

Pattern separation, completion, and categorisation in the hippocampus and neocortex



Edmund T. Rolls*

Oxford Centre for Computational Neuroscience, Oxford, England, United Kingdom
University of Warwick, Department of Computer Science, Coventry CV4 7AL, England, United Kingdom

ARTICLE INFO

Article history:

Received 1 May 2015

Revised 2 July 2015

Accepted 11 July 2015

Available online 17 July 2015

Keywords:

Hippocampus

Attractor network

Competitive network

Pattern association network

Episodic memory

Pattern separation

Completion

Categorisation

ABSTRACT

The mechanisms for pattern completion and pattern separation are described in the context of a theory of hippocampal function in which the hippocampal CA3 system operates as a single attractor or autoassociation network to enable rapid, one-trial, associations between any spatial location (place in rodents, or spatial view in primates) and an object or reward, and to provide for completion of the whole memory during recall from any part. The factors important in the pattern completion in CA3 and also a large number of independent memories stored in CA3 include: a sparse distributed representation, representations that are independent due to the randomizing effect of the mossy fibres, heterosynaptic long-term depression as well as long-term potentiation in the recurrent collateral synapses, and diluted connectivity to minimize the number of multiple synapses between any pair of CA3 neurons which otherwise distort the basins of attraction. Recall of information from CA3 is implemented by the entorhinal cortex perforant path synapses to CA3 cells, which in acting as a pattern associator allow some pattern generalization. Pattern separation is performed in the dentate granule cells using competitive learning to convert grid-like entorhinal cortex firing to place-like fields, and in the dentate to CA3 connections that have diluted connectivity. Recall to the neocortex is achieved by a reverse hierarchical series of pattern association networks implemented by the hippocampo-cortical backprojections, each one of which performs some pattern generalization, to retrieve a complete pattern of cortical firing in higher-order cortical areas. New results on competitive networks show which factors contribute to their ability to perform pattern separation, pattern clustering, and pattern categorisation, and how these apply in different hippocampal and neocortical systems.

© 2015 Elsevier Inc. All rights reserved.

1. Pattern separation and pattern completion in the hippocampal system

1.1. Introduction

There is great interest in how pattern separation and pattern completion in the hippocampus contribute to its functions in memory and spatial function (Giocomo, Moser, & Moser, 2011; Hasselmo & Wyble, 1997; Hunsaker & Kesner, 2008, 2013; Jezek, Henriksen, Treves, Moser, & Moser, 2011; Kesner, 2007, 2013; Kesner, Morris, & Weeden, 2012; Kesner & Rolls, 2015; Leutgeb, Leutgeb, Moser, & Moser, 2007; McHugh et al., 2007; Nakashiba et al., 2012; Nakazawa et al., 2002, 2003; Rolls & Kesner, 2006;

Rolls & Treves, 1998; Wills, Lever, Cacucci, Burgess, & O'Keefe, 2005).

The aim of Section 1 of this paper is to describe some of the different types of pattern separation and pattern completion in the hippocampal system, and the mechanisms that implement them, in part as an Introduction to the special issue of the *Neurobiology of Learning and Memory* (2016) on pattern separation and pattern completion. It is important to appreciate that there are different mechanisms each of which contributes to pattern separation or pattern completion in the hippocampal system, for this helps not only in the understanding of how the hippocampal system operates, but also helps in the evaluation of the effects of changes that influence each of these mechanisms. These different mechanisms are separated into different subsections of this paper, so that the operation and contributions of each mechanism can be clarified and evaluated. The different mechanisms for pattern separation and pattern completion are considered in the context of a theory of hippocampal function (Rolls, 2008, 2010b). More comprehensive

* Address: Oxford Centre for Computational Neuroscience, Oxford, England, United Kingdom.

E-mail address: Edmund.Rolls@oxcns.org

URL: <http://www.oxcns.org>

descriptions of this theory of hippocampal function, and of differences between the primate and rodent hippocampal neuronal representations and the implications for understanding human memory, are provided elsewhere (Kesner & Rolls, 2015; Rolls, 2008, 2010b, 2013; Rolls & Kesner, 2006; Rolls & Xiang, 2006). The theory presented here (Rolls' theory of hippocampal function) has been developed through many stages (Kesner & Rolls, 2015; Rolls, 1987, 1989a, 1989b, 1989d, 1990a, 1990b, 1991, 1995, 1996b, 2008, 2010b; Rolls & Deco, 2010; Rolls & Kesner, 2006; Rolls & Treves, 1998; Treves & Rolls, 1991, 1992, 1994), has as a predecessor developments made by David Marr (Marr, 1971) in understanding cortical structure and computation (though he never identified the CA3 system as an autoassociation network), and has benefitted greatly from collaborations with many whose names appear below in the citations, including Alessandro Treves and Simon Stringer.¹ The operation of pattern association networks (also known as heteroassociation networks because one pattern is associated with another pattern), autoassociation networks, and competitive networks has been described elsewhere (Hertz, Krogh, & Palmer, 1991; Rolls, 2008, 2016; Rolls & Treves, 1998).

In Section 2 of this paper, I address how, somewhat in contrast, important computational functions of the cerebral neocortex involve pattern categorisation, in which similar patterns are treated as being of the same object, and different patterns are treated as being of different objects.

In Section 3 of this paper, new and original computational analyses are described of these different processes of pattern separation, pattern completion, and pattern categorisation, which are fundamental to understanding the operation of the neocortex as well as hippocampal cortex (Rolls, 2016). These new results are shown in Figs. 5–8.

¹ Historical note: The theory was originally developed as described next. 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. Marr did not apply this to particular networks within the hippocampus, and dealt with binary neurons and binary synapses which utilised heavily the properties of the binomial distribution. 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, and used different techniques from theoretical physics. Early work of Gardner-Medwin (1976) showed how progressive recall could operate in a network of binary neurons with binary synapses. The analysis of these autoassociation or attractor networks was developed by Kohonen (1977, 1984), Kohonen, Oja, and Lehtio (1981) and Hopfield (1982), and the value of sparse representations was quantified by Treves and Rolls (1991). Marr's model was assessed by Willshaw and Buckingham (1990), who found that the hippocampal model described by Marr was not dependent on a recurrent collateral effect. Rolls presented a theory of the hippocampus 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 fibres could act to set up different representations for each memory to be stored in the CA3 cells. He suggested that the CA1 cells operate as a recoder for the information recalled from the CA3 cells to a partial memory cue, so that the recalled information would be represented more efficiently to enable recall, via the backprojection synapses, of activity in the neocortical areas similar to that which had been present during the original episode. This appeared in print as (Rolls, 1987), and further details were made explicit by Rolls (1989a, 1989b, 1989c, 1989d). 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, 1989b, 1989c, 1989d) papers and quantitatively by Treves and Rolls (1992, 1994)), nor how retrieval to the neocortex of hippocampal memories could be produced, for which a theory was developed by Rolls (1987, 1989a, 1989b, 1989c, 1989d) and made quantitative by Treves and Rolls (1994). McNaughton and Morris (1987) suggested that the CA3 network might be an autoassociation network, and that the mossy fibre to CA3 connections might implement 'detonator' synapses. However, the concepts that the diluted mossy fibre 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, McNaughton and O'Reilly (McClelland, McNaughton, & O'Reilly, 1995), and many others, are described below.

A subtheme of this paper is that diluted connectivity in all the types of network described here, pattern association, autoassociation, and competitive networks, provides an important component in how they operate in the brain.

1.2. Background to the approach to hippocampal function

1.2.1. Event and episodic memory

The focus is on a fundamental property of episodic memory, the ability to store and retrieve the memory of a particular single event involving an association between items such as the place and the object or reward seen at that place. Episodic memory in the sense of a series of linked events requires this type of event memory, and could be implemented by linking together a series of events.

An event consists of a set of items that occur together, such as seeing a particular object or person's face in a particular place. An everyday example might be remembering where one was for dinner, who was present, what was eaten, what was discussed, and the time at which it occurred. The spatial context is almost always an important part of an episodic memory (Dere, Easton, Nadel, & Huston, 2008), and it may be partly for this reason that episodic memory is linked to the functions of the hippocampal system, which is involved in spatial processing and memory. The ability to recall a whole memory from a partial cue is an important property of episodic memory, and is referred to as completion.

1.2.2. Systems-level functions and connections of the primate hippocampus

Any theory of the hippocampus must state at the systems level what is computed by the hippocampus. Some of the relevant evidence about the functions of the hippocampus in memory comes from the effects of damage to the hippocampus, the responses of neurons in the hippocampus during behaviour, and the systems-level connections of the hippocampus, described in more detail elsewhere (Kesner & Rolls, 2015; Rolls, 2008, 2010b; Rolls & Kesner, 2006; Rolls & Xiang, 2006). Many of the memory functions 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 is suggested 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, and for completion to then occur in recall. An important property of this autoassociation and completion is that completion of a whole memory can occur from any part. This 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, 2008, 2014; Rolls & Treves, 1998).

Information stored in the hippocampus will need to be retrieved and affect other parts of the brain in order to be used. The information about episodic events recalled from the hippocampus could be used to help form semantic memories (Rolls, 1989b, 1989d, 1990a; Treves & Rolls, 1994). For example, remembering many particular journeys could help to build a geographic cognitive map in the neocortex. The hippocampus and neocortex would thus be complementary memory systems, with the hippocampus being used for rapid, "on the fly", unstructured storage of information involving activity potentially arriving from many areas of the neocortex; while the neocortex would gradually build and adjust on the basis of much accumulating information, often recalled from the hippocampal unstructured store, the semantic

representation (McClelland, McNaughton, & O'Reilly, 1995; Moscovitch et al., 2005; Rolls, 1989b; Treves & Rolls, 1994). The theory shows how information could be retrieved within the hippocampus, and how this retrieved information could enable the activity in neocortical areas that was present during the original storage of the episodic event to be reinstated, thus implementing recall, by using hippocampo-neocortical backprojections as described elsewhere (Rolls, 1995, 1996b, 2008, 2010b; Treves & Rolls, 1994) (see Fig. 1).

To understand the functions of the primate hippocampus in event or episodic memory, it is necessary to understand which other parts of the brain it receives information from. Does it for example receive object as well as spatial information as indicated by its anatomical connectivity? The primate hippocampus receives inputs via the entorhinal cortex (area 28) and the highly developed parahippocampal gyrus (areas TF and TH) as well as the perirhinal cortex from the ends of many processing streams of the cerebral association cortex, including the visual and auditory temporal lobe association cortical areas, the prefrontal cortex, and the parietal cortex (Amaral, 1987; Amaral, Price, Pitkanen, & Carmichael, 1992; Lavenex, Suzuki, & Amaral, 2004; Rolls, 2008; Rolls & Kesner, 2006; Suzuki & Amaral, 1994b; Van Hoesen, 1982; Witter, Wouterlood, Naber, & Van Haften, 2000) (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 & Price, 1995; Pitkanen, Kelly, & Amaral, 2002; Stefanacci, Suzuki, & Amaral, 1996; Suzuki & Amaral, 1994a).

The primary output from the hippocampus to neocortex originates in CA1 and projects to subiculum, entorhinal cortex, and parahippocampal structures (areas TF–TH) as well as prefrontal including orbitofrontal cortex (Delatour & Witter, 2002; van Haften, Baks-te-Bulte, Goede, Wouterlood, & Witter, 2003; Van Hoesen, 1982; Witter, 1993) (see Fig. 1), though there are other outputs (Kesner & Rolls, 2015; Rolls, 2015b; Rolls & Kesner, 2006). These are the pathways that are likely to be involved in the retrieval of information from the hippocampus back to the neocortex.

The theory is a quantitative theory, and the numbers of synapses on the different types of neuron is an important feature of the circuitry emphasized next.

1.2.3. Hippocampal circuitry

Hippocampal circuitry (Amaral, 1993; Amaral & Witter, 1989; Andersen, Morris, Amaral, Bliss, & O'Keefe, 2007; Kondo, Lavenex, & Amaral, 2009; Lavenex et al., 2004; Naber, Lopes da Silva, & Witter, 2001; Storm-Mathiesen, Zimmer, & Ottersen, 1990; Witter, 2007; Witter, Wouterlood, et al., 2000) is illustrated in Fig. 1.

Projections from the entorhinal cortex layer 2 reach the granule cells (of which there are 10^6 in the rat) in the dentate gyrus (DG), via the perforant path (pp) (Witter, 1993). The granule cells project to CA3 cells via the mossy fibres (mf), which provide a *sparse* but powerful connection to the 3×10^5 CA3 pyramidal cells in the rat. Each CA3 cell receives approximately 46 mossy fibre inputs, so that the sparseness (or dilution) of this connectivity is thus 0.000046. By contrast, there are many more – possibly weaker – direct perforant path inputs also from layer 2 of the entorhinal

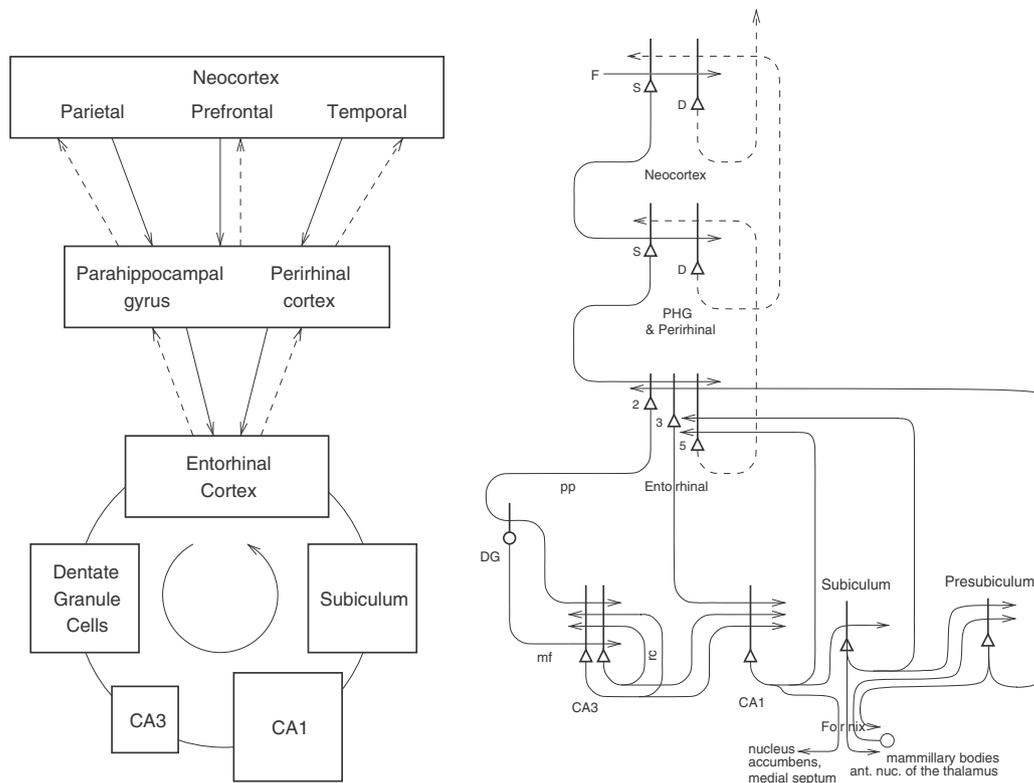


Fig. 1. Forward connections (solid lines) from areas of cerebral association neocortex via the parahippocampal gyrus and perirhinal cortex, and entorhinal cortex, to the hippocampus; and backprojections (dashed lines) via the hippocampal CA1 pyramidal cells, subiculum, and parahippocampal gyrus to the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

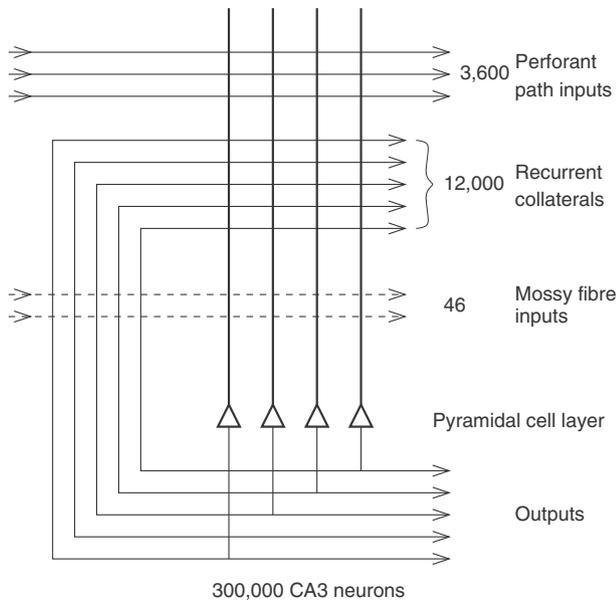


Fig. 2. The numbers of connections from three different sources onto each CA3 cell from three different sources in the rat. (After Rolls & Treves, 1998; Treves & Rolls, 1992.)

cortex onto each CA3 cell, in the rat of the order of 3600. The largest number of synapses (about 1.2×10^4 in the rat) on the dendrites of CA3 pyramidal cells is, however, provided by the (recurrent) axon collaterals of CA3 cells themselves (rc) (see Fig. 2). It is remarkable that the recurrent collaterals are distributed to other CA3 cells largely throughout the hippocampus (Amaral, Ishizuka, & Claiborne, 1990; Amaral & Witter, 1989, 1995; Ishizuka, Weber, & Amaral, 1990; Witter, 2007), so that effectively the CA3 system provides a single network, with a connectivity of approximately 2% between the different CA3 neurons given that the connections are bilateral. The CA3–CA3 recurrent collateral system is even more extensive in macaques than in rats (Kondo et al., 2009). 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).

1.3. CA3 as an autoassociation or attractor memory: pattern completion

1.3.1. 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, 1989, 2003; Morris et al., 2003; Nakazawa, McHugh, Wilson, & Tonegawa, 2004; Nakazawa et al., 2003; Wang & Morris, 2010). On the basis of the evidence summarized above, Rolls (1987, 1989a, 1989b, 1989d, 1990a, 1990b, 1991) and others (Levy, 1989; McNaughton, 1991; McNaughton & 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 recall 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 & Deco, 2002, 2010; Rolls &

Treves, 1998) including *Memory, Attention, and Decision-Making* (Rolls, 2008).)

The architecture of an autoassociation network is effectively that of the recurrent collateral synapses shown in Fig. 2, and the learning rule for the change in the synaptic weights is as shown in Eq. (1) (Rolls, 2008; Rolls & Treves, 1998).

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

where k is a constant, r_i is the activation of the dendrite (the post-synaptic term), r'_j is the presynaptic firing rate, and δw_{ij} is the change in the synaptic weight w_{ij} . (w_{ij} refers to the j 'th synapse onto the i 'th neuron). An introduction to autoassociation, competitive, and pattern association networks is provided in the Appendices of *Memory, Attention and Decision-Making: A Unifying Computational Neuroscience Approach* (Rolls, 2008); and *Cerebral Cortex: Principles of Operation* (Rolls, 2016), which also provides tutorial Matlab simulations of these networks.

The hypothesis is that because the CA3 operates effectively as a single network, it can allow arbitrary associations between inputs originating from very different parts of the cerebral cortex to be formed. These might involve associations between information originating in the temporal visual cortex about the presence of an object, and information originating in the parietal cortex about where it is. I note that although there is some spatial gradient in the CA3 recurrent connections, so that the connectivity is not fully uniform (Ishizuka et al., 1990; Witter, 2007), nevertheless the network will still have the properties of a single interconnected autoassociation network allowing associations between arbitrary neurons to be formed, given the presence of many long-range connections which overlap from different CA3 cells, and the ability of attractor networks to operate with diluted connectivity shown in our computational studies prompted by this issue (Rolls, 1995, 2012a; Rolls & Webb, 2012; Treves, 1990; Treves & Rolls, 1991). 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 fibre distribution, clear in the rat, are much more subtle in the nonhuman primate brain (Kondo et al., 2009). The implication is that in primates, the CA3 network operates even more as a single network than in rodents.

Crucial issues include how many memories could be stored in this system (to determine whether the autoassociation hypothesis leads to a realistic estimate of the number of memories that the hippocampus could store); whether the whole of a memory could be completed from any part; whether the autoassociation memory can act as a short term memory, for which the architecture is inherently suited; and whether the system could operate with spatial representations, which are essentially continuous because of the continuous nature of space. These and related issues are considered in the remainder of Section 1.3 and in more detail elsewhere (Rolls, 2008; Rolls & Kesner, 2006).

1.3.2. Storage capacity

We have performed quantitative analyses of the storage and retrieval processes in the CA3 network (Rolls, 2012a; Rolls & Webb, 2012; Treves & Rolls, 1991, 1992; Webb, Rolls, Deco, & Feng, 2011). We have extended previous formal models of autoassociative memory (see Amit, 1989) by analysing a network with graded response units, so as to represent more realistically the continuously variable rates at which neurons fire, and with incomplete connectivity (Rolls, Treves, Foster, & Perez-Vicente, 1997; Rolls & Webb, 2012; Treves, 1990; Treves & Rolls, 1991; Webb et al., 2011). 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 (RC) synapses on to each neuron, by a factor that increases roughly with the inverse of the sparseness a of the neuronal representation. [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.] The neuronal population sparseness a of the representation can be measured by extending the binary notion of the proportion of neurons that are firing to any one stimulus or event as

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

where r_i is the firing rate of the i 'th neuron in the set of N neurons. The sparseness ranges from $1/N$, when only one of the neurons responds to a particular stimulus (a local or grandmother cell representation), to a value of 1.0, attained when all the neurons are responding to a given stimulus. Considering the maximum number p_{\max} of firing patterns, we find that

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

where k is a factor that depends weakly on the detailed structure of the rate distribution, on the connectivity pattern, etc., but is roughly in the order of 0.2–0.3 (Treves & Rolls, 1991). For example, for $C^{\text{RC}} = 12,000$ and $a = 0.02$, p_{\max} is calculated to be approximately 36,000. This analysis emphasizes the utility of having a sparse representation in the hippocampus, for this enables many different memories to be stored. [The sparseness a in this equation is strictly the population sparseness (Franco, Rolls, Aggelopoulos, & Jerez, 2007; Treves & 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 neuron sparseness or selectivity a^s would be measured by the distribution of firing rates to a set of stimuli, which would take a long time. The selectivity or sparseness a^s of a single neuron measured across a set of stimuli often takes a similar value to the population sparseness a^p in the brain, and does so if the tuning profiles of the neurons to the set of stimuli are uncorrelated (Franco et al., 2007). These concepts are elucidated by Franco et al. (2007).] (I note that the sparseness estimates obtained by measuring early gene changes, which are effectively population sparsenesses, would be expected to depend greatly on the range of environments or stimuli in which these were measured. If the environment was restricted to one stimulus, this would reflect the population sparseness. If the environment was changing, the measure from early gene changes would be rather undefined, as all the populations of neurons activated in an undefined number of testing situations would be likely to be activated.)

In order for most associative networks to store information efficiently, heterosynaptic Long Term Depression (as well as LTP) is required (Fazeli & Collingridge, 1996; Rolls, 2008; Rolls & Deco, 2002; Rolls & Treves, 1990, 1998; Treves & Rolls, 1991). Simulations that are fully consistent with the analytic theory are provided by Rolls (1995, 2012a), Simmen, Treves, and Rolls (1996) and Rolls, Treves, et al. (1997).

A number of points that arise, including measurement of the total amount of information (in bits per synapse) that can be retrieved from the network, the computational definition of a memory, the computational sense in which CA3 is an attractor network, and the possible computational utility of memory reconsolidation, are treated elsewhere (Rolls, 2008; Rolls & Kesner, 2006). Here I note that given that the memory capacity of the hippocampal CA3 system is limited, it is necessary to have some form of forgetting in this store, or other mechanism to ensure that its

capacity is not exceeded. (Exceeding the capacity can lead to a loss of much of the information retrievable from the network.) Heterosynaptic LTD could help this *forgetting*, by enabling new memories to overwrite old memories (Rolls, 1996a, 2008). The limited capacity of the CA3 system does also provide one of the arguments that some transfer of information from the hippocampus to neocortical memory stores may be useful (see Treves & 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.

1.3.3. Recall and completion

A fundamental property of the autoassociation model of the CA3 recurrent collateral network is that the recall can be symmetric, that is, the whole of the memory can be retrieved and completed from any part (Rolls, 2008; Rolls & Kesner, 2006; Rolls & Treves, 1998). For example, in an object–place autoassociation memory, an object could be recalled from a place retrieval cue, and vice versa. In a test of this, Day, Langston, and Morris (2003) trained rats in a study phase to learn in one trial an association between two flavours of food and two spatial locations. During a recall test phase they were presented with a flavour which served as a cue for the selection of the correct location. They found that injections of an NMDA receptor blocker (AP5) or AMPA/kainate receptor 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 learning one-trial odour–place associations, and that recall can be performed without further involvement of NMDA receptors.

Evidence that the CA3 system is not necessarily required during recall in a reference memory spatial task, such as the water maze spatial navigation for a single spatial location task, is that CA3 lesioned rats are not impaired during recall of a previously learned water maze task (Brun et al., 2002; Florian & Rouillet, 2004). However, if completion from an incomplete cue is needed, then CA3 NMDA receptors are necessary (presumably to ensure satisfactory CA3–CA3 learning) even in a reference memory task (Gold & Kesner, 2005; Nakazawa et al., 2002). Thus, the CA3 system appears to be especially needed in rapid, one-trial object–place recall, and when completion from an incomplete cue is required (see further Section 1.5).

Especially important though in assessing the implications of all such tests 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 behavioural 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.

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).

1.3.4. Continuous, spatial, patterns and CA3 representations

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

A ‘Continuous Attractor’ neural network (CANN) can maintain the firing of its neurons to represent any location along a continuous physical dimension such as spatial view, spatial position, and head direction. 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 or packet or bump 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, 2008), and have the same architecture. 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. 3. For comparison, autoassociation networks normally have discrete (separate) patterns (each pattern implemented by the firing of a particular subset of the neurons), with no continuous distribution of the patterns throughout the space (see Fig. 3). 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 towards one of the previously learned attractor states, and thus implements the recall of a particular previously learned pattern from an incomplete or noisy (distorted) version of one of the previously learned patterns.

Space is continuous, and object representations are discrete. If these representations are to be combined in for example an object–place memory, then we need to understand the operation of networks that combine these representations. Rolls, Stringer, and Trappenberg (2002) have shown that attractor networks can store both continuous patterns and discrete patterns (as illustrated in Fig. 3), 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. We showed this by storing associated continuous and

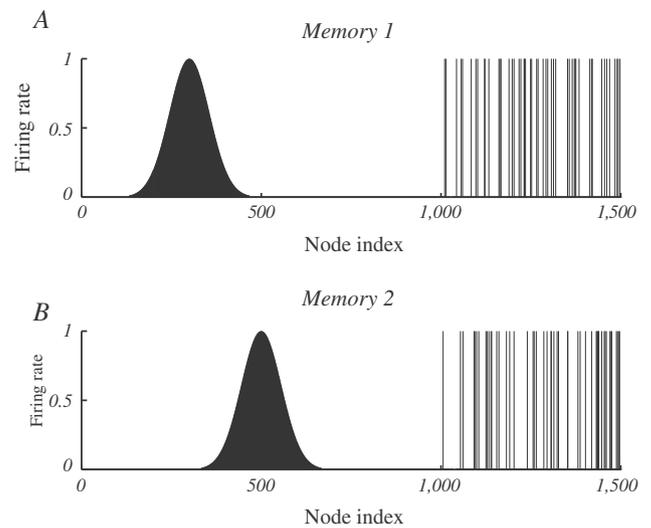


Fig. 3. The types of firing patterns stored in continuous attractor networks are illustrated for the patterns present on neurons 1–1000 for Memory 1 (when the firing is that produced when the spatial state represented is that for location 300), and for Memory 2 (when the firing is that produced when 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.

discrete representations in the same single attractor network, and then showing that the representation in the continuous space could be retrieved by the discrete object that was associated with that spatial position; and that the representation of the discrete object could be retrieved by providing the position in the continuous representation of space.

If spatial representations are stored in the hippocampus, the important issue arises in terms of understanding memories that include a spatial component or context of how many such spatial representations could be stored in a continuous attractor network. The very interesting result is that because there are in general low correlations between the representations of places in different maps or charts (where each map or chart might be of one room or locale), very many different maps or charts can be simultaneously stored in a continuous attractor network (Battaglia & Treves, 1998). Although it has been suggested that entorhinal cortex grid cell representations of location might be associated directly with objects (Hasselmo, 2009), the hippocampal CA3 cell representation of place in rodents or spatial view in primates would seem to be much more suitable for association with objects, because the complex conjunction of cells firing in the entorhinal cortex to represent a place is not sufficiently sparse and discrete.

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, a property of some spatial representations is that they can be updated by self-motion, idiothetic, input, and mechanisms have been proposed for how this could be achieved (Rolls & Stringer, 2005; Samsonovich & McNaughton, 1997; Stringer & Rolls, 2006; Stringer, Rolls, & Trappenberg, 2005;

Stringer, Rolls, et al., 2002; Stringer, Trappenberg, et al., 2002; Walters, Stringer, & Rolls, 2013), including in the entorhinal cortex grid cell system (Burak & Fiete, 2009; Burgess, Barry, & O'Keefe, 2007; Fuhs & Touretzky, 2006; Giocomo et al., 2011; Kropff & Treves, 2008; Zilli, 2012). The ways in which path integration could be implemented in recurrent networks such as the CA3 system in the hippocampus or in related systems are described elsewhere (McNaughton, Battaglia, Jensen, Moser, & Moser, 2006; Samsonovich & McNaughton, 1997; Stringer, Rolls, et al., 2002; Stringer, Trappenberg, et al., 2002), and have been applied to primate spatial view cells by Rolls and colleagues (Rolls & Stringer, 2005; Stringer, Rolls, & Trappenberg, 2004; Stringer et al., 2005). Cognitive maps (O'Keefe & Nadel, 1978) can be understood by the operations of these attractor networks, and how they are updated by learning and by self-motion (Rolls, 2008). 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 & Treves, 2013; Stella, Cerasti, & Treves, 2013).

1.3.5. 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 & Rolls, 1992) that an input of the perforant path type, alone, is unable to direct efficient information storage. Such an input is too weak, it turns out, to drive the firing of the cells, as the “dynamics” of the network is dominated by the randomizing effect of the recurrent collaterals. On the other hand, an autoassociative memory network needs afferent inputs to apply the retrieval cue to the network. We have shown (Treves & Rolls, 1992) that the perforant path system is likely to be the one involved in relaying the cues that initiate retrieval in CA3. The concept is that to initiate retrieval, a numerically large input (the perforant path system, see Fig. 2) 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, 2008; Treves & 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, 2008; Rolls & 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, 2008, 2014), 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 fibre 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, 2008; Treves & Rolls, 1992).

1.3.6. The dilution of the CA3 recurrent collateral connectivity enhances memory storage capacity and pattern completion

Fig. 2 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 that 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 & Treves, 1990; Rolls & Webb, 2012; Treves, 1990, 1991; Treves & 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. 2). 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, 2008). 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) (Banks, Sziklas, Sodums, & Jones-Gotman, 2012; Barkas, Henderson, Hamilton, Redhead, & Gray, 2010; Bohbot, Allen, & Nadel, 2000; Bonelli et al., 2010; Rolls, 2015c; 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 & 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., 2002, 2003, 2004; Rondi-Reig, Libbey, Eichenbaum, & Tonegawa, 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, 2008, 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).

1.3.7. 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, Rolls, Albantakis, & Romo, 2013; Faisal, Selen, & Wolpert, 2008; Rolls & 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 & 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 & Brunel, 1997; Miller & Wang, 2006; van Vreeswijk & Sompolinsky, 1998). 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 & Rolls, 2006; Rolls & 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, 2008). 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, 2014; Rolls & 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 & 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 & Webb, 2012).

Representations in the neocortex and in the hippocampus are often distributed with graded firing rates in the neuronal populations (Rolls & Treves, 2011). The firing rate probability distribution of each neuron to a set of stimuli is often exponential or gamma (Rolls & 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, Rolls, & Robertson, 1999; McNaughton, Barnes, & O'Keefe, 1983; O'Keefe & Speakman, 1987; O'Keefe, 1979; Robertson, Rolls, & Georges-François, 1998; Rolls, 2008; Rolls, Robertson, & Georges-François, 1997; Rolls & Treves, 2011; Rolls, Treves, Robertson, Georges-François, & Panzeri, 1998). 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).

1.4. 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; McNaughton & Morris, 1987; Rolls & 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 (Amit, 1989; Hopfield, 1982; Rolls & Treves, 1998).

We have proposed that there are several mechanisms that help to achieve this pattern separation, namely the mossy fibre pattern separation effect produced by the small number of connections received by a CA3 neuron from mossy fibres which dominate the CA3 cell firing; the expansion recoding, and the sparse representation provided by the dentate granule cells that form the mossy fibre synapses; and the sparseness of the CA3 cell representation. The expansion recoding (described by Rolls (2008, 2016)), although suggested for this entorhinal cortex to dentate connectivity (McNaughton & Morris, 1987), is modest, in that there are approximately 660,000 neurons in the rat entorhinal cortex (Merrill, Chiba, & Tuszyński, 2001), projecting to approximately 10^6 dentate granule cells. Neurogenesis of dentate granule cells is a fifth potential contributor to achieving pattern separation of CA3 cell firing. The five factors are described next. Before this, 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 fibre system to produce new random sets of neurons activated (Rolls, 2008).

1.4.1. Pattern separation and the sparse connectivity of the mossy fibre inputs to CA3 cells

We hypothesize that the mossy fibre inputs force efficient information storage by virtue of their strong and sparse influence on the CA3 cell firing rates (Rolls, 1987, 1989b, 1989d; Treves & Rolls, 1992). (The strong effects likely to be mediated by the mossy fibres were also emphasized by McNaughton and Morris (1987) and McNaughton and Nadel (1990).) We (Rolls and Treves) (Rolls, 1987, 1989b, 1989d, 1990b, 2008; Rolls & Treves, 1998; Treves & Rolls, 1992) hypothesize that the mossy fibre input appears to be particularly appropriate in several ways. First, the fact that mossy fibre synapses are large and located very close to the soma makes them relatively powerful in activating the postsynaptic cell. Second, the firing activity of dentate granule cells appears to be very sparse (Jung & 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 sparse firing activity in CA3 by a threshold effect. The hypothesis is that the mossy fibre sparse connectivity solution performs the appropriate function to enable learning to operate correctly in CA3 (Cerasti & Treves, 2010; Treves & 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 & Rolls, 1992) (see further Section 1.3.2.6).

The particular property of the small number of mossy fibre connections onto a CA3 cell, approximately 46 (see Fig. 2), 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, 1989d, 2008; Rolls & Kesner, 2006; Rolls & Treves, 1998; Treves & 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, 2008; Rolls & Treves, 1998; Treves & Rolls, 1991).

The requirement for a small number of mossy fibre connections onto each CA3 neuron applies not only to discrete (Treves & 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 fibres that may be active at any one time to select a set of random neurons in the CA3 (Cerasti & Treves, 2010). Any learning may help by reducing the accuracy required for a particular number of mossy fibre connections to be specified genetically onto each CA3 neuron. The optimal number of mossy fibres for the best information transfer from dentate granule cells to CA3 cells is in the order of 35–50 (Cerasti & Treves, 2010; Treves & Rolls, 1992). The mossy fibres also make connections useful for feedforward inhibition in CA3 (Acsady, Kamondi, Sik, Freund, & Buzsáki, 1998), which is likely to be useful to help in the sparse representations being formed in CA3.

On the basis of these and other points, we predicted that the mossy fibres may be necessary for new learning in the hippocampus, but may not be necessary for the recall of existing memories from the hippocampus (Rolls, 2008; Rolls & Treves, 1998; Treves & Rolls, 1992). Experimental evidence consistent with this prediction about the role of the mossy fibres in learning has been found in rats with disruption of the dentate granule cells (Lassalle, Bataille, & Halley, 2000) (Section 1.5).

We (Rolls & Kesner, 2006) have hypothesized that non-associative plasticity of mossy fibres (see Brown, Ganong, Kairiss, Keenan, & Kelso, 1989; Brown, Kairiss, & Keenan, 1990) might have a useful effect in enhancing the signal-to-noise ratio, in that a consistently firing mossy fibre would produce nonlinearly amplified currents in the postsynaptic cell, which would not happen with an occasionally firing fibre (Treves & 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, Wittner, & Buzsáki, 2002).

As acetylcholine turns down the efficacy of the recurrent collateral synapses between CA3 neurons (Giocomo & Hasselmo, 2007; Hasselmo, Schnell, & Barkai, 1995), then cholinergic activation also might help to allow external inputs from the mossy fibres 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 behavioural activation (Giocomo & Hasselmo, 2007; Hasselmo et al., 1995), when presumably it may be particularly relevant to allocate some of the limited memory capacity to new memories.

1.4.2. Pattern separation and the sparseness of the firing of the dentate granule cell input via the mossy fibres to CA3 cells

The firing activity of dentate granule cells appears to be very sparse (Jung & McNaughton, 1993; Leutgeb et al., 2007) and this, together with the small number of dentate mossy fibre connections on each CA3 cell, produces a sparse signal, which can then be transformed into sparse firing activity in CA3 by a threshold

effect. The pattern separation mechanisms that enable the dentate to provide a sparse firing input to CA3 are described below.

1.4.3. Pattern separation and the large number of dentate granule cells providing inputs via the mossy fibres to CA3 cells

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

1.4.4. Sparseness of the CA3 cell representation and pattern separation

The firing of CA3 cells is relatively sparse, and this helps to decorrelate different population vectors of CA3 cell firing for different memories. (Sparse representations are more likely to be decorrelated with each other (Rolls, 2008).) 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 & Leutgeb, 2007). In the macaque hippocampus, in which spatial view cells are found (Georges-François et al., 1999; Robertson et al., 1998; Rolls, Robertson, et al., 1997; Rolls et al., 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, 2008; Rolls & 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, 2008; Rolls & 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, 2008; Rolls & 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.

1.4.5. Neurogenesis of dentate granule cells to provide new representations in CA3 uncorrelated with previous CA3 representations

If adult neurogenesis in the dentate gyrus does prove to be functionally relevant, its computational role could be to facilitate pattern separation for new patterns, by providing new dentate granule cells with new sets of random connections to CA3 neurons (Becker, 2005; Rolls, 2010b). Consistent with the dentate spatial pattern separation hypothesis (Rolls, 1989b, 1989d, 1996b, 2008; Treves & 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).

1.4.6. The direct perforant path to CA3 cell input: poor at pattern separation and forcing a new memory pattern into CA3 cell firing

It has been suggested that the feedforward connectivity from the entorhinal cortex via the perforant path to the CA3 neurons may act as a feedforward pattern association network that is more important than the CA3–CA3 recurrent collateral autoassociation system (Cheng, 2013). The quantitative properties of pattern association networks are described elsewhere (Rolls, 2008, 2016; Rolls & 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 do generalize*. (As noted above, pattern association networks do not perform pattern completion, in that the unconditioned stimulus cannot recall the conditioned stimulus.) The analyses described in these sources shows that the capacity of pattern association networks (the maximum number of memories that can be stored and retrieved, here denoted by p_{\max}) is approximately

$$p_{\max} \approx \frac{C^{\text{PA}}}{[a_o \ln(1/a_o)]} \quad (4)$$

where C^{PA} is the number of feedforward associatively modifiable connections per neuron, and a_o is the sparseness of the representation in the output neurons of the pattern associator (Rolls, 2008). Given that there are fewer feedforward (perforant path) synaptic connections onto CA3 neurons (3600) than recurrent synaptic connections between CA3 neurons (12,000 in the rat) (see Fig. 2), then the capacity of the feedforward system would be considerably smaller than that of the recurrent collateral CA3–CA3 system. (It is noted that the a_o of Eq. (4) would be the same number as the a of Eq. (3), as that is just the sparseness of the firing of the population of CA3 neurons. The number of perforant path synapses is sufficiently large that it can generalize given even a partial retrieval pattern, so that the CA3–CA3 connections can then complete the retrieval, given that the recall signal for the perforant path pattern associator is proportional to the square root of the number of perforant path synapses, as shown by Eq. (17) of Treves and Rolls (1992).) The feedforward hypothesis (Cheng, 2013) thus has a strong argument against it of storage capacity, which would be much less (approximately 3600/12,000) than that of the CA3–CA3 recurrent collateral system operating as an autoassociation memory. Another disadvantage of the feedforward hypothesis is that the attractor properties of the CA3–CA3 connections would be lost, and these potentially contribute to holding one or more items simultaneously active in short-term memory (Rolls, 2008; Rolls, Dempere-Marco, & Deco, 2013), and providing a basis for temporal order memory as described in Section 1.9. Another disadvantage is that we have been able to show (Treves & 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. Another disadvantage of the feedforward hypothesis is that a pattern associator may not with an incomplete cue be able to recall the complete pattern that was stored, whereas an attractor network has the property that it can fall into an attractor basin that can reflect perfect retrieval of the complete memory (Rolls, 2008; Rolls & Treves, 1998).

1.5. Pattern separation performed by dentate granule cells

The theory is that the dentate granule cell stage of hippocampal processing which precedes the CA3 stage acts as a competitive network in a number of ways to produce during learning the sparse yet efficient (i.e. non-redundant) representation in CA3 neurons that is required for the autoassociation implemented by CA3 to

perform well (Rolls, 1989b, 1989d, 1990b; Rolls & Kesner, 2006; Rolls, Stringer, & Elliot, 2006; Treves & 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; Rolls & Kesner, 2006; Rolls et al., 2006; Treves & Rolls, 1992), enabling the hippocampus to store different memories of even similar events, and this prediction has been confirmed (Gilbert, Kesner, & Lee, 2001; Goodrich-Hunsaker, Hunsaker, & Kesner, 2008; Kesner et al., 2012; Leutgeb & Leutgeb, 2007; McHugh et al., 2007; Rolls, 2008; Rolls & Kesner, 2006). Consistently with this evidence for pattern separation by dentate granule cells, in rats small changes in the shape of the environment in which rats are exploring can substantially alter the activity patterns among place-modulated granule cells (Leutgeb et al., 2007).

As just described, the dentate granule cells could be important in helping to build and prepare spatial representations for the CA3 network. The actual representation of space in the primate hippocampus includes a representation of spatial view (Georges-François et al., 1999; Robertson et al., 1998; Rolls, Robertson, et al., 1997; Rolls & Xiang, 2006; Rolls et al., 1998), 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, Rolls, & Stringer, 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 deg, 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 deg, a combination of visual features formed over such a wide visual angle would effectively define a position in space that is a place (de Araujo et al., 2001).

The entorhinal cortex contains grid cells, which have high firing in the rat in a two-dimensional spatial grid as a rat traverses an environment, with larger grid spacings in the ventral entorhinal cortex (Fyhn, Molden, Witter, Moser, & Moser, 2004; Hafting, Fyhn, Molden, Moser, & Moser, 2005; Moser, Rowland, & Moser, 2015). This may be a system optimized for path integration (McNaughton et al., 2006) which may self-organize during locomotion with longer time constants producing more widely spaced grids in the ventral entorhinal cortex (Kropff & Treves, 2008). How are the grid cell representations, which would not be suitable for association of an object or reward with a place to form an episodic memory, transformed into a place representation that would be appropriate for this type of episodic memory? I have proposed that this could be implemented by a competitive network (Rolls, 2008) in the dentate gyrus which operates to form place cells, implemented by each dentate granule cell learning to respond to particular combinations of entorhinal cortex cells firing, where each combination effectively specifies a place, and this has been shown to be feasible computationally (Rolls et al., 2006). The sparse representations in the dentate gyrus, implemented by the mutual inhibition through inhibitory interneurons and competitive learning, help to implement this 'pattern separation' effect (Rolls, 1989b, 1989d, 2008; Rolls & Treves, 1998). The investigations showed that learning in the perforant path to dentate granule cell representation, and the sparse representation in the dentate

granule cells, are both important in the formation of place-like fields in dentate granule cells from the grid cells in the entorhinal cortex (Georges-François et al., 1999; Robertson et al., 1998; Rolls, Robertson, et al., 1997; Rolls et al., 1998). To illustrate this, Fig. 4 shows from these simulations the responses of the simulated grid cells (a, b), the dentate receptive fields formed by feedforward connections and a sparse representation in the dentate gyrus (c, d), and the dentate receptive fields formed when Hebbian synaptic modification and training is included in the feedforward connections to implement competitive learning (e, f). It is only with the full competitive learning that the dentate receptive fields self-organised to become small place-like receptive fields (Rolls et al., 2006) similar to those found in the rat dentate granule cells.

In primates, there is now evidence that there is a 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, Jutras, & Buffalo, 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, 2008; Rolls, Robertson, et al., 1997; Rolls & Xiang, 2006; Rolls et al., 1998). Further support for this type of representation of space being viewed 'out there' rather than only where one is located as for rat place cells is that cells in the human hippocampus have been found that respond to landmarks (or spatial views), with others coding for location (Ekstrom et al., 2003; Rolls, 2015c).

1.6. CA1 cells and pattern completion prior to hippocampo-directed recall to the neocortex

The CA3 cells connect to the CA1 cells by the Schaeffer collateral synapses. The associative modifiability in this connection helps the full information present in CA3 to be retrieved in the CA1 neurons (Rolls, 1995; Schultz & Rolls, 1999; Treves, 1995; Treves & Rolls, 1994). Part of the hypothesis is that the separate sub-parts of an episodic memory, which must be represented separately in CA3 to allow for completion, can be combined together by competitive learning in CA1 to produce an efficient retrieval cue for the recall via the backprojection pathways to the neocortex of memories stored in the neocortex (Rolls, 1989a, 1989b, 1995, 1996b; Treves & Rolls, 1994). Associative recall by pattern association in the CA3 to CA1 feedforward connections is a prominent property which implements generalization, so that if completion in CA3 was not perfect, CA1 could generalize to the closest CA1 representation (Rolls, 1995, 2008; Schultz, Panzeri, Rolls, & Treves, 2000, chap. 14).

1.7. Backprojections to the neocortex, and memory retrieval from the hippocampus to provide a complete neocortical memory representation

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 (Kesner & Rolls, 2015; Rolls, 1995, 1996b, 2008, 2010b; Treves & 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

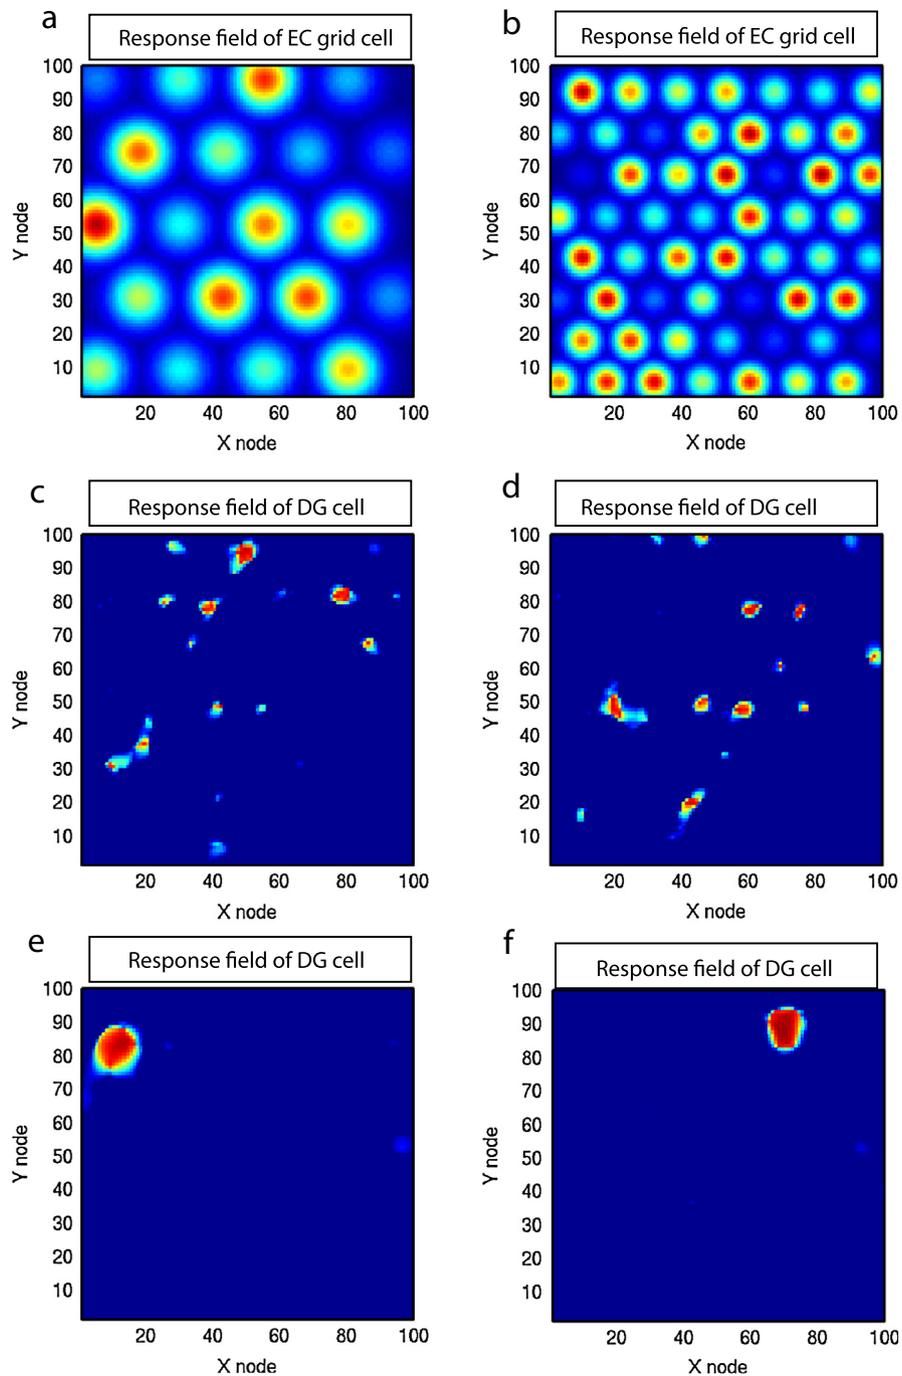


Fig. 4. Simulation of competitive learning in the dentate gyrus to produce place cells from the entorhinal cortex grid cell inputs. (a and b) Firing rate profiles of two entorhinal cortex (EC) grid cells with frequencies of 4 and 7 cycles. (c and d) Firing rate profiles of two dentate gyrus (DG) cells 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.)

be produced in CA1. The CA1 neurons would then activate, via their termination in the deep layers of the entorhinal cortex, at least the pyramidal cells in the deep layers of the entorhinal cortex (see Fig. 1). These entorhinal cortex layer 5 neurons would then, by virtue of their backprojections (Lavenex & Amaral, 2000; Witter, Naber, et al., 2000) 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, 1989b, 1989d). 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, 2014, 2015b)), and also areas of unimodal association cortex (e.g. inferior temporal visual cortex). The backprojections, by recalling previous episodic events, could provide information useful to the neocortex in the building of new representations in the multimodal and unimodal association cortical areas, which by building new long-term and structured representations can be considered as a form of memory consolidation

(Rolls, 1989a, 1989b, 1989d, 1990a, 1990b, 2008), or in organizing actions.

The hypothesis of the architecture with which this would be achieved is shown in Fig. 1. The feedforward connections from association areas of the cerebral neocortex (solid lines in Fig. 1), show major convergence as information is passed to CA3, with the CA3 autoassociation network having the smallest number of neurons at any stage of the processing. The backprojections allow for divergence back to neocortical areas. The way in which I suggest that the backprojection synapses are set up to have the appropriate strengths for recall is as follows (Kesner & Rolls, 2015; Rolls, 1989a, 1989b, 1989d; Treves & Rolls, 1994). During the setting up of a new episodic memory, there would be strong feedforward activity progressing towards the hippocampus. During the episode, the CA3 synapses would be modified, and via the CA1 neurons and the subiculum, a pattern of activity would be produced on the backprojecting synapses to the entorhinal cortex. Here the backprojecting synapses from active backprojection axons onto pyramidal cells being activated by the forward inputs to entorhinal cortex would be associatively modified. A similar process would be implemented at preceding stages of neocortex, that is in the parahippocampal gyrus/perirhinal cortex stage, and in association cortical areas, as shown in Fig. 1. The operation of the whole system from the entorhinal cortex through the hippocampus and back to implement memory recall in the entorhinal cortex to test the theory was simulated by Rolls (1995). Acetylcholine may modulate learning vs recall in the backprojection synapses from CA1 to the entorhinal cortex (Hasselmo & Wyble, 1997).

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, McNaughton, and O'Reilly (1995), that recall might be useful to the neocortex to help it build new semantic memories, which might inherently be a slow process and is not part of the theory of recall.

A plausible requirement for a successful hippocampo-directed recall operation, is that the signal generated from the hippocampally retrieved pattern of activity, and carried backwards towards neocortex, remain undegraded when compared to the noise due, at each stage, to the interference effects caused by the concurrent storage of other patterns of activity on the same backprojecting synaptic systems. That requirement is equivalent to that used in deriving the storage capacity of such a series of heteroassociative memories, and it was shown by Treves and Rolls (1991, 1994) 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 $C^{\text{HBP}} = C^{\text{RC}} a_{\text{nc}} / a_{\text{CA3}}$, where C^{RC} is the number of recurrent collaterals to any given cell in CA3, the average sparseness of a representation is a_{nc} , and a_{CA3} is the sparseness of memory representations there in CA3.

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. 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 & Rolls, 1994) (cf. Fig. 1).

The theory of recall by the backprojections thus provides a quantitative account of why the cerebral cortex has as many backprojection as forward projection connections (Kesner & Rolls, 2015; Rolls, 2008, 2016; Treves & Rolls, 1994).

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, 2014). 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 multi-stage 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.

Further aspects of the operation of the backprojecting systems are described elsewhere (Kesner & Rolls, 2015; Rolls, 2008).

1.8. Tests of pattern separation and pattern completion

There is now a large literature on tests of pattern separation and pattern completion in the hippocampus (Giocomo et al., 2011;

Hunsaker & Kesner, 2008, 2013; Jezek et al., 2011; Kesner, 2007, 2013; Kesner & Rolls, 2015; Kesner et al., 2012; Leutgeb et al., 2007; McHugh et al., 2007; Nakashiba et al., 2012; Nakazawa et al., 2002, 2003; Rolls & Kesner, 2006; Wills, Lever, Cacucci, Burgess, & O'Keefe, 2005), and a brief summary of some of the findings is provided next. An important point is that the theory (Kesner & Rolls, 2015; Rolls, 1987, 1989a, 1989b, 1989d, 1990a, 1990b, 1991, 1995, 1996b, 2008, 2010b, 2013; Rolls & Deco, 2010; Rolls & Kesner, 2006; Rolls & Treves, 1998; Treves & Rolls, 1991, 1992, 1994) is a quantitative theory of hippocampal function, and addresses how pattern separation and pattern completion are important in enabling the hippocampal system to operate up to capacity, which is in the order of tens of thousands of different memories. Some predictions from the theory may only hold when the system is well loaded, that is tested when the system is operating with thousands of memories, for then the pattern separation will be important. It is possible to test the predictions in simulations, where the system can be trained up to capacity (Rolls, 1995, 2012a; Rolls, Treves, et al., 1997). In vivo, it may be useful to test the storage and recall of as many memories as possible, and in addition testing animals kept in environments where memories of the hippocampal type are needed may also help to test hypotheses in situations where the hippocampus has been at least moderately well loaded with many different memories.

1.8.1. Dentate granule cells

The theory predicts that pattern separation is performed by competitive learning by the dentate granule cells. Evidence consistent with this has been found neurophysiologically in the small sparsely encoded place fields of dentate neurons (Jung & McNaughton, 1993; Leutgeb & Leutgeb, 2007) and their reflection in CA3 neurons (Leutgeb & Leutgeb, 2007). Further, and consistent with the theory, it has been shown that selective dentate lesions in rats (Gilbert & Kesner, 2003; Gilbert et al., 2001; Goodrich-Hunsaker et al., 2008; Hunsaker & Kesner, 2013; Kesner, 2013; Rolls, 2008; Rolls & Kesner, 2006) or dentate granule cell NMDA receptor knockouts in mice (McHugh et al., 2007) impair spatial, object–place (or reward–place: remembering where to find a reward) association tasks especially when the places are close together and require pattern separation before storage in CA3.

1.8.2. Mossy fibre inputs to CA3 and learning

The theory predicts that the dentate granule cell mossy fibre system of inputs to the CA3 neurons is necessary to store spatial memories, but not to recall them. Lassalle et al. (2000) have obtained evidence consistent with this in rats with damage to the mossy fibre system (Lassalle et al., 2000), and there is further evidence consistent with this (Daumas, Ceccom, Halley, Frances, & Lassalle, 2009; Kesner & Rolls, 2015; Lee & Kesner, 2004; Rolls & Kesner, 2006).

1.8.3. Perforant path inputs to CA3 and recall

The theory predicts that the direct perforant path input from the entorhinal cortex to the CA3 cells (which bypasses the dentate granule cells) is involved in the recall of memory from the CA3 system, and Lee and Kesner (2004) have obtained evidence consistent with this in a Hebb–Williams maze recall task.

1.8.4. CA3 and pattern completion

The theory predicts that the CA3 system is especially important in object–place or reward–place tasks in which associations must be formed between any spatial location and any object (referred to as *arbitrary associations*). There is much evidence from subregion analyses involving disruption of CA3 that CA3 is necessary for arbitrary associations between places and objects or rewards (Gilbert &

Kesner, 2003; Hunsaker & Kesner, 2013; Kesner & Rolls, 2015; Rolls & Kesner, 2006). Similar impairments were obtained following deletion of CA3 NMDA receptors in mice in the acquisition of an odour–context paired associate learning task (Rajji, Chapman, Eichenbaum, & Greene, 2006). If place or time is not a component, associative tasks such as odour–object association are not impaired (Kesner & Rolls, 2015; Rolls & Kesner, 2006) underlining the fact that the hippocampus is especially involved in episodic types of associative memory which typically involve place and/or time.

The theory predicts that the CA3 is especially important in object–place or reward–place *completion* tasks, in which associations must be completed from a part of the whole. It has been shown that if completion from an incomplete cue is needed, then CA3 NMDA receptors are necessary (presumably to ensure satisfactory CA3–CA3 learning) even in a reference memory task (Gold & Kesner, 2005; Hunsaker & Kesner, 2013; Nakazawa et al., 2002).

The theory predicts that the CA3 system is especially needed in *rapid, one-trial object–place, learning and recall*. It has been shown that hippocampal NMDA receptors (necessary for Long Term Potentiation to occur) are needed for one-trial flavour–place association learning, and that hippocampal AMPA/kainate receptors are sufficient for the recall, though the hippocampal subregion involved was not tested (Day et al., 2003). In subregion studies, Kesner and colleagues have shown that CA3 lesions produce chance performance on a one-trial object–place recall task (Kesner, Hunsaker, & Warthen, 2008) and other object–spatial tasks (Kesner & Rolls, 2001, 2015; Rolls & Kesner, 2006). For example, CA3 lesions produced chance performance on both a one-trial object–place recall and place–object recall task (Kesner et al., 2008). This is evidence that CA3 supports arbitrary associations as well as episodic memory based on 1-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, episodic) learning that is impaired (Kesner et al., 2008). CA3 NMDA receptors are as predicted by the theory necessary for rapid/one-trial spatial learning, as shown by a mouse knockout study by Nakazawa, Tonegawa and colleagues (Nakazawa et al., 2003, 2004; Tonegawa, Nakazawa, & Wilson, 2003). We have shown that hippocampal CA3 neurons reflect the computational processes necessary for one-trial object–place event memory, used as a model for episodic memory (Rolls & Xiang, 2006).

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, that hippocampal neurons will fall in 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, Lever, Cacucci, Burgess, & O'Keefe, 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 is an indication, predicted previously (Hasselmo & Wyble, 1997; Rolls & Treves, 1998), 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.

In a direct neurophysiological test of the CA3 pattern completion hypothesis, simultaneous recordings were made from CA3 and DG neurons of behaving rats when local and global reference frames were placed in conflict. CA3 showed a coherent population response to the conflict (pattern completion), even though its DG

inputs were severely disrupted (pattern separation) (Neunuebel & Knierim, 2014). The results thus confirm hallmark predictions of this computational theory of hippocampal memory processing.

Evidence that the firing of hippocampal pyramidal cells in macaques is more sparse than in neocortical areas is described in Section 1.4.4. This is consistent with the premium placed in the hippocampus for storing and retrieving large numbers of independent memories.

The theory predicts that if primates including humans can form an episodic memory in which objects or people are seen at particular locations even though the observer viewing the space has never been to those locations ‘out there’ in space, there should be a neural system in CA3 that can support such associations between places ‘out there’ in a scene and objects. Exactly this is provided by the spatial view neurons Rolls and colleagues have discovered that are present in CA3 (Georges-François et al., 1999; Robertson et al., 1998; Rolls & Xiang, 2005; Rolls, Xiang, & Franco, 2005; Rolls, Robertson, et al., 1997; Rolls et al., 1998, 2006). Place cells will not do for this type of episodic memory (Kesner & Rolls, 2015; Rolls, 2010b, 2013).

1.8.5. Recall via CA1 to neocortex: a reverse hierarchy of pattern associators each performing pattern completion

The theory shows quantitatively, analytically, how memories could be retrieved from the hippocampus to the neocortex (Treves & Rolls, 1994), and this has been shown by simulation of the multistage hippocampal system including the entorhinal cortex, dentate, CA3, CA1, and return to the entorhinal cortex to recall the memory to be quantitatively realistic (Rolls, 1995).

It has been shown that after learning in hippocampal-dependent tasks, neocortical representations may change (Schwindel & McNaughton, 2011). Although this has been interpreted as the transfer of memories from the hippocampus to the neocortex (Schwindel & 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 of neuronal firing changes and influences another vector of neuronal firing through fixed connections.)

It has also been suggested that the transfer of information from the hippocampus to the neocortex occurs especially during sleep (Marr, 1971; McClelland, McNaughton, & O’Reilly, 1995; Schwindel & McNaughton, 2011; Wilson & McNaughton, 1994). My own view is that during waking would be the best time to retrieve a memory from the hippocampus to the neocortex by using the hippocampus to retrieve the complete episodic memory from a fragment. The retrieval would reinstate the neocortical activity present when the event was originally learned. The retrieved information now present in the neocortex could then be used to build new semantic memories, for example a narrative account of all the events that took place on one’s fifth birthday party. During waking the building of semantic representations could be guided and organised by rational thought into useful semantic representations. To do this during sleep would run the risk of forming bizarre semantic representations of the type that we dream about during the unguided noise-driven stochastic firing during sleep (Rolls, 2008; Rolls & Deco, 2010). Further, the active recall during waking of memories from the hippocampus means that mainly relevant or useful memories would be retrieved from the hippocampus (not useless memories such as where one parked one’s bicycle two weeks ago), and only these memories would tend to become incorporated into useful long-term semantic representations, allowing memories not retrieved from the hippocampus to be overwritten by new memories in the process of forgetting

that involves using CA3 sets of neurons chosen at random for new episodic memories (Rolls, 2008).

Many further tests of the theory are described elsewhere (Hunsaker & Kesner, 2013; Kesner & Rolls, 2015; Kesner et al., 2012; Rolls, 2008, 2010b; Rolls & Kesner, 2006).

1.9. Temporal order encoding and hippocampal function

There has for some time been evidence that the hippocampus plays a role in temporal order memory, even when there is no spatial component (Hoge & Kesner, 2007; Kesner, Gilbert, & Barua, 2002; Rolls & Kesner, 2006). In humans, the hippocampus becomes activated when the temporal order of events is being processed (Lehn et al., 2009). An approach is now being developed on how temporal order memory could be implemented in the hippocampus (Rolls, 2010b, 2013; Rolls & Deco, 2010), as follows, and temporal pattern separation may be understood with this approach.

The approach is based on recent neurophysiological evidence of MacDonald, Eichenbaum and colleagues (Macdonald, Lepage, Eden, & Eichenbaum, 2011) and others (Kraus, Robinson, White, Eichenbaum, & Hasselmo, 2013; Pastalkova, Itskov, Amarasingham, & Buzsaki, 2008) showing that neurons in the rat hippocampus have firing rates that reflect which temporal part of the task is current. For example, a sequence of different neurons is activated at successive times during a time delay period. The tasks used included an object–odour paired associate non-spatial task with a 10 s delay period between the visual stimulus and the odour. 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 odour), 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. (Macdonald et al., 2011; Rolls & Deco, 2010). This provides the foundation for a new computational theory of temporal order memory within the hippocampus (and also the prefrontal cortex) which I outline next, and which 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.

First, because some neurons fire at different times in a trial of a temporal order memory task or delay task, the time in a trial at which an object (e.g. a visual stimulus or odour) was presented could become encoded in the hippocampus by an association implemented in the CA3 recurrent collaterals between the neurons that represent the object (already known to be present in the hippocampus for tasks for which the hippocampus is required (Rolls & Xiang, 2006; Rolls et al., 2005)) and the ‘time encoding’ neurons in the hippocampus (Macdonald et al., 2011). This would allow associations for the time at which the object was present to be formed.

Second, 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.

Third, we need a theory of what the origin is of the temporal effect whereby different hippocampal (or potentially prefrontal

cortex) neurons fire in different parts of a trial or delay period. The properties of the 'time encoding neurons' (Macdonald et al., 2011; Rolls & Deco, 2010) are key here, and we need to understand how they are generated. Are they generated within the hippocampus, or elsewhere, and in any case, what is the mechanism by which different neurons have high firing rates at different times in a trial? The fundamentally new approach to hippocampal function I am taking here is that rate encoding is being used, that is, the firing rates of different neurons are high at different times within a trial (Macdonald et al., 2011; Rolls & Deco, 2010)). This is a radically different approach to order encoding than that based on phenomena such as theta and gamma oscillations that has been investigated by Lisman and colleagues (Lisman & Redish, 2009).

We can consider three hypotheses about how the firing of the 'time encoding' hippocampal neurons is produced. All utilize slow transitions between attractor states that can be a property of noisy attractor networks. The first hypothesis is that an attractor network with realistic dynamics (modelled at the integrate-and-fire level with a dynamical implementation of the neuronal membrane and synaptic current dynamics, and with synaptic or neuronal adaptation) can implement a sequence memory (Deco & Rolls, 2005). The hypothesis is that there are several different attractors, and that there are weak connections between the different attractors (Deco & Rolls, 2005; Kleinfeld & Sompolinsky, 1988; Sompolinsky & Kanter, 1986). In the model, adaptation produces effects whereby whatever sequence (order of stimuli) is presented on an individual trial, that order can be replayed in the same sequence because as one attractor state dies as a result of the adaptation, the next attractor to emerge from the spontaneous firing because of the spiking-related noise is the one that has been active least recently, as it is the one that is least adapted (Deco & Rolls, 2005). The whole system operates at a rather slow timescale for the transitions between the attractors partly because of the time for the noise to drive the system from one attractor state to another, and the slow time course of the adaptation (Deco & Rolls, 2005; Rolls & Deco, 2010). This implements a type of order memory.

The second hypothesis is analogous, and is also implemented in a recurrently connected system such as the hippocampal CA3 system or local recurrent circuits in the neocortex (Rolls & Deco, 2010). This second theory is that again there are several attractors, but that each attractor is connected by slightly stronger forward than reverse synaptic weights to the next. In previous work, we have shown that with an integrate-and-fire implementation with spiking noise this allows slow transitions from one attractor state to the next (Deco, Ledberg, Almeida, & Fuster, 2005; Deco & Rolls, 2003). During learning of the synaptic weights in the network, adaptation might lead to each 'time encoding' population of neurons responding for only a limited period, helping to produce multiple sequentially activated populations of time encoding neurons (Macdonald et al., 2011; Rolls & Deco, 2010). In this scenario, an associative pool of neurons is unlikely to be helpful, and stronger forward than reverse synaptic weights between different attractors each consisting of a different population of 'time encoding' neurons would be the essence. It will be of interest to investigate whether this system, because of the noise, is limited to transitions between up to perhaps 7 ± 2 different sequential firing rate states with different neuronal subpopulations for each state, and thus provides an account for the limit of the magical number 7 ± 2 on short-term memory and related types of processing (Miller, 1956), and for the recency part of short-term memory in which the items are naturally recalled in the order in which they were presented. This is the most likely model at present of short-term memory and its natural propensity to store and to recall items in the order in which they were received (Rolls & Deco, 2010).

A variation on this implementation that I have proposed would be to have short-term attractor memories with different time constants (for example of adaptation), but all started at the same time (Rolls & Deco, 2010). This could result in some attractors starting early in the sequence and finishing early, and with other attractors starting up a little later, but lasting for much longer in time. The neurons recorded in the rat (Macdonald et al., 2011) are not inconsistent with this possibility. This type of time-encoding representation could also be used to associate with items, to implement an item-order memory.

Other models of time cells incorporate oscillatory interference (Hasselmo & Stern, 2014) or an inverse Laplace transform (Howard et al., 2014).

It is thus suggested that temporal order memory could be implemented in the hippocampus in this way, and could make an important contribution to episodic memory in which several events linked in the correct order might form an episode. The theory shows how items in a particular temporal order could be separated from each other, a property we have referred to as the temporal pattern separation effect (Rolls & Kesner, 2006). The natural implementation of such temporal order memory would be in the hippocampal CA3–CA3 recurrent collateral network. It is therefore somewhat of a puzzle that some of the evidence implicates the CA1 region in temporal order memory (Hunsaker, Lee, & Kesner, 2008; Kesner et al., 2012; Rolls & Kesner, 2006). This issue remains to be clarified. In any case, important in this temporal pattern separation would be the sparseness of the representation in the attractor, for if the representation became less sparse, this would impair the ability of the attractor network to maintain a long temporal sequence of different attractor states.

2. Categorisation in the neocortex

We have seen that in the hippocampus a key computational foundation is preparing patterns to be relatively orthogonal to each other using pattern separation, storing the components of a single memory in the CA3 autoassociation network, completing a memory using a partial recall cue in the CA3 autoassociation network, and then retrieving the whole memory back to the cerebral cortex. The emphasis is thus on accurate storage and recall, without altering the representations.

In contrast, in the neocortex, the foundations of its operation include computing new representations, used for example in perception, in semantic memory, and in much cognitive function including language. New representations are computed in the cortex that go far beyond what is represented in the periphery (Rolls, 2016), including for example invariant representations of objects (Rolls, 2012b). One key component is placing similar exemplars, for example similar visual stimuli, into the same category. One example is the categorisation of many types of vehicle as cars, and another is placing many shapes of potato into the category of potato. A second key component is placing somewhat different exemplars into different distinct categories, for example cars and trucks into separate categories, or potatoes and turnips into different categories. This process is referred to as categorisation, and does not require a teacher, but is based on similarities and differences in the input space. Categorisation is used to refer to this type of unsupervised learning. The essence is that somewhat similar patterns must be categorised into the same category, and somewhat different patterns into different categories. This can be performed by a competitive network, in ways that are analysed in Section 3 (Rolls, 2008, 2016).

Patterns can also be classified into the same or different classes using a teacher in supervised learning. Networks that can perform this include pattern association networks, in which there is one

input that specifies the output of the network, and a second input (sometimes referred to as a conditioned stimulus) that is associated with the firing produced by the first (unconditioned) stimulus using associative synaptic modification (Rolls, 2008). One example is the pairing of a visual input (the conditioned input) with a taste (the unconditioned input) to produce an output that classifies the visual stimuli based on their associations with different tastes. In this case, the taste inputs define the classes into which the visual stimuli must be classified by pattern association learning. Similar networks can perform a similar classification function if they compute an error between the firing produced by the input and the firing specified by a teacher, and use the error to modify the synaptic weights, and are described as one-layer delta rule perceptrons, where the delta rule refers to the error correcting rule (Rolls, 2008). Networks of this type may be implemented in other brain systems such as the cerebellum, in which the parallel fibres represent the input, the Purkinje cells the output, and the climbing fibres the error-correcting teacher (Albus, 1971; Ito, 1984, 2006; Marr, 1969; Rolls & Treves, 1998). An extension is a multilayer backpropagation of error network which has a similar teacher for the output layer, and calculates corrections needed for correction of synaptic connections at earlier levels of the network. Backpropagation of errors in such multilayer networks is possible in a computer, but there is no simple way that the errors could be backpropagated correctly to earlier layers in a biologically plausible way (O'Reilly & Munakata, 2000; Rolls, 2008, 2016).

3. Analyses of pattern separation and categorisation in competitive networks

The new analyses described in this section highlight in a more precise computational way how networks with the architecture of a competitive network can perform pattern separation, and categorisation, depending on the ways in which they are set up. The different modes of operation are related to particular systems in the hippocampus and neocortex.

3.1. The architecture and operation of competitive networks

The basic architecture of a competitive network is shown in Fig. 5a. It is a one-layer network with a set of inputs that make

modifiable excitatory synapses w_{ij} with the output neurons. The output cells compete with each other (for example by mutual inhibition through inhibitory neurons) in such a way that the most strongly activated neuron or neurons win the competition, and are left firing strongly. The synaptic weights, w_{ij} , are initialized to random values before learning starts. If some of the synapses are missing, that is if there is randomly diluted connectivity, that is not a problem for such networks, and can even help them (see below). The operation is as follows:

1. Apply an input vector of firing rates x_j on the inputs to the network, and calculate the activation h_i of each neuron i

$$h_i = \sum_j x_j w_{ij}$$

where the sum is over the C input axons, indexed by j . (It is useful to normalize the length of each input firing rate vector. In the brain, a scaling effect is likely to be achieved both by feedforward inhibition, and by feedback inhibition among the set of input cells (in a preceding network) that give rise to the axons conveying the inputs.)

The output firing y_i^1 is a function of the activation of the neuron

$$y_i^1 = f(h_i)$$

The function f can be linear, linear threshold, binary, sigmoid, monotonically increasing, etc., as described elsewhere (Rolls, 2008, 2016).

2. Allow competitive interaction between the output neurons by a mechanism such as lateral or mutual inhibition (possibly with self-excitation), to produce a contrast-enhanced version of the firing rate vector

$$y_i = g(y_i^1)$$

Function g is typically a non-linear operation, and in its most extreme form may be a winner-take-all function, in which after the competition one neuron may be firing, and the others not firing. Algorithms that produce softer competition without a single winner to produce a distributed representation are described elsewhere (Rolls, 2008, 2016).

3. Apply an associative Hebb-like learning rule

$$\delta w_{ij} = \alpha y_i x_j$$

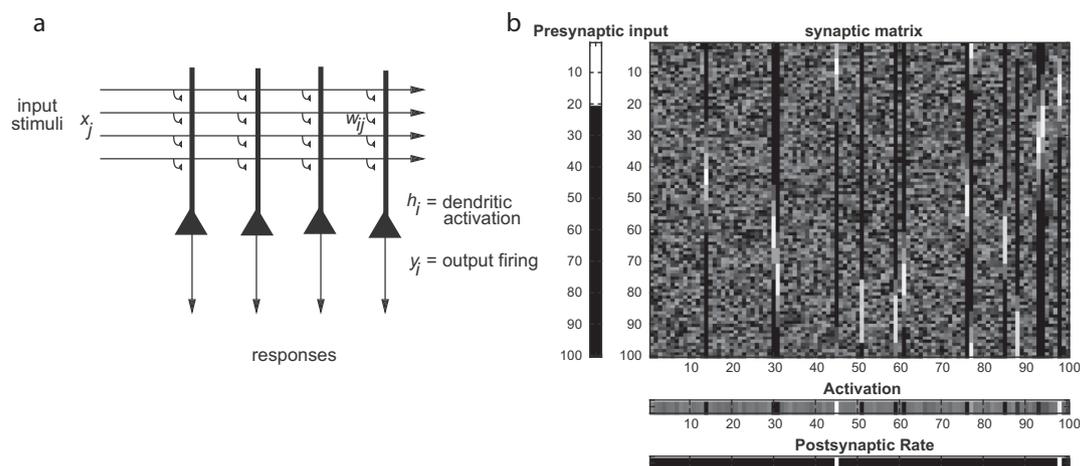


Fig. 5. (a) The architecture of a competitive network. (b) Simulation of a competitive network. The synaptic matrix is shown after training with 20 presynaptic input patterns or stimuli, each 20 elements long, and each shifted down 5 elements with respect to the previous pattern. In the synaptic matrix, white corresponds to the maximum synaptic weight, and black to 0. The network is shown during testing with presynaptic input pattern 1, which activates strongly 2 of the neurons as shown because of the pattern of their increased synaptic weights after training. After competition, two of the postsynaptic neurons have high firing rates (white), corresponding to a sparseness of 0.02 in this competitive network with $N = 100$ output neurons, and $C = 100$ synaptic inputs onto each neuron. The synaptic weight vector for each neuron corresponds to one vertical column in the synaptic matrix. The neurons left with random synaptic weights have not learned to this set of patterns, and remain available for further presynaptic input patterns or stimuli. These new simulations were written in Matlab.

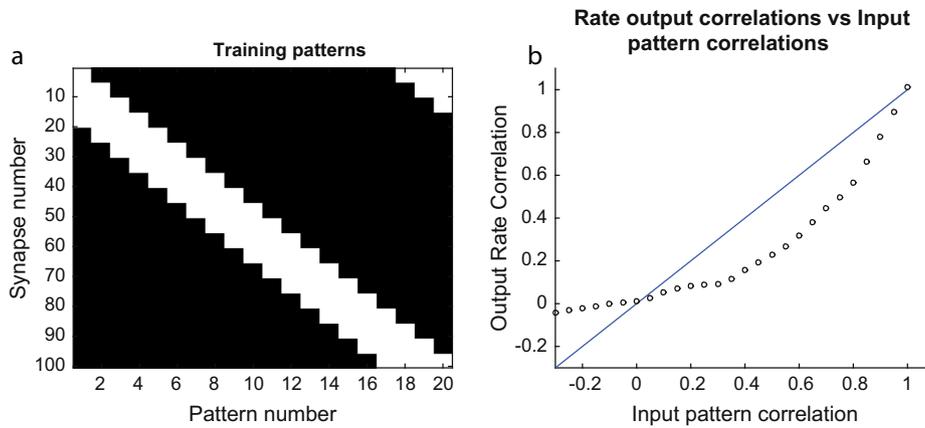


Fig. 6. Competitive net with diluted connectivity of 2 and no learning. (a) The 20 basic stimuli, each a vertical column with 20 values of 1 (white) and the remainder of the 100 values 0 (black). Each stimulus overlaps in 15 positions with the preceding stimulus in the series. After this stage, some bits were flipped to produce a wider range of correlations between different stimuli (see text). (b) The correlations between the output firing rate vectors as a function of the correlations between the different input patterns. The diagonal line with a slope of 1 shows where correlations between the output firing rate vectors would fall if there was no pattern separation of the input pattern vectors.

where δw_{ij} is the change in the synaptic weight, and α is a learning rate constant.

4. Normalize the length of the synaptic weight vector on each dendrite to prevent the same few neurons always winning the competition:

$$\sum_j (w_{ij})^2 = 1$$

(Biologically plausible implementations of this involve heterosynaptic long-term depression during the synaptic modification (Oja, 1982; Rolls, 2008, 2016).)

5. Repeat steps 1–4 for each different input stimulus a number of times in random sequence.

We now consider how this network operates with different sets of some key parameters, including whether the learning described in step 3 occurs, and whether the network has diluted connectivity, that is whether there are fewer synapses onto each neuron than there are inputs j to the network.

Simple simulations are described to illustrate the properties of the network operating with these different parameters. Small simulations are illustrated to enable the effects found to be easily demonstrated, but similar effects can be found when the networks are scaled up. Simple Matlab programs to illustrate some of the properties of competitive networks and also pattern association and autoassociation networks have been made available (Rolls, 2016). The networks described here had $N = 100$ neurons; $C = N$ inputs per neuron if there was no dilution of the connectivity, and half that number if the dilution was 2; a learning rate α of 0.1 if learning was allowed and otherwise zero; 30 epochs of training; length normalization of the synaptic weight vectors; a binary activation function with the firing of a neuron either 0 or 1 (with similar effects found with a threshold-linear activation function); and the threshold set so that the sparseness a of the output representation achieved was 0.02, where the sparseness in this case is the proportion of neurons firing with a high firing rate of 1. The input stimuli consisted of a set of 20 stimuli each 20 elements long (i.e. with a sparseness of 0.2), with each overlapping in 5 elements with the preceding stimulus, as illustrated in Fig. 6a. The stimuli were modified to provide a richer set of different correlations between each of the input stimuli by flipping a random number of elements between 1 and 4 to the opposite state. (That is for each

stimulus between 1 and 4 of the elements were set from 1 to 0, with the same number set from 0 to 1.)

The results of a simulation to illustrate the operation of a competitive network with complete (undiluted) connectivity are illustrated in Fig. 5b. The synaptic matrix is shown after training with 20 presynaptic input patterns or stimuli, each 20 elements long, and each shifted down 5 elements with respect to the previous pattern, as illustrated in Fig. 6a. In the synaptic matrix, white corresponds to the maximum synaptic weight, and black to 0. The network is shown during testing with presynaptic input pattern 1, which activated strongly two of the neurons as shown because of the pattern of their increased synaptic weights. After competition, two of the postsynaptic neurons had high firing rates (white), corresponding to a sparseness of 0.02 in this competitive network with $N = 100$ output neurons, and $C = 100$ synaptic inputs onto each neuron. The synaptic weight vector for each neuron corresponds to one vertical column in the synaptic matrix. The neurons left with random synaptic weights have not learned to this set of patterns, and remain available for further presynaptic input patterns or stimuli. The synaptic weight vectors show that each neuron with modified synapses has learned to respond to part of one or more of the input patterns, and that this is how the network has learned to respond differently to each of the different inputs.

3.2. Competitive networks without associative learning but with diluted connectivity

This provides a model of the mossy fibre inputs to the hippocampal CA3 neurons.

The results of a simulation with this scenario are shown in Fig. 6b. The dilution was 2, and the learning rate was 0, and the synapses that were present had the same positive value. Fig. 6b shows that the correlations between the firing rate vectors (each vector the firing rates of the 100 output neurons) were lower than the correlations between the inputs of the corresponding stimulus patterns that produced the output firing rates. That is, the computation that was achieved was some pattern separation or orthogonalization of the inputs in transforming them to outputs. This was confirmed by the results of an analysis that showed that 90–100% of the output vectors to the different stimuli (on different runs) were less correlated than 0.8.

If there was no dilution, or almost no dilution, of the connectivity, then the outputs to the different patterns were of course the same or very similar, that is, no pattern separation or

orthogonalization was produced. They were very close to the diagonal line in Fig. 6b, which shows the situation in which the correlations between the different outputs are the same as those between the different inputs.

The simulation thus shows that pattern separation can be produced by diluted connectivity in a competitive network. (For completeness, it is noted that an analogous type of pattern separation can be produced by a fully connected competitive net without learning in which the synaptic weights are set to random values with a uniform distribution.)

The mode of operation considered in this section is the condition considered analytically by Treves and Rolls (1992) for discrete input patterns (e.g. random patterns of 0s and 1s as the input vectors) that might represent objects, and by Cerasti and Treves (2010) for input patterns with Gaussian firing rate distributions that might represent a position in continuous space. This was in the context of how the dentate gyrus projections via the mossy fibres to the CA3 neurons, which have diluted connectivity and no associative learning, might operate as a system to activate relatively orthogonal representations in CA3 cell firing for different inputs. This pattern separation and orthogonalization is important for the CA3 autoassociation system, so that each episodic memory is as different as possible from all the other episodic memories stored in CA3, to minimize interference between them, and to maximize the number of memories that can be stored and correctly retrieved (Kesner & Rolls, 2015; Rolls, 1989b; Treves & Rolls, 1992).

3.3. Competitive networks with learning and with diluted connectivity

If in the same network we allow associative learning, then the result is that the synaptic weight vectors of some of the neurons come to point towards a cluster of similar stimuli. This results in the important new property of categorisation, in which the outputs for similar stimuli are more similar than the inputs. This is the opposite of pattern separation. This is illustrated in Fig. 7 in the same network with a diluted connectivity of 2 and with the same stimulus input patterns, but now with the learning between the inputs and the output neurons enabled (with the standard learning rate of 0.1), and with the synaptic weights initialised to random values with a uniform distribution. Fig. 7b shows that pattern separation is achieved if the inputs are less correlated than approximately 0.5, and that in contrast the outputs are more similar (as shown by the correlation between the firing rate vectors on the

ordinate) if the inputs are more correlated than approximately 0.5. The combination of these two processes, pattern separation for different patterns and pattern clustering for similar patterns, that is implemented is the categorisation that can be performed by a competitive network, as illustrated for example in Fig. 4.3 of Rolls and Treves (1998). Processes such as these have been considered previously (Guzowski, Knierim, & Moser, 2004; O'Reilly & McClelland, 1994), but the aim of the present analysis is to characterise the different categorisation and pattern separation processes performed by different types of competitive networks.

It should be noted that this process is quite distinct from the process of pattern completion, in which the whole of a pattern can be retrieved from any part (Rolls, 2008).

Fig. 7a illustrates the pattern separation that is achieved, by showing that the correlations of the output firing rate vectors for most of the patterns are low. (In this run, only the outputs to patterns 16–17 and 20–1 were highly correlated with each other, resulting in 80% correct pattern separation. These pairs of input patterns were closer together than the other pairs as a result of the small amount of randomness introduced into the patterns by flipping 1–4 of the bits.)

The competitive network can thus be described as generalizing to similar patterns, in that it places similar patterns into the same category. This generalization was found to untrained as well as trained patterns provided that the patterns were similar. It has been noted that some generalization in the hippocampal system may be desirable, and a connectionist model of this was designed (Kumaran & McClelland, 2012). The generalization described here for competitive networks is relevant to this need. Indeed the dentate granule cells and the CA1 network are hypothesized to operate as competitive networks with this property (Kesner & Rolls, 2015; Rolls, 1987, 1989b, 2008; Rolls & Treves, 1998; Schultz & Rolls, 1999).

3.4. Competitive networks with learning and with full (undiluted) connectivity

If the same network is set up with full (undiluted) connectivity, associative learning, and random initialization of the synaptic weights, the clustering of similar patterns still occurs as a result of the learning, but the network may perform less pattern separation. This is illustrated in Fig. 8 in the same network with a connectivity of 1 and with the same stimulus input patterns, with the learning between the inputs and the output neurons enabled (with

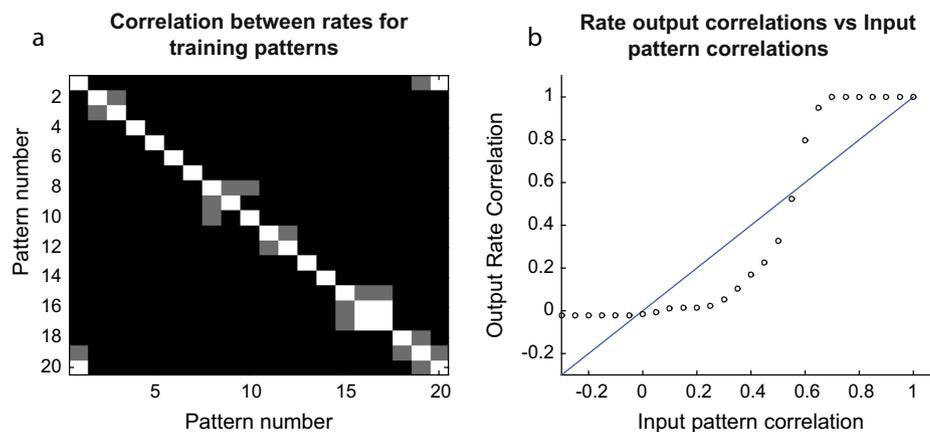


Fig. 7. Competitive net with diluted connectivity of 2 and learning. (a) The correlations between the output firing rate vectors produced by each of the input stimulus patterns. White represents a correlation of 1, and black represents a correlation of approximately 0. (b) The correlations between the output firing rate vectors as a function of the correlations between the different input patterns. The diagonal line with a slope of 1 shows where correlations between the output firing rate vectors would fall if there was no pattern separation of the input pattern vectors. Each point in the graph represents the mean value across the presentation of many stimuli within the same correlation bin.

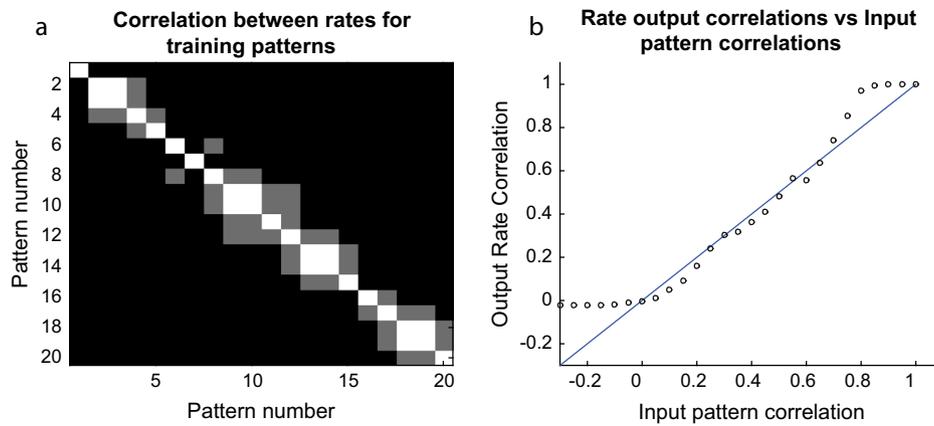


Fig. 8. Competitive net with undiluted connectivity and with learning. (a) The correlations between the firing rate vectors between produced by each of the patterns. White represents a correlation of 1, and black represents a correlation of approximately 0. (b) The correlations between the output firing rate vectors as a function of the correlations between the different input patterns. The diagonal line with a slope of 1 shows where correlations between the output firing rate vectors would fall if there was no pattern separation of the input pattern vectors.

the standard learning rate of 0.1), and with the synaptic weights initialised to random values with a uniform distribution. Fig. 8b shows that the outputs are more similar (as shown by the correlation between the firing rate vectors on the ordinate) if the inputs are more correlated than approximately 0.7, and this is the clustering together of similar patterns. However, almost no pattern separation is achieved: if the inputs are less correlated than approximately 0.7, then the correlations between the firing rate output vectors produced are almost the same as the correlations between the input stimulus vectors. (The output correlations fall close to the diagonal line in Fig. 8b for correlations less than 0.7.)

This finding establishes the point that diluted connectivity can improve the pattern separation that can be performed by competitive networks. The magnitude of the effect depends on several factors, such as the size of the network, and the correlations between the inputs, which will be harder to pattern separate without the dilution, because similar patterns will tend to activate the same neurons during the learning. Indeed, the mechanism by which the dilution works is that it forces different neurons to be capable of pointing in only certain directions in the input space of patterns, so that some neurons can be thought of as inherently more likely to learn to some input patterns than to other input patterns.

Fig. 8a illustrates that pattern separation with full connectivity is less good than with the diluted connectivity, by showing that the correlations of the output firing rate vectors for several pairs of the patterns were high. (In this run, the outputs to patterns 2–3, 9–10, 13–14 and 18–19 were highly correlated with each other, resulting in only 60% correct pattern separation.)

4. Discussion

The new analyses of the operation of competitive networks described in Section 3 of this paper help to elucidate the operation of mechanisms in the hippocampus and neocortex for pattern separation and pattern clustering.

The effects of diluted connectivity in networks without associative synaptic modification described in Section 3.2 is relevant to how the dentate gyrus projections via the mossy fibres to the CA3 neurons, which have diluted connectivity and no associative learning, might operate as a system to activate relatively orthogonal representations in CA3 cell firing for different inputs. This pattern separation and orthogonalization is important for the CA3 autoassociation system, so that each episodic memory is as different as possible from all the other episodic memories stored in CA3, to minimize interference between them and to maximize the

number of different memories that can be stored (Kesner & Rolls, 2015; Rolls, 1989b; Treves & Rolls, 1992). The dentate to CA3 mossy fibre pathway may be one of the few major cortical pathways that does not show associative synaptic modification, for in this dentate to CA3 system the main computation appears to be the production of pattern separation so that CA3 representations are relatively orthogonal to each other. In this situation, a single attractor state can be set up in CA3 by associative synaptic modification in the CA3–CA3 recurrent collateral connections that is particular for each episodic memory (Rolls, 2008). The theory is that the direct perforant path input to the CA3 cells becomes associatively modified during learning, and later during recall provides the partial retrieval cue for completion to occur in CA3, which then connects to CA1 on the path back to neocortex for the whole of the memory to be retrieved in each neocortical area that provided part of the original input to the episodic memory (Kesner & Rolls, 2015; Rolls, 1989b, 2010b; Treves & Rolls, 1992). In this situation, the dentate to CA3 projection via the mossy fibres may be one of the few places in the brain where information is not stored by a change of synaptic weight to store memories or build new perceptual or cognitive representations, as its function is to orthogonalize new patterns, and not to act as the retrieval cue for previously formed memories. In this situation, the lack of associative synaptic learning may in fact be useful, for as shown in Section 3.3, if learning is present, this will tend to cluster similar patterns, whereas the goal is to orthogonalize all patterns, even similar ones, in the hippocampal system, to maximize the number of different episodic memories that can be stored and correctly retrieved. The presence of neurogenesis in this dentate system (Section 1.4.5), but not in most other cortical pathways, can be understood in this framework.

The effect of clustering of nearby patterns, and separation of less correlated patterns, in competitive networks with associative synaptic modification is well known (Rolls, 2008; Rolls & Treves, 1998). In performing this function, making the output representation more sparse, by pointing the synaptic weight vectors of neurons towards the centres of clusters of patterns, also occurs. This sparsification can be a useful step in many functions, including memory storage in which sparse representations enable more memories to be stored (Rolls & Treves, 1990; Treves & Rolls, 1991). This type of computation of a sparse representation involving learning may be important in the learning that allows medial entorhinal cortex grid cells (Moser et al., 2015) to produce sparse place-like representations in dentate granule cells (Jung & McNaughton, 1993; Leutgeb et al., 2007), for which the learning

in a dentate granule cell competitive network appears to be beneficial (Rolls et al., 2006).

What has been less clearly formulated previously is the beneficial role of dilution of connectivity in competitive networks. As shown here this can help to separate the outputs to patterns that are highly correlated. This will become especially important in highly loaded systems with highly correlated patterns.

Another beneficial role of diluted connectivity in competitive networks is to increase the stability of the categorisation if new patterns that drift continuously across the input space are introduced to the system. If learning is still on, this can lead to drifting, across the output neurons of the representations. This effect can be reduced by the diluted connectivity that is found in all cortical networks (Rolls, 2008, 2016), for this makes individual neurons in a competitive network specialized to point with their synaptic weight vectors towards only parts of the input space of patterns, so that neurons cannot remap across the whole input space (Rolls, 2008). Another factor that may help competitive networks in the brain is the sensitive or critical early period during development in which early cortical layers may learn new categories for perceptual inputs, after which these areas become much less modifiable, so providing stability in the inputs to higher cortical levels in a hierarchy (Rolls, 2008, 2012b, 2016).

The importance of diluted connectivity in competitive networks that is important in pattern separation and categorisation adds to the evidence that diluted connectivity is important in other types of cortical network. In particular, it has been shown that randomly diluted connectivity between neurons in an autoassociation network such as that in the CA3 region of the hippocampus or in most neocortical areas helps to maintain the capacity of the network for storing large numbers of memories close to its high theoretical limit by minimizing the numbers of more than one synapse between any pair of neurons in the autoassociation network (Rolls, 2012a). If there was on average 1 synapse between every pair of neurons, then many pairs would have more than one synapse, and this distorts the basins of attraction of the network very greatly, so that many fewer memories can be stored (Rolls, 2012a). In the contrasting situation with diluted connectivity with a value of 0.1, or even 0.04 as it appears to be in the CA3–CA3 connectivity, the probability of the number of multiple connections between every pair of neurons is low, and the memory capacity remains close to its theoretical value (Rolls, 2012a), in which the leading term is the number of connections C onto any one neuron in the attractor network (Rolls, 2012a; Treves & Rolls, 1991). The background to this point is that there are insufficient genes to specify the connections between every pair of neurons in the brain, so it is likely that connections are made on the basis of genetically specified connectivities between classes of neurons, for example with CA3 neurons connecting to CA3 neurons in the CA3 recurrent collateral system, but after that largely at random. In that situation, diluted connectivity is a way to minimize the probability of multiple connections between pairs of neurons, and as we have seen here, can be beneficial in competitive networks, as well as in attractor networks (Rolls, 2012a). Consistent with this approach, we have produced a list of the properties of neurons that might be specified by genes to provide the basis for the architecture and function of different classes of network, including attractor, pattern association, and competitive networks, and have shown that a genetic algorithm operating with these rules can be shaped by genetic evolution mechanisms to build these three classes of neuronal network based on the problem to be solved by the individuals in each generation (Rolls & Stringer, 2000).

To complete the trio, in pattern association networks it has been shown that multiple connections at random between pairs of neurons can by distorting the effects of the recall cue distort the retrieval of associations from the pattern association network

(Rolls, 2015a). This is relevant to the recall of information from the hippocampus, for which the theory as described in Section 1.8.5 is that there are multiple stages of pattern association involved in recall to the neocortex (Rolls, 1989b, 2008; Treves & Rolls, 1994) (see Fig. 1). In this situation, any distortion at every stage would increase with the multiple stages if there were many double or more synapses onto any one neuron from a sending neuron in this pathway, and this provides an account (Rolls, 2015a) of why the connectivity in the backprojection pathways to the neocortex from the hippocampus to the neocortex is diluted (Rolls, 2008).

In summary, in Section 1 of this paper, the computational hypotheses and some of the evidence on which they are based and which test the hypotheses that particular systems in the hippocampus are involved in pattern separation and pattern completion are described. In Section 2 the distinction that categorisation is a separate process with different functionality, frequently of importance in the neocortex, was introduced. In Section 3, new analyses were described which elucidate the differences in the functional architecture of competitive networks that can enable them to perform only pattern separation, as is believed to happen in the dentate-mossy fibre-CA3 system, or categorisation as in the neocortex. Finally, the importance of diluted connectivity for the operation of competitive networks was elucidated, and this adds to the evidence that diluted connectivity is advantageous (though for different reasons) in autoassociation and pattern association networks. This fits with the diluted connectivity found in both the hippocampus and neocortex, and to the importance of competitive, autoassociation, and pattern association networks in cortical function (Rolls, 2016).

Acknowledgments

Different parts of the research described here were supported by Programme Grants from the Medical Research Council; by a Human Frontier Science Program Grant; by an EEC BRAIN Grant; by the MRC Oxford Interdisciplinary Research Centre in Cognitive Neuroscience; 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 which is incorporated in some of the ideas presented here on the hippocampus with many colleagues, including Alessandro Treves, Simon Stringer, Ray Kesner, Robert Robertson, Pierre Georges-François, Shane O'Mara, and Alain Berthoz, and their contributions are sincerely acknowledged. Discussions with Alessandro Treves on some of the simulations described in Section 3 are warmly acknowledged. Pdfs of some of the papers cited are available at www.oxcns.org.

References

- 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. *Journal of Neuroscience*, *18*, 3386–3403.
- Albus, J. S. (1971). A theory of cerebellar function. *Mathematical Biosciences*, *10*, 25–61.
- Amaral, D. G. (1993). Emerging principles of intrinsic hippocampal organisation. *Current Opinion in Neurobiology*, *3*, 225–229.
- Amaral, D. G., Ishizuka, N., & Claiborne, B. (1990). Neurons, numbers and the hippocampal network. *Progress in Brain Research*, *83*, 1–11.
- Amaral, D. G. (1987). Memory: Anatomical organization of candidate brain regions. In V. B. Mountcastle (Ed.), *Handbook of physiology. Section 1, the nervous system* (pp. 211–294). Washington, DC: American Physiological Society.
- Amaral, D. G., Price, J. L., Pitkanen, A., & Carmichael, S. T. (1992). Anatomical organization of the primate amygdaloid complex. In J. P. Aggleton (Ed.), *The amygdala* (pp. 1–66). New York: Wiley-Liss.
- 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 G. Paxinos (Ed.), *The rat nervous system* (pp. 443–493). San Diego: Academic Press.

- Amari, S. (1977). Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological Cybernetics*, 27, 77–87.
- Amit, D. J. (1989). *Modeling brain function*. Cambridge: Cambridge University Press.
- Amit, D. J., & Brunel, N. (1997). Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cerebral Cortex*, 7, 237–252.
- Andersen, P., Morris, R. G. M., Amaral, D. G., Bliss, T. V. P., & O'Keefe, J. (2007). *The hippocampus book*. London: Oxford University Press.
- 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 & Behavior*, 25, 42–49.
- 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. *Behavioural Brain Research*, 208, 535–544.
- Battaglia, F. P., & Treves, A. (1998). Attractor neural networks storing multiple space representations: A model for hippocampal place fields. *Physical Review E*, 58, 7738–7753.
- Becker, S. (2005). A computational principle for hippocampal learning and neurogenesis. *Hippocampus*, 15, 722–738.
- Bohbot, V. D., Allen, J. J., & Nadel, L. (2000). Memory deficits characterized by patterns of lesions to the hippocampus and parahippocampal cortex. *Annals of the New York Academy of Sciences*, 911, 355–368.
- Bonelli, S. B., Powell, R. H., Yogarajah, M., Samson, R. S., Symms, M. R., Thompson, P. J., et al. (2010). Imaging memory in temporal lobe epilepsy: Predicting the effects of temporal lobe resection. *Brain*, 133, 1186–1199.
- Brown, T. H., Ganong, A. H., Kairiss, E. W., Keenan, C. L., & Kelso, S. R. (Eds.). (1989). *Long-term potentiation in two synaptic systems of the hippocampal brain slice*. San Diego: Academic Press.
- Brown, T. H., Kairiss, E. W., & Keenan, C. L. (1990). Hebbian synapses: Biophysical mechanisms and algorithms. *Annual Review of Neuroscience*, 13, 475–511.
- Brun, V. H., Otnass, M. K., Molden, S., Steffenach, H. A., Witter, M. P., Moser, M. B., et al. (2002). Place cells and place recognition maintained by direct entorhinal-hippocampal circuitry. *Science*, 296, 2243–2246.
- Burak, Y., & Fiete, I. R. (2009). Accurate path integration in continuous attractor network models of grid cells. *PLoS Computational Biology*, 5, e1000291.
- Burgess, N., Barry, C., & O'Keefe, J. (2007). An oscillatory interference model of grid cell firing. *Hippocampus*, 17, 801–812.
- Carmichael, S. T., & Price, J. L. (1995). Limbic connections of the orbital and medial prefrontal cortex in macaque monkeys. *Journal of Comparative Neurology*, 346, 403–434.
- Cerasti, E., & Treves, A. (2010). How informative are spatial CA3 representations established by the dentate gyrus? *PLoS Computational Biology*, 6, e1000759.
- Cerasti, E., & Treves, A. (2013). The spatial representations acquired in CA3 by self-organizing recurrent connections. *Frontiers in Cellular Neuroscience*, 7, 112.
- Cheng, S. (2013). The CRISP theory of hippocampal function in episodic memory. *Frontiers in Neural Circuits*, 7, 88.
- Clelland, C. D., Choi, M., Romberg, C., Clemenson, G. D., Jr., Fragniere, A., Tyers, P., et al. (2009). A functional role for adult hippocampal neurogenesis in spatial pattern separation. *Science*, 325, 210–213.
- Daumas, S., Ceccom, J., Halley, H., Frances, B., & Lassalle, J. M. (2009). Activation of metabotropic glutamate receptor type 2/3 supports the involvement of the hippocampal mossy fiber pathway on contextual fear memory consolidation. *Learning & Memory*, 16, 504–507.
- Day, M., Langston, R., & Morris, R. G. (2003). Glutamate-receptor-mediated encoding and retrieval of paired-associate learning. *Nature*, 424, 205–209.
- de Araujo, I. E. T., Rolls, E. T., & Stringer, S. M. (2001). A view model which accounts for the spatial fields of hippocampal primate spatial view cells and rat place cells. *Hippocampus*, 11, 699–706.
- Deco, G., Ledberg, A., Almeida, R., & Fuster, J. (2005). Neural dynamics of cross-modal and cross-temporal associations. *Experimental Brain Research*, 166, 325–336.
- Deco, G., & Rolls, E. T. (2003). Attention and working memory: A dynamical model of neuronal activity in the prefrontal cortex. *European Journal of Neuroscience*, 18, 2374–2390.
- Deco, G., & Rolls, E. T. (2005). Sequential memory: A putative neural and synaptic dynamical mechanism. *Journal of Cognitive Neuroscience*, 17, 294–307.
- Deco, G., & Rolls, E. T. (2006). Decision-making and Weber's Law: A neurophysiological model. *European Journal of Neuroscience*, 24, 901–916.
- Deco, G., Rolls, E. T., Albantakis, L., & Romo, R. (2013). Brain mechanisms for perceptual and reward-related decision-making. *Progress in Neurobiology*, 103, 194–213.
- Delatour, B., & Witter, M. P. (2002). Projections from the parahippocampal region to the prefrontal cortex in the rat: Evidence of multiple pathways. *European Journal of Neuroscience*, 15, 1400–1407.
- Dere, E., Easton, A., Nadel, L., & Huston, J. P. (Eds.). (2008). *Handbook of episodic memory*. Amsterdam: Elsevier.
- Ekstrom, A. D., Kahana, M. J., Caplan, J. B., Fields, T. A., Isham, E. A., Newman, E. L., et al. (2003). Cellular networks underlying human spatial navigation. *Nature*, 425, 184–188.
- Faisal, A. A., Selen, L. P., & Wolpert, D. M. (2008). Noise in the nervous system. *Nature Reviews Neuroscience*, 9, 292–303.
- Fazeli, M. S., & Collingridge, G. L. (Eds.). (1996). *Cortical plasticity: LTP and LTD*. Oxford: Bios.
- Florian, C., & Roulet, P. (2004). Hippocampal CA3-region is crucial for acquisition and memory consolidation in Morris water maze task in mice. *Behavioural Brain Research*, 154, 365–374.
- Franco, L., Rolls, E. T., Aggelopoulos, N. C., & Jerez, J. M. (2007). Neuronal selectivity, population sparseness, and ergodicity in the inferior temporal visual cortex. *Biological Cybernetics*, 96, 547–560.
- Fuhs, M. C., & Touretzky, D. S. (2006). A spin glass model of path integration in rat medial entorhinal cortex. *Journal of Neuroscience*, 26, 4266–4276.
- Fyhn, M., Molden, S., Witter, M. P., Moser, E. I., & Moser, M. B. (2004). Spatial representation in the entorhinal cortex. *Science*, 305, 1258–1264.
- Gardner-Medwin, A. R. (1976). The recall of events through the learning of associations between their parts. *Proceedings of the Royal Society of London B*, 194, 375–402.
- Georges-François, P., Rolls, E. T., & Robertson, R. G. (1999). Spatial view cells in the primate hippocampus: Allocentric view not head direction or eye position or place. *Cerebral Cortex*, 9, 197–212.
- Gilbert, P. E., & Kesner, R. P. (2003). Localization of function within the dorsal hippocampus: The role of the CA3 subregion in paired-associate learning. *Behavioral Neuroscience*, 117, 1385–1394.
- 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. *Molecular Neurobiology*, 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. *Behavioral Neuroscience*, 122, 16–26.
- Guzowski, J. F., Knierim, J. J., & Moser, E. I. (2004). Ensemble dynamics of hippocampal regions CA3 and CA1. *Neuron*, 44, 581–584.
- 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.
- Hasselmo, M. E. (2009). A model of episodic memory: Mental time travel along encoded trajectories using grid cells. *Neurobiology of Learning and Memory*, 92, 559–573.
- 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. *Journal of Neuroscience*, 15, 5249–5262.
- Hasselmo, M. E., & Stern, C. E. (2014). Theta rhythm and the encoding and retrieval of space and time. *Neuroimage*, 85(Pt 2), 656–666.
- 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. *Behavioural Brain Research*, 89, 1–34.
- Henze, D. A., Wittner, L., & Buzsáki, G. (2002). Single granule cells reliably discharge targets in the hippocampal CA3 network in vivo. *Nature Neuroscience*, 5, 790–795.
- Hertz, J., Krogh, A., & Palmer, R. G. (1991). *An introduction to the theory of neural computation*. Wokingham: Addison-Wesley.
- Hoge, J., & Kesner, R. P. (2007). Role of CA3 and CA1 subregions of the dorsal hippocampus on temporal processing of objects. *Neurobiology of Learning and Memory*, 88, 225–231.
- Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Science of the United States of America*, 79, 2554–2558.
- Howard, M. W., MacDonald, C. J., Tiganj, Z., Shankar, K. H., Du, Q., Hasselmo, M. E., et al. (2014). A unified mathematical framework for coding time, space, and sequences in the hippocampal region. *Journal of Neuroscience*, 34, 4692–4707.
- 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. *Neuroscience and Biobehavioral Reviews*, 37, 36–58.
- Hunsaker, M. R., Lee, B., & Kesner, R. P. (2008). Evaluating the temporal context of episodic memory: The role of CA3 and CA1. *Behavioural Brain Research*, 188, 310–315.
- Ishizuka, N., Weber, J., & Amaral, D. G. (1990). Organization of intrahippocampal projections originating from CA3 pyramidal cells in the rat. *Journal of Comparative Neurology*, 295, 580–623.
- Ito, M. (1984). *The cerebellum and neural control*. New York: Raven Press.
- Ito, M. (2006). Cerebellar circuitry as a neuronal machine. *Progress in Neurobiology*, 78, 272–303.
- Jackson, M. B. (2013). Recall of spatial patterns stored in a hippocampal slice by long-term potentiation. *Journal of Neurophysiology*, 110, 2511–2519.
- 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.
- Jung, M. W., & McNaughton, B. L. (1993). Spatial selectivity of unit activity in the hippocampal granular layer. *Hippocampus*, 3, 165–182.
- Kesner, R. P. (2007). Behavioral functions of the CA3 subregion of the hippocampus. *Learning & Memory*, 14, 771–781.
- Kesner, R. P. (2013). An analysis of the dentate gyrus function. *Behavioural Brain Research*, 254, 1–7.

- 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. *Behavioral Neuroscience*, *116*, 286–290.
- 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. *Behavioral Neuroscience*, *122*, 1217–1225.
- 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 T. R. Zentall & E. A. Wasserman (Eds.), *Oxford handbook of comparative cognition* (pp. 322–346). Oxford: Oxford University Press.
- Kesner, R. P., & Rolls, E. T. (2001). Role of long term synaptic modification in short term memory. *Hippocampus*, *11*, 240–250.
- Kesner, R. P., & Rolls, E. T. (2015). A computational theory of hippocampal function, and tests of the theory: New developments. *Neuroscience and Biobehavioral Reviews*, *48*, 92–147.
- Killian, N. J., Jutras, M. J., & Buffalo, E. A. (2012). A map of visual space in the primate entorhinal cortex. *Nature*, *491*, 761–764.
- Kleinfeld, D., & Sompolinsky, H. (1988). Associative neural network model for the generation of temporal patterns. Theory and application to central pattern generators. *Biophysical Journal*, *54*, 1039–1051.
- Kohonen, T. (1977). *Associative memory: A system theoretical approach*. New York: Springer.
- Kohonen, T. (1984). *Self-organization and associative memory*. Berlin: Springer-Verlag.
- Kohonen, T., Oja, E., & Lehtio, P. (1981). Storage and processing of information in distributed memory systems. In G. E. Hinton & J. A. Anderson (Eds.), *Parallel models of associative memory* (pp. 129–167). Hillsdale: Lawrence Erlbaum.
- Kondo, H., Lavenex, P., & Amaral, D. G. (2009). Intrinsic connections of the macaque monkey hippocampal formation: II. CA3 connections. *Journal of Comparative Neurology*, *515*, 349–377.
- Kraus, B. J., Robinson, R. J., 2nd, White, J. A., Eichenbaum, H., & Hasselmo, M. E. (2013). Hippocampal “time cells”: Time versus path integration. *Neuron*, *78*, 1090–1101.
- Kropff, E., & Treves, A. (2008). The emergence of grid cells: Intelligent design or just adaptation? *Hippocampus*, *18*, 1256–1269.
- Kumaran, D., & McClelland, J. L. (2012). Generalization through the recurrent interaction of episodic memories: A model of the hippocampal system. *Psychological Review*, *119*, 573–616.
- 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. *Neurobiology of Learning and Memory*, *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. *Journal of Comparative Neurology*, *472*, 371–394.
- Lee, I., & Kesner, R. P. (2004). 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.
- 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. *Journal of Neuroscience*, *29*, 3475–3484.
- Leutgeb, S., & Leutgeb, J. K. (2007). Pattern separation, pattern completion, and new neuronal codes within a continuous CA3 map. *Learning & Memory*, *14*, 745–757.
- Leutgeb, J. K., Leutgeb, S., Moser, M. B., & Moser, E. I. (2007). Pattern separation in the dentate gyrus and CA3 of the hippocampus. *Science*, *315*, 961–966.
- Levy, W. B. (1989). A computational approach to hippocampal function. In R. D. Hawkins & G. H. Bower (Eds.), *Computational models of learning in simple neural systems* (pp. 243–305). San Diego: Academic Press.
- Lisman, J., & Redish, A. D. (2009). Prediction, sequences and the hippocampus. *Philosophical Transactions of the Royal Society of London, Series B: Biological Sciences*, *364*, 1193–1201.
- Lynch, M. A. (2004). Long-term potentiation and memory. *Physiological Reviews*, *84*, 87–136.
- Macdonald, C. J., Lepage, K. Q., Eden, U. T., & Eichenbaum, H. (2011). Hippocampal “time cells” bridge the gap in memory for discontinuous events. *Neuron*, *71*, 737–749.
- Markov, N. T., Vezoli, J., Chameau, P., Falchier, A., Quilodran, R., Huissoud, C., et al. (2014). Anatomy of hierarchy: Feedforward and feedback pathways in macaque visual cortex. *Journal of Comparative Neurology*, *522*, 225–259.
- Marr, D. (1969). A theory of cerebellar cortex. *Journal of Physiology*, *202*, 437–470.
- Marr, D. (1971). Simple memory: A theory for archicortex. *Philosophical Transactions of the Royal Society of London, Series B: Biological Sciences*, *262*, 23–81.
- 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. *Psychological Review*, *102*, 419–457.
- McHugh, T. J., Jones, M. W., Quinn, J. J., Balthasar, N., Coppari, R., Elmquist, J. K., et al. (2007). Dentate gyrus NMDA receptors mediate rapid pattern separation in the hippocampal network. *Science*, *317*, 94–99.
- McNaughton, B. L. (1991). Associative pattern completion in hippocampal circuits: New evidence and new questions. *Brain Research Reviews*, *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. *Experimental Brain Research*, *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’. *Nature Reviews Neuroscience*, *7*, 663–678.
- McNaughton, B. L., & Morris, R. G. M. (1987). Hippocampal synaptic enhancement and information storage within a distributed memory system. *Trends in Neurosciences*, *10*, 408–415.
- McNaughton, B. L., & Nadel, L. (1990). Hebb–Marr networks and the neurobiological representation of action in space. In M. A. Gluck & D. E. Rumelhart (Eds.), *Neuroscience and connectionist theory* (pp. 1–64). Hillsdale, NJ: Erlbaum.
- Merrill, D. A., Chiba, A. A., & Tuszyński, M. H. (2001). Conservation of neuronal number and size in the entorhinal cortex of behaviorally characterized aged rats. *Journal of Comparative Neurology*, *438*, 445–456.
- Miller, G. A. (1956). The magic number seven, plus or minus two: Some limits on our capacity for the processing of information. *Psychological Reviews*, *63*, 81–93.
- Miller, P., & Wang, X. J. (2006). Inhibitory control by an integral feedback signal in prefrontal cortex: A model of discrimination between sequential stimuli. *Proceedings of the National Academy of Science of the United States of America*, *103*, 201–206.
- Morris, R. G. (2003). Long-term potentiation and memory. *Philosophical Transactions of the Royal Society of London, Series B: Biological Sciences*, *358*, 643–647.
- Morris, R. G. M. (1989). Does synaptic plasticity play a role in information storage in the vertebrate brain? In R. G. M. Morris (Ed.), *Parallel distributed processing: Implications for psychology and neurobiology* (pp. 248–285). Oxford: Oxford University Press.
- Morris, R. G., Moser, E. I., Riedel, G., Martin, S. J., Sandin, J., Day, M., et al. (2003). Elements of a neurobiological theory of the hippocampus: The role of activity-dependent synaptic plasticity in memory. *Philosophical Transactions of the Royal Society of London, Series B: Biological Sciences*, *358*, 773–786.
- Moscovitch, M., Rosenbaum, R. S., Gilboa, A., Addis, D. R., Westmacott, R., Grady, C., et al. (2005). Functional neuroanatomy of remote episodic, semantic and spatial memory: A unified account based on multiple trace theory. *Journal of Anatomy*, *207*, 35–66.
- Moser, M. B., Rowland, D. C., & Moser, E. I. (2015). Place cells, grid cells, and memory. *Cold Spring Harbor Perspectives in Biology*, *7*, a021808.
- 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.
- Nakashiba, T., Cushman, J. D., Pelkey, K. A., Renaudineau, S., Buhl, D. L., McHugh, T. J., et al. (2012). Young dentate granule cells mediate pattern separation, whereas old granule cells facilitate pattern completion. *Cell*, *149*, 188–201.
- Nakazawa, K., McHugh, T. J., Wilson, M. A., & Tonegawa, S. (2004). NMDA receptors, place cells and hippocampal spatial memory. *Nature Reviews Neuroscience*, *5*, 361–372.
- Nakazawa, K., Quirk, M. C., Chitwood, R. A., Watanabe, M., Yeckel, M. F., Sun, L. D., et al. (2002). Requirement for hippocampal CA3 NMDA receptors in associative memory recall. *Science*, *297*, 211–218.
- Nakazawa, K., Sun, L. D., Quirk, M. C., Rondi-Reig, L., Wilson, M. A., & Tonegawa, S. (2003). Hippocampal CA3 NMDA receptors are crucial for memory acquisition of one-time experience. *Neuron*, *38*, 305–315.
- Neunuebel, J. P., & Knierim, J. J. (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, J., & Speakman, A. (1987). Single unit activity in the rat hippocampus during a spatial memory task. *Experimental Brain Research*, *68*, 1–27.
- Oja, E. (1982). A simplified neuron model as a principal component analyser. *Journal of Mathematical Biology*, *15*, 267–273.
- O’Keefe, J. (1979). A review of the hippocampal place cells. *Progress in Neurobiology*, *13*, 419–439.
- O’Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map*. Oxford: Clarendon Press.
- 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., & Munakata, Y. (2000). *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain*. Cambridge, MA: MIT Press.
- Pastalkova, E., Itskov, V., Amarasingham, A., & Buzsáki, G. (2008). Internally generated cell assembly sequences in the rat hippocampus. *Science*, *321*, 1322–1327.
- 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.
- Rajji, T., Chapman, D., Eichenbaum, H., & Greene, R. (2006). The role of CA3 hippocampal NMDA receptors in paired associate learning. *Journal of Neuroscience*, *26*, 908–915.
- Robertson, R. G., Rolls, E. T., & Georges-François, P. (1998). Spatial view cells in the primate hippocampus: Effects of removal of view details. *Journal of Neurophysiology*, *79*, 1145–1156.
- Rolls, E. T. (1990b). Theoretical and neurophysiological analysis of the functions of the primate hippocampus in memory. *Cold Spring Harbor Symposia in Quantitative Biology*, *55*, 995–1006.
- Rolls, E. T. (1991). Functions of the primate hippocampus in spatial and non-spatial memory. *Hippocampus*, *1*, 258–261.
- Rolls, E. T. (1995). A model of the operation of the hippocampus and entorhinal cortex in memory. *International Journal of Neural Systems*, *6*, 51–70.

- 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. (2008). *Memory, attention, and decision-making: A unifying computational neuroscience approach*. Oxford: Oxford University Press.
- Rolls, E. T. (2010a). Attractor networks. *WIREs Cognitive Science*, 1, 119–134.
- Rolls, E. T. (2010b). A computational theory of episodic memory formation in the hippocampus. *Behavioural Brain Research*, 215, 180–196.
- Rolls, E. T. (2012a). Advantages of dilution in the connectivity of attractor networks in the brain. *Biologically Inspired Cognitive Architectures*, 1, 44–54.
- Rolls, E. T. (2012b). Invariant visual object and face recognition: Neural and computational bases, and a model, VisNet. *Frontiers in Computational Neuroscience*, 6(35), 1–70.
- Rolls, E. T. (2013). A quantitative theory of the functions of the hippocampal CA3 network in memory. *Frontiers in Cellular Neuroscience*, 7, 98.
- Rolls, E. T. (2014). *Emotion and decision-making explained*. Oxford: Oxford University Press.
- Rolls, E. T. (2015a). Diluted connectivity in pattern association networks facilitates the recall of information from the hippocampus to the neocortex. *Progress in Brain Research*, 219, 21–32.
- Rolls, E. T. (2015b). Limbic systems for emotion and for memory, but no single limbic system. *Cortex*, 62, 119–157.
- Rolls, E. T. (2015c). The neuronal representation of information in the human brain. *Review. Brain*, in press.
- Rolls, E. T. (2016). *Cerebral cortex: Principles of operation*. Oxford: Oxford University Press.
- Rolls, E. T. (1989b). Functions of neuronal networks in the hippocampus and neocortex in memory. In J. H. Byrne & W. O. Berry (Eds.), *Neural models of plasticity: Experimental and theoretical approaches* (pp. 240–265). San Diego: Academic Press.
- Rolls, E. T. (1987). Information representation, processing and storage in the brain: Analysis at the single neuron level. In J.-P. Changeux & M. Konishi (Eds.), *The neural and molecular bases of learning* (pp. 503–540). Chichester: Wiley.
- Rolls, E. T. (1989a). Functions of neuronal networks in the hippocampus and cerebral cortex in memory. In R. M. J. Cotterill (Ed.), *Models of brain function* (pp. 15–33). Cambridge: Cambridge University Press.
- Rolls, E. T., & Deco, G. (2002). *Computational neuroscience of vision*. Oxford: Oxford University Press.
- Rolls, E. T., & Deco, G. (2010). *The noisy brain: Stochastic dynamics as a principle of brain function*. Oxford: Oxford University Press.
- Rolls, E. T., Dempere-Marco, L., & Deco, G. (2013). Holding multiple items in short term memory: A neural mechanism. *PLoS ONE*, 8, e61078.
- Rolls, E. T. (1989d). The representation and storage of information in neuronal networks in the primate cerebral cortex and hippocampus. In R. Durbin, C. Miall, & G. Mitchison (Eds.), *The computing neuron* (pp. 125–159). Wokingham, England: Addison-Wesley.
- Rolls, E. T. (1996a). Roles of long term potentiation and long term depression in neuronal network operations in the brain. In M. S. Fazel & G. L. Collingridge (Eds.), *Cortical plasticity* (pp. 223–250). Oxford: Bios.
- Rolls, E. T., & Kesner, R. P. (2006). A computational theory of hippocampal function, and empirical tests of the theory. *Progress in Neurobiology*, 79, 1–48.
- Rolls, E. T. (1989c). Parallel distributed processing in the brain: Implications of the functional architecture of neuronal networks in the hippocampus. In R. G. M. Morris (Ed.), *Parallel distributed processing: Implications for psychology and neurobiology* (pp. 286–308). Oxford: Oxford University Press.
- Rolls, E. T. (1990a). Functions of the primate hippocampus in spatial processing and memory. In D. S. Olton & R. P. Kesner (Eds.), *Neurobiology of comparative cognition* (pp. 339–362). Hillsdale, NJ: L. Erlbaum.
- Rolls, E. T., Robertson, R. G., & Georges-François, P. (1997). Spatial view cells in the primate hippocampus. *European Journal of Neuroscience*, 9, 1789–1794.
- Rolls, E. T., & Stringer, S. M. (2000). On the design of neural networks in the brain by genetic evolution. *Progress in Neurobiology*, 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 Networks*, 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: Computation in Neural Systems*, 17, 447–465.
- Rolls, E. T., Stringer, S. M., & Trappenberg, T. P. (2002). A unified model of spatial and episodic memory. *Proceedings of the Royal Society of London 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*, 1, 407–421.
- Rolls, E. T., & Treves, A. (1998). *Neural networks and brain function*. Oxford: Oxford University Press.
- Rolls, E. T., & Treves, A. (2011). The neuronal encoding of information in the brain. *Progress in Neurobiology*, 95, 448–490.
- Rolls, E. T., Treves, A., Foster, D., & Perez-Vicente, C. (1997). Simulation studies of the CA3 hippocampal subfield modelled as an attractor neural network. *Neural Networks*, 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. *Journal of Neurophysiology*, 79, 1797–1813.
- Rolls, E. T., & Webb, T. J. (2012). Cortical attractor network dynamics with diluted connectivity. *Brain Research Reviews*, 1434, 212–225.
- Rolls, E. T., & Xiang, J.-Z. (2005). Reward-spatial view representations and learning in the hippocampus. *Journal of Neuroscience*, 25, 6167–6174.
- Rolls, E. T., & Xiang, J.-Z. (2006). Spatial view cells in the primate hippocampus, and memory recall. *Reviews in the Neurosciences*, 17, 175–200.
- Rolls, E. T., Xiang, J.-Z., & Franco, L. (2005). Object, space and object-space representations in the primate hippocampus. *Journal of Neurophysiology*, 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. *Proceedings of the National Academy of Science of the United States of America*, 98, 3543–3548.
- Samsonovich, A., & McNaughton, B. L. (1997). Path integration and cognitive mapping in a continuous attractor neural network model. *Journal of Neuroscience*, 17, 5900–5920.
- Schultz, S., Panzeri, S., Rolls, E. T., & Treves, A. (2000). Quantitative model analysis of a Schaffer collateral model. In P. Hancock, P. Földiák, & R. Baddeley (Eds.), *Information theory and the brain* (pp. 257–272). Cambridge, UK: Cambridge University Press.
- 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. *Progress in Brain Research*, 193, 163–177.
- Sidhu, M. K., Stretton, J., Winston, G. P., Bonelli, S., Centeno, M., Vollmar, C., et al. (2013). A functional magnetic resonance imaging study mapping the episodic memory encoding network in temporal lobe epilepsy. *Brain*, 136, 1868–1888.
- Simmen, M. W., Treves, A., & Rolls, E. T. (1996). Pattern retrieval in threshold-linear associative nets. *Network*, 7, 109–122.
- Sompolinsky, H., & Kanter, I. I. (1986). Temporal association in asymmetric neural networks. *Physical Review Letters*, 57, 2861–2864.
- 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. *Journal of Comparative Neurology*, 375, 552–582.
- Stella, F., Cerasti, E., & Treves, A. (2013). Unveiling the metric structure of internal representations of space. *Frontiers in Neural Circuits*, 7, 81.
- Storm-Mathieson, J., Zimmer, J., & Ottersen, O. P. (Eds.). (1990). *Understanding the brain through the hippocampus*. Oxford: Elsevier.
- Stringer, S. M., & Rolls, E. T. (2002). Invariant object recognition in the visual system with novel views of 3D objects. *Neural Computation*, 14, 2585–2596.
- Stringer, S. M., & Rolls, E. T. (2006). Self-organizing path integration using a linked continuous attractor and competitive network: Path integration of head direction. *Network: Computation in Neural Systems*, 17, 419–445.
- 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 Networks*, 17, 5–27.
- Stringer, S. M., Rolls, E. T., & Trappenberg, T. P. (2005). Self-organizing continuous attractor network models of hippocampal spatial view cells. *Neurobiology of Learning and Memory*, 83, 79–92.
- Stringer, S. M., Rolls, E. T., Trappenberg, T. P., & Araujo, I. E. T. (2002). Self-organizing continuous attractor networks and path integration. Two-dimensional models of place cells. *Network: Computation in Neural Systems*, 13, 429–446.
- Stringer, S. M., Trappenberg, T. P., Rolls, E. T., & Araujo, I. E. T. (2002). Self-organizing continuous attractor networks and path integration: One-dimensional models of head direction cells. *Network: Computation in Neural Systems*, 13, 217–242.
- Suzuki, W. A., & Amaral, D. G. (1994a). Perirhinal and parahippocampal cortices of the macaque monkey – Cortical afferents. *Journal of Comparative Neurology*, 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. *Journal of Neuroscience*, 14, 1856–1877.
- Taylor, J. G. (1999). Neural “bubble” dynamics in two dimensions: Foundations. *Biological Cybernetics*, 80, 393–409.
- Tonegawa, S., Nakazawa, K., & Wilson, M. A. (2003). Genetic neuroscience of mammalian learning and memory. *Philosophical Transactions of the Royal Society of London, Series B: Biological Sciences*, 358, 787–795.
- Treves, A. (1990). Graded-response neurons and information encodings in autoassociative memories. *Physical Review A*, 42, 2418–2430.
- Treves, A. (1991). Dilution and sparse coding in threshold-linear nets. *Journal of Physics A*, 24, 327–335.
- Treves, A. (1995). Quantitative estimate of the information relayed by Schaffer collaterals. *Journal of Computational Neuroscience*, 2, 259–272.
- Treves, A., & Rolls, E. T. (1991). What determines the capacity of autoassociative memories in the brain? *Network*, 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.
- 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 in Neuroscience*, 5, 345–350.
- van Vreeswijk, C., & Sompolinsky, H. (1998). Chaotic balanced state in a model of cortical circuits. *Neural Computation*, 10, 1321–1371.

- Walters, D. M., Stringer, S. M., & Rolls, E. T. (2013). Path integration of head direction: Updating a packet of neural activity at the correct speed using axonal conduction delays. *PLoS ONE*, 8, e58330.
- Wang, X. J. (2002). Probabilistic decision making by slow reverberation in cortical circuits. *Neuron*, 36, 955–968.
- Wang, S. H., & Morris, R. G. (2010). Hippocampal-neocortical interactions in memory formation, consolidation, and reconsolidation. *Annual Review of Psychology*, 61(49–79), C41–C44.
- 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.
- 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., & Buckingham, J. T. (1990). An assessment of Marr's theory of the hippocampus as a temporary memory store. *Philosophical Transactions of the Royal Society B*, 329, 205–215.
- Wilson, M. A., & McNaughton, B. L. (1994). Reactivation of hippocampal ensemble memories during sleep. *Science*, 265, 676–679.
- 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 connective heterogeneity. *Learning & Memory*, 14, 705–713.
- Witter, M. P., Naber, P. A., van Haften, T., Machielsen, W. C., Rombouts, S. A., Barkhof, F., et al. (2000). Cortico-hippocampal communication by way of parallel parahippocampal-subicular pathways. *Hippocampus*, 10, 398–410.
- Witter, M. P., Wouterlood, F. G., Naber, P. A., & Van Haften, T. (2000). Anatomical organization of the parahippocampal-hippocampal network. *Annals of the New York Academy of Sciences*, 911, 1–24.
- Zhang, K. (1996). Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: A theory. *Journal of Neuroscience*, 16, 2112–2126.
- Zilli, E. A. (2012). Models of grid cell spatial firing published 2005–2011. *Front Neural Circuits*, 6, 16.