

A unified model of spatial and episodic memory

Edmund T. Rolls*, Simon M. Stringer and Thomas P. Trappenberg

Department of Experimental Psychology, University of Oxford, South Parks Road, Oxford OX1 3UD, UK

Medial temporal lobe structures including the hippocampus are implicated by separate investigations in both episodic memory and spatial function. We show that a single recurrent attractor network can store both the discrete memories that characterize episodic memory and the continuous representations that characterize physical space. Combining both types of representation in a single network is actually necessary if objects and where they are located in space must be stored. We thus show that episodic memory and spatial theories of medial temporal lobe function can be combined in a unified model.

Keywords: continuous attractor network; discrete attractor network; spatial memory; place cells; spatial view cells; hippocampus

1. INTRODUCTION

Evidence implicating the hippocampus and connected structures in the medial temporal lobe in episodic memory is that in humans with bilateral damage to this region anterograde amnesia is produced, with the memory for particular events and episodes being particularly impaired (Squire 1992; Rempel-Clower *et al.* 1996). We use the term episodic memory to capture the memory of a single event that might occur when a previous particular occasion is recalled. The event (something that happens at a particular time) might include both spatial and non-spatial components, including for example where the event occurred, who was present, and which objects were seen. (We use the term episodic memory to refer to the memory of a particular event, and are not concerned in this paper with describing linked temporal sequences of events.) The hippocampus is also implicated in spatial memory. For example, damage to the hippocampal system in monkeys produces deficits in learning about where objects are and where responses must be made (Rolls 1996; Gaffan 1998), and in rats produces spatial learning deficits (Martin *et al.* 2000). Neurophysiologically, hippocampal neurons in rats respond to the place where the animal is located (O'Keefe 1990; Kubie & Muller 1991; Wilson & McNaughton 1993), and in primates to the place being viewed (Rolls *et al.* 1997; Rolls 1999). It has thus been a long-standing question about whether the hippocampus and nearby temporal lobe structures are involved in episodic memory or spatial function. In this paper we show that this question can be resolved by revealing that a single neural network can implement both episodic and spatial memory.

The simplest form of network that can store discrete (separate) patterns that characterize the single events, or objects that are the parts of an episodic memory and can retrieve the whole memory from any of its parts, is an autoassociation network, illustrated in figure 1 (Kohonen 1977; Hopfield 1982). In this type of network, each memory is stored by using an associative (Hebbian) learning rule which can be implemented by long-term potentiation

to modify the synaptic connections in the recurrent collateral synapses for the neurons that are firing to represent the pattern. It is suggested that a memory of this type, implemented by the associatively modifiable synapses in the recurrent connections between hippocampal CA3 cells, could implement episodic memories of discrete events or objects (Rolls 1987, 1989*a,b*; Treves & Rolls 1992, 1994).

The simplest memory model that can implement the spatial representations provided by place cells in the rat hippocampus and by spatial view cells in the primate hippocampus, and enables the firing to be maintained in darkness as found experimentally, is a continuous attractor network (Amari 1977; Samsonovich & McNaughton 1997; Battaglia & Treves 1998; Taylor 1999; Tsodyks 1999; Stringer *et al.* 2002*b*). This network has the same architecture and rules of operation as the discrete attractor network of episodic memory, but uses continuous patterns which are needed to reflect the continuous nature of physical space.

In this paper we analyse an attractor network that can store both continuous (e.g. spatial) patterns and discrete (e.g. single object) patterns. We show in this paper that both types of pattern (continuous spatial and discrete, as illustrated in figure 2) can be stored in the same network; that the packet of activity representing the current location can still be moved in the continuous space; that the place can be used as a retrieval cue to recall the object at that place; and that the object can be used as a retrieval cue to recall the place of the object. Moreover, we argue that both types of memory pattern must coexist in the same network in the brain if associations are to be made between patterns which are inherently continuous (e.g. spatial) and inherently separate (e.g. different objects). This approach unifies the spatial and episodic approaches to medial temporal lobe function, showing that both types of representation must be stored in the same network if associations between them are to be learned and retrieved. Previous models of hippocampal function have focused primarily on networks that can store discrete patterns suitable for episodic memory (Marr 1971; Treves & Rolls 1992, 1994; McClelland *et al.* 1995; Rolls 1996), or on continuous patterns suitable for spatial memory (Samsonovich & McNaughton 1997; Redish & Touretzky

* Author for correspondence (edmund.rolls@psy.ox.ac.uk).

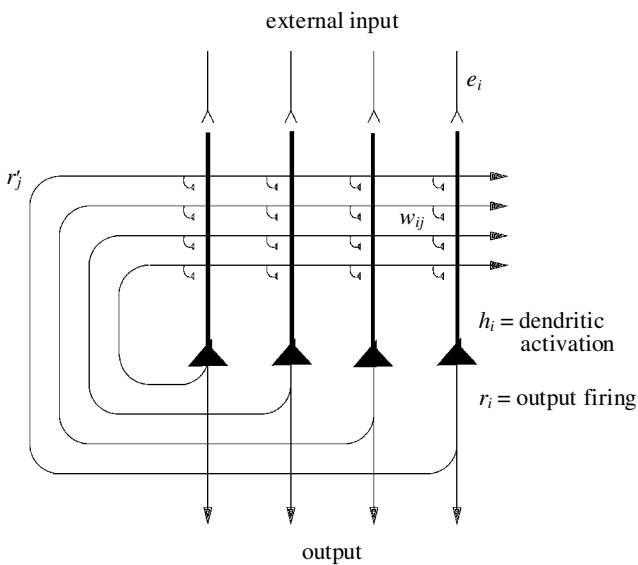


Figure 1. The architecture of an attractor network. A set of neurons indexed by i receives external inputs e_i . The firing of each neuron is sent by recurrent collateral connections to make synapses w_{ij} on the dendrites of other neurons in the population. The symbol h_i represents the activation of node i by the external inputs and by the recurrent synapses w_{ij} , and $r_i (= r'_i)$ the corresponding firing rate.

1998; Tsodyks 1999; Kali & Dayan 2000; Stringer *et al.* 2002a).

2. METHODS AND RESULTS

(a) *The model*

The model we propose that could combine continuous and discrete patterns in the same single network to allow associations between any of the patterns is as follows. This model instantiates the hypothesis we propose and consists of a single attractor network of the type illustrated in figure 1. There are recurrent excitatory modifiable connections between the neurons. The synaptic connectivity could be complete or the connectivity could be partial (diluted). Inputs e_i are introduced to the network by unmodifiable connections.

Two types of input pattern are used. For neurons trained with spatial patterns, each neuron has Gaussian firing as a function of location, consistent with place cells in rats and spatial-view cells in primates. More precisely, the firing rate r_i of a spatial cell is a Gaussian function of the displacement of the agent from the optimal firing location of the cell in the space,

$$r_i = \exp(-\delta^2/2\sigma^2), \quad (2.1)$$

where δ is the difference between the actual spatial location x of the agent and the optimal spatial location x_i for cell i . Different spatial neurons have their response centred at different locations in the continuous space. An illustration of the type of tuning of individual spatial cells is provided in figure 2.

For neurons tuned to represent discrete objects/events, each object is represented by a set of active neurons. The distributed representation could be fully distributed (with, for example, half of the object neurons in the active state, and half in the inactive state for each object), or sparse.

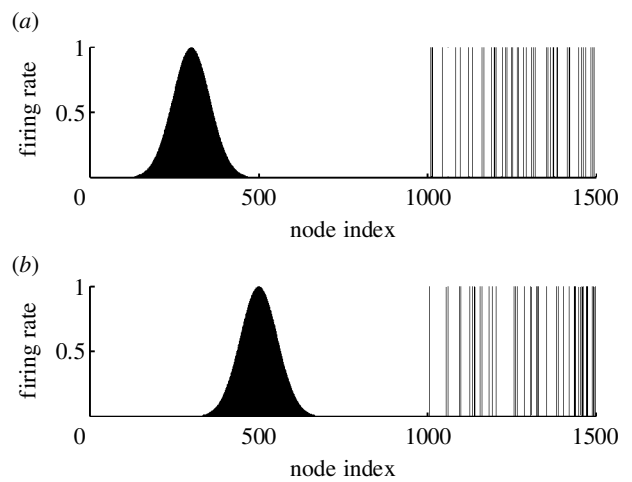


Figure 2. The firing patterns for two different memories (1 and 2) of nodes 1–1000, which represent the spatial firing and nodes 1001–1500, which represent the discrete events that occur at the location. The continuous nature of the spatial representation results from the fact that each neuron has a Gaussian firing rate that peaks at its optimal location. (a) The firing rates of the set of active neurons when the animal is located at the place at which neuron 300 fires maximally (memory 1). (b) The activity of the spatial and event neurons when the animal is at location 500 (memory 2). During learning, memory 1 is stored in the synaptic weights, then memory 2, etc., and each memory contains a part which is continuously distributed to represent physical space, and a part which represents a discrete event.

For sparse representations, the sparseness a can be defined as (Treves & Rolls 1991; Rolls & Treves 1998)

$$a = \frac{\left(\sum_i r_i/N\right)^2}{\sum_i (r_i^2)/N}, \quad (2.2)$$

where N is the number of neurons in the network. (For binary neurons where 1 is the active state and 0 the inactive, the sparseness evaluates to 0.1 if 10% of the neurons are active.) For the simulations, to ensure that the representations of each object were discrete, we made the representations orthogonal.

Identical training is used for the two types of pattern. The recurrent collateral synaptic weights between the nodes are modified with an associative (Hebbian) rule

$$\delta w_{ij} = k r_i r_j, \quad (2.3)$$

where δw_{ij} is the change in the synaptic weight w_{ij} from node j to node i , r_i is the firing rate of the postsynaptic node i , r_j is the firing rate of the presynaptic node j , and k is the learning rate. We note that this very simple and biologically plausible learning rule (equation (2.3)) is different from the Hebbian covariance rule

$$\delta w_{ij} = k(r_i - \langle r_i \rangle)(r_j - \langle r_j \rangle) \quad (2.4)$$

commonly employed in discrete attractor network memory models. The rule we use is more biologically plausible in that it does not postulate that the synaptic weights increase if both the pre- and the postsynaptic elements are inactive. However, we note that the covariance rule

becomes close to equation (2.3) as the sparse limit is approached ($a \rightarrow 0$). We demonstrate in the simulations described below that equation (2.3) together with global inhibition is indeed sufficient to enable the network to store and retrieve both object and spatial memories. (We note that in the brain there are decremting synaptic strength processes that may act to limit the total strength reached by synapses (Rolls & Deco 2002).)

The firing rate of each neuron in the network during retrieval is calculated as follows. The activation h_i of the i th neuron is calculated as

$$h_i = \sum_j w_{ij} r_j + I_i^{\text{EXT}} - \sum_j w^{\text{INH}} r_j, \quad (2.5)$$

where r_j the firing rate of the j th input neuron, w_{ij} is an excitatory (positive) synaptic weight from node j to node i trained with equation (2.3), and I_i^{EXT} is the synaptic current produced by an external input to the network (e_i in figure 1). The last term in the equation is subtractive inhibition, where w^{INH} represents the effects of a set of inhibitory interneurons with unmodifiable synaptic weights which receive inputs from the excitatory neurons in the network. This inhibition keeps the operation of the network in a stable regime. The resulting activations of the excitatory neurons are converted into firing r_i by a typical nonlinear activation function

$$r_i = g(h_i) \quad (2.6)$$

incorporating some type of neuronal firing threshold, e.g. a sigmoid, binary threshold or threshold linear activation function.

This completes the description of the model. We next demonstrate, in the simulations described below, the operation of a network that simulates the model, and use this to illustrate some of the properties of the model. We then go on to consider some of the generic properties of the model, such as its storage capacity for memories that consist of combinations of discrete and continuous patterns.

(b) Simulation and properties of the model

(i) The network simulated

The 1500 neurons or nodes of the network used in the simulation were divided into two classes of cells with different response properties, of the type illustrated in figure 2. Nodes 1–1000 provided a spatial representation, with each neuron tuned to have Gaussian firing as a function of location, consistent with place cells in rats and spatial-view cells in primates. More precisely, the firing rate of the spatial cells was a Gaussian function of the displacement of the agent from the optimal firing location of the cell in the space, as defined in equation (2.1). The standard deviation σ of the Gaussian was 55 nodes. In the hippocampus, nearby neurons do not typically represent nearby locations, but for purposes of illustration the neurons have been arranged in figure 2 so neurons that represent neighbouring locations are neighbours in the diagram.

Nodes 1001–1500 represented discrete objects or events. Each ‘object’ was represented by a set of 50 of these 500 cells in the active (1) state, with the remaining cells in the inactive (0) state; that is, the sparseness of the (orthogonal) representations was 0.1 (see equation (2.2)). Two examples of the resulting stimulus patterns that were

presented during learning to the single attractor network are illustrated in figure 2. In memory pattern 1 (figure 2a), the spatial nodes represent location 300, and 50 randomly selected nodes in the set 1001–1500 represent an object or event in that location. The set of neurons that are active when the animal is at, or looking at, the location where neuron 500 fires optimally is shown in figure 2b.

To train the network with the patterns, the recurrent collateral synaptic weights between the nodes were trained with the associative Hebbian rule shown in equation (2.3). The network was trained at 50 equispaced locations, at 10 of which discrete objects were present. (Although we could have trained with spatial patterns at every one of the 1000 spatial locations, we have found that the properties of continuous attractor networks that we have described (Stringer *et al.* 2002b) can all be obtained with training on a subset of the spatial locations. This interesting fact is due to the relatively wide spatial tuning of each cell, so that after training, the whole net is not very far from having a level energy landscape and can display the properties described previously. Indeed, this property of continuous attractor networks helps these networks to be biologically plausible, in that training at every possible location is not needed, and this is an interesting aspect of the results described here. To ensure that the weights produced by the training with spatial patterns did not dominate the network due to there being more synaptic updates with the spatial than discrete patterns, the learning rate k in equation (2.3) was set to be 1 for the continuous patterns, and *ca.* 5 for the discrete patterns.)

The time evolution of the system was simulated by sigmoidal nodes with the continuous dynamics of leaky integrator neurons,

$$\tau \frac{dh_i(t)}{dt} = -h_i(t) + \frac{\phi_0}{C} \sum_j (w_{ij} - w^{\text{INH}}) r_j(t) + I_i^{\text{EXT}} \quad (2.7)$$

and

$$r_i(t) = \frac{1}{[1 + \exp(-2\beta(h_i(t) - \alpha))]} \quad (2.8)$$

where h_i is the activity of node i , r_i the corresponding firing rate, I_i^{EXT} is the synaptic current produced by an external input (e_i in figure 1), and the matrix element w_{ij} is an excitatory (positive) synaptic weight from node j to node i trained with equation (2.3). The inhibition parameter w^{INH} describing the effect of inhibitory interneurons was set to $w^{\text{INH}} = 49$ in the simulations. The value of the time constant was set to $\tau = 10$ ms. The slope of the sigmoidal neuronal activation function was fixed to the value $\beta = 0.035$. C is the number of connections received by each neuron in the recurrent collateral network, and ϕ_0 is a constant.

To help stabilize the activity packet in the network, we incorporated a voltage-dependent nonlinearity in the postsynaptic neuron, which can help to maintain firing in a recurrent network and which could be implemented by N-methyl-D-aspartate (NMDA) receptors (Lisman *et al.* 1998; Stringer *et al.* 2002b). The effect was implemented by altering the threshold value in the activation function by changing α in equation (2.8) from 0 to -10 when the firing rate r_i of the postsynaptic cell was 50% or more of its maximum value.




Figure 3. The performance during and after external stimulation of the trained network. (a) Retrieval in the event nodes, shown by the correlation between the recalled state of firing and the stored state for each of the trained discrete events represented by the activity of nodes 1001–1500. (b) The firing of each node (indicated by the colour, where red represents the highest firing) is shown as a function of time in ms. (The time constant of the dynamics of the neurons was 10 ms.) Nodes 1–1000 reflect the spatial component of the memory, and nodes 1001–1500 represent the discrete event that is a part of each memory. Retrieval cues were applied as shown during the solid bars at the base of the diagram (see description in text). The spatial nodes were arranged topologically in the diagram for clarity. The actual topology in the hippocampus is represented by the strength of the connections between neurons, rather than by their physical proximity to each other.

(ii) Recall of event memory from spatial or object retrieval cues

For all the experiments described, the network was first trained as described above. In the first set of four experiments, the results of which are illustrated in figure 3, we tested the performance of the memory network for various combinations of location and object retrieval cues as follows. From time 0 until time $t = 100$ ms, an external location retrieval cue centred at node 300 was applied to the initially quiescent network. This external spatial input produced (within 5 ms) strong firing in the neurons that represent continuous space, and this firing was maintained even after the retrieval cue was withdrawn at time $t = 100$ ms, illustrating the persistent spatial memory implemented by the continuous attractor properties of the network. At the same time that the continuous attractor firing developed, the discrete memory for the event that occurred at that location was retrieved, as indicated by the set of nodes in the range 1001–1500 that became active. The measure of retrieval, the Pearson correlation between the retrieved discrete event and the event originally stored at that location, became high (close to 1.0; figure 3a) at the same time, illustrating correct retrieval of the discrete event associated with the location. This memory of the event continued after the spatial retrieval cue was removed at 100 ms.

Between 200 and 300 ms the network was provided

with a new spatial retrieval cue centred at location 800. The location represented by the packet of activity in the network changed to represent the new location (taking 7 ms), and at the same time the correct discrete event associated with that location was retrieved. The continuous spatial and discrete event memories were both maintained after the spatial retrieval cue was removed at time 300 ms.¹

Between 400 and 500 ms an event retrieval cue was applied to nodes 1001–1500. The retrieval cue was 50% of the event originally learned, i.e. only 25 of the 50 nodes that were the original training pattern used to represent that event were activated by the retrieval cue. It is shown in figure 3 that the correct discrete event representation was retrieved soon after 400 ms, demonstrating the important property of completion of a whole memory from a partial retrieval cue. At the same time, the correct spatial location where that event had occurred during training was retrieved, as shown by the packet of activity present in nodes 1–1000. This demonstrates correct retrieval of the location by the object. Both the correct location and the correct event were maintained after the event retrieval cue was removed at time 500 ms.

Between 600 and 700 ms the network was provided with a new spatial retrieval cue centred at location 300 together with an inconsistent event retrieval cue, the incomplete event retrieval cue used in the previous experiment (corresponding to the event associated during learning with location 500). The firing of the spatial cells in the network soon changed and reflected the external spatial retrieval cue. However, the object cells continued firing to the object corresponding to location 500, as this object or event was being presented to the network. Thus when the spatial and event inputs are inconsistent with previous learning, the inputs dominate the firing of the neurons. However, at time 700 ms both the spatial and the memory inputs were removed, and the network then moved into a state which was consistent with its memory, and the particular state it happened to fall into was the one in which the memory was of spatial location 300 and of the object originally trained at that location. This mode of operation is appropriate in that incoming sensory inputs dominate, and not memories, and this matches human perceptual and memory performance.

One of the interesting points made by the results shown in figure 3 is that in such spatial and object memory systems, individual neurons can assume states which reflect either the spatial input or the object input. Indeed, depending on the relative strengths of these two inputs (which is likely to be influenced by attention and the task being performed), a cell might appear to be more spatial, or more object-related. The implication of this is that when analysing the activity of single neurons in such a system (e.g. in the brain) concerned with object and spatial memories, cells that reflect object and event information will almost always be influenced in addition by spatial information. This is consistent with what is usually found from single neurons recorded in the hippocampus and closely connected areas, which rarely reflect only object/event information, and typically are also influenced by spatial factors (Rolls *et al.* 1989; Rolls & Treves 1998; Rolls 1999).

In a separate experiment we tested the object retrieval




Figure 4. Object retrieval of the network for different spatial retrieval cues centred around a spatial node at which an object had been trained. The network was trained with object memories centred at only 10 equispaced locations of the 1000 possible locations. The object training locations were thus 100 spatial locations apart. The correlation between the network states and the 10 different object patterns that had been trained are shown for locations on different sides of a training location.




Figure 5. The relationship between the final spatial state of the network (ordinate) after a spatial retrieval cue was shown at different distances from a location at which an object had been trained. In the clamped mode, which is when the spatial retrieval cue is being applied (dashed line), the final position of the network is relatively close to the spatial retrieval cue. When the spatial retrieval cue has been removed and the network has settled (solid line), then the final spatial location tends, for nodes close to a location at which an object has been trained, to settle at that node (the horizontal part of the solid line). For a retrieval cue a little further away from an object training location, the network remains stable at the node where the spatial retrieval pattern was presented (the diagonal part of the solid line).

of the network for different spatial retrieval cues centred around every node in the network. (Recall that the network was trained with object memories centred at only 10 equispaced locations of the 1000 possible locations. The object training locations were thus 100 spatial locations apart.) The correlation between the network states and the 10 different object patterns that had been trained are shown for locations on different sides of a training location in figure 4. In each instance the object that was retrieved by the spatial cue was that of the object in the closest spatial training location. Thus the network generalizes correctly in retrieving the object that was seen at the closest training location to the testing location.

(iii) *Operation with clamped and unclamped retrieval cues, with and without NMDA-based stabilization*

Having shown that the network implements object memory correctly and can retrieve a spatial memory from an object cue and vice versa, we now show that the network does retain some of the properties of a continuous attractor, such as maintaining a bubble of activity close to any location in which a retrieval cue has been presented when the location retrieval cue has been removed. The results of this investigation are shown in figure 5. The figure shows that in the clamped mode, which is when the spatial retrieval cue is being applied (dashed line), the final position of the network is relatively close to the spatial retrieval cue. The figure also shows that when the retrieval cue has been removed and the network has settled (solid line), then the final spatial location tends, for nodes close to a location at which an object has been trained, to settle at that node (the almost horizontal part of the solid line).

For a retrieval cue a little further away from an object training location, the network remains stable at the node where the spatial retrieval pattern was presented (the diagonal part of the solid line). Thus for locations where a discrete object has been trained, the continuous spatial property of the energy landscape has been made a bit uneven, so that nearby spatial location firing tends to drift towards the location at which an object has been trained. This result thus shows what is reasonable theoretically, that discrete patterns in an otherwise continuous landscape do alter the ability to support a stable attractor at every spatial location, but that this effect can be mild, with regions of normal spatial attractor operation between the object training locations (the diagonal parts of the solid lines in figure 5).

The results shown in figure 5 were obtained with mild NMDA receptor based stabilization ($\alpha = -30$ in equation (2.8) when the firing rate of the neuron was 50% or more of its maximum value). When no NMDA receptor stabilization was used, the width of the region within which the objects influenced the continuity of the spatial attractor increased from approximately ± 20 nodes to approximately ± 35 nodes. If the stabilization is increased ($\alpha = -100$), then the size of the basin of attraction created by the discrete pattern in the continuous landscape decreases, but the packet then becomes more stable, and is harder to move with a retrieval cue.

(c) Memory capacity of the model

The issue arises of the capacity of such a combined network. First, the capacity of a discrete attractor network with sparse patterns is in the order

$$p_{\max} \approx \frac{C}{a \ln(1/a)^k}, \quad (2.9)$$

where C is the number of connections per neuron from the recurrent collateral synapses, and k is a factor that depends weakly on the detailed structure of the rate distribution, the connectivity pattern, etc., but is roughly in the order of 0.2–0.3 (Treves & Rolls 1991). Such a network effectively stores a fixed amount of information per synapse, and more patterns can therefore be stored if less information is required to represent each pattern (Treves 1990, 1991; Treves & Rolls 1991).

Second, the capacity (or resolution) of a continuous attractor network can be considered to be specified in an analogous information-based way, where a^{SP} represents the sparseness of firing of a packet of neuronal activity in the spatial map or chart (in that the information per synapse is again approximately a constant) (Battaglia & Treves 1998; Tsodyks 1999).²

These considerations suggest that the (information) capacity of the combined network will simply reflect the contributions to the sparseness of the two sets of representations. Each set contributes to setting the overall sparseness of the representation of any one pattern of neuronal firing reflecting what is stored at any one time. Thus the number of memories that can be stored can be very large given the large number of connections onto each neuron from other neurons in attractor networks in the brain. For example, hippocampal CA3 neurons in the rat have approximately 12 000 connections from other CA3 neurons, leading to an estimated storage capacity which is well above 12 000 discrete memories if the patterns are sparse (Treves & Rolls 1991, 1994; Rolls & Treves 1998).

3. DISCUSSION

The results described here show that it is possible to store memories represented by discrete attractors, and spatial memories represented by continuous attractors, in the same network, and to use the learned associations to retrieve an object (episodic event) memory from a place, and a place from an object memory. Moreover, we propose that it is actually essential for the spatial and the event representations to be stored in the same single network, for otherwise arbitrary associations between events and places could not be stored and retrieved. Given that spatial representations imply continuous distributions because of the physical continuity of space, we believe that a network which stores continuous and discrete representations will inevitably be required in the brain, so that events and places can be associated, and indeed so that episodic memories, which usually consist of a spatial and a non-spatial component, can be formed.

A neurophysiological complement to the theory presented here is that recordings have established that hippocampal CA3 neurons provide both continuous spatial representations (place cells in rats (O'Keefe 1990; Kubie & Muller 1991; Wilson & McNaughton 1993), and spatial view cells in macaques (Rolls *et al.* 1997; Rolls

1999)), and also information about discrete objects and events.

The spatial representations are of an underlying physical space in that they can be updated by idiothetic (self-motion) cues made as an animal traverses the environment (Quirk *et al.* 1990; Markus *et al.* 1994; Robertson *et al.* 1998). Non-spatial information (which is a typical component of episodic memories) is represented in that in the macaque hippocampus, the activity of some neurons does reflect the object that was shown in an object–place memory task (Rolls *et al.* 1989; Rolls & Treves 1998; Rolls 1999), or a conditional spatial response task where the behavioural response depends on the object shown (Miyashita *et al.* 1989), and in that in the rat some neurons respond to non-spatial information in memory tasks (Wiener *et al.* 1989; Otto & Eichenbaum 1992; Hampson *et al.* 1999; Eichenbaum & Harris 2000). Both spatial and non-spatial information is likely to reach the entorhinal cortex, given its connections (Rolls 1996). From these projections reach the dentate granule cells, which operate, it is suggested, as a pattern separator (Rolls 1989*a,b*; Treves & Rolls 1992), allowing both spatial and non-spatial information (and combinations of the two) to reach the hippocampal CA3 cells, which could then store representations based on both spatial and non-spatial information as described here. We note that the direct entorhinal cortex to CA3 cell connections may be especially suitable for recalling memories from the CA3 network (Treves & Rolls 1992). Once the combined spatial and non-spatial memories have been recalled in CA3, the backprojections to the neocortex via CA1 and the entorhinal cortex could allow the separate neocortical spatial and non-spatial neuronal systems, active originally when each memory was stored, to become active again in the full memory retrieval operation (Treves & Rolls 1994; Rolls 1996).

In conclusion we have shown that the same process in a single network could store and retrieve spatial and non-spatial representations. The results described here show that such a network can operate with the properties of both continuous attractor and discrete attractor networks to provide a basis for remembering episodic events with spatial and non-spatial components, and can recall spatial from event representations, and vice versa. The relative preponderance of spatial cells in the rat hippocampus, and the inclusion of more non-spatial events in primate and human hippocampal processing, may reflect at least in part the great development of the visual object representation system in the inferior temporal visual cortical areas of primates (Rolls 2000; Rolls & Deco 2002).

This research was supported by the Medical Research Council, grant PG9826105, by the Human Frontier Science Program, and by the MRC Interdisciplinary Research Centre for Cognitive Neuroscience.

ENDNOTES

¹During the application of the spatial retrieval cue, it was found in the simulation that the activity of the object cells decreased. This is consistent with global inhibition within the hippocampus operating to maintain the sparseness of the representation approximately constant (Rolls & Treves 1998).

²In particular, the number of different charts that can be stored in a continuous attractor network is

$$p_{\max} \approx \frac{C}{\ln(1/a^{\text{SP}})^k}.$$

REFERENCES

- Amari, S. 1977 Dynamics of pattern formation in lateral-inhibition type neural fields. *Biol. Cybernet.* **27**, 77–87.
- Battaglia, F. P. & Treves, A. 1998 Attractor neural networks storing multiple space representations: a model for hippocampal place fields. *Phys. Rev. E* **58**, 7738–7753.
- Eichenbaum, H. & Harris, K. 2000 Toying with memory in the hippocampus. *Nature Neurosci.* **3**, 205–206.
- Gaffan, D. 1998 Idiothetic input into object–place configuration as the contribution to memory of the monkey and human hippocampus: a review. *Exp. Brain Res.* **123**, 201–209.
- Hampson, R. E., Simalal, J. D. & Deadwyler, S. A. 1999 Distribution of spatial and nonspatial information in dorsal hippocampus. *Nature* **402**, 610–614.
- Hopfield, J. J. 1982 Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl Acad. Sci. USA* **79**, 2554–2558.
- Kali, S. & Dayan, P. 2000 The involvement of recurrent connections in area CA3 in establishing the properties of place fields: a model. *J. Neurosci.* **20**, 7463–7477.
- Kohonen, T. 1977 *Associative memory: a system theoretical approach*. New York: Springer.
- Kubie, J. L. & Muller, R. U. 1991 Multiple representations in the hippocampus. *Hippocampus* **1**, 240–242.
- Lisman, J. E., Fellous, J. M. & Wang, X. J. 1998 A role for NMDA-receptor channels in working memory. *Nature Neurosci.* **1**, 273–275.
- McClelland, J. L., McNaughton, B. L. & O'Reilly, R. C. 1995 Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychol. Rev.* **102**, 419–457.
- Markus, E. J., Barnes, C. A., McNaughton, B. L., Gladden, V. L. & Skaggs, W. 1994 Spatial information content and reliability of hippocampal CA1 neurons: effects of visual input. *Hippocampus* **4**, 410–421.
- Marr, D. 1971 Simple memory: a theory for archicortex. *Phil. Trans. R. Soc. Lond. B* **262**, 23–81.
- Martin, S. J., Grimwood, P. D. & Morris, R. G. 2000 Synaptic plasticity and memory: an evaluation of the hypothesis. *A. Rev. Neurosci.* **23**, 649–711.
- Miyashita, Y., Rolls, E. T., Cahusac, P. M. B., Niki, H. & Feigenbaum, J. D. 1989 Activity of hippocampal neurons in the monkey related to a conditional spatial response task. *J. Neurophysiol.* **61**, 669–678.
- O'Keefe, J. 1990 A computational theory of the cognitive map. *Prog. Brain Res.* **83**, 301–312.
- Otto, T. & Eichenbaum, H. 1992 Neuronal activity in the hippocampus during delayed non-match to sample performance in rats: evidence for hippocampal processing in recognition memory. *Hippocampus* **2**, 323–334.
- Quirk, G. L., Muller, R. U. & Kubie, J. L. 1990 The firing of hippocampal place cells in the dark depends on the rat's recent experience. *J. Neurosci.* **10**, 2008–2017.
- Redish, A. D. & Touretzky, D. S. 1998 The role of the hippocampus in solving the Morris water maze. *Neural Comput.* **10**, 73–111.
- Rempel-Clower, N. L., Zola, S. M., Squire, L. R. & Amaral, D. G. 1996 Three cases of enduring memory impairment after bilateral damage limited to the hippocampal formation. *J. Neurosci.* **16**, 5233–5255.
- Robertson, R. G., Rolls, E. T. & Georges-Francois, P. 1998 Spatial view cells in the primate hippocampus: effects of removal of view details. *J. Neurophysiol.* **79**, 1145–1156.
- Rolls, E. T. 1987 Information representation, processing and storage in the brain: analysis at the single neuron level. In *The neural and molecular bases of learning* (ed. J.-P. Changeux & M. Konishi), pp. 503–540. Chichester, UK: Wiley.
- Rolls, E. T. 1989a Functions of neuronal networks in the hippocampus and neocortex in memory. In *Neural models of plasticity: experimental and theoretical approaches*, ch. 13 (ed. J. Byrne & W. Berry), pp. 240–265. San Diego, CA: Academic.
- Rolls, E. T. 1989b Parallel distributed processing in the brain: implications of the functional architecture of neuronal networks in the hippocampus. In *Parallel distributed processing: implications for psychology and neurobiology*, ch. 12 (ed. R. G. M. Morris), pp. 286–308. Oxford University Press.
- Rolls, E. T. 1996 A theory of hippocampal function in memory. *Hippocampus* **6**, 601–620.
- Rolls, E. T. 1999 Spatial view cells and the representation of place in the primate hippocampus. *Hippocampus* **9**, 467–480.
- Rolls, E. T. 2000 Functions of the primate temporal lobe cortical visual areas in invariant visual object and face recognition. *Neuron* **27**, 205–218.
- Rolls, E. T. & Deco, G. 2002 *Computational neuroscience of vision*. Oxford University Press.
- Rolls, E. T. & Treves, A. 1998 *Neural networks and brain function*. Oxford University Press.
- Rolls, E. T., Miyashita, Y., Cahusac, P. M. B., Kesner, R. P., Niki, H., Feigenbaum, J. & Bach, L. 1989 Hippocampal neurons in the monkey with activity related to the place in which a stimulus is shown. *J. Neurosci.* **9**, 1835–1845.
- Rolls, E. T., Robertson, R. G. & Georges-François, P. 1997 Spatial view cells in the primate hippocampus. *Eur. J. Neurosci.* **9**, 1789–1794.
- Samsonovich, A. & McNaughton, B. 1997 Path integration and cognitive mapping in a continuous attractor neural network model. *J. Neurosci.* **17**, 5900–5920.
- Squire, L. R. 1992 Memory and the hippocampus: a synthesis from findings with rats, monkeys and humans. *Psychol. Rev.* **99**, 195–231.
- Stringer, S. M., Rolls, E. T., Trappenberg, T. P. & de Araujo, I. E. T. 2002a Self-organizing continuous attractor networks and path integration: two-dimensional models of place cells. *Network: Comput. Neural Syst.* (In the press.)
- Stringer, S. M., Trappenberg, T. P., Rolls, E. T. & de Araujo, I. E. T. 2002b Self-organizing continuous attractor networks and path integration: one-dimensional models of head direction cells. *Network: Comput. Neural Syst.* (In the press.)
- Taylor, J. G. 1999 Neural 'bubble' dynamics in two dimensions: foundations. *Biol. Cybernet.* **80**, 393–409.
- Treves, A. 1990 Graded-response neurons and information encodings in autoassociative memories. *Phys. Rev. A* **42**, 2418–2430.
- Treves, A. 1991 Dilution and sparse encoding in threshold-linear nets. *J. Phys. A* **23**, 1–9.
- 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.
- Tsodyks, M. 1999 Attractor neural network models of spatial maps in the hippocampus. *Hippocampus* **9**, 481–489.
- Wiener, S. I., Paul, C. A. & Eichenbaum, H. 1989 Spatial and behavioural correlates of hippocampal neuronal activity. *J. Neurosci.* **9**, 2737–2763.
- Wilson, M. A. & McNaughton, B. L. 1993 Dynamics of the hippocampal ensemble code for space. *Science* **261**, 1055–1058.

As this paper exceeds the maximum length normally permitted, the authors have agreed to contribute to production costs.