



## CONTRIBUTED ARTICLE

## Simulation Studies of the CA3 Hippocampal Subfield Modelled as an Attractor Neural Network

EDMUND T. ROLLS,\* ALESSANDRO TREVES,† DAVID FOSTER AND CONRADO PEREZ-VICENTE‡

University of Oxford

(Received 10 November 1995; accepted 21 October 1996)

**Abstract**—Real neuronal networks in the brain utilize networks of neurons with graded not binary firing rates. A theoretical analysis of the operation of autoassociative networks with neurons with graded firing rates has therefore been developed. The present simulation study was performed in order to investigate the operation of such a network with values for the asymmetric diluted neuronal connectivity typical of some brain regions such as the hippocampus, which are outside the range to which the theoretical analysis strictly applies. We report that, in line with theoretical predictions, the amount of information that can be retrieved is relatively independent of the resolution of the stored patterns (binary, ternary, decimal, or fifty-fold). The implication of this is that if the network stores many graded patterns, which it can, then the retrieval quality of each of the patterns becomes low. The implications of this trade-off between the number of patterns stored and the retrieval quality of each pattern when graded firing rates are stored for understanding the operation of networks in the hippocampus are considered. © 1997 Elsevier Science Ltd. All rights reserved.

### 1. INTRODUCTION

The presence of recurrent connectivity amongst neurons in the CA3 subfield region of the hippocampus, along with the physiological evidence for the existence of mechanisms of long-term synaptic modification, has led to the consideration of this part of the brain as an attractor network, or autoassociative memory (Marr, 1971, Rolls, 1987, 1989a, 1989b, 1989c, 1990a, 1990b, 1990c, 1991; Rolls & O'Mara, 1993, Treves & Rolls, 1991, 1992, 1994, Rolls & O'Mara, 1993; Rolls & Treves, 1998). The hypothesis developed in this work is that the CA3 region is part of a system involving the hippocampus and neocortex that is involved in the formation and retrieval of episodic memories. Retrieval of recent episodic memories from the hippocampus may be useful when the neocortex is building more structured, semantic memories (Treves & Rolls, 1994).

Systems of large number of neurons acting together in a dynamically recurrent fashion have been amendable to analysis using techniques based on forming analogies with physical systems of large number of interacting particles. Physicists have long been used to the idea of looking at the long-term evolution of macroscopic variables in a system where details of the activity of individuals are not known. Just such a situation is considered in the mean field theory of attractor neural nets. The asymptotic behaviour of an attractor neural net can be derived analytically for fixed point attractors, and estimates can be made for the storage capacity of the net in terms of pattern and information retrieval (Hopfield, 1982; Amit, 1989).

Since the beginnings of the involvement of statistical physics in neural network theory, advances have been made towards biological relevance. One area of interest has been to consider attractor networks with graded firing rates, rather than binary firing rates, because it is found that neurons in the brain have continuously variable firing rates (e.g. in the visual cortex Rolls and Tovee, 1994; in the primate hippocampus, Rolls and O'Mara, 1993). Hopfield (1984) extended his analysis to fully connected networks with sigmoid activation functions. However, the approach from statistical mechanics deals most naturally with fully connected networks, because then there is a formal energy minimum; but biologically

Acknowledgements: This research was supported by Medical Research Council grant PG8513790, The Human Frontier-Science Program, and the European Community Human Capital and Mobility Programme.

\*Please address correspondence to: Edmund T. Rolls, University of Oxford, Department of Experimental Psychology, South Parks Road, Oxford, OX1 3UD, U.K. Tel. 01865 271348.

† Present address: SISSA, Cognitive Neuroscience, Trieste, Italy.

‡ Present address: Department of Fundamental Physics, University of Barcelona, Spain.

realistic networks typically have diluted connectivity, and in addition sparse representations. Because the CA3 network in the hippocampus has a diluted connectivity (approximately 2%), and has a sparse representation with only a small proportion of the neurons active at any one time, as well as neurons with continuously variable firing rates, we have developed an analysis applicable to these biologically relevant conditions of storage and retrieval in autoassociative neural networks (Treves, 1990, Treves and Rolls, 1991, 1992).<sup>1</sup> This has been possible using for the analysis networks of neurons with threshold linear activation functions. In these analyses, the connectivity was either complete, as above, or had to be assumed to be very diluted.<sup>2</sup> The biologically relevant case, however, is that of a moderately (and asymmetrically) diluted connected network (2–4% connectivity, see Rolls, 1990a, Treves and Rolls, 1992, 1994). To obtain results closer to this case, it is necessary to conduct simulation studies. Simmen et al. (Simmen et al., 1995) have simulated an attractor network with threshold linear neurons and an asymmetric connectivity of 20%. The network was able to retrieve stored patterns with results very close to the capacity predicted analytically. However, in that paper, only binary and ternary patterns, not the more biologically plausible more graded patterns, were considered. The purpose of the investigation described here is to extend the simulation results to moderately connected networks with neurons with graded, not just binary (or ternary), patterns of firing rates.

It is of considerable interest, theoretically as well as in the context of the hippocampus, to extend the results to networks with graded firing rates, for in part of a simulation of the whole hippocampus, Rolls (1995) found with threshold linear neurons that if the firing rates in CA3 were allowed to be continuously graded, then, relative to the binary condition, there was a large deterioration in the retrieval quality of the patterns, as well as some decrease in the number of patterns that could be stored. (The retrieval quality was measured by the correlation of the pattern retrieved using a partial cue to the pattern which was learned). Rolls (1995) suggested that if the function of the hippocampus was to store and retrieve correctly as many memories as possible, as might well be the case for a structure involved in storing episodic memories, then it could be advantageous to use relatively binary encoding of patterns in structures such as CA3, and to forgo graded firing rates, which though potentially capable of storing more information per pattern, do so only when the memory loading (the number of patterns stored in the network) is low. The specific aim of the

present investigation was therefore to analyse attractor networks in which the degree of grading of the encoding was specifically controlled, and the quality of the retrieval, as well as the number of patterns for which some retrieval could be shown, was investigated. Four different kinds of patterns were investigated: patterns with two-fold (binary), three-fold (ternary), ten-fold (decimal), and fifty-fold structure. In addition, also involving an extension to previous work, we focussed on the completion of incomplete patterns by the auto-associative network of (cf Simmen et al., 1995), as completion of past episodes from remembered fragments may be an important property of the hippocampus in memory (Rolls, 1989a, 1989b, 1989c).

## 2. THE NETWORK

The network can be described under four headings which correspond to the four stages in which the simulation of the network operates. The formal specification of the operation of the network is the same as that of the network analysed by Treves (1990), except where indicated. Firstly, the patterns that the net is to be trained on are created from the distributions specified in the theoretical analysis described by Treves (1990). Secondly, the weights are set according to a Hebbian covariance rule. Thirdly, the weight matrix is ablated, that is a proportion of its elements are probabilistically set to zero, to achieve an effective degree of recurrent connectivity. Fourthly, the net undergoes testing with incomplete persistent external cues until the state has settled into retrieval or otherwise.

### 2.1. Pattern Generation

The memory patterns were generated according to distributions which satisfied the general requirement that  $\langle \eta \rangle = \langle \eta^2 \rangle = a$  ( $\eta$  is the firing rate), i.e. that the first and second moments of the distribution be constrained to certain values, setting the sparseness  $a$  of the distribution. The sparseness is defined as:  $a = \langle \eta \rangle^2 / \langle \eta^2 \rangle$ . The distributions were as follows:

Binary:

$$P\hat{\eta} = (1 - a)\delta(\eta) + a\delta(\eta - 1)$$

Ternary<sup>3</sup>:

$$P\hat{\eta} = (1 - \frac{4}{3}a)\delta(\eta) + a\delta(\eta - \frac{1}{2}) + \frac{a}{3}\delta(\eta - \frac{3}{2})$$

10- and 50-fold: The continuous exponential-like distribution

$$P\hat{\eta} = (1 - 2a)\delta(\eta) + \lambda 4a \exp^{-2\eta}$$

<sup>1</sup> In the rat CA3 region, the sparseness parameter  $a$  defined below is approximately 0.02 (Barnes et al., 1990, Rolls and O'Mara, 1993).

<sup>2</sup> Analytical results are also available for intermediate degrees of dilution, but only with the additional constraint of symmetry in the connections (O'Kane and Treves, unpublished).

<sup>3</sup> There are, of course, many possible choices for a distribution  $P_{\eta}$  which consists of three  $\delta$  functions which satisfy the required constraints on the moments set out above. We choose the particular form described next because this has been used in related analytic investigations by Treves (1990).

which satisfies the required constraints on the moments, was evaluated at 10 equispaced values ( $\eta = \eta_0, \eta_0 + 1/3, \dots, \eta_0 + 3$ ) and, respectively, at 50 equispaced values ( $\eta = \eta_0, \eta_0 + 1/15, \dots, \eta_0 + 49/15$ ) with  $\eta_0 = 0$  and  $\lambda = 1$  adjusted to make sure the discretized distributions also satisfy the constraints.  $\delta(x)$  is Dirac's delta function, and it allocates to the quiescent state the probability remaining after the probabilities of the different firing levels (all proportional to  $a$ ) have been determined. Note that for 10- and 50-fold distributions the weight of the delta function is augmented by the contribution, at  $\eta_0 = 0$ , of the first discrete firing level.

The simulation allowed the sparseness  $a$  of the coding of the patterns generated to be specified. The issue of sparseness is important, for it is an important parameter setting the number of patterns that can be stored in the network (Treves & Rolls, 1991). The sparseness of the retrieved patterns was measured, to ensure that the network was operating in such a way that the sparseness of the retrieved patterns was close to that of the stored patterns. The gain factor  $g$  in the threshold linear activation function was a parameter which altered the retrieved sparseness.

## 2.2. Learning

The learning mechanism is a form of Hebbian covariance synaptic modification, a one-step application of a simple rule which takes account of simple pairwise covariance relationships within each pattern. The exact rule is as follows. Note that the form of the covariance rule is commutative with respect to units  $i$  and  $j$ , therefore forcing a fully connected net with such a rule to have symmetric weights.

$$J_{ij}^c = \frac{1}{Na^2} \sum_{\mu=1}^p (\eta_i^\mu - a)(\eta_j^\mu - a)$$

where  $J_{ij}^c$  is the weight between units  $i$  and  $j$ .  $\eta_i^\mu$  represents the firing rate of unit  $i$  within pattern  $\mu$ . This is a simple covariance rule, and  $a$  represents, along with the sparseness, the mean activation of the net. As noted above, the sparseness of the representation is a property which is a characteristic of the neuronal firing in each brain area.

## 2.3. Connectivity

One of the central aims of this study was to investigate net performance at moderate levels of diluted connectivity, and thus the weight matrix as prescribed by the above learning rule was modified. The connectivity required was interpreted by the program as a probability that each element of the weight matrix retained its prescribed value; thus with a complementary probability, each element was 'ablated', i.e. set to zero.

This was achieved in one of two ways, depending upon whether this dilution of connectivity was to be applied symmetrically or asymmetrically. In the asymmetric case, each element was simply considered individually.

In the symmetric case, only one half of the weight matrix, split by the leading diagonal, was considered. As each element was probabilistically considered, so its mirror element across the diagonal was also treated. In this report, results are given for asymmetric dilution only, because this is the biologically plausible case. The performance of networks with diluted symmetric connectivity was rather better than the performance with the asymmetric connectivity described here. This is because the symmetric case is more likely to produce a stable state.

The total number of units is  $N$ , each unit receives on average  $C$  inputs, and  $p = \alpha C$  is the number of patterns stored in the net. The critical loading of the net, when it fails to operate as a memory, is denoted as  $\alpha_c$ .

## 2.4. Testing Recall by the Net

During recall, the activity of each neuron in the network was asynchronously updated according to a rule which, by analogy with the theoretical analysis, considered a local field  $h_i$  at each unit  $i$  consisting of an external field, internal field, and threshold term, as follows

$$h_i = \sum_{(j \neq i)} J_{ij}^c V_j + b \left( \sum_j \frac{V_j}{N} \right) + \sum_{\mu} s^{\mu} \frac{\eta_i^{\mu}}{a}$$

$V_j$  represents the output of neuron  $j$ , and  $s^{\mu}$  represents the relative strength of pattern  $\mu$ , see below.

The external field (the last term in the above equation) is equivalent to the clamping, persistent external cue, which is believed to be provided by the direct perforant path, afferents into CA3 from entorhinal cortex (Treves & Rolls, 1992). The ratio between the average number of perforant path synapses per CA3 cell and that of the recurrent collaterals is in this model allowed to determine their relative influence on the firing of CA3 cells. Anatomical evidence available from the rat suggests that the ratio of the external input (the retrieval cue) to the internal recall provided by the recurrent collaterals should be in the order of 0.25 (see Treves and Rolls, 1992), and we set  $s^{\mu}$  to produce this ratio (for example when the retrieval cue had a correlation of 0.5 with the originally learned pattern).

The internal field (the first term in the above equation) is equivalent to the recurrent activation provided by the recurrent collaterals in CA3. This is implemented through a standard autoassociation update rule involving weighted inputs from each of the other units. As explained above, this is qualified by the connectivity enforced through zero weights.

The threshold term  $b$  is equivalent to the effect of the inhibitory mechanisms at work in the CA3 system. The inhibitory mechanism included here is a function of the mean activity of the entire system, and it is subtractive, of the form

$$\begin{aligned} b &= f(\bar{V}) \\ &= \kappa(\bar{a} - \bar{V})^3 \end{aligned}$$




FIGURE 1. Recall performance is only weakly dependent on the value of the gain parameter,  $g$ , over a certain range. Twenty trials were performed at five correlation levels, for each gain value for binary patterns. The recall performance measure used was an estimate of the area above the line of zero completion  $y = x$  shown in Figure 3 (see text).

The value  $\kappa$  was set to 10,000. (At lower values,  $\alpha_c$  is artificially low, and at very high values completion was limited.) We ensured that the retrieval sparseness was close to 0.1 (which was the sparseness of the stored patterns), by adjusting  $\tilde{a}$  to a lower value than  $a$ . (If we had allowed the sparseness of the recalled patterns to become different from 0.1, then this would have altered the recall properties of the system being studied. If, as in the original theory (Treves, 1990), the retrieval sparseness is allowed to decrease, then  $\alpha_c$  becomes higher. We focus here on the condition in which the sparseness is the same during storage and retrieval.) A gain parameter which itself depends on the activity in the net is tantamount to a more realistic inhibitory mechanism, e.g. divisive inhibition, but as threshold setting is just a generic operation, it was not investigated in detail here.

The activation function of the neurons is a threshold

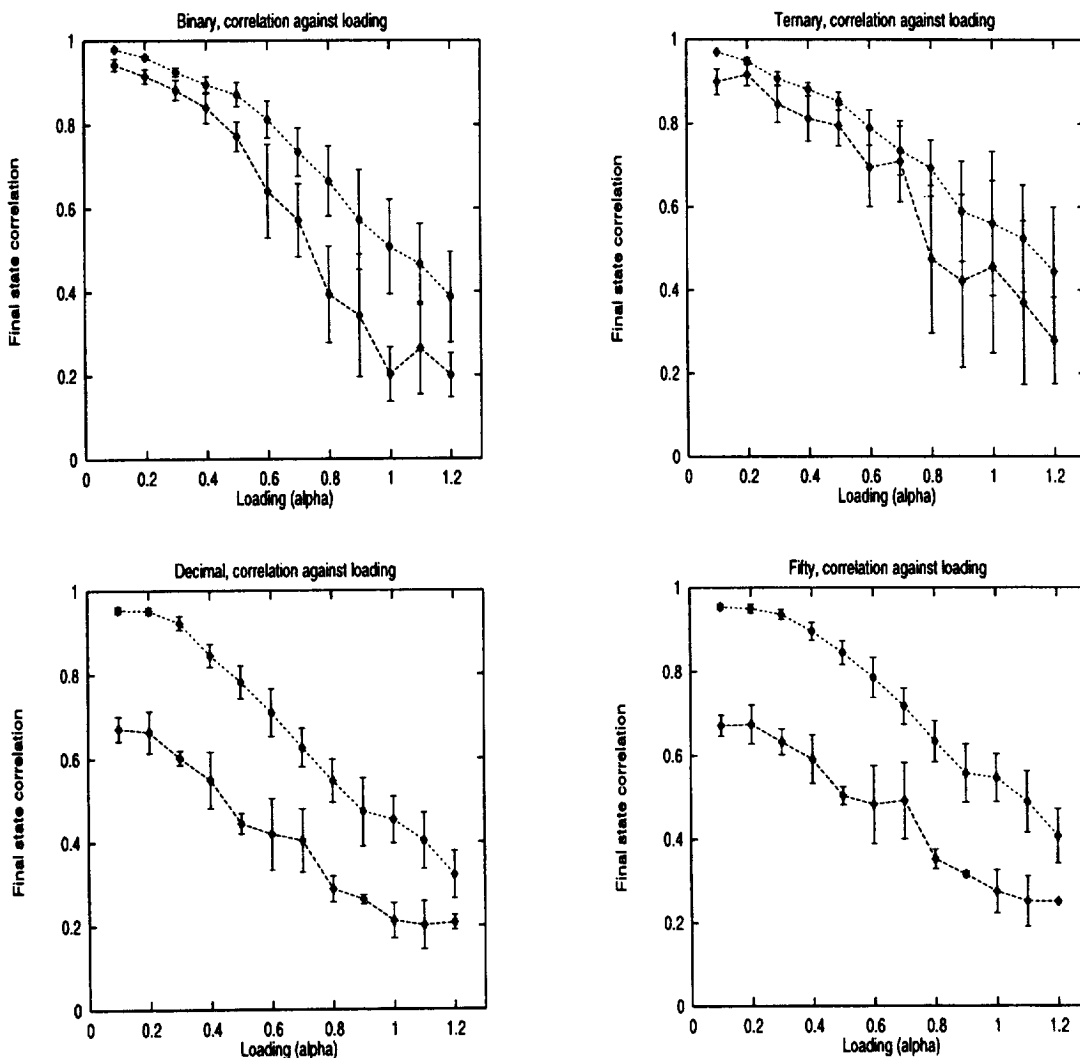
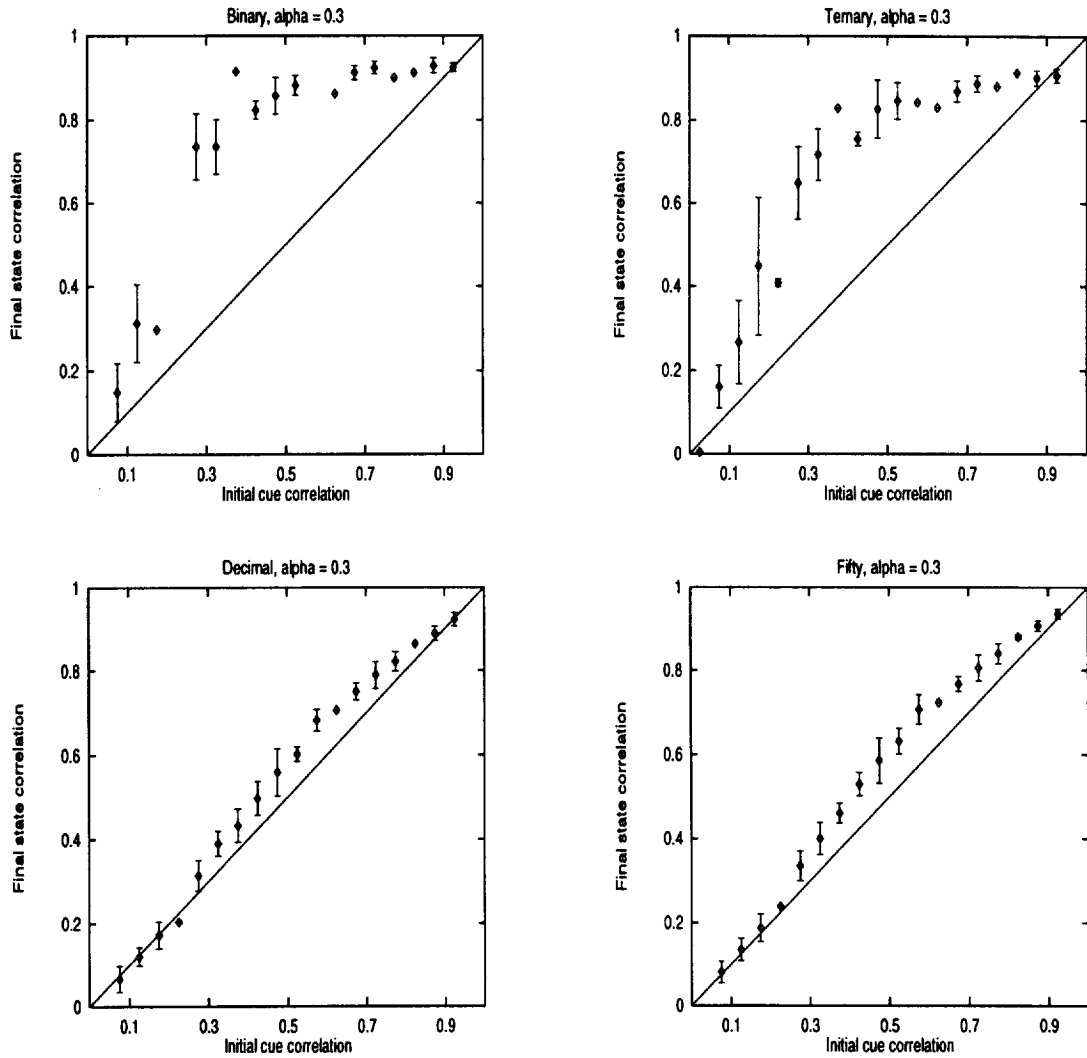


FIGURE 2. This figure shows the correlation of the final state of the network after recall with the stored pattern, as a function of loading, for the four levels of pattern resolution studied. The two lines in each graph correspond to two cue levels: the heavy dashed line is for a cue correlated 0.5–0.55 with the original stored pattern, and the light dashed line is for a cue correlation of 0.9–0.95. The loading,  $\alpha$ , is  $P/C$ . Note that here and in the following graphs, the error bars represent the standard deviations.



**FIGURE 3.** This figure shows pattern completion across the four pattern types, at a loading of  $\alpha = 0.3$ . The abscissa shows the correlation of the retrieval cue with the stored pattern. The ordinate shows the correlation of the retrieved pattern with the stored pattern. Each point represents the average of 5 different patterns retrieved.

linear function of the local field  $h_i$  with a gain factor  $g$ , as described by Treves (1990).

$$V = \begin{cases} g(h - T_{\text{thr}}), & h > T_{\text{thr}} \\ 0 & , \quad h < T_{\text{thr}} \end{cases}$$

The recall of the net was measured by the correlation of the retrieved pattern with that stored, when incomplete retrieval cues were used. The performance of the network was also measured by the information retrievable from the network in bits per synapse about the set of stored patterns, as follows

$$I = \frac{\alpha}{\log 2} \sum_{\kappa=0}^m \sum_{l=0}^n c_{\kappa}^l \log \frac{c_{\kappa}^l}{c_{\kappa} c^l} \quad (1)$$

where we refer to Treves (1990) for further details. Briefly, for each element of the retrieved network state and the corresponding stored pattern, the firing was first discretized into bins, and then the expression above was evaluated. In the above,  $c_{\kappa}$  is the probability that the

pattern element is in the  $k$ th bin of  $m$  bins,  $c^l$  is the probability that the retrieved element is in the  $l$ th bin of  $n$  bins, and  $c_{\kappa}^l$  is the probability that the retrieved element is in the  $l$ th bin, and the pattern element is in the  $k$ th bin. In our implementation of this calculation, due to practical limitations, we binned both the patterns and network states into 15 bins. Note that the factor  $\alpha$  means that the result is in bits per synapse, which is proportional to the total information stored in, and retrievable from, the whole network.

## 2.5. Parameters

The network functioned with a set of parameters chosen to be biologically relevant. Where the parameters are not in correspondence with measurable quantities, they were optimized to the values required for the theory to apply (Treves, 1990). This subsection details some of these parameters, and the reason for their choice.

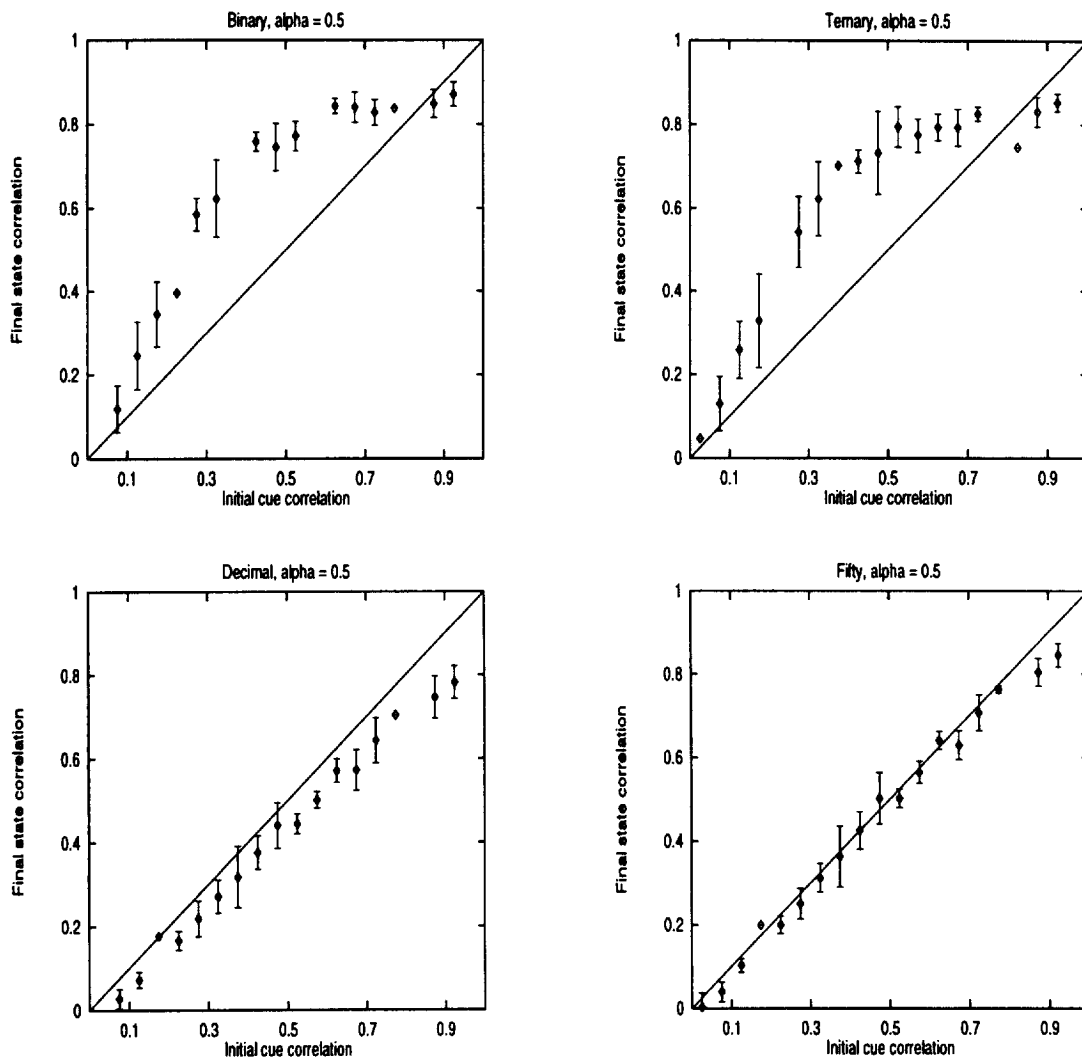


FIGURE 4. This figure shows pattern completion across the four pattern types, at a loading of  $\alpha = 0.5$ .

The total number of units in the net was set to 2000 for the results reported. This represented a compromise between size, where too small a net might lead to small-size effects, and available storage space on the computer. A limited set of further results showed similar performance for nets with 4000 neurons, and with 1000 neurons the mean results were comparable, but this decrease in size increased the noisiness of individual runs.

The connectivity was kept at 0.2, that is, 400 connections per neuron. The sparseness of coding parameter of the stored patterns  $a$  was kept at 0.1.

The gain parameter, equivalent to the gradient of the linear threshold function producing the output of each unit from its incoming local field, was tested at a number of different values, and found to be optimal in the region of  $g = 0.5$  for binary neurons (see Figure 1). Repeating the search for other pattern types (3-, 10- and 50-fold), the parameter was set close to its optimal value for each type. The actual values of  $g$  used were 0.5 for binary patterns, 0.45 for ternary, and 0.28 for decimal and fifty-fold patterns.

The loading  $\alpha$  was expressed as the ratio  $P/C$ , where  $P$  is the number of patterns stored, and  $C$  is the number of connections per neuron, set to 400. The loading was varied between 0.1 and 1.2 to investigate the maximum value of the storage capacity  $\alpha_c$ , in terms of patterns, as well as to investigate the effect of over-loading. Thus the number of patterns was varied from 40 to 480. The net was allowed to iterate for a maximum of 30 epochs.

### 3. RESULTS

Figure 2 shows the (Pearson product-moment) correlation of the final state of the net with the stored pattern that is being retrieved, as a function of loading. The storage capacity, defined as the critical value  $\alpha_c = \frac{P_{MAX}}{C}$  is approximately indicated in Figure 2 by the point where the correlation of the retrieved pattern with the original is not better than that of the retrieval cue with the original. It is clear that retrieval begins to fail at similar loadings between 0.4 and 0.7 for all the pattern resolutions.

Figure 3 shows the correlation between the final state

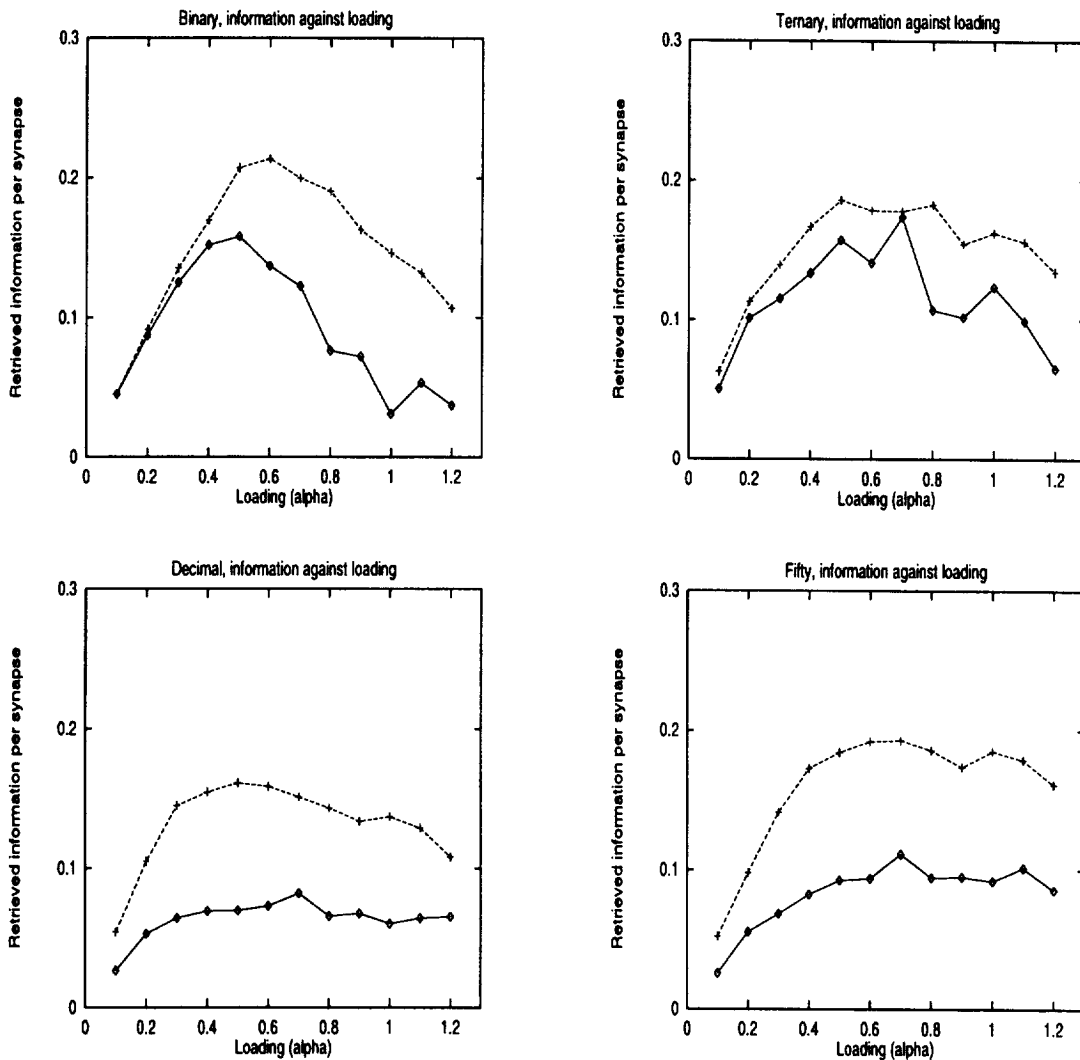


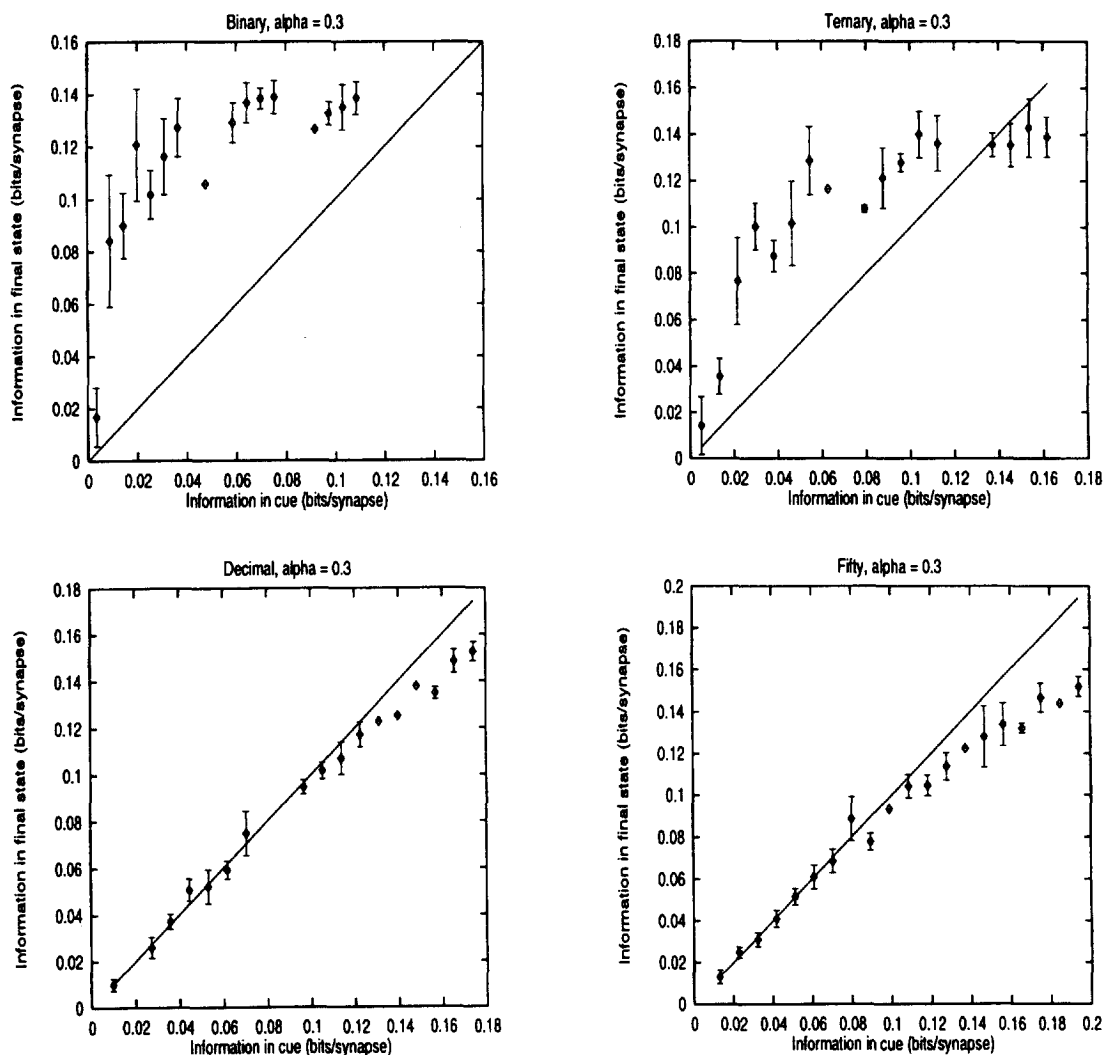
FIGURE 5. This figure shows the information retrieved in bits per synapse, as a function of loading, for the four types of patterns. The two lines in each graph correspond to two cue levels; the heavy dashed line is for a cue correlation of 0.5–0.55 with the stored patterns; the light dashed line is for a cue correlation of 0.9–0.95. The loading,  $\alpha$ , is  $P/C$ .

of the network during retrieval and the stored pattern as a function of the correlation between the initial retrieval cue and the stored pattern of which it is a part, at a given loading of  $\alpha = 0.3$ . For each plot, the line  $y = x$  is shown. This represents the line of zero retrieval, where the cue is effectively reproduced without any pattern completion. If a point is above this line, completion has occurred, and the cue presented can be said to have been within the basin of attraction of the stored pattern. Where a point is below this line, the final state has deteriorated even with respect to the cue. Thus these graphs show pattern completion, and are informative with respect to both the basins of attraction, and the issue of retrieval quality. Retrieval quality, which may be thought of as the extent to which partial cues undergo completion, producing a retrieved state that has a high correlation with the stored pattern, is shown to become worse as the pattern resolution increases from the binary and ternary cases to the decimal and 50-fold cases. One

can see that the binary and ternary cases allow for good completion from a quite incomplete cue, thus exhibiting large basins of attraction. In contrast, the retrieval quality for the decimal and 50-fold patterns becomes poor.

Figure 4 shows the same type of graphs as Figure 3, but for a loading of  $\alpha = 0.5$ , closer to the theoretical capacity. The behaviour is consistent with the previous figure, with in this case even poorer quality of retrieval with the 10-fold and fifty-fold patterns. Indeed, with 10-fold and 50-fold pattern resolutions, the retrieved state was just worse than the retrieval cue itself.

Figures 5, 6 and 7 show the performance of the network measured by the average information stored in the network by each synapse for different pattern resolutions. Figure 5 is the information analysis that corresponds to Figure 2. In all cases, the information per synapse increases as the number of patterns stored increases, until some information capacity limit is reached, beyond which no further information can be retrieved as the



**FIGURE 6.** This figure shows information in the final state of the network as a function of the information in the retrieval cue, in bits per synapse. The loading,  $\alpha$ , was 0.3.

number of patterns is increased. An important point here is that comparable information is stored for all pattern types, as predicted (Treves, 1990, Treves & Rolls, 1991). Beyond the point where the information per synapse reaches its maximum, in some cases, as expected with overloading, the retrieved information actually deteriorates. The retrieved information is higher with the cue correlated 0.9 with the stored patterns as compared to the cues correlated 0.5, partly because the external cue helps to stabilize the firing of the network. To show whether the information in the retrieved pattern (the final state of the network, which is what is plotted in Figure 5) is greater than that in the retrieval cue itself, we show this relation in Figure 6 (for  $\alpha = 0.3$ ) and Figure 7 (for  $\alpha = 0.5$ ). Figure 6 (which corresponds to Figure 3) shows that useful retrieval is achieved by the network for binary and ternary patterns, but that even at this quite low loading of 0.3, the information in the final state of the network is

not better than that in the retrieval cue for 10- and 50-fold patterns. This trend is even greater with a loading  $\alpha$  of 0.5, as shown in Figure 7.

Taken together, the results summarized in Figures 2–4 show the following.  $\alpha_c$  has been reached by around 0.7, a region which agrees with theoretical estimates (Treves & Rolls, 1991). ( $\alpha_c$  is taken to be the point in these figures at which the final state of the network during retrieval is not better than the retrieval cue itself.) However, at each loading, the decimal and 50-fold cases are poor in terms of retrieval quality, with only little completion found. On the other hand, as shown in Figure 5, the actual amount of information retrieved is similar for the different pattern resolutions. Considering that higher resolution implies more information in the stored pattern, there appears to be a trade-off, set by the total information that can be retrieved, between retrieval quality and the information load on each pattern.



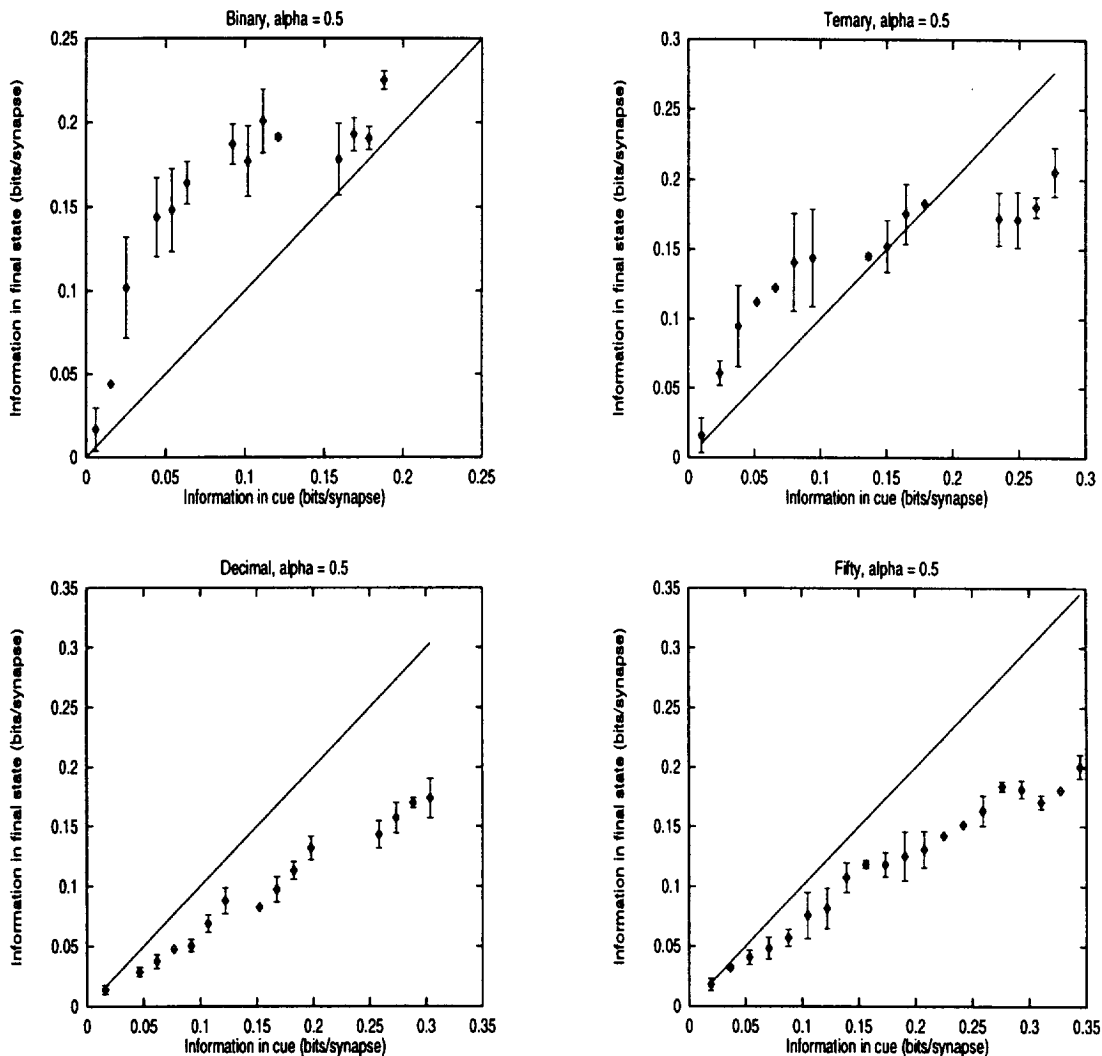


FIGURE 7. This figure shows information in the final state of the network as a function of the information in the retrieval cue, in bits per synapse. The loading,  $\alpha$ , was 0.5.

#### 4. DISCUSSION

The results described here on the storage capacity, that is the number of patterns that can be stored for a given number of connections per neuron, are in line with theoretical analyses which show that  $\alpha_c$  is largely comparable for different pattern types (Treves, 1990, Treves & Rolls, 1991, Fig. 5a). In more detail, however, the analyses predict higher capacities for more graded patterns, provided the sparseness of retrieved patterns is not constrained. The capacity is in fact maximised when, with graded patterns, the intermediate firing levels are thresholded out, producing sparser patterns at retrieval than the original ones. In the present simulations, instead, the sparseness at retrieval is constrained by adjusting the thresholds to be similar to that of the original patterns, and as a result the number of patterns that can be retrieved is not larger for more graded patterns, as shown in Figure 2. Moreover, the results shown in Figure 2 are directly relevant to the biologically

plausible case of moderate connectivity, rather than the full or very diluted connectivity for which the theoretical analyses were originally made (see Treves & Rolls, 1991).

The results described here on the decrease in the quality of the retrieval which occurs with more and more graded patterns, shown in Figures 3 and 4, provide a useful quantitative description of what was suggested to be the case in the hippocampal simulations described by Rolls (Rolls, 1995). The most straightforward point that can be made to account for this is the following. Theoretical analyses, and the results now provided in Figure 5, show that the total amount of information that can be stored and retrieved from such a network (or the amount of information stored per synapse) is relatively independent of the resolution of the stored patterns (Treves & Rolls, 1991, Fig. 5; Treves, 1990, Fig. 2). Given that more information is loaded onto more graded patterns, this means that the quality of the

retrieval will decrease with more and more graded patterns. (Alternatively, the network parameters can be set to produce somewhat better retrieval quality, but then  $\alpha_c$  decreases.) Another way in which this effect can be understood is in terms of the way in which the threshold linear neuronal network operates. The optimal capacity  $\alpha_c$  is achieved for non-binary patterns when, as mentioned above, the threshold is set between the intermediate and the strongest activations (unless the sparseness of the retrieved patterns is controlled). This however means not only that the retrieved patterns end up sparser than those stored, leading to an apparent increase in  $\alpha_c$ , but also that the quality of the recall deteriorates. In the simulations performed by Rolls (1995), the sparseness of the retrieved patterns was carefully controlled, and the effect on the quality of retrieval of graded pattern encodings was very apparent. In the present simulations, we checked that the sparseness of the retrieved patterns was approximately that of the stored patterns, and this was the case (indicating that the gain parameter  $g$  and the inhibitory parameters were set as we had intended).

The exact point at which the network performance became poor as loading was increased was found during the simulations we performed to depend on a number of factors. One was the ratio of the external activation of each neuron. As this was increased beyond about 0.25, then the final retrieval state of course reflected the retrieval cue more, and this could lead to an apparent higher  $\alpha_c$  because the network was being stabilized and helped by the strong external cue. Conversely, if the external cue was much less than 0.25, then the external retrieval cue was insufficient to lead to good retrieval with this asymmetric diluted connectivity.

The trade-off described here between higher information load and poorer retrieval quality, which occurs with more graded patterns, must be considered when defining whether binary or graded encoding is more appropriate in autoassociation networks. In other words, the potential for loading more information on to each graded pattern remains unutilizable by these simple networks. Since neither the number of patterns nor the total amount of information that can be retrieved increase significantly with the use of more graded patterns, one can say that the higher resolution available is just wasted by the network. In the context of storing as much as possible about as many recent episodic memories as possible, in a structure such as the hippocampal CA3 network, there may be even some advantage to using relatively coarsely graded (almost binary) encoding of patterns. This will maximize the number of different memories that can be stored and retrieved correctly. Moreover, the implications of the results described here are not limited by the fact that a correlation measure is more sensitive for graded than for binary patterns, for the operation of the next stage of processing in the hippocampus, the CA1 cells, was found to be limited by the poor retrieval quality of CA3 in the simulation of the

hippocampus described by Rolls (1995). The noisiness or poor predictability of the retrieval described here may also limit the value of using very graded encodings of firing rate in other autoassociative memory systems required to operate at high loading.

One aspect of the patterns of firing of hippocampal neurons may contribute to the interpretation of their firing rates as being relatively binary, without a great range of different firing rates needing to be stored. It is that in addition to the representation being sparse, so that the probability that CA3 neurons fire may be low (e.g. 0.02, see Barnes et al., 1990; Rolls and O'Mara, 1993), the firing rates of hippocampal cells with spatial (place in rats, or view in monkeys, see Rolls and O'Mara, 1993) responses tends to be low, with rates in the range 0–10 spikes/s being typical in the rat, and 0–20 spikes/s in the monkey. Given that the integration time of the biophysical neuronal processes involved in synaptic modification may be in the order of 50 ms, this would mean that there is insufficient time in such short periods for accurate measures of firing rate to be reflected in what is stored, given that only one or two spikes are likely to occur in such a short period. In contrast, the rates with which neurons fire are very much faster in cerebral neocortical areas involved in encoding visual stimuli such as faces (Rolls & Tovee, 1994), and this may be to allow very rapid information processing to occur (see Rolls, 1994; Rolls & Tovee, 1994). This suggests that the constraint in what is thought to be a memory storage system such as the hippocampus may be towards *storing* and *retrieving* large numbers of different memories, and large amounts of information, for which sparse patterns, and low firing rates, may be more appropriate, given the findings described here. If the constraint were more that of *representing* large numbers of patterns, or large amounts of information, the advantages would be with less sparseness and more graded encodings, which may be utilized in some neocortical areas (see Rolls & Treves, 1998).

In conclusion, the most important results described here are that in autoassociative networks with neurons with graded firing rates the retrieval quality decreases markedly when the pattern grading is increased from ternary to 10-fold and 50-fold. The effect was not found by Simmen et al. (1995), who did not test for graded patterns greater than ternary. The effect was observed initially in the simulations of the hippocampus described by Rolls (1995). The result does have implications for how autoassociative networks in the brain may operate.

## REFERENCES

- Amit, D.J. (1989). *Modelling brain function*. New York: Cambridge University Press.
- Barnes, C.A., McNaughton, B.L., Mizumori, S.J., & Lim, L.H. (1990). Comparison of spatial and temporal characteristics of neuronal activity in sequential stages of hippocampal processing. *Progress in Brain Research*, 83, 287–300.

- Hopfield, J.J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Science USA*, 79, 2554–2558.
- Hopfield, J.J. (1984). Neurons with graded response have collective properties like those of two-state neurons. *Proceedings of the National Academy of Science USA*, 81, 3088–3092.
- Marr, D. (1971). Simple memory: a theory for archicortex. *Philosophical Transactions of the Royal Society*, B262, 24–81.
- Rolls, E.T. (1987). Information representation, processing and storage in the brain: analysis at the single neuron level. In J.-P. Changeux and 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 neocortex in memory. In J.H. Byrne and W.O. Berry (Eds.), *Neural models of plasticity: Experimental and theoretical approaches* (Ch. 13, pp. 240–265). San Diego: Academic Press.
- Rolls, E.T. (1989b). The representation and storage of information in neuronal networks in the primate cerebral cortex and hippocampus. In R. Durbin, C. Miall and G. Mitchison (Eds.), *The computing neuron* (Ch. 8, pp. 125–159). Wokingham, England: Addison-Wesley.
- Rolls, E.T. (1989c). 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. (1990a). 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. (1990b). Functions of the primate hippocampus in spatial processing and memory. In D.S. Olton and R.P. Kesner (Eds.), *Neurobiology of comparative cognition* (Ch. 12, pp. 339–362). Hillsdale, N.J.: Lawrence Erlbaum.
- Rolls, E.T. (1990c). A theory of emotion, and its application to understanding the neural basis of emotion. *Cognition and Emotion*, 4, 161–190.
- Rolls, E.T. (1991). Functions of the primate hippocampus in spatial and non-spatial memory. *Hippocampus*, 1, 258–261.
- Rolls, E.T. (1994). Brain mechanisms for invariant visual recognition and learning. *Behavioural Processes*, 33, 113–138.
- Rolls, E.T. (1995). A model of the operation of the hippocampus and entorhinal cortex in memory. *International Journal of Neural Systems*, 6, Supplement, 51–70.
- Rolls, E.T. & O'Mara, S. (1993). Neurophysiological and theoretical analysis of how the hippocampus functions in memory. In T. Ono, L.R. Squire, M. Raichle, D. Perrett, M. Fukuda (Eds.), *Brain mechanisms of perception: From neuron to behavior* (Ch. 17, pp. 276–300).
- Rolls, E.T., & Tovee, M.J. (1994). Processing speed in the cerebral cortex and the neurophysiology of visual masking. *Proceedings of the Royal Society*, B, 257, 9–15.
- Rolls, E.T. & Treves, A. (1998). *Neural networks and brain function*. Oxford: Oxford University Press.
- Simmen, M.W., Treves, A., & Rolls, E.T. (1995). Pattern retrieval in threshold-linear associative nets. *Network*, 7, 109–122.
- Treves, A. (1990). Graded-response neurons and information encodings in autoassociative memories. *Physical Review A*, 42, 2418–2430.
- 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.