (e.g. 5 min, 24 h) time delays, rats with CA1 lesions were tested in a modified Hebb–Williams maze. The results indicate that CA1 lesioned rats are impaired in retrieval with a 24 h delay (across-day tests), but have no difficulty in encoding new information (i.e. within-day tests). Using a spatial contextual fear conditioning paradigm, rats with dorsal CA1 lesions are relative to controls impaired in retention (intermediate-term memory) of conditioning when tested 24 h following acquisition (Lee and Kesner, 2004a). Using the same paradigm Hall et al. (2000) have demonstrated a relationship between contextual fear memory and the expression of LTP-related immediate early genes (i.e. *BDNF* and *zif268*) in the CA1 30 min after the exposure to a contextual fear-conditioning situation. Hall et al. (2001) also demonstrated that 30 min following a test for contextual retrieval administered 24 h after fear conditioning resulted in an increase in *zif 268* detected in CA1. Strekalova et al. (2003) showed that following a single footshock, a test for contextual retrieval 48 hrs later revealed 90 min later an induction of c-Fos and JunB early genes in dorsal CA1. No changes were observed in the absence of a test of contextual retrieval. In addition, Gall et al. (1998) demonstrated that *c-fos* gene expression was elevated selectively more in CA1 compared to CA3 only after overtraining in an olfactory discrimination task, whereas the same gene expression in CA3 compared to that in CA1 was selectively enhanced after the initial acquisition of the olfactory discrimination task (see also Hess et al., 1995a, 1995b). In another study Bertaina-Anglade et al. (2000) showed that learning of an appetitive operant conditioning task in mice resulted in increased c-Fos expression in CA1 primarily as a function of improvement in performance between Day 1 and Day 2. As these early genes containing the cAMP response element (CRE) are essential for gene transcription and protein synthesis after induction of LTP, the enhanced expression of *BDNF* and *zif268*, c-Fos, and JunB also suggests that CA1 is a key subregion for the establishment of intermediate-term memory using cellular consolidation processes.

Support for assigning the above mentioned results to the processing of intermediate memory within CA1 comes from the observation that early gene changes in CA1 while observed on day 1 are not observed with longer term retention test at 5 days (Hall et al., 2001) or 28 days (Bertaina-Anglade et al., 2000).

In addition, it has been shown in a water maze study that post-training lesions of the CA1 24 h after learning disrupted subsequent recall, whereas the same lesion made 3 weeks later did not produce a disruptive effect (Remondes and Schuman, 2004). In another study injections of propranolol into CA1 to block beta adrenergic receptors 5 min after contextual fear conditioning disrupted subsequent retention, whereas the same injections had no effect when administered 6 h later (Ji et al., 2003).

How could memories become hippocampal-independent after long time delays? One possibility is that there is active recall to the neocortex which can then use episodic information to build new semantic memories, as described in Section 2.1.2 (McClelland et al., 1995; Rolls, 1989b,d, 1990a; Treves and Rolls, 1994). In this case, the hippocampus could be said to play a role in the consolidation of memories in the neocortex. We note that in another sense a consolidation process is required in the neocortex during the original learning of the episodic memory, in that we hypothesize that associative connections between hippocampal backprojections and neocortical neurons activated by forward inputs are formed (and presumably consolidated) during the original learning, so that later a hippocampal signal can recall neocortical representations. Another possibility is that there are more passive effects of what has been learned in the hippocampus on the neocortex, occurring for example during sleep (Wilson, 2002; Wilson and McNaughton, 1994). Indeed, given that the hippocampus is massively connected to the neocortex by backprojections as described above, it is very likely that any alteration of the functional connectivity within the hippocampus, for example the coactivity of place neurons representing nearby places in an environment which has just been learned, will have an influence on the correlations found between neocortical neurons also representing the places of events (Wilson, 2002; Wilson and McNaughton, 1994). In this case, there could be synaptic modification in the neocortex between newly correlated neuronal activity, and this would effectively produce a form of new learning ("consolidation") in the neocortex. Another sense in which consolidation might be used is to refer to the findings considered in this section that there is a time-dependency of the learning in CA1 that seems special, enabling CA1 to contribute to the formation of memories ("consolidation") with long time delays, and also, perhaps with a different time course, to their subsequent forgetting or reconsolidation. The term "consolidation" could thus be used to refer to many different types of computational process, which should be distinguished.

3.3.5. Interactions and dissociations between CA3 and CA1

The dominant view of the relationship between CA3 and CA1 and short-term and intermediate-term memory is that they operate as a feed-forward sequential processing system. The more recent data, however, suggest that, for certain tasks, there are dissociations between short-term vs. intermediate-term memory and between the involvements of the CA3 vs. CA1 subregions. Yet for a different set of tasks both the CA1 and CA3 interact in processing of short-term and intermediate-term memory. Any parallel processing between the CA3 and CA1 would imply possible independence for short-term and intermediate-term memory.

One result that has been obtained when one compares the relationships between CA3 and CA1 is that there is a deficit following dysfunction of the CA3 subregion, but no concomitant deficit following dysfunction of the CA1 subregion. For example, (A) short-term memory in a delayed (10 s) nonmatching-to-sample for a spatial location task is disrupted by lesions in CA3 or AP5 injections into the CA3, but not into the CA1 subregion (Lee and Kesner, 2002, 2003), (B) CA3, but not CA1, lesions impair within-day learning (encoding) in a Hebb–Williams maze (Jerman et al., 2006), (C) After a change in cue configuration the CA3 place fields shifted their locations backwards prominently for the first time, but not thereafter compared to the place fields in CA1 (Lee et al., 2004a) and when Arc is enhanced in CA3, but not CA1 when a rat is exposed to a single lap in a rectangular track (Miyashita et al., 2009). This double dissociation in the time course of plasticity between CA1 and CA3 place fields suggests that CA3 reacts first to any changed components in the environment, presumably to incorporate the novel components into an existing system or build a new representation of the environment if changes are significant.

These results suggest that short-term and intermediate-term memory can operate independently of each other. Is there an anatomical basis for this apparent parallel operation between CA3 and CA1 and short-term and intermediate term episodic memory? Given that CA1 represents the primary output from the hippocampus, especially in the light of the evidence that that CA3 does not have direct axonal projections to the subiculum or entorhinal cortex, how can information be transmitted to other neural regions outside the hippocampus once CA1 is ablated? This is important since other studies have also reported a lack of or transient deficit in spatial memory following selective damage to CA1 (Mizumori et al., 1990). One possibility is that outputs that originate from CA3 project via the fimbria directly to the vertical limb of the diagonal band of Broca and medial septum (MS) or indirectly via the lateral septum to MS (Gaykema et al., 1991; Raisman et al., 1966; Swanson and Cowan, 1979; Wyss et al., 1980). The vertical limb of the diagonal band of Broca and MS in turn provides cholinergic and GABA-ergic inputs back to the hippocampus, especially the CA3 subregion. These connections have been characterized anatomically via degeneration, fluorescent tracer, and immunohistochemical experiments (Gaykema et al., 1991; Raisman et al., 1966; Swanson and Cowan, 1979; Wyss et al., 1980). The connections via the fimbria to the vertical limb of the diagonal band of Broca and MS and especially onto the cholinergic neurons provide an anatomical locus for modulating cholinergic inputs into the hippocampus and especially CA3. An alternate output pathway from the MS to the hippocampus is provided by projections to subiculum and entorhinal cortex, as well as the mammillary bodies. In turn, the mammillary bodies project to the anterior thalamus which projects to entorhinal cortex via the pre- and para-subiculum (Wyss et al., 1980). There are, therefore, multiple pathways whereby the outputs from CA3 in the fimbria can send information to the entorhinal cortex and subsequently the hippocampus.

In order to examine the hypothesis that the CA3 efferents via the fimbria are involved in acquiring tasks that require learning over multiple trials, the ventrolateral half of the fimbria was transected under electrophysiological control with a retractable wire knife. The rats were then tested in the acquisition of a non-match to sample 1-trial spatial location task on an 8-arm maze with 10 s delays. Results indicate that fimbria lesioned rats were impaired in acquiring the task similar to CA3 lesioned rats. Furthermore, pairing a CA1 lesion with a fimbria transection mimics the effect of a CA3 lesion, but with longer lasting deficits. In addition, neurophysiological and acetylcholinesterase and BDA (an anterograde tracer) results suggest that the output from dorsal CA3 to the medial septum via the fimbria was damaged without removing the input connections back to dorsal CA3 (Hunsaker et al., 2007c, 2008c). Furthermore, the same fimbria transection also mimics CA3 lesions in that there is a deficit in the encoding associated with within day learning, but not between day retrieval, in the Hebb–Williams maze and an impairment in detection of spatial novelty (Hunsaker et al., 2007c, 2008c).

One explanation that could provide a mechanism for understanding the fimbria cut induced acquisition deficit is based on the importance of the MS to CA3 connection (Swanson and Cowan, 1977), since even a small disruption to the modulatory cholinergic system could disrupt the ability of CA3 neurons to encode new information (Hasselmo et al., 1995). It is assumed that the fimbria output to the MS is excitatory, so that a cut in the fimbria output would result in reduced activation of cholinergic fiber projections, especially to CA3. The consequence of a reduction in the cholinergic influence on the hippocampus could be an increase in the effective strength of the reciprocal, recurrent CA3 connections, which would impair new learning because existing associations in the CA3 to CA3 network would tend to dominate the CA3 cell firing (Hasselmo et al., 1995). In this sense, low acetylcholine is appropriate for recall but not for new learning; with low acetylcholine, memories already in CA3 may interfere with setting up new representations in CA3 to be learned. In addition, low acetylcholine after a fimbria cut could impair new learning because long-term potentiation in the CA3 to CA3 synapses is reduced (Hasselmo et al., 1995). These acquisition deficits may be especially evident when the task is difficult to learn and requires multiple trials, and interference in encoding of new spatial information becomes more pronounced. Consistent with these hypotheses are observations of Rogers and Kesner (2003, 2004), where they examined the effects of scopolamine injections into dorsal CA3 on the acquisition of the modified Hebb–Williams maze as well as the acquisition of contextual delay fear conditioning. Scopolamine acted to inhibit or reduce encoding or acquisition over multiple trials, but had no real effect on retrieval. The same encoding/retrieval effect was observed for spatial contextual fear conditioning (Rogers and Kesner, 2003). In another study, Marques Pereira et al. (2005) showed that injections of 192 IgG-saporin in the hippocampus, which produced a severe depletion of cholinergic cells in the MS, was sufficient to disrupt the acquisition of the standard Hebb–Williams maze. These data also fit with modeling data indicating that MS innervates CA3 more strongly than CA1 (Hasselmo et al., 1995) which again emphasizes the role of CA3 in encoding. Another possibility that needs to be considered involves the output route from MS which projects directly to the subiculum and entorhinal cortex or directly to the mammillary bodies and then via the anterior thalamus and pre- and para-subiculum to entorhinal cortex and then hippocampus (Vann and Aggleton, 2004).

On a different dimension, differences between CA3 and CA1 might be based on the role of CA3 in processing predominantly spatial information and the CA1 in processing predominantly temporal information. For example, A) lesions of the CA3, but not the CA1, subregion impair novelty detection of a spatial location (Lee et al., 2005), B) lesions of CA3, but not CA1 impairs acquisition of object–place and object–odor associations (Gilbert and Kesner, 2003a). A summary of the dissociations between CA3 and CA1 is provided in Table 1.

A second type of result that has been obtained when one compares the relationships between CA3 and CA1 is that there is a deficit following dysfunction of the CA1 subregion, but no concomitant deficit following dysfunction of the CA3 subregion. For example, A) intermediate-term memory in a delayed (5 min) nonmatching-to-sample for a spatial location task is disrupted by lesions in CA3 or AP5 injections into the CA1, but not into the CA3 subregion (Lee and Kesner, 2002, 2003), B) CA1, but not CA3, lesions impair retrieval of information acquired in a Hebb–Williams maze (Jerman et al., 2006; Vago et al., 2007), C) after a change in cue configuration the CA1 place fields shifted their locations backwards 2 days later compared to the place fields in CA3 (Lee et al., 2004a). These results also suggest that intermediate-term and short-term episodic memory can operate independent of each other. Is there a different anatomical basis for this apparent parallel operation between CA3 vs. CA1 and short-term vs. intermediate term memory? In all these cases when there is a deficit following CA1 lesions, but not CA3 lesions, the possibility exists that the deficit is due to a faulty input from the direct perforant pathway to CA1, since the Schaffer collateral input is intact. Based on the idea that in vitro dopamine injections into the CA1 region inhibit the direct perforant path projection to the CA1 region without affecting the Schaffer collateral projection into the CA1 region (Otmakhova and Lisman, 1999), rats were injected with apomorphine (a nonselective dopamine agonist) or vehicle control in the CA1 region in the delayed (5 min) nonmatching-to-sample for a spatial location task, in the Hebb–Williams maze task, and the contextual fear conditioning task. The results show that apomorphine injections into the CA1 disrupt performance in all these tasks. Furthermore, apomorphine injections into CA1 do not

Table 1
CA_3–CA_1 dissociations.

Task	Method	CA_3	CA_1	References
Short-term memory for spatial information	AP5 injection or lesion (ibotenic acid)	Deficit	No deficit	Lee and Kesner (2002, 2003)
Hebb–Williams maze	Lesion (ibotenic acid)	Deficit in encoding	No deficit in encoding	Kesner et al. (2002)
Place fields shift following a change in cue configuration	Electrophysiological recording	Short-term memory deficit	Short-term memory no deficit	Lee et al. (2004a)
Acquisition of object–place and object–odor associations	Lesion (ibotenic acid)	Deficit	No deficit	Gilbert and Kesner (2003a,b)
Novelty detection of a spatial change	Lesion (ibotenic acid)	Deficit	No Deficit	Lee et al. (2005)

Table 2
CA_1–CA_3 dissociations continued.

Task	Method	CA_3	CA_1	References
Intermediate memory for spatial locations	AP5 injection or lesion (ibotenic acid)	No deficit	Deficit	Lee and Kesner (2002, 2003)
Hebb–Williams maze	Lesion (ibotenic acid)	No deficit in retrieval	Deficit in retrieval	Jerman et al. (2006), Vago et al. (2007)
Place fields shift following a change in cue configuration	Electrophysiological recording	Long-term memory no deficit	Long-term memory deficit	Lee et al. (2004a)
Temporal order memory for visual objects	Lesion (ibotenic acid)	No deficit	Deficit	Hoge and Kesner (2007)
Acquisition of object–trace–odor associations	Lesion (ibotenic acid)	No deficit	Deficit	Kesner et al. (2005)

disrupt encoding in the Hebb–Williams maze and transfer of short-term memory for a spatial location in a nonmatching-to-sample task to a novel environment (a new maze) which are processes that are not dependent on the CA1 region (Vago et al., 2007; Vago and Kesner, 2008). Apomorphine injections into the CA1 subregion produces the same pattern of deficits as are produced by CA1, but not CA3, lesions. This result implies that in situations where there is a CA1, but no CA3, lesion effect, the direct perforant path input into the CA1 region represents the main input for the integrity of intact performance perhaps based on intermediate-term memory processing.

On a different dimension, differences between CA3 and CA1 might be based on the role of CA3 in processing predominantly spatial information and the CA1 in processing predominantly temporal information. For example, A) lesions of the CA1, but not the CA3, subregion impair temporal order memory for visual objects (Hoge and Kesner, 2007), and B) lesions of CA1, but not CA3, impairs acquisition of object–trace–odor associations (Kesner et al., 2005). For a summary of the dissociations between CA1 and CA3 see Table 2.

A third type of result that has been obtained when one compares the relationships between CA3 and CA1 is that one can observe a deficit following dysfunction of either the CA3 or CA1 subregions. For example, A) both lesions of CA1 or CA3 disrupt temporal order memory for spatial locations (Hunsaker and Kesner, 2008), B) both lesions of CA3 or CA1 impair the acquisition of an object–trace–place task (Hunsaker et al., 2006), C) both CA3 or CA1 lesions disrupt sequential performance in a runway task (Hunsaker and Kesner, 2008), and D) both CA3 or CA1 lesions impair temporal pattern completion for a known sequence of spatial information (Hoang and Kesner, 2008). When both a spatial and temporal memory are required to perform a memory task, then there is the implication that both regions are working cooperatively. For a summary of the potential interactions between CA1 and CA3 see Table 3.

A summary of some of the differences and similarities between the effects of DG and CA3 lesions is shown in Table 4.

In summary, the anatomical organization in terms of afferent and efferent connections of the different subregions of the hippocampus as well as intrinsic operations associated with spatial vs. temporal processing of information of each subregion can account for dissociations as well as associations between CA3 and CA1 in supporting short-term and intermediate-term memory, respectively. Thus short-term episodic memory mediated by CA3 and intermediate-term episodic memory mediated by CA1 can operate independent of each other for one set of tasks and interact with each other for a different set of tasks.

4. Comparison with other theories of hippocampal function

The model described here is quantitative, and supported by both formal analyses and quantitative simulations. Many of the points made, such as on the number of memories that can be stored in autoassociative networks, the utility of sparse representations, and the dynamics of the operation of networks with recurrent connections, are quite general, and will apply to networks in a number of different brain areas. With respect to the hippocampus, the theory specifies the maximum number of memories that could be stored in it, and this has implications for how it could be used biologically. It indicates that if this number is approached, it will be useful to have a mechanism for recalling information from the hippocampus for incorporation into memories elsewhere. With respect to recall, the theory provides a quantitative account for why there are as many backprojections as forward projections in the cerebral cortex. Overall, the theory provides an explanation for *how* this part of the brain could work, and even if this theory needs to be revised, it is suggested that a fully quantitative theory along the lines proposed which is based on the evidence available from a wide range

Table 3
CA_3–CA_1 associations.

Task	Method	CA_3	CA_1	References
Temporal order memory for spatial locations	Lesion (ibotenic acid)	Deficit	Deficit	Hunsaker and Kesner (2008)
Acquisition of place–trace–object task	Lesion (ibotenic acid)	Deficit	Deficit	Hunsaker et al. (2006)
Sequential learning in a runway	Lesion (ibotenic acid)	Deficit	Deficit	Hunsaker et al. (2008b)
Temporal pattern completion for spatial information	Lesion (ibotenic acid)	Deficit	Deficit	Hoang and Kesner (2008)

of techniques will be essential before we can say that we understand *how* a part of the brain operates. This theory is almost unique in being a quantitative theory of hippocampal function.

Hypotheses have been described about how a number of different parts of hippocampal and related circuitry might operate. Although these hypotheses are consistent with a theory of how the hippocampus operates, some of these hypotheses could be incorporated into other views or theories. In order to highlight the differences between alternative theories, and in order to lead to constructive analyses that can test them, the theory described above is compared with other theories of hippocampal function in this section. Although the differences between the theories are highlighted in this section, the overall view 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 geometric 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 as inputs information 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.[7] 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., 2000, 1994) 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 elsewhere (Rolls, 1996a, 2008a; Rolls and Treves, 1998). 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, we 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 Grusser and Landis, 1991; 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 (Rolls, 1987, 1989a,b,d, 2008a, 2010b, 2013b,c; Treves and Rolls, 1992, 1994) at the systems level, except that it takes a stronger position on the gradient of retrograde amnesia, emphasises that recall from

[7] John O'Keefe was one of the recipients of the Nobel Prize for Physiology or Medicine in 2014, after this paper was written. His work on the discovery of hippocampal place cells in rats (O'Keefe and Dostrovsky, 1971) was cited, and he is congratulated. Indeed, the announcement for the award of the Nobel Prize described this system as a 'component of a positioning system, an "inner GPS" in the brain'. John O'Keefe has continued to emphasize the role of the hippocampus and rodent place cells in navigation (O'Keefe, 1990; Hartley et al., 2014). Rolls' discoveries and theory are thus somewhat different, in that Rolls has shown that spatial view cells may be especially relevant to the operation of the hippocampus in primates including humans; and in that the roles of the hippocampal system in memory are emphasized.

Table 4
DG–CA_3 dissociations and similarities.

Task	Method	DG	CA_3	References
Spatial pattern separation	Lesion	Partial deficit	Total deficit	Gilbert and Kesner (2006), Gilbert et al. (2001)
Detection of a geometric context change	Lesion	Deficit	No deficit	Hunsaker et al. (2008a)
Temporal pattern separation for spatial locations	Lesion	No deficit	Deficit	Gilbert and Kesner (2006), Gilbert et al. (2001)
Hebb–Williams maze	Lesion	Deficit in encoding not retrieval	Deficit in encoding not retrieval	Jerman et al. (2006)

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 (see Fig. 1). It also emphasizes the important point that the hippocampal and neocortical memory systems may be quite different, with the hippocampus specialized for the rapid learning of single events or episodes, and the neocortex for the slower learning of semantic representations which may necessarily benefit from the many exemplars needed to shape the semantic representation. In the formulation by McClelland et al. (1995), the entorhinal cortex connections via the perforant path onto the CA1 cells are non-modifiable (in the short term), and allow a representation of neocortical long-term memories to activate the CA1 cells. The new information learned in an episode by the CA3 system is then linked to existing long-term memories by the CA3 to CA1 rapidly modifiable synapses. All the connections from the CA1 back via the subiculum, entorhinal cortex, parahippocampal cortex etc. to the association neocortex are held to be unmodifiable in the short term, during the formation of an episodic memory. The formal argument that leads us to suggest that the backprojecting synapses are associatively modifiable during the learning of an episodic memory is similar to that which we have used to show that for efficient recall, the synapses which initiate recall in the CA3 system (identified above with the perforant path projection to CA3) must be associatively modifiable if recall is to operate efficiently (see Treves and Rolls, 1992). The present theory holds that it is possible that for several stages back into neocortical processing, the backprojecting synapses should be associatively modifiable, with a similar time course to the time it takes to learn a new episodic memory. It may well be that at earlier stages of cortical processing, for example from inferior temporal visual cortex to V4, and from V4 to V2, the backprojections are relatively more fixed, being formed (still associatively) during early developmental plasticity or during the formation of new long-term semantic memory structures. Having such relatively fixed synaptic strengths in these earlier cortical backprojection systems could ensure that whatever is recalled in higher cortical areas, such as objects, will in turn recall relatively fixed and stable representations of parts of objects or features. Given that the functions of backprojections may include many top-down processing operations, including attention (Deco and Rolls, 2005a) and priming, it may be useful to ensure that there is consistency in how higher cortical areas affect activity in earlier "front-end" or preprocessing cortical areas. Indeed, the current theory shows that at least one backprojection stage the hippocampo-cortical connections must be associatively modifiable during the learning of an episodic memory, but does not require the associative backprojection learning to occur at all backprojection stages during the learning of the episodic memory. Crucial stages might include CA1 to subiculum or to entorhinal cortex (see Fig. 1), or entorhinal cortex to parahippocampal gyrus and perirhinal cortex. It would be interesting to test this using local inactivation of NMDA receptors at different stages of the backprojection system to determine where this impairs the learning of for example place–object recall, i.e. a task that is likely to utilize the hippocampo-cortical backprojection pathways.

If a model of the hippocampal/neocortical memory system could store only a small number of patterns, it would not be a good model of the real hippocampal/neocortical memory system in the brain. Indeed, this appears to be a major limitation of another model presented by Alvarez and Squire (1994). The model specifies that the hippocampus helps the operation of a neocortical multimodal memory system in which all memories are stored by associative recurrent collaterals between the neocortical neurons. Although the idea worked in the model with twenty neurons and two patterns to be learned (Alvarez and Squire, 1994), the whole idea is computationally not feasible, because the number of memories that can be stored in a single autoassociative network of the type described is limited by the number of inputs per neuron from other neurons, not by the number of neurons in the network (Treves and Rolls, 1991, 1994). This would render the capacity of the whole neocortical multimodal (or amodal) memory store very low (in the order of the number of inputs per neuron from the other neurons, that is in the order of 5000–10,000) (cf. O'Kane and Treves, 1992). This example makes it clear that it is important to take into account analytic and quantitative approaches when investigating memory systems. The current work is an attempt to do this.

The theory of Lisman and colleagues (2005) of the memory for sequences has been described and evaluated in Section 2.2.3.7. 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 III. The proposal that acetylcholine could be important during encoding by facilitating CA3–CA3 LTP, and should be lower during retrieval (Hasselmo et al., 1995), is a useful concept.

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

Teyler and DiScenna (1986) proposed that the hippocampus stores pointers to neocortical areas, and uses those pointers to recall memories in the cortex. The present theory is different, for it emphasizes that spatial and timing representations are present in the hippocampus, and can be combined with object information to store episodic memories in the hippocampus by forming associations between these components. However, the theory also shows how recall of these episodic memories within the hippocampus can then using backprojections to neocortical areas allow neocortical representations present at the time of the encoding to be retrieved, and then potentially stored in the neocortex in a schema or semantic memory. The actual retrieval process to neocortex is implemented by using associatively modified synapses in at least one stage of the backprojection system, and this provides a deaddressing method for the hippocampus to retrieve a

content-addressable neocortical memory. This deaddressing process for recall to the neocortex that is part of the current theory can be likened to the use of a pointer. The present theory (Rolls, 2008a, 2010b) is we believe unique in that it enables both the hippocampal storage process, and the recall to the neocortex, to be understood quantitatively.

The aim of this comparison of the present theory with other theories has been to highlight differences between the theories, to assist in the future assessment of the utility and the further development of each of the theories.

Acknowledgements

(ETR) Different parts of the research described here were supported by Programme Grants from the Medical Research Council; by a Human Frontier Science program grant; by an EEC BRAIN grant; by the MRC Oxford Interdisciplinary Research Centre in Cognitive Neuroscience; by the Oxford McDonnell-Pew Centre in Cognitive Neuroscience; and by the Oxford Centre for Computational Neuroscience. The author has performed the experimental and theoretical work that is incorporated in some of the ideas presented here on the hippocampus with many colleagues, including Alessandro Treves, Simon Stringer, Robert Robertson, Pierre Georges-François, Shane O'Mara, Ray Kesner, Yasushi Miyashita, Janet Feigenbaum, Ivan de Araujo, Simon Schultz, Stefano Panzeri, David Foster, Peter Cahusac, Hiroaki Niki, and Alain Berthoz, and their contributions are sincerely acknowledged.

(RPK) This research was supported by NSF Grant IBN-0135273 and NIH Grant 5R01MH065314-02 awarded to R.P. Kesner, who acknowledges the excellent collaborations with former doctoral students Paul Gilbert, Inah Lee, and Ryan Hunsaker.

References

Abbott, L.F., Blum, K.I., 1996. Functional significance of long-term potentiation for sequence learning and prediction. Cereb. Cortex 6, 406–416.
Abbott, L.F., Nelson, S.B., 2000. Synaptic plasticity: taming the beast. Nat. Neurosci. 3, 1178–1183.
Acsady, L., Kamondi, A., Sik, A., Freund, T., Buzsaki, G., 1998. GABAergic cells are the major postsynaptic targets of mossy fibers in the rat hippocampus. J. Neurosci. 18, 3386–3403.
Aggelopoulos, N.C., Franco, L., Rolls, E.T., 2005. Object perception in natural scenes: encoding by inferior temporal cortex simultaneously recorded neurons. J. Neurophysiol. 93, 1342–1357.
Ahn, J.R., Lee, I., 2014. Intact CA3 in the hippocampus is only sufficient for contextual behavior based on well-learned and unaltered visual background. Hippocampus 24, 1081–1093.
Aimone, J.B., Deng, W., Gage, F.H., 2010. Adult neurogenesis: integrating theories and separating functions. Trends Cogn. Sci. 14, 325–337.
Aimone, J.B., Gage, F.H., 2011. Modeling new neuron function: a history of using computational neuroscience to study adult neurogenesis. Eur. J. Neurosci. 33, 1160–1169.
Aimone, J.B., Wiles, J., Gage, F.H., 2006. Potential role for adult neurogenesis in the encoding of time in new memories. Nat. Neurosci. 9, 723–727.
Aimone, J.B., Wiles, J., Gage, F.H., 2009. Computational influence of adult neurogenesis on memory encoding. Neuron 61, 187–202.
Alvarez, P., Squire, L.R., 1994. Memory consolidation and the medial temporal lobe: a simple network model. Proc. Natl. Acad. Sci. U. S. A. 91, 7041–7045.
Amaral, D.G., 1986. Amygdalohippocampal and amygdalocortical projections in the primate brain. In: Schwarz, R., Ben-Ari, Y. (Eds.), Excitatory Amino Acids and Epilepsy. Plenum, New York, pp. 3–17.
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., Price, J.L., Pitkanen, A., Carmichael, S.T., 1992. Anatomical organization of the primate amygdaloid complex. In: Aggleton, J.P. (Ed.), The Amygdala. Wiley-Liss, New York, pp. 1–66.
Amaral, D.G., 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.
Amit, D.J., Brunel, N., 1997. Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. Cereb. Cortex 7, 237–252.
Andersen, P., Morris, R.G.M., Amaral, D.G., Bliss, T.V.P., O'Keefe, J., 2007. The Hippocampus Book. Oxford University Press, London.
Anderson, M.I., Jeffery, K.J., 2003. Heterogeneous modulation of place cell firing by changes in context. J. Neurosci. 23, 8827–8835.
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.
Banks, W.P., 1978. Encoding and processing of symbolic information in comparative judgements. In: Bower, G.H. (Ed.), The Psychology of Learning and Motivation: Advances in Theory and Research. Academic Press, New York, pp. 101–159.
Banta Lavenex, P., Lavenex, P., 2009. Spatial memory and the monkey hippocampus: not all space is created equal. Hippocampus 19, 8–19.
Barkas, L.J., Henderson, J.L., Hamilton, D.A., Redhead, E.S., Gray, W.P., 2010. Selective temporal resections and spatial memory impairment: cue dependent lateralization effects. Behav. Brain Res. 208, 535–544.
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.
Baxter, M.G., Murray, E.A., 2001. Opposite relationship of hippocampal and rhinal cortex damage to delayed nonmatching-to-sample deficits in monkeys. Hippocampus 11, 61–71.
Baylis, G.C., Rolls, E.T., 1987. Responses of neurons in the inferior temporal cortex in short term and serial recognition memory tasks. Exp. Brain Res. 65, 614–622.
Bertaina-Anglade, V., Tramu, G., Destrade, C., 2000. Differential learning-stage dependent patterns of c-Fos protein expression in brain regions during the acquisition and memory consolidation of an operant task in mice. Eur. J. Neurosci. 12, 3803–3812.
Bilkey, D.K., Goddard, G.V., 1985. Medial septal facilitation of hippocampal granule cell activity is mediated by inhibition of inhibitory interneurones. Brain Res. 361, 99–106.
Bonelli, S.B., Powell, R.H., Yogarajah, M., Samson, R.S., Symms, M.R., Thompson, P.J., Koepp, M.J., Duncan, J.S., 2010. Imaging memory in temporal lobe epilepsy: predicting the effects of temporal lobe resection. Brain 133, 1186–1199.
Brown, M.W., Aggleton, J.P., 2001. Recognition memory: what are the roles of the perirhinal cortex and hippocampus? Nat. Rev. 2, 51–61.
Brown, M.W., Xiang, J.Z., 1998. Recognition memory: neuronal substrates of the judgement of prior occurrence. Prog. Neurobiol. 55, 149–189.
Brown, T.H., Ganong, A.H., Kairiss, E.W., Keenan, C.L., Kelso, S.R., 1989. Long-term potentiation in two synaptic systems of the hippocampal brain slice. In: Byrne, J.H., Berry, W.O. (Eds.), Neural Models of Plasticity. Academic Press, San Diego, pp. 266–306.
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.
Buckley, M.J., Gaffan, D., 2006. Perirhinal cortical contributions to object perception. Trends Cogn. Sci. 10, 100–107.
Buckmaster, P.S., Schwartzkroin, P.A., 1994. Hippocampal mossy cell function: a speculative view. Hippocampus 4, 393–402.
Buhl, E.H., Halasy, K., Somogyi, P., 1994. Diverse sources of hippocampal unitary inhibitory postsynaptic potentials and the number of synaptic release sites. Nature 368, 823–828.
Bunsey, M., Eichenbaum, H., 1996. Conservation of hippocampal memory function in rats and humans. Nature 379, 255–257.
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.
Burke, S.N., Wallace, J.L., Hartzell, A.L., Nematollahi, S., Plange, K., Barnes, C.A., 2011. Age-associated deficits in pattern separation functions of the perirhinal cortex: a cross-species consensus. Behav. Neurosci. 125, 836–847.
Burke, S.N., Wallace, J.L., Nematollahi, S., Uprety, A.R., Barnes, C.A., 2010. Pattern separation deficits may contribute to age-associated recognition impairments. Behav. Neurosci. 124, 559–573.
Burwell, R.D., Witter, M.P., Amaral, D.G., 1995. Perirhinal and postrhinal cortices of the rat: a review of the neuroanatomical literature and comparison with findings from the monkey brain. Hippocampus 5, 390–408.

Bussey, T.J., Saksida, L.M., Murray, E.A., 2002. Perirhinal cortex resolves feature ambiguity in complex visual discriminations. Eur. J. Neurosci. 15, 365–374.

Bussey, T.J., Saksida, L.M., Murray, E.A., 2005. The perceptual-mnemonic/feature conjunction model of perirhinal cortex function. Q. J. Exp. Psychol. B 58, 269–282.

Cahusac, P.M.B., Rolls, E.T., Miyashita, Y., Niki, H., 1993. Modification of the responses of hippocampal neurons in the monkey during the learning of a conditional spatial response task. Hippocampus 3, 29–42.

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.

Carr, V.A., Rissman, J., Wagner, A.D., 2010. Imaging the human medial temporal lobe with high-resolution fMRI. Neuron 65, 298–308.

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.

Cerasti, E., Treves, A., 2013. The spatial representations acquired in CA3 by self-organizing recurrent connections. Front. Cell. Neurosci. 7, 112.

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.

Chapuis, J., Wilson, D.A., 2012. Bidirectional plasticity of cortical pattern recognition and behavioral sensory acuity. Nat. Neurosci. 15, 155–161.

Cheng, S., 2013. The CRISP theory of hippocampal function in episodic memory. Front. Neural Circ. 7, 88.

Chiba, A.A., Kesner, R.P., Reynolds, A.M., 1994. Memory for spatial location as a function of temporal lag in rats: role of hippocampus and medial prefrontal cortex. Behav. Neural Biol. 61, 123–131.

Cho, Y.H., Kesner, R.P., 1995. Relational object association learning in rats with hippocampal lesions. Behav. Brain Res. 67, 91–98.

Cleland, T.A., Morse, A., Yue, E.L., Linster, C., 2002. Behavioral models of odor similarity. Behav. Neurosci. 116, 222–231.

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.

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.

Debiec, J., LeDoux, J.E., Nader, K., 2002. Cellular and systems reconsolidation in the hippocampus. Neuron 36, 527–538.

Deco, G., Rolls, E.T., 2003. Attention and working memory: a dynamical model of neuronal activity in the prefrontal cortex. Eur. J. Neurosci. 18, 2374–2390.

Deco, G., Rolls, E.T., 2004. A neurodynamical cortical model of visual attention and invariant object recognition. Vision Res. 44, 621–644.

Deco, G., Rolls, E.T., 2005a. Attention, short-term memory, and action selection: a unifying theory. Prog. Neurobiol. 76, 236–256.

Deco, G., Rolls, E.T., 2005b. Neurodynamics of biased competition and cooperation for attention: a model with spiking neurons. J. Neurophysiol. 94, 295–313.

Deco, G., Rolls, E.T., 2005c. Sequential memory: a putative neural and synaptic dynamical mechanism. J. Cogn. Neurosci. 17, 294–307.

Deco, G., Rolls, E.T., 2006. Decision-making and Weber's Law: a neurophysiological model. Eur. J. Neurosci. 24, 901–916.

Deco, G., Rolls, E.T., Albantakis, L., Romo, R., 2013. Brain mechanisms for perceptual and reward-related decision-making. Prog. Neurobiol. 103, 194–213.

Deco, G., Rolls, E.T., Zihl, J., 2005. A neurodynamical model of visual attention. In: Itti, L., Rees, G., Tsotos, J. (Eds.), Neurobiology of Attention. Elsevier, San Diego, CA, pp. 593–599 (Chapter 597).

DeCoteau, W.E., Kesner, R.P., 2000. A double dissociation between the rat hippocampus and medial caudoputamen in processing two forms of knowledge. Behav. Neurosci. 114, 1096–1108.

Dees, R.L., Kesner, R.P., 2013. The role of the dorsal dentate gyrus in object and object-context recognition. Neurobiol. Learn. Mem. 106, 112–117.

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.

Dellu, F., Fauchey, V., Le Moal, M., Simon, H., 1997. Extension of a new two-trial memory task in the rat: influence of environmental context on recognition processes. Neurobiol. Learn. Mem. 67, 112–120.

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.

Dix, S.L., Aggleton, J.P., 1999. Extending the spontaneous preference test of recognition: evidence of object-location and object-context recognition. Behav. Brain Res. 99, 191–200.

Dudchenko, P.A., Wood, E.R., Eichenbaum, H., 2000. Neurotoxic hippocampal lesions have no effect on odor span and little effect on odor recognition memory but produce significant impairments on spatial span, recognition, and alternation. J. Neurosci. 20, 2964–2977.

Eacott, M.J., Norman, G., 2004. Integrated memory for object, place, and context in rats: a possible model of episodic-like memory. J. Neurosci. 24, 1948–1953.

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

Eichenbaum, H., Cohen, N.J., 2001. From Conditioning to Conscious Recollection: Memory Systems of the Brain. Oxford University Press, New York.

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.

Emerich, D.F., Walsh, T.J., 1989. Selective working memory impairments following intradentate injection of colchicine: attenuation of the behavioral but not the neuropathological effects by gangliosides GM1 and AGF2. Physiol. Behav. 45, 93–101.

Ennaceur, A., Delacour, J., 1988. A new one-trial test for neurobiological studies of memory in rats. 1: Behavioral data. Behav. Brain Res. 31, 47–59.

Estes, W.K., 1986. Memory for temporal information. In: Michon, J.A., Jackson, J.L. (Eds.), Time, Mind and Behavior. Springer-Verlag, New York, pp. 151–168.

Faisal, A.A., Selen, L.P., Wolpert, D.M., 2008. Noise in the nervous system. Nat. Rev. Neurosci. 9, 292–303.

Fazeli, M.S., Collingridge, G.L., 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.

Ferbinteanu, J., Holsinger, R.M., McDonald, R.J., 1999. Lesions of the medial or lateral perforant path have different effects on hippocampal contributions to place learning and on fear conditioning to context. Behav. Brain Res. 101, 65–84.

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.

Földiák, P., 1991. Learning invariance from transformation sequences. Neural Comput. 3, 194–200.

Fortin, N.J., Agster, K.L., Eichenbaum, H.B., 2002. Critical role of the hippocampus in memory for sequences of events. Nat. Neurosci. 5, 458–462.

Foster, T.C., Castro, C.A., McNaughton, B.L., 1989. Spatial selectivity of rat hippocampal neurons: dependence on preparedness for movement. Science 244, 1580–1582.

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.

Franzius, M., Sprekeler, H., Wiskott, L., 2007. Slowness and sparseness lead to place, head-direction, and spatial-view cells. PLoS Comput. Biol. 3, e166.

Fried, I., MacDonald, K.A., Wilson, C.L., 1997. Single neuron activity in human hippocampus and amygdala during recognition of faces and objects. Neuron 18, 753–765.

Fujimichi, R., Naya, Y., Koyano, K.W., Takeda, M., Takeuchi, D., Miyashita, Y., 2010. Unitized representation of paired objects in area 35 of the macaque perirhinal cortex. Eur. J. Neurosci. 32, 659–667.

Fyhn, M., Molden, S., Witter, M.P., Moser, E.I., Moser, M.B., 2004. Spatial representation in the entorhinal cortex. Science 305, 1258–1264.

Gaffan, D., 1993. Additive effects of forgetting and fornix transection in the temporal gradient of retrograde amnesia. Neuropsychologia 31, 1055–1066.

Gaffan, D., 1994. Scene-specific memory for objects: a model of episodic memory impairment in monkeys with fornix transection. J. Cogn. Neurosci. 6, 305–320.

Gaffan, D., Harrison, S., 1989a. A comparison of the effects of fornix section and sulcus principalis ablation upon spatial learning by monkeys. Behav. Brain Res. 31, 207–220.

Gaffan, D., Harrison, S., 1989b. Place memory and scene memory: effects of fornix transection in the monkey. Exp. Brain Res. 74, 202–212.

Gaffan, D., Saunders, R.C., 1985. Running recognition of configural stimuli by fornix transected monkeys. Q. J. Exp. Psychol. 37B, 61–71.

Gall, C.M., Hess, U.S., Lynch, G., 1998. Mapping brain networks engaged by, and changed by, learning. Neurobiol. Learn. Mem. 70, 14–36.

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.

Gaykema, R.P., van der Kuil, J., Hersh, L.B., Luiten, P.G., 1991. Patterns of direct projections from the hippocampus to the medial septum-diagonal band complex: anterograde tracing with *Phaseolus vulgaris* leucoagglutinin combined with immunohistochemistry of choline acetyltransferase. Neuroscience 43, 349–360.

Ge, S., Yang, C.H., Hsu, K.S., Ming, G.L., Song, H., 2007. A critical period for enhanced synaptic plasticity in newly generated neurons of the adult brain. Neuron 54, 559–566.

Gelbard-Sagiv, H., Mukamel, R., Harel, M., Malach, R., Fried, I., 2008. Internally generated reactivation of single neurons in human hippocampus during free recall. Science 322, 96–101.

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., Brushfield, A.M., 2009. The role of the CA3 hippocampal subregion in spatial memory: a process oriented behavioral assessment. Prog. Neuro-Psychopharmacol. Biol. Psychiatry 33, 774–781.

Gilbert, P.E., Kesner, R.P., 2002. The amygdala but not the hippocampus is involved in pattern separation based on reward value. Neurobiol. Learn. Mem. 77, 338–353.

Gilbert, P.E., Kesner, R.P., 2003a. Localization of function within the dorsal hippocampus: the role of the CA3 subregion in paired-associate learning. Behav. Neurosci. 117, 1385–1394.

Gilbert, P.E., Kesner, R.P., 2003b. Recognition memory for complex visual discriminations is influenced by stimulus interference in rodents with perirhinal cortex damage. Learn. Mem. (Cold Spring Harbor, NY) 10, 525–530.
Gilbert, P.E., Kesner, R.P., 2006. The role of the dorsal CA3 hippocampal subregion in spatial working memory and pattern separation. Behav. Brain Res. 169, 142–149.
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., 2008. The interactions and dissociations of the dorsal hippocampus subregions: how the dentate gyrus, CA3, and CA1 process spatial information. Behav. Neurosci. 122, 16–26.
Grabenhorst, F., Rolls, E.T., 2010. Attentional modulation of affective vs sensory processing: functional connectivity and a top-down biased activation theory of selective attention. J. Neurophysiol. 104, 1649–1660.
Grabenhorst, F., Rolls, E.T., 2011. Value, pleasure, and choice in the ventral prefrontal cortex. Trends Cogn. Sci. 15, 56–67.
Grusser, O.-J., Landis, T., 1991. Visual Agnosias. Macmillan, London.
Gulyas, A.I., Miles, R., Hajos, N., Freund, T.F., 1993. Precision and variability in postsynaptic target selection of inhibitory cells in the hippocampal CA3 region. Eur. J. Neurosci. 5, 1729–1751.
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.
Hall, J., Thomas, K.L., Everitt, B.J., 2000. Rapid and selective induction of BDNF expression in the hippocampus during contextual learning. Nat. Neurosci. 3, 533–535.
Hall, J., Thomas, K.L., Everitt, B.J., 2001. Cellular imaging of zif268 expression in the hippocampus and amygdala during contextual and cued fear memory retrieval: selective activation of hippocampal CA1 neurons during the recall of contextual memories. J. Neurosci. 21, 2186–2193.
Handelmann, G.E., Olton, D.S., 1981. Spatial memory following damage to hippocampal CA3 pyramidal cells with kainic acid: impairment and recovery with preoperative training. Brain Res. 217, 41–58.
Hannesson, D.K., Howland, G.J., Phillips, A.G., 2004. Interaction between perirhinal and medial prefrontal cortex is required for temporal order but not recognition memory for objects in rats. J. Neurosci. 24, 4596–4604.
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., 1999. Neuromodulation: acetylcholine and memory consolidation. Trends Cogn. Sci. 3, 351–359.
Hasselmo, M.E., 2005. What is the function of hippocampal theta rhythm? – Linking behavioral data to phasic properties of field potential and unit recording data. Hippocampus 15, 936–949.
Hasselmo, M.E., Bodelon, C., Wyble, B.P., 2002a. A proposed function for hippocampal theta rhythm: separate phases of encoding and retrieval enhance reversal of prior learning. Neural Comput. 14, 793–817.
Hasselmo, M.E., Bower, J.M., 1993. Acetylcholine and memory. Trends Neurosci. 16, 218–222.
Hasselmo, M.E., Eichenbaum, H.B., 2005. Hippocampal mechanisms for the context-dependent retrieval of episodes. Neural Netw. 18, 1172–1190.
Hasselmo, M.E., Fehlau, B.P., 2001. Differences in time course of ACh and GABA modulation of excitatory synaptic potentials in slices of rat hippocampus. J. Neurophysiol. 86, 1792–1802.
Hasselmo, M.E., Giocomo, L.M., 2006. Cholinergic modulation of cortical function. J. Mol. Neurosci. MN 30, 133–135.
Hasselmo, M.E., Hay, J., Ilyn, M., Gorchetchnikov, A., 2002b. Neuromodulation, theta rhythm and rat spatial navigation. Neural Netw. 15, 689–707.
Hasselmo, M.E., McGaughy, J., 2004. High acetylcholine levels set circuit dynamics for attention and encoding and low acetylcholine levels set dynamics for consolidation. Prog. Brain Res. 145, 207–231.
Hasselmo, M.E., Schnell, E., 1994. Laminar selectivity of the cholinergic suppression of synaptic transmission in rat hippocampal region CA1: computational modeling and brain slice physiology. J. Neurosci. 14, 3898–3914.
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.
Hasselmo, M.E., Wyble, B.P., 1997. Free recall and recognition in a network model of the hippocampus: simulating effects of scopolamine on human memory function. Behav. Brain Res. 89, 1–34.
Hasselmo, M.E., Wyble, B.P., Wallenstein, G.V., 1996. Encoding and retrieval of episodic memories: role of cholinergic and GABAergic modulation in the hippocampus. Hippocampus 6, 693–708.
Henze, D.A., Wittner, L., Buzsaki, G., 2002. Single granule cells reliably discharge targets in the hippocampal CA3 network in vivo. Nat. Neurosci. 5, 790–795.
Hertz, J., Krogh, A., Palmer, R.G., 1991. An Introduction to the Theory of Neural Computation. Addison-Wesley, Wokingham.
Hess, U.S., Lynch, G., Gall, C.M., 1995a. Changes in c-fos mRNA expression in rat brain during odor discrimination learning: differential involvement of hippocampal subfields CA1 and CA3. J. Neurosci. 15, 4786–4795.
Hess, U.S., Lynch, G., Gall, C.M., 1995b. Regional patterns of c-fos mRNA expression in rat hippocampus following exploration of a novel environment versus performance of a well-learned discrimination. J. Neurosci. 15, 7796–7809.
Hirabayashi, T., Takeuchi, D., Tamura, K., Miyashita, Y., 2013. Functional microcircuit recruited during retrieval of object association memory in monkey perirhinal cortex. Neuron 77, 192–203.
Hoang, L.T., Kesner, R.P., 2008. Dorsal hippocampus, CA3, and CA1 lesions disrupt temporal sequence completion. Behav. Neurosci. 122, 9–15.
Hoge, J., Kesner, R.P., 2007. Role of CA3 and CA1 subregions of the dorsal hippocampus on temporal processing of objects. Neurobiol. Learn. Mem. 88, 225–231.
Hölscher, C., Jacob, W., Mallot, H.A., 2003a. Reward modulates neuronal activity in the hippocampus of the rat. Behav. Brain Res. 142, 181–191.
Hölscher, C., Rolls, E.T., 2002. Perirhinal cortex neuronal activity is actively related to working memory in the macaque. Neural Plas. 9, 41–51.
Hölscher, C., Rolls, E.T., Xiang, J.-Z., 2003b. Perirhinal cortex neuronal activity related to long-term familiarity memory in the macaque. Eur. J. Neurosci. 18, 2037–2046.
Hopfield, J.J., 1982. Neural networks and physical systems with emergent collective computational abilities. Proc. Natl. Acad. Sci. U. S. A. 79, 2554–3255.
Huerta, P.T., Sun, L.D., Wilson, M.A., Tonegawa, S., 2000. Formation of temporal memory requires NMDA receptors within CA1 pyramidal neurons. Neuron 25, 473–480.
Hunsaker, M.R., Allan, K.D., Kesner, R.P., 2007a. Role of dCA3 efferents via the fimbria in the acquisition of a delay nonmatch to place task. Hippocampus 17, 494–502.
Hunsaker, M.R., Kesner, R.P., 2008. Evaluating the differential roles of the dorsal dentate gyrus, dorsal CA3, and dorsal CA1 during a temporal ordering for spatial locations task. Hippocampus 18, 955–964.
Hunsaker, M.R., Kesner, R.P., 2013. The operation of pattern separation and pattern completion processes associated with different attributes or domains of memory. Neurosci. Biobehav. Rev. 37, 36–58.
Hunsaker, M.R., Lee, B., Kesner, R.P., 2008a. Evaluating the temporal context of episodic memory: the role of CA3 and CA1. Behav. Brain Res. 188, 310–315.
Hunsaker, M.R., Mooy, G.G., Swift, J.S., Kesner, R.P., 2007b. Dissociations of the medial and lateral perforant path projections into dorsal DG, CA3, and CA1 for spatial and nonspatial (visual object) information processing. Behav. Neurosci. 121, 742–750.
Hunsaker, M.R., Rogers, J.L., Kesner, R.P., 2007c. Behavioral characterization of a transection of dorsal CA3 subcortical efferents: comparison with scopolamine and physostigmine infusions into dorsal CA3. Neurobiol. Learn. Mem. 88, 127–136.
Hunsaker, M.R., Rosenberg, J.S., Kesner, R.P., 2008b. The role of the dentate gyrus, CA3a,b, and CA3c for detecting spatial and environmental novelty. Hippocampus 18, 1064–1073.
Hunsaker, M.R., Thorup, J.A., Welch, T., Kesner, R.P., 2006. The role of CA3 and CA1 in the acquisition of an object–trace–place paired-associate task. Behav. Neurosci. 120, 1252–1256.
Hunsaker, M.R., Tran, G.T., Kesner, R.P., 2008c. A double dissociation of subcortical hippocampal efferents for encoding and consolidation/retrieval of spatial information. Hippocampus 18, 699–709.
Hunsaker, M.R., Tran, G.T., Kesner, R.P., 2009. A behavioral analysis of the role of CA3 and CA1 subcortical efferents during classical fear conditioning. Behav. Neurosci. 123, 624–630.
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.
Itskov, P.M., Vinnik, E., Diamond, M.E., 2011. Hippocampal representation of touch-guided behavior in rats: persistent and independent traces of stimulus and reward location. PLoS ONE 6, e16462.
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.
Jakab, R.L., Leranth, C., 1995. Septum. In: Paxinos, G. (Ed.), The Rat Nervous System. , 2nd ed. Academic Press, San Diego.
Jarrard, E.L., 1993. On the role of the hippocampus in learning and memory in the rat. Behav. Neural Biol. 60, 9–26.
Jarrard, L.E., Davidson, T.L., 1990. Acquisition of concurrent conditional discriminations in rats with ibotenate lesions of hippocampus and subiculum. Psychobiology 18, 68–73.
Jay, T.M., Witter, M.P., 1991. Distribution of hippocampal CA1 and subicular efferents in the prefrontal cortex of the rat studied by means of anterograde transport of *Phaseolus vulgaris*-leucoagglutinin. J. Comp. Neurol. 313, 574–586.
Jeffery, K.J., 2011. Place cells, grid cells, attractors, and remapping. Neural Plas. 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.

Jeltsch, H., Bertrand, F., Lazarus, C., Cassel, J.C., 2001. Cognitive performances and locomotor activity following dentate granule cell damage in rats: role of lesion extent and type of memory tested. Neurobiol. Learn. Mem. 76, 81–105.

Jensen, O., Idiart, M.A., Lisman, J.E., 1996. Physiologically realistic formation of autoassociative memory in networks with theta/gamma oscillations: role of fast NMDA channels. Learn. Mem. (Cold Spring Harbor, NY) 3, 243–256.

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, NY) 3, 264–278.

Jensen, O., Lisman, J.E., 2005. Hippocampal sequence-encoding driven by a cortical multi-item working memory buffer. Trends Neurosci. 28, 67–72.

Jerman, T., Kesner, R.P., Hunsaker, M.R., 2006. Disconnection analysis of CA3 and DG in mediating encoding but not retrieval in a spatial maze learning task. Learn. Mem. (Cold Spring Harbor, NY) 13, 458–464.

Jezek, K., Henriksen, E.J., Treves, A., Moser, E.I., Moser, M.-B., 2011. Theta-paced flickering between place-cell maps in the hippocampus. Nature 278, 246–249.

Ji, J.Z., Wang, X.M., Li, B.M., 2003. Deficit in long-term contextual fear memory induced by blockade of beta-adrenoceptors in hippocampal CA1 region. Eur. J. Neurosci. 17, 1947–1952.

Jung, M.W., McNaughton, B.L., 1993. Spatial selectivity of unit activity in the hippocampal granular layer. Hippocampus 3, 165–182.

Kempermann, G., Wiskott, L., Gage, F.H., 2004. Functional significance of adult neurogenesis. Curr. Opin. Neurobiol. 14, 186–191.

Kesner, R.P., 1998. Neural mediation of memory for time: role of hippocampus and medial prefrontal cortex. Psychol. Bull. Rev. 5, 585–596.

Kesner, R.P., 2007. Behavioral functions of the CA3 subregion of the hippocampus. Learn. Mem. (Cold Spring Harbor, NY) 14, 771–781.

Kesner, R.P., 2013. An analysis of the dentate gyrus function. Behav. Brain Res. 254, 1–7.

Kesner, R.P., Gilbert, P.E., 2006. The role of the medial caudate nucleus, but not the hippocampus, in a matching-to sample task for a motor response. Eur. J. Neurosci. 23, 1888–1894.

Kesner, R.P., Gilbert, P.E., Barua, L.A., 2002. The role of the hippocampus in memory for the temporal order of a sequence of odors. Behav. Neurosci. 116, 286–290.

Kesner, R.P., Hui, X., Sommer, T., Wright, C., Barrera, V.R., Fanselow, M.S., 2014. The role of postnatal neurogenesis in supporting remote memory and spatial metric processing. Hippocampus 24, 1663–1671.

Kesner, R.P., Hunsaker, M.R., Gilbert, P.E., 2005. The role of CA1 in the acquisition of an object–trace–odor paired associate task. Behav. Neurosci. 119, 781–786.

Kesner, R.P., Hunsaker, M.R., Warthen, M.W., 2007. CA3 Mediates Arbitrary Associations and Episodic Memory Processing Based on Performance of Reciprocal Object Place Recall Tasks. Society for Neuroscience Abstract.

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., Hunsaker, M.R., Ziegler, W., 2011. The role of the dorsal and ventral hippocampus in olfactory working memory. Neurobiol. Learn. Mem. 96, 361–366.

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., Morris, A.M., Weeden, C.S.S., 2012. Spatial, temporal, and associative behavioral functions associated with different subregions of the hippocampus. In: Zentall, T.R., Wasserman, E.A. (Eds.), Oxford Handbook of Comparative Cognition. Oxford University Press, Oxford, pp. 322–346.

Kesner, R.P., Rolls, E.T., 2001. Role of long term synaptic modification in short term memory. Hippocampus 11, 240–250.

Kesner, R.P., Warthen, D.K., 2010. Implications of CA3 NMDA and opiate receptors for spatial pattern completion in rats. Hippocampus 20, 550–557.

Killian, N.J., Jutras, M.J., Buffalo, E.A., 2012. A map of visual space in the primate entorhinal cortex. Nature 491, 761–764.

Kirkwood, A., Dudek, S.M., Gold, J.T., Aizenman, C.D., Bear, M.F., 1993. Common forms of synaptic plasticity in the hippocampus and neocortex in vitro. Science 260, 1518–1521.

Kleinfeld, D., 1986. Sequential state generation by model neural networks. Proc. Natl. Acad. Sci. U. S. A. 83, 9469–9473.

Knierim, J.J., Neunuebel, J.P., Deshmukh, S.S., 2014. Functional correlates of the lateral and medial entorhinal cortex: objects, path integration and local-global reference frames. Philos. Trans. R. Soc. Lond. 369, 20130369.

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, 6th ed. Worth, New York.

Komorowski, R.W., Manns, J.R., Eichenbaum, H., 2009. Robust conjunctive item-place coding by hippocampal neurons parallels learning what happens where. J. Neurosci. 29, 9918–9929.

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.

Kraus, B.J., Brandon, M.P., Ronbinson, 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.

Kreiman, G., Koch, C., Fried, I., 2000. Category-specific visual responses of single neurons in the human medial temporal lobe. Nat. Neurosci. 3, 946–953.

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.

Lee, I., Hunsaker, M.R., Kesner, R.P., 2005. The role of hippocampal subregions in detecting spatial novelty. Behav. Neurosci. 119, 145–153.

Lee, I., Kesner, R.P., 2002. Differential contribution of NMDA receptors in hippocampal subregions to spatial working memory. Nat. Neurosci. 5, 162–168.

Lee, I., Kesner, R.P., 2003. Differential roles of dorsal hippocampal subregions in spatial working memory with short versus intermediate delay. Behav. Neurosci. 117, 1044–1053.

Lee, I., Kesner, R.P., 2004a. Differential contributions of dorsal hippocampal subregions to memory acquisition and retrieval in contextual fear-conditioning. Hippocampus 14, 301–310.

Lee, I., Kesner, R.P., 2004b. Encoding versus retrieval of spatial memory: double dissociation between the dentate gyrus and the perforant path inputs into CA3 in the dorsal hippocampus. Hippocampus 14, 66–76.

Lee, I., Rao, G., Knierim, J.J., 2004a. A double dissociation between hippocampal subfields: differential time course of CA3 and CA1 place cells for processing changed environments. Neuron 42, 803–815.

Lee, I., Yoganarasimha, D., Rao, G., Knierim, J.J., 2004b. Comparison of population coherence of place cells in hippocampal subfields CA1 and CA3. Nature 430, 456–459.

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.

Leutgeb, S., Leutgeb, J.K., 2007. Pattern separation, pattern completion, and new neuronal codes within a continuous CA3 map. Learn. Mem. (Cold Spring Harbor, NY) 14, 745–757.

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., 1996. A sequence predicting CA3 is a flexible associator that learns and uses context to solve hippocampal-like tasks. Hippocampus 6, 579–590.

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.

Levy, W.B., Wu, X., Baxter, R.A., 1995. Unification of hippocampal function via computational/encoding considerations. Int. J. Neural Syst. 6, 71–80.

Li, H., Matsumoto, K., Watanabe, H., 1999. Different effects of unilateral and bilateral hippocampal lesions in rats on the performance of radial maze and odor-paired associate tasks. Brain Res. Bull. 48, 113–119.

Li, X.G., Somogyi, P., Ylinen, A., Buzsaki, G., 1994. The hippocampal CA3 network: an in vivo intracellular labeling study. J. Comp. Neurol. 339, 181–208.

Lisman, J., Redish, A.D., 2009. Prediction, sequences and the hippocampus. Philos. Trans. R. Soc. Lond. 364, 1193–1201.

Lisman, J.E., 1999. Relating hippocampal circuitry to function: recall of memory sequences by reciprocal dentate–CA3 interactions. Neuron 22, 233–242.

Lisman, J.E., Idiart, M.A., 1995. Storage of 7 ± 2 short-term memories in oscillatory subcycles. Science 267, 1512–1515.

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.

Lorente de No, R., 1934. Studies on the structure of the cerebral cortex. II. Continuation of the study of the Ammonic system. Bd. 46, Heft 2 3, 113–177.

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

Maaswinkel, H., Jarrard, L.E., Whishaw, I.Q., 1999. Hippocampectomized rats are impaired in homing by path integration. Hippocampus 9, 553–561.

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.

Madsen, J., Kesner, R.P., 1995. The temporal-distance effect in subjects with dementia of the Alzheimer type. Alzheimer Dis. Assoc. Disord. 9, 94–100.

Maguire, E.A., 2014. Memory consolidation in humans: new evidence and opportunities. Exp. Physiol. 99, 471–486.

Malkova, L., Bachevalier, J., Mishkin, M., Saunders, R.C., 2001. Neurotoxic lesions of perirhinal cortex impair visual recognition memory in rhesus monkeys. Neuroreport 12, 1913–1917.

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.
Markov, N.T., Vezoli, J., Chameau, P., Falchier, A., Quilodran, R., Huissoud, C., Lamy, C., Misery, P., Giroud, P., Ullman, S., Barone, P., Dehay, C., Knoblauch, K., Kennedy, H., 2014. Anatomy of hierarchy: feedforward and feedback pathways in macaque visual cortex. J. Comp. Neurol. 522, 225–259.
Markram, H., Pikus, D., Gupta, A., Tsodyks, M.V., 1998. Information processing with frequency-dependent synaptic connections. Neuropharmacology 37, 489–500.
Markus, E.J., Qin, Y.L., Leonard, B., Skaggs, W., McNaughton, B.L., Barnes, C.A., 1995. Interactions between location and task affect the spatial and directional firing of hippocampal neurons. J. Neurosci. 15, 7079–7094.
Marques Pereira, P., Cosquer, B., Schimchowitsch, S., Cassel, J.C., 2005. Hebb–Williams performance and scopolamine challenge in rats with partial immunotoxic hippocampal cholinergic deafferentation. Brain Res. Bull. 64, 381–394.
Marr, D., 1971. Simple memory: a theory for archicortex. Philos. Trans. R. Soc. Lond. B 262, 23–81.
Marshuetz, C., 2005. Order information in working memory: an integrative review of evidence from brain and behavior. Psychol. Bull. 131, 323–339.
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.
Martinez, C.O., Do, V.H., Martinez Jr., J.L., Derrick, B.E., 2002. Associative long-term potentiation (LTP) among extrinsic afferents of the hippocampal CA3 region in vivo. Brain Res. 940, 86–94.
Matsumura, N., Nishijo, H., Tamura, R., Eifuku, S., Endo, S., Ono, T., 1999. Spatial- and task-dependent neuronal responses during real and virtual translocation in the monkey hippocampal formation. J. Neurosci. 19, 2318–2393.
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.
McDonald, R.J., White, N.M., 1995. Hippocampal and nonhippocampal contributions to place learning in rats. Behav. Neurosci. 109, 579–593.
McEchron, M.D., Tseng, W., Disterhoft, J.F., 2003. Single neurons in CA1 hippocampus encode trace interval duration during trace heart rate (fear) conditioning in rabbit. J. Neurosci. 23, 1535–1547.
McGaugh, J.L., 2000. Memory – a century of consolidation. Science 287, 248–251.
McLamb, R.L., Mundy, W.R., Tilson, H.A., 1988. Intradentate colchicine impairs acquisition of a two-way active avoidance response in a Y-maze. Neurosci. Lett. 94, 338–342.
McLennan, H., Miller, J.J., 1974. The hippocampal control of neuronal discharges in the septum of the rat. J. Physiol. 237, 607–624.
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., Battaglia, F.P., Jensen, O., Moser, E.I., Moser, M.-B., 2006. Path integration and the neural basis of the 'cognitive map'. Nat. Rev. Neurosci. 7, 663–678.
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.
McNaughton, N., Miller, J.J., 1986. Collateral specific long term potentiation of the output of field CA3 of the hippocampus of the rat. Exp. Brain Res. (Experimentelle Hirnforschung) 62, 250–258.
Miller, E.K., Li, L., Desimone, R., 1998. Activity of neurons in anterior inferior temporal cortex during a short-term memory task. J. Neurosci. 13, 1460–1478.
Miller, P., Wang, X.J., 2006. Inhibitory control by an integral feedback signal in prefrontal cortex: a model of discrimination between sequential stimuli. Proc. Natl. Acad. Sci. U. S. A. 103, 201–206.
Minai, A.A., Levy, W.B., 1993. Sequence learning in a single trial. Int. Neural Netw. Soc. World Congr. Neural Netw. 2, 505–508.
Miyashita, T., Kubik, S., Haghighi, N., Steward, O., Guzowski, J.F., 2009. Rapid activation of plasticity-associated gene transcription in hippocampal neurons provides a mechanism for encoding of one-trial experience. J. Neurosci. 29, 898–906.
Miyashita, Y., Kameyama, M., Hasegawa, I., Fukushima, T., 1998. Consolidation of visual associative long-term memory in the temporal cortex of primates. Neurobiol. Learn. Mem. 1, 197–211.
Miyashita, Y., Okuno, H., Tokuyama, W., Ihara, T., Nakajima, K., 1996. Feedback signal from medial temporal lobe mediates visual associative mnemonic codes of inferotemporal neurons. Cogn. Brain Res. 5, 81–86.
Miyashita, Y., Rolls, E.T., Cahusac, P.M., Niki, H., Feigenbaum, J.D., 1989. Activity of hippocampal formation neurons in the monkey related to a conditional spatial response task. J. Neurophysiol. 61, 669–678.
Mizumori, S.J., Perez, G.M., Alvarado, M.C., Barnes, C.A., McNaughton, B.L., 1990. Reversible inactivation of the medial septum differentially affects two forms of learning in rats. Brain Res. 528, 12–20.
Monaghan, D.T., Cotman, C.W., 1985. Distribution of N-methyl-D-aspartate-sensitive L-[3H]glutamate-binding sites in the rat brain. J. Neurosci. 5, 2909–2919.
Moore, D.B., Lee, P., Paiva, M., Walker, D.W., Heaton, M.B., 1998. Effects of neonatal ethanol exposure on cholinergic neurons of the rat medial septum. Alcohol (Fayetteville, NY) 15, 219–226.
Morris, A.M., Churchwell, J.C., Kesner, R.P., Gilbert, P.E., 2012. Selective lesions of the dentate gyrus produce disruptions in place learning for adjacent spatial locations. Neurobiol. Learn. Mem. 97, 326–331.
Morris, A.M., Curtis, B.J., Churchwell, J.C., Maasberg, D.W., Kesner, R.P., 2013a. Temporal associations for spatial events: the role of the dentate gyrus. Behav. Brain Res. 256, 250–256.
Morris, A.M., Weeden, C.S., Churchwell, J.C., Kesner, R.P., 2013b. The role of the dentate gyrus in the formation of contextual representations. Hippocampus 23, 162–168.
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, E.I., 2004. Hippocampal place cells demand attention. Neuron 42, 183–185.
Moser, M.B., Moser, E.I., 1998. Functional differentiation in the hippocampus. Hippocampus 8, 608–619.
Moyer Jr., J.R., Deyo, R.A., Disterhoft, J.F., 1990. Hippocampectomy disrupts trace eye-blink conditioning in rabbits. Behav. Neurosci. 104, 243–252.
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.
Muller, R.U., Ranck Jr., J.B., Taube, J.S., 1996. Head direction cells: properties and functional significance. Curr. Opin. Neurobiol. 6, 196–206.
Mumby, D.G., 2001. Perspectives on object-recognition memory following hippocampal damage: lessons from studies in rats. Behav. Brain Res. 127, 159–181.
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.
Murray, E.A., Gaffan, D., Mishkin, M., 1993. Neural substrates of visual stimulus-stimulus association in rhesus monkeys. J. Neurosci. 13, 4549–4561.
Myers, C.E., Ermita, B.R., Harris, K., Hasselmo, M., Solomon, P., Gluck, M.A., 1996. A computational model of cholinergic disruption of septohippocampal activity in classical eyeblink conditioning. Neurobiol. Learn. Mem. 66, 51–66.
Myers, C.E., Scharfman, H.E., 2009. A role for hilar cells in pattern separation in the dentate gyrus: a computational approach. Hippocampus 19, 321–337.
Myers, C.E., Scharfman, H.E., 2011. Pattern separation in the dentate gyrus: a role for the CA3 backprojection. Hippocampus 21, 1190–1215.
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.
Nakamura, N.H., Flasbeck, V., Maingret, N., Kitsukawa, T., Sauvage, M.M., 2013. Proximodistal segregation of nonspatial information in CA3: preferential recruitment of a proximal CA3-distal CA1 network in nonspatial recognition memory. J. Neurosci. 33, 11506–11514.
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.
Nanry, K.P., Mundy, W.R., Tilson, H.A., 1989. Colchicine-induced alterations of reference memory in rats: role of spatial versus non-spatial task components. Behav. Brain Res. 35, 45–53.
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.
Neunuebel, J.P., Knierim, J.J., 2014. CA3 retrieves coherent representations from degraded input: direct evidence for CA3 pattern completion and dentate gyrus pattern separation. Neuron 81, 416–427.
O'Keefe, Speakman, J.A., 1987. Single unit activity in the rat hippocampus during a spatial memory task. Exp. Brain Res. 68, 1–27.
O'Brien, N., Lehmann, H., Lecluse, V., Mumby, D.G., 2006. Enhanced context-dependency of object recognition in rats with hippocampal lesions. Behav. Brain Res. 170, 156–162.
O'Kane, D., Treves, A., 1992. Why the simplest notion of neocortex as an autoassociative memory would not work. Network (Bristol, Engl.) 3, 379–384.
O'Keefe, J., 1979. A review of the hippocampal place cells. Prog. Neurobiol. 13, 419–439.

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'Mara, S.M., Rolls, E.T., Berthoz, A., Kesner, R.P., 1994. Neurons responding to whole-body motion in the primate hippocampus. J. Neurosci. 14, 6511–6523.
O'Reilly, R.C., McClelland, J.L., 1994. Hippocampal conjunctive encoding, storage, and recall: avoiding a trade-off. Hippocampus 4, 661–682.
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.
Ono, T., Nakamura, K., Nishijo, H., Eifuku, S., 1993. Monkey hippocampal neurons related to spatial and nonspatial functions. J. Neurophysiol. 70, 1516–1529.
Otmakhova, N.A., Lisman, J.E., 1999. Dopamine selectively inhibits the direct cortical pathway to the CA1 hippocampal region. J. Neurosci. 19, 1437–1445.
Otto, T., Eichenbaum, H., 1992. Complementary roles of the orbital prefrontal cortex and the perirhinal-entorhinal cortices in an odor-guided delayed-nonmatching-to-sample task. Behav. Neurosci. 106, 762–775.
Panzeri, S., Rolls, E.T., Battaglia, F., Lavis, R., 2001. Speed of feedforward and recurrent processing in multilayer networks of integrate-and-fire neurons. Network: Comput. Neural Syst. 12, 423–440.
Parisi, G., 1986. A memory which forgets. J. Phys. A: Math. Gen. 19, L617.
Parkinson, J.K., Murray, E.A., Mishkin, M., 1988. A selective mnemonic role for the hippocampus in monkeys: memory for the location of objects. J. Neurosci. 8, 4059–4167.
Petrides, M., 1985. Deficits on conditional associative-learning tasks after frontal- and temporal-lobe lesions in man. Neuropsychologia 23, 601–614.
Piterkin, P., Cole, E., Cossette, M.P., Gaskin, S., Mumby, D.G., 2008. A limited role for the hippocampus in the modulation of novel-object preference by contextual cues. Learn. Mem. (Cold Spring Harbor, NY) 15, 785–791.
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.
Poucet, B., 1989. Object exploration, habituation, and response to a spatial change in rats following septal or medial frontal cortical damage. Behav. Neurosci. 103, 1009–1016.
Preston, A.R., Eichenbaum, H., 2013. Interplay of hippocampus and prefrontal cortex in memory. Curr. Biol. 23, R764–R773.
Raisman, G., Cowan, W.M., Powell, T.P., 1966. An experimental analysis of the efferent projection of the hippocampus. Brain 89, 83–108.
Ranck, J.B.J., 1985. Head direction cells in the deep cell layer of dorsal presubiculum in freely moving rats. In: Buzsaki, G., Vanderwolf, C.H. (Eds.), Electrical Activity of the Archicortex. Akademiai Kiado, Budapest, pp. 217–220.
Rawlins, J.N., Feldon, J., Gray, J.A., 1979. Septo-hippocampal connections and the hippocampal theta rhythm. Exp. Brain Res. (Experimentelle Hirnforschung) 37, 49–63.
Rawlins, J.N.P., 1985. Associations across time: the hippocampus as a temporary memory store. Behav. Brain Sci. 8, 479–496.
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., http://dx.doi.org/10.1016/j.nlm.2014.05.009.
Remondes, M., Schuman, E.M., 2004. Role for a cortical input to hippocampal area CA1 in the consolidation of a long-term memory. Nature 431, 699–703.
Renart, A., Moreno, R., de la Rocha, J., Parga, N., Rolls, E.T., 2001. A model of the IT-PF network in object working memory which includes balanced persistent activity and tuned inhibition. Neurocomputing 38-40, 1525–1531.
Risold, P.Y., Swanson, L.W., 1997. Connections of the rat lateral septal complex. Brain Res. 24, 115–195.
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.
Robertson, R.G., Rolls, E.T., Georges-François, P., Panzeri, S., 1999. Head direction cells in the primate pre-subiculum. Hippocampus 9, 206–219.
Rogers, J.L., Hunsaker, M.R., Kesner, R.P., 2006. Effects of ventral and dorsal CA1 subregional lesions on trace fear conditioning. Neurobiol. Learn. Mem. 86, 72–81.
Rogers, J.L., Kesner, R.P., 2003. Cholinergic modulation of the hippocampus during encoding and retrieval. Neurobiol. Learn. Mem. 80, 332–342.
Rogers, J.L., Kesner, R.P., 2004. Cholinergic modulation of the hippocampus during encoding and retrieval of tone/shock-induced fear conditioning. Learn. Mem. (Cold Spring Harbor, NY) 11, 102–107.
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., 1992. Neurophysiological mechanisms underlying face processing within and beyond the temporal cortical visual areas. Philos. Trans. R. Soc. Lond. B 335, 11–21.
Rolls, E.T., 1994. Neurophysiological and neuronal network analysis of how the primate hippocampus functions in memory. In: Delacour, J. (Ed.), The Memory System of the Brain. World Scientific, London, pp. 713–744.
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., 2000a. Functions of the primate temporal lobe cortical visual areas in invariant visual object and face recognition. Neuron 27, 205–218.
Rolls, E.T., 2000b. Memory systems in the brain. Annu. Rev. Psychol. 51, 599–630.
Rolls, E.T., 2003. Consciousness absent and present: a neurophysiological exploration. Prog. Brain Res. 144, 95–106.
Rolls, E.T., 2004. The functions of the orbitofrontal cortex. Brain Cogn. 55, 11–29.
Rolls, E.T., 2005. Emotion Explained. Oxford University Press, Oxford.
Rolls, E.T., 2008a. Memory, Attention, and Decision-Making: A Unifying Computational Neuroscience Approach. Oxford University Press, Oxford.
Rolls, E.T., 2008b. Top-down control of visual perception: attention in natural vision. Perception 37, 333–354.
Rolls, E.T., 2010a. Attractor networks. WIREs Cogn. Sci. 1, 119–134.
Rolls, E.T., 2010b. 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. Architect. 1, 44–54.
Rolls, E.T., 2012b. Invariant visual object and face recognition: neural and computational bases, and a model, VisNet. Front. Comput. Neurosci. 35 (6), 1–70.
Rolls, E.T., 2013a. A biased activation theory of the cognitive and attentional modulation of emotion. Front. Hum. Neurosci. 7 (74), 1–15.
Rolls, E.T., 2013b. The mechanisms for pattern completion and pattern separation in the hippocampus. Front. Syst. Neurosci. 7, 74.
Rolls, E.T., 2013c. A quantitative theory of the functions of the hippocampal CA3 network in memory. Front. Cell Neurosci. 7, 98.
Rolls, E.T., 2014a. Emotion and Decision-Making Explained. Oxford University Press, Oxford.
Rolls, E.T., 2014b. Limbic systems for emotion and for memory, but no single limbic system. Cortex, http://dx.doi.org/10.1016/j.cortex.2013.12.005.
Rolls, E.T., 2015. Diluted connectivity in pattern association networks facilitates the recall of information from the hippocampus to the neocortex. Prog. Brain Res. (in press).
Rolls, E.T., Deco, G., 2002. Computational Neuroscience of Vision. Oxford University Press, Oxford.
Rolls, E.T., Deco, G., 2006. Attention in natural scenes: neurophysiological and computational bases. Neural Netw. 19, 1383–1394.
Rolls, E.T., Deco, G., 2010. The Noisy Brain: Stochastic Dynamics as a Principle of Brain Function. Oxford University Press, Oxford.
Rolls, E.T., Franco, L., Stringer, S.M., 2005a. The perirhinal cortex and long-term familiarity memory. Q. J. Exp. Psychol. B 58, 234–245.
Rolls, E.T., Grabenhorst, F., 2008. The orbitofrontal cortex and beyond: from affect to decision-making. Prog. Neurobiol. 86, 216–244.
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., O'Mara, S.M., 1995. View-responsive neurons in the primate hippocampal complex. Hippocampus 5, 409–424.
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., Stringer, S.M., 2000. On the design of neural networks in the brain by genetic evolution. Prog. Neurobiol. 61, 557–579.
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. Network. 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, Engl.) 1, 407–421.
Rolls, E.T., Treves, A., 1998. Neural Networks and Brain Function. Oxford University Press, Oxford.
Rolls, E.T., Treves, A., 2011. The neuronal encoding of information in the brain. Prog. Neurobiol. 95, 448–490.
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., 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., Webb, T.J., 2012. Cortical attractor network dynamics with diluted connectivity. Brain Res. 1434, 212–225.
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., 2005b. Object, space and object–space representations in the primate hippocampus. J. Neurophysiol. 94, 833–844.
Rondi-Reig, L., Libbey, M., Eichenbaum, H., Tonegawa, S., 2001. CA1-specific N-methyl-D-aspartate receptor knockout mice are deficient in solving a nonspatial transverse patterning task. Proc. Natl. Acad. Sci. U. S. A. 98, 3543–3548.
Rupniak, N.M.J., Gaffan, D., 1987. Monkey hippocampus and learning about spatially directed movements. J. Neurosci. 7, 2331–2337.
Samsonovich, A., McNaughton, B.L., 1997. Path integration and cognitive mapping in a continuous attractor neural network model. J. Neurosci. 17, 5900–5920.
Saper, C.B., Swanson, L.W., Cowan, W.M., 1979. An autoradiographic study of the efferent connections of the lateral hypothalamic area in the rat. J. Comp. Neurol. 183, 689–706.
Scharfman, H.E., Myers, C.E., 2012. Hilar mossy cells of the dentate gyrus: a historical perspective. Front Neural Circ. 6, 106.
Schmidt-Hieber, C., Jonas, P., Bischofberger, J., 2004. Enhanced synaptic plasticity in newly generated granule cells of the adult hippocampus. Nature 429, 184–187.
Schmidt, B., Marrone, D.F., Markus, E.J., 2012. Disambiguating the similar: the dentate gyrus and pattern separation. Behav. Brain Res. 226, 56–65.
Schultz, S., Rolls, E.T., 1999. Analysis of information transmission in the Schaffer collaterals. Hippocampus 9, 582–598.
Schwindel, C.D., McNaughton, B.L., 2011. Hippocampal-cortical interactions and the dynamics of memory trace reactivation. Prog. Brain Res. 193, 163–177.
Shapiro, M.L., Olton, D., 1994. Hippocampal function and interference. In: Schachter, D.L., Tulving, E. (Eds.), Memory Systems. MIT Press, London.
Sidhu, M.K., Stretton, J., Winston, G.P., Bonelli, S., Centeno, M., Vollmar, C., Symms, M., Thompson, P.J., Koepp, M.J., Duncan, J.S., 2013. A functional magnetic resonance imaging study mapping the episodic memory encoding network in temporal lobe epilepsy. Brain 136, 1868–1888.
Simmen, M.W., Rolls, E.T., Treves, A., 1996a. On the dynamics of a network of spiking neurons. In: Eekman, F.H., Bower, J.M. (Eds.), Computations and Neuronal Systems: Proceedings of CNS 1995. Kluwer, Boston.
Simmen, M.W., Treves, A., Rolls, E.T., 1996b. Pattern retrieval in threshold-linear associative nets. Network (Bristol, Engl.) 7, 109–122.
Skaggs, W.E., Knierim, J.J., Kudrimoti, H.S., McNaughton, B.L., 1995. A model of the neural basis of the rat's sense of direction. In: Tesauro, G., Touretzky, D.S., Leen, T.K. (Eds.), Advances in Neural Information Processing Systems. MIT Press, Cambridge, MA, pp. 173–180.
Smith, G.K., Kesner, R.P., Korenberg, J.R., 2014. Dentate gyrus mediates cognitive function in the Ts65Dn/DnJ mouse model of Down syndrome. Hippocampus 24, 354–362.
Smith, M.L., Milner, B., 1981. The role of the right hippocampus in the recall of spatial location. Neuropsychologia 19, 781–793.
Sompolinsky, H., Kanter, I.I., 1986. Temporal association in asymmetric neural networks. Phys. Rev. Lett. 57, 2861–2864.
Spanswick, S.C., Sutherland, R.J., 2010. Object/context-specific memory deficits associated with loss of hippocampal granule cells after adrenalectomy in rats. Learn. Mem. (Cold Spring Harbor, NY) 17, 241–245.
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.
Stella, F., Cerasti, E., Treves, A., 2013. Unveiling the metric structure of internal representations of space. Front. Neural Circ. 7, 81.
Storm-Mathiesen, J., Zimmer, J., Ottersen, O.P., 1990. Understanding the Brain Through the Hippocampus, Progress in Brain Research. Elsevier, Oxford.
Strekalova, T., Zorner, B., Zacher, C., Sadovska, G., Herdegen, T., Gass, P., 2003. Memory retrieval after contextual fear conditioning induces c-Fos and JunB expression in CA1 hippocampus. Genes Brain Behav. 2, 3–10.
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., 2005. Self-organizing continuous attractor network models of hippocampal spatial view cells. Neurobiol. Learn. Mem. 83, 79–92.
Stringer, S.M., Rolls, E.T., Trappenberg, T.P., 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.
Sutherland, R.J., Whishaw, I.Q., Kolb, B., 1983. A behavioural analysis of spatial localization following electrolytic, kainate- or colchicine-induced damage to the hippocampal formation in the rat. Behav. Brain Res. 7, 133–153.
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.
Swanson, L.W., Cowan, W.M., 1977. An autoradiographic study of the organization of the efferent connections of the hippocampal formation in the rat. J. Comp. Neurol. 172, 49–84.
Swanson, L.W., Cowan, W.M., 1979. The connections of the septal region in the rat. J. Comp. Neurol. 186, 621–655.
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.
Takehara-Nishiuchi, K., 2014. Entorhinal cortex and consolidated memory. Neurosci. Res. 84C, 27–33.
Takeuchi, T., Duszkiewicz, A.J., Morris, R.G., 2014. The synaptic plasticity and memory hypothesis: encoding, storage and persistence. Philos. Trans. R. Soc. Lond. 369, 20130288.
Tanila, H., 1999. Hippocampal place cells can develop distinct representations of two visually identical environments. Hippocampus 9, 235–246.
Taube, J.S., Goodridge, J.P., Golob, E.J., Dudchenko, P.A., Stackman, R.W., 1996. Processing the head direction signal: a review and commentary. Brain Res. Bull. 40, 477–486.
Taube, J.S., Muller, R.U., Ranck, J.B.J., 1990. Head-direction cells recorded from the postsubiculum in freely moving rats 1: description and quantitative analysis. J. Neurosci. 10, 420–435.
Taylor, J.G., 1999. Neural bubble dynamics in two dimensions: foundations. Biol. Cybern. 80, 393–409.
Teyler, T.J., DiScenna, P., 1986. The hippocampal memory indexing theory. Behav. Neurosci. 100, 147–154.
Tilson, H.A., Rogers, B.C., Grimes, L., Harry, G.J., Peterson, N.J., Hong, J.S., Dyer, R.S., 1987. Time-dependent neurobiological effects of colchicine administered directly into the hippocampus of rats. Brain Res. 408, 163–172.
Treves, A., 1990. Graded-response neurons and information encodings in autoassociative memories. Phys. Rev A 42, 2418–2430.
Treves, A., 1991. Dilution and sparse coding in threshold-linear nets. J. Phys. A 24, 327–335.
Treves, A., 1993. Mean-field analysis of neuronal spike dynamics. Network (Bristol, Engl.) 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, Engl.) 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.
Vago, D.R., Bevan, A., Kesner, R.P., 2007. The role of the direct perforant path input to the CA1 subregion of the dorsal hippocampus in memory retention and retrieval. Hippocampus 17, 977–987.
Vago, D.R., Kesner, R.P., 2008. Disruption of the direct perforant path input to the CA1 subregion of the dorsal hippocampus interferes with spatial working memory and novelty detection. Behav. Brain Res. 189, 273–283.
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.
Vann, S.D., Aggleton, J.P., 2004. The mammillary bodies: two memory systems in one? Nat. Rev. 5, 35–44.
Vazdarjanova, A., Guzowski, J.F., 2004. Differences in hippocampal neuronal population responses to modifications of an environmental context: evidence for distinct, yet complementary, functions of CA3 and CA1 ensembles. J. Neurosci. 24, 6489–6496.
Vinogradova, O.S., Brazhnik, E.S., Kichigina, V.F., Stafekhina, V.S., 1996. Modulation of the reaction of hippocampal neurons to sensory stimuli by cholinergic substances. Neurosci. Behav. Physiol. 26, 113–124.
Wallace, D.G., Whishaw, I.Q., 2003. NMDA lesions of Ammon's horn and the dentate gyrus disrupt the direct and temporally paced homing displayed by rats exploring a novel environment: evidence for a role of the hippocampus in dead reckoning. Eur. J. Neurosci. 18, 513–523.
Wallenstein, G.V., Hasselmo, M.E., 1997. GABAergic modulation of hippocampal population activity: sequence learning, place field development, and the phase precession effect. J. Neurophysiol. 78, 393–408.
Wallis, G., Rolls, E.T., Földiák, P., 1993. Learning invariant responses to the natural transformations of objects. Int. Joint Conf. Neural Netw. 2, 1087–1090.
Walsh, T.J., Schulz, D.W., Tilson, H.A., Schmechel, D.E., 1986. Colchicine-induced granule cell loss in rat hippocampus: selective behavioral and histological alterations. Brain Res. 398, 23–36.
Wang, S.H., Morris, R.G., 2010. Hippocampal–neocortical interactions in memory formation, consolidation, and reconsolidation. Annu. Rev. Psychol. 61, 49–79, C41–C44.
Wang, X.J., 2002. Probabilistic decision making by slow reverberation in cortical circuits. Neuron 36, 955–968.
Webb, T., Rolls, E.T., Deco, G., Feng, J., 2011. Noise in attractor networks in the brain produced by graded firing rate representations. PLoS ONE 6, e23630.
Weeden, C.S., Hu, N.J., Ho, L.U., Kesner, R.P., 2014. The role of the ventral dentate gyrus in olfactory pattern separation. Hippocampus 24, 553–559.
Weiss, C., Bouwmeester, H., Power, J.M., Disterhoft, J.F., 1999. Hippocampal lesions prevent trace eyeblink conditioning in the freely moving rat. Behav. Brain Res. 99, 123–132.
Whishaw, I.Q., Hines, D.J., Wallace, D.G., 2001. Dead reckoning (path integration) requires the hippocampal formation: evidence from spontaneous exploration and spatial learning tasks in light (allothetic) and dark (idiothetic) tests. Behav. Brain Res. 127, 49–69.
Wills, T.J., Lever, C., Cacucci, F., Burgess, N., O'Keefe, J., 2005. Attractor dynamics in the hippocampal representation of the local environment. Science 308, 873–876.
Willshaw, D.J., 1981. Holography, associative memory, and inductive generalization. In: Hinton, G.E., Anderson, J.A. (Eds.), Parallel Models of Associative Memory. Lawrence Erlbaum, Hillsdale, pp. 103–124.
Willshaw, D.J., Buckingham, J.T., 1990. An assessment of Marr's theory of the hippocampus as a temporary memory store. Philos. Trans. R. Soc. B 329, 205–215.
Wilson, F.A.W., Riches, I.P., Brown, M.W., 1990. Hippocampus and medial temporal cortex: neuronal responses related to behavioural responses during the performance of memory tasks by primates. Behav. Brain Res. 40, 7–28.
Wilson, M.A., 2002. Hippocampal memory formation, plasticity, and the role of sleep. Neurobiol. Learn. Mem. 78, 565–569.
Wilson, M.A., McNaughton, B.L., 1993. Dynamics of the hippocampal ensemble code for space. Science 261, 1055–1058.
Wilson, M.A., McNaughton, B.L., 1994. Reactivation of hippocampal ensemble memories during sleep. Science 265, 676–679.
Winocur, G., Moscovitch, M., 2011. Memory transformation and systems consolidation. J. Int. Neuropsychol. Soc.: JINS 17, 766–780.
Winters, B.D., Bussey, T.J., 2005. Transient inactivation of perirhinal cortex disrupts encoding, retrieval, and consolidation of object recognition memory. J. Neurosci. 25, 52–61.
Wirth, S., Yanike, M., Frank, L.M., Smith, A.C., Brown, E.N., Suzuki, W.A., 2003. Single neurons in the monkeys hippocampus and learning of new associations. Science 300, 1578–1581.
Wiskott, L., Sejnowski, T.J., 2002. Slow feature analysis: unsupervised learning of invariances. Neural Comput. 14, 715–770.
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. (Cold Spring Harbor, NY) 14, 705–713.
Witter, M.P., Groenewegen, H.J., Lopes da Silva, F.H., Lohman, A.H., 1989a. Functional organization of the extrinsic and intrinsic circuitry of the parahippocampal region. Prog. Neurobiol. 33, 161–253.
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., 1989b. 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., Agster, K.M., Eichenbaum, H., 2004. One-trial odor-reward association: a form of event memory not dependent on hippocampal function. Behav. Neurosci. 118, 526–539.
Wood, E.R., Dudchenko, P.A., Eichenbaum, H., 1999. The global record of memory in hippocampal neuronal activity. Nature 397, 613–616.
Wu, X., Baxter, R.A., Levy, W.B., 1996. Context codes and the effect of noisy learning on a simplified hippocampal CA3 model. Biol. Cybern. 74, 159–165.
Wyss, J.M., Swanson, L.W., Cowan, W.M., 1980. The organization of the fimbria, dorsal fornix and ventral hippocampal commissure in the rat. Anat. Embryol. (Berl.) 158, 303–316.
Xavier, G.F., Oliveira-Filho, F.J., Santos, A.M., 1999. Dentate gyrus-selective colchicine lesion and disruption of performance in spatial tasks: difficulties in place strategy because of a lack of flexibility in the use of environmental cues? Hippocampus 9, 668–681.
Xiang, J.-Z., Brown, M.W., 1998. Differential neuronal encoding of novelty, familiarity and recency in regions of the anterior tempral lobe. Neuropharmacology 37, 657–676.
Yassa, M.A., Stark, C.E., 2011. Pattern separation in the hippocampus. Trends Neurosci. 34, 515–525.
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 Circ. 6, 16.
Zola-Morgan, S., Squire, L.R., Amaral, D.G., Suzuki, W.A., 1989. Lesions of perirhinal and parahippocampal cortex that spare the amygdala and hippocampal formation produce severe memory impairment. J. Neurosci. 9, 4355–4370.
Zola-Morgan, S., Squire, L.R., Ramus, S.J., 1994. Severity of memory impairment in monkeys as a function of locus and extent of damage within the medial temporal lobe memory system. Hippocampus 4, 483–494.