Information About Spatial View in an Ensemble of Primate Hippocampal Cells

EDMUND T. ROLLS, ALESSANDRO TREVES, ROBERT G. ROBERTSON, PIERRE GEORGES-FRANÇOIS, AND STEFANO PANZERI
Department of Experimental Psychology, University of Oxford, Oxford OX1 3UD, United Kingdom

Rolls, Edmund T., Alessandro Treves, Robert G. Robertson, Pierre Georges-François, and Stefano Panzeri. Information about spatial view in an ensemble of primate hippocampal cells. *J. Neurophysiol.* 79: 1797–1813, 1998. Hippocampal function was analyzed by making recordings from hippocampal neurons in monkeys actively walking in the laboratory. ''Spatial view'' cells, which respond when the monkey looks at a part of the environment, were analyzed. To assess quantitatively the information about the spatial environment represented by these cells, we applied information theoretic techniques to their responses. The average information provided by these cells about which location the monkey was looking at was 0.32 bits, and the mean across cells of the maximum information conveyed about which location was being looked at was 1.19 bits, measured in a period of 0.5 s. There were 16 locations for this analysis, each being one-quarter of one of the walls of the room. It also was shown that the mean spontaneous rate of firing of the neurons was 0.1 spikes/s, that the mean firing rate in the center of the spatial field of the neurons was 13.2 spikes/s, and that the mean sparseness of the representation measured in a 25-ms period was 0.04 and in a 500-ms time period was 0.19. (The sparseness is approximately equivalent to the proportion of the 25- or 500-ms periods in which the neurons showed one or more spikes.) Next it was shown that the mean size of the view fields of the neurons was 0.9 of a wall. In an approach to the issue of how an ensemble of neurons might together provide more precise information about spatial location than a single neuron, it was shown that in general the neurons had different centers for their view fields. It then was shown that the information from an ensemble of these cells about where in space is being looked at increases approximately linearly with the number of cells in the ensemble. This indicates that the number of places that can be represented increases approximately exponentially with the number of cells in the population. It is concluded that there is an accurate representation of space ''out there'' in the primate hippocampus. This representation of space out there would be an appropriate part of a primate memory system involved in memories of where in an environment an object was seen, and more generally in the memory of particular events or episodes, for which a spatial component normally provides part of the context.

INTRODUCTION

Damage to the temporal lobe that includes the hippocampal formation or to one of its main connection pathways, the fornix, produces amnesia (see Gaffan 1994; Scoville and Milner 1957; Squire and Knowlton 1994). One of the memory deficits in amnesic humans is a major impairment in remembering not just what objects have been seen recently but also where they have been seen (Smith and Milner 1981). This type of memory is the type of memory used for example in remembering where one's keys have been left. In experimental studies in monkeys to define the crucial structures to which damage produces memory impairments, it has been shown that hippocampal or fornix damage produces deficits in learning about where objects have been seen, in object-place memory tasks (Angeli et al. 1993; Gaffan 1994; Parkinson et al. 1988).

To analyze how the hippocampus operates to help implement this type of memory, Rolls and colleagues have recorded from single neurons in the hippocampus while monkeys perform object-place memory tasks in which they must remember where on a video monitor a picture has been shown. They found that ~10% of hippocampal neurons responded when images were shown in some positions on the screen (Rolls et al. 1989). Moreover, they showed that the representation was in allocentric (world) rather than egocentric (related to the body) coordinates, in that the spatial fields of these neurons remained in the same position on the video monitor even when the whole monitor was moved relative to the monkey's body axis (Feigenbaum and Rolls 1991).

A theory that the hippocampus is a computer for spatial navigation, computing bearings and distances to the next place, has been built on the basis of the properties of rat hippocampal place cells (Burgess et al. 1994). In contrast to the findings in primates, the spatial representation provided by hippocampal neurons in rats appears to be related to the place where the rat currently is located. That is, individual hippocampal neurons in rats respond when the rat is in one place in a test environment (O'Keefe and Speakman 1987). Because it is not clear whether the primate hippocampus should be considered a spatial computer, with perhaps place cells like those of rats (Ono et al. 1993), or is instead a structure involved in storing memories, including those with a spatial component such as where an object has been seen, we recorded from single hippocampal neurons while monkeys actively locomoted in a rich spatial environment. We set up the recording situation to allow perambulation by the monkey, because it is only during active locomotion that the place fields of rat hippocampal neurons become evident (Foster et al. 1989). We used a rich testing environment, as compared with a cue-controlled environment with only a few spatial cues, to maximize the possibility that many cells with spatial response properties would be found. In one previous study, without active locomotion and with a cue-controlled environment, we found a small number of hippocampal cells that responded to spatial views of the environment, but no cells with response fields that defined the place where the monkey was located (Rolls and O'Mara 1995). However, that study was not with active locomotion nor with a spatially rich environment. In a previous study with active

locomotion in the same rich spatial environment used here, the open laboratory, we found spatial view cells that responded when the monkey looked at one part of the environment but not when it looked at another (Rolls et al. 1997a). These responses occurred relatively independently of where the monkey was in the testing environment, provided that it was looking toward a particular part of the environment. Eye position recordings with the monkey stationary confirmed that these neurons fired when the monkey looked at a particular part of the spatial environment and not in relation to where it was (Rolls et al. 1997a). For these reasons, the cells were named ''spatial view'' and not ''place'' cells. It also has been shown that these neurons respond in relation to where the monkey is looking in space and not to head direction per se or to eye gaze angle per se (Georges-François et al. 1998).

The new investigation described here is designed to analyze the spatial properties of these cells further by comparing for a population of these cells where in space the view field is centered, measuring the width of each view field, and quantifying how much information is obtained about spatial view from the responses of these cells. The information theoretic approach used for measuring the information available in the responses of single hippocampal neurons was based on that used for single neurons in the inferior temporal visual cortex (Rolls et al. 1997c) and the orbitofrontal olfactory cortex (Rolls et al. 1996). Of particular interest also was how the information increases as more cells are added to the ensemble. An attractive property of distributed encoding is that the information available from an ensemble can scale linearly with the size of the ensemble. If this were true of the representation of spatial view by primate hippocampal neurons, this would mean that the firing of even relatively few neurons in a sparse representation would provide considerable information about spatial view. As it has been suggested that a sparse code in the hippocampus might enable it to store many different memories (each one for example about where in space a different object was located) (Rolls 1989; Treves and Rolls 1991, 1994), it was of great interest to try to estimate how much information might be available (potentially for storage) when a sparse ensemble of hippocampal cells was active (that is, an ensemble with few neurons active). A first step in assessing this was to analyze whether each cell was tuned to a different part of the environment. Only if the cells coded for different parts of the environment, would the information rise rapidly (linearly) with the number of cells in the ensemble. If the neurons were just replicates of each other, so that distributing the information only served to suppress noise through massive redundancy, the signal-to-noise ratio would tend to rise in proportion to the square root of the number of cells in the ensemble, and the information would tend to rise only logarithmically with that number. Then we applied the information theoretic approach described by Rolls, Treves and Tovee (1997b) to estimating the information available from the ensemble. In this paper, we introduce an additional procedure that can be used when the ensemble of cells is small.

METHODS

Neurophysiological recordings

Single neurons were recorded with glass-insulated tungsten microelectrodes with methods that have been described previously (Rolls et al. 1989) in two rhesus monkeys. During the recordings, each monkey (*Macaca mulatta*) was free to roam a 2.7 × 2.7 m area in an open 4 ×4 m laboratory in a chair on four wheels, which allowed it to face forward. Small pieces of food were placed in three of the four cups (c1–c4) shown in Fig. 1 from time to time during the experiment and also were scattered sometimes on the floor to ensure that the monkey explored the environment fully. Three of the cups c1–c4 were provided with food to encourage the monkey to learn about the places of food in the spatial environment. Eye position was measured to an accuracy of 1° with the search coil technique, with the field coils attached to the walker to which the head also was attached. The angle visible to the monkey by eye movements was ~35° left and right and 35° up and down, with respect to head direction. The head direction and position in the room were measured using a video tracking device (Datawave) with the camera in the ceiling tracking two light-emitting diodes placed in line 25 cm apart above the head on the top of the chair. We wrote software to provide the position of the monkey's head in the room, the head direction, and the eye position (i.e., the horizontal and vertical angles of the eye in the orbit) every 67 ms, and from these, the gaze direction (i.e., the direction of the eyes in world coordinates) and thus the position on the wall of the room at which the monkey was looking were determined. Each action potential was recorded to an accuracy of 1 ms. The Datawave spike cutting software was used to ensure that the spikes of well-isolated neurons were analyzed. Software was written to measure the firing rate of the neuron whenever the monkey was looking at a position in space. The algorithm took a fixed length record (usually 500-ms long) whenever the eyes were steadily fixating a position in the room during the recording and calculated the firing rate together with where the monkey was looking during that record. (The computer determined that the eyes were fixating a location by taking into account both the eye gaze angle and the head direction and position.) If there was no eye movement, the next record was taken immediately after the preceding one. The algorithm allowed a delay in neuronal data collection after a steady eye position. (If the neuron started to respond 100 ms after the monkey moved his eyes to an effective location in space, this lag could be set to 100 ms. In practice, the lag was set for all neurons to the small value of 50 ms.) From all such records containing a firing rate and where the monkey was looking during the record, it was possible to plot diagrams of the firing rate of the cell when different locations were being viewed. [The records were binned typically into 64 bins horizontally (16 for each wall) and 16 vertically, and smoothed.] It was possible to measure the neuronal responses either while the monkey was walking round the room or when it was stationary. In the experiments described here, it was sometimes advantageous for the monkey to be stationary facing in a particular direction for a number of seconds. This was facilitated by slipping a panel into the bottom of the walker for the monkey to stand on instead of the floor. The monkey of course still could actively explore his environment by making eye movements in this condition. As described previously (Rolls et al. 1997a), the neuronal responses when the monkey looked at a particular position in space while it was walking were very similar to those while it was still.

The neurons were selected to be similar to those described previously as having spatial view-related responses, that is they responded when the monkey looked at a given position in space, relatively independently of where the monkey was (Rolls et al. 1997a). The responses of each neuron were recorded for several minutes during which the monkey looked at all the walls of the environment and moved round the environment. From the hundreds of 0.5-s samples of firing rate (each with the eyes still), graphs were made to show the firing rate as a function of where the monkey was looking on the four walls of the room. From these graphs, the width of the half-maximal response was measured. For the purposes of analyzing the information available from the cell

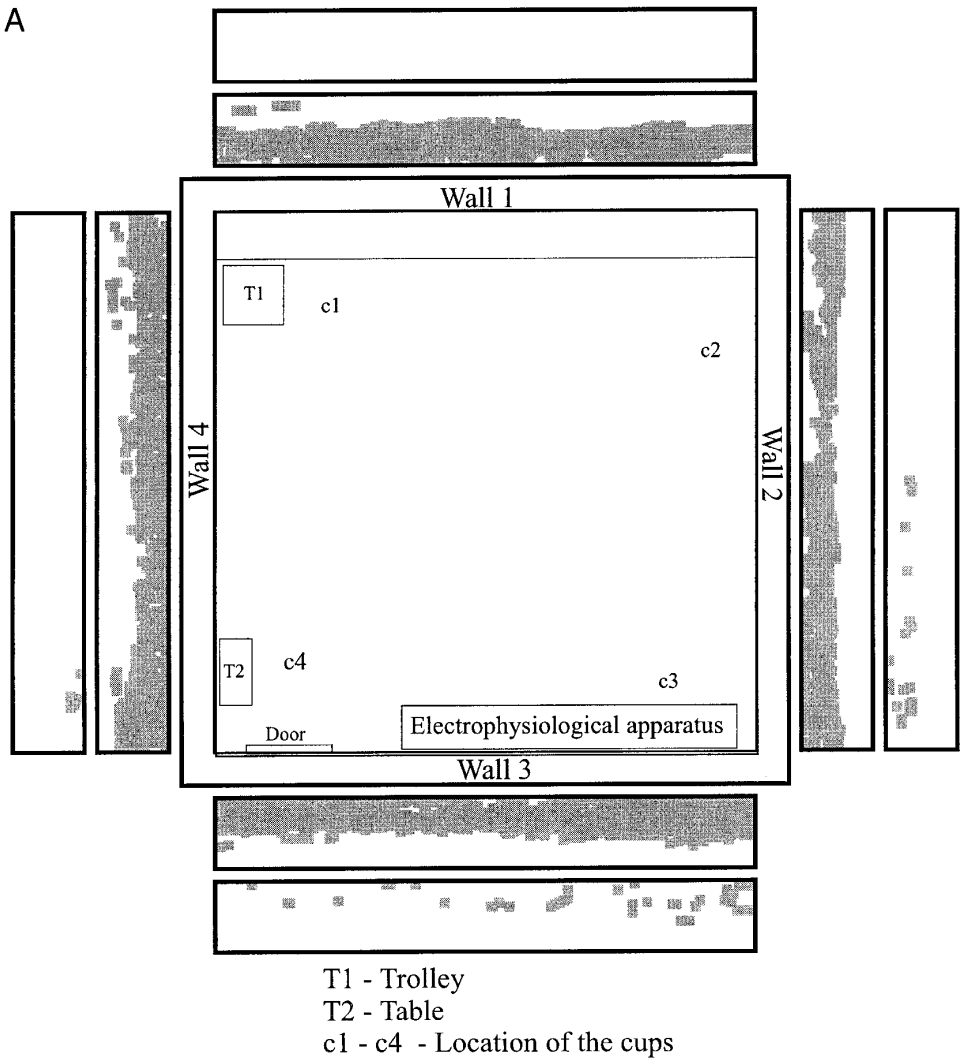

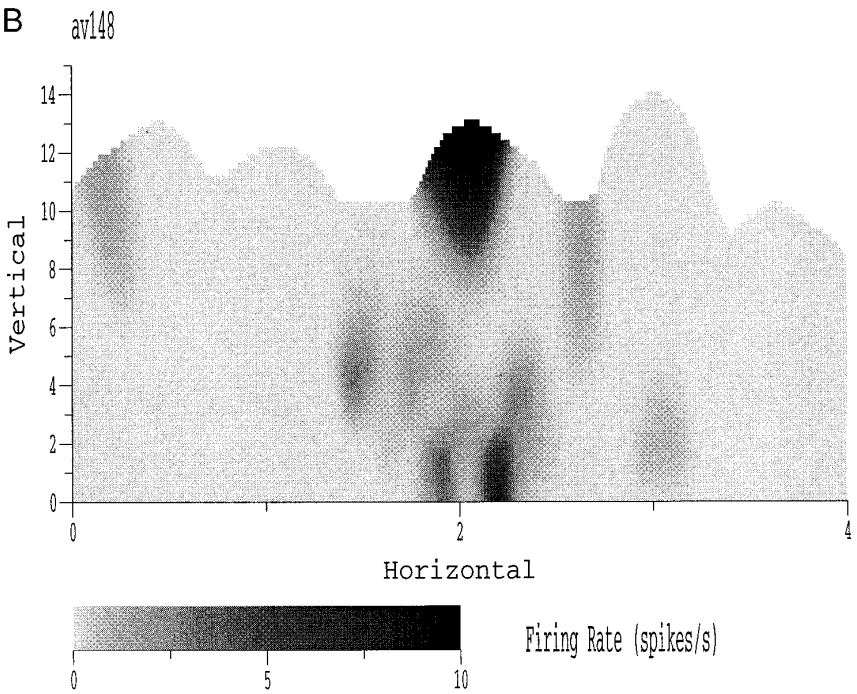

FIG. 1. Cell av148. *A*: view field of a hippocampal pyramidal cell during free exploration of the 2.7 × 2.7 m area. A spot on the *inner* set of 4 rectangles, each of which represents one of the walls of the room, indicates where on a wall the monkey had been looking every 25 ms during the 7-min recording session. Bottom of the wall is represented closest to the center of the diagram. A spot on the outer set of 4 rectangles indicates where the cell fired during the 7-min recording session. One spot is shown for each action potential. *B*: for the same cell the firing rate (gray scale calibration, *bottom*) projected onto a representation of the 4 walls of the room.

about where on the walls of the environment the monkey was looking, we binned that data as follows. We divided the four walls into eight half-walls horizontally and two half-walls vertically and then had available a large set of firing rates for each of the 16 quarter walls of the room (left and right, upper, and lower). From these rates, we used the techniques described later to analyze the information available both from each neuron taken alone and from an ensemble of neurons about where the monkey was looking. For the purposes of the exposition, we describe the walls as ''stimuli,'' and rephrase the analysis as estimating how much information is available from the neuronal response of the cell in any one 0.5-s epoch about which wall (or which part of a wall) the monkey is looking at. We note that dividing the walls into 16 stimuli means that the information required to decode correctly where the monkey is looking is 4 bits. Provided that this ''ceiling'' is not reached by the information available from one cell or the ensemble of cells, it is not necessary to divide the space into more stimuli, in that not much more information would be measured in the neuronal responses, as is evident also from the typical widths of the spatial view fields. This binning was used for the information analysis.

When we calculated the sparseness of the representation as described next, the original 64 × 16 resolution of the spatial view measurement was used because for sparseness, large numbers of samples in each bin are not necessary and a better estimate of the sparseness is obtained with a large number of bins. We note that the binning for position on the wall works best if the size of the spatial view field on the wall remains constant even when the monkey is at different distances from the wall. Our data do suggest that the angular width of the receptive field decreases as the distance of the animal from the wall increases (see e.g., Fig. 2 in Rolls et al. 1997a). Although this finding is consistent with an approximate constancy in size of the spatial view field as a function of the distance of the animal from the wall, a quantitative determination of this issue requires the collection of more data with a precise sampling of the view field over a large number of different spatial positions. However, even if the spatial view field is not perfectly constant in size, the result on the information analysis would be only the injection of a little additional noise into the data (resulting in a small underestimate of the true information) because the spatial bins used for the calculation of the information were quite large (one-quarter of the wall of the room).

Sparseness of the representation

The sparseness, a, of the representation of a set of stimuli (spatial locations in this case) provided by these neurons can be defined and was calculated as

$$a = \left(\sum_{s=1,S} (r_s/S) \right)^2 \Big/ \sum_{s=1,S} (r_s^2/S)$$

where r_s is the firing rate to the sth stimulus in the set of S stimuli. The sparseness has a maximal value of 1.0. This is a measure of the extent of the tail of the distribution, in this case, of the firing rates of the neuron to each stimulus. A low value indicates that there is a long tail to the distribution, equivalent in this case to only a few stimuli with high firing rates. If these neurons were binary (either responding with a high firing rate, or not responding), then a value of 0.2 would indicate that 20% of the stimuli produced high firing rates in a neuron and 80% produced no response. In the more general case of a continuous distribution of firing rates, the sparseness measure, a, still provides a quantitative measure of the length of the tail of the firing rate distribution (Treves and Rolls 1991). This measure of the sparseness of the representation of a set of stimuli by a single neuron has a number of advantages, detailed by Rolls and Tovee (1995). One is that it is the same measure of sparseness that has proved to be useful and tractable in formal analyses of the capacity of neural networks that use an approach derived from theoretical physics (see Rolls and Treves 1990; Treves 1990; Treves and Rolls 1991). A second advantage is that it can be applied to neurons that have continuously variable (graded) firing rates and not just to firing rates with a binary distribution (e.g., 0 or 100 spikes/s) (Treves and Rolls 1991). A third is that it makes no assumption about the form of the firing rate distribution (e.g., binary, ternary, exponential, etc.) and can be applied to different firing rate distributions (Treves and Rolls 1991). Fourth, it makes no assumption about the mean and the variance of the firing rate. Fifth, the measure does not make any assumption about the number of stimuli in the set and can be used with different numbers of test stimuli. Its maximal value is always 1.0, corresponding to the situation when a neuron responds equally to all the stimuli in a set of stimuli. The use of this measure of sparseness in neurophysiological investigations has the advantage that the neurophysiological findings then provide one set of the parameters useful in understanding theoretically (Rolls and Treves 1990; Treves and Rolls 1991, 1994) how the system operates.

For the purpose of calculating the sparseness a, the spatial locations were the 64 × 16 bins. A rate for a bin was used only if there was ≥1 s of data when the monkey was looking at that particular location. In typical experiments, there was sufficient data for the sparseness to be calculated over 100–300 such stimuli. Obviously the spatial resolution of the binning is limited by the recording time (or number of 500-ms periods) available, because finer bins would necessarily sometimes be empty. One can realize easily that taking a coarser binning on the same data produces higher values for a or apparently more distributed representations. Because of this overestimation effect, we chose to use relatively many bins to compute a, even if that meant relying on as few as two 500-ms samples per bin. Note that sparseness is a measure not strongly affected by limited sampling. Because the information measures described in the following text are, instead, strongly sensitive to limited sampling, to compute them we had to limit spatial resolution to that achieved with only 16 bins.

Information available in the responses of single neurons

The principles of the information theoretic analysis for single neurons were similar to those developed by Richmond and Optican (Optican and Richmond 1987; Richmond and Optican 1987) except that we applied a novel correction procedure for the limited number of trials. The analytic correction procedure we use was developed by Treves and Panzeri (Panzeri and Treves 1996; Treves and Panzeri 1995) to which we refer for a detailed discussion, and its efficacy in eliminating the limited sampling bias recently was compared with that of an alternative empiric procedure by Golomb et al. (1997). As in Rolls et al. (1997c), a novel aspect of the data analysis described here is that we investigated how much information was available about each stimulus in the set. The information theoretic analyses described and used here were based on the information available from the firing rate measured in 500-, 100-, and 25-ms periods when the eyes were steadily fixating a position in the room.

If each stimulus, s, were to evoke its own response, r (or its own set of unique responses), then on measuring r one would ascertain s, and thus gain $I(s) = -\log_2 P(s)$ bits of information, where $P(s)$ is the a priori probability of occurrence of a particular stimulus (in this case, a location in space) s. If instead, as happens in general, the same response sometimes can be shared, with different probabilities, by several stimuli, the probabilistic stimulus-response relation will be expressed by a table of probabilities $P(s, r)$ or, equivalently, of conditional probabilities $P(r|s) = P(s, r)/P(s)$. The information about s gained by knowing r can be evaluated from the formula

$$I(s, R) = \sum_r P(r|s) \log_2 \frac{P(s|r)}{P(s)} \qquad (1)$$

[This can be regarded as the difference between the original uncertainty about s (or a priori entropy) and the residual uncertainty after r is known, and attains its maximum value $I(s) = -\log_2 P(s)$ only if the probabilistic relation reduces to the deterministic one $P(s|r) = 1$ for $s = s(r)$, and $P(s|r) = 0$ otherwise.]

Averaging over different stimuli s in the set of stimuli S one obtains the average information gain about the set of stimuli S present in the neuronal spike data R (where R denotes the set of responses r) as

$$I(S, R) = \sum_s P(s)I(s, R) = \sum_{s,r} P(s, r) \log_2 \frac{P(s, r)}{P(s)P(r)} \qquad (2)$$

In the results, we show both $I(s, R)$, the information available in the responses of the cell about each individual stimulus s; and $I(S, R)$, the average information across all stimuli that is provided about which of the set of stimuli was presented.

In evaluating the information content from the data recorded, the neuronal responses were simply quantified by the number of spikes within any 500-ms time period, as stated above, that is we used a unidimensional measure based on a firing rate measurement. Both the set of stimuli S and the set of responses R in general could be continua (and the information I in the relation between the two still would be well defined because of the finite resolution with which responses can help discriminate among stimuli). However, in practice, to evaluate I, it is better to discretize both stimuli and responses to ensure adequate sampling of the spaces, and the number of discrete bins in each space must not be too high for limited sampling effects, even after the correction procedure we apply, not to bias information estimates based on limited numbers of trials (Treves and Panzeri 1995). In our analysis, S is discretized into 16 spatial bins as explained above, and there is no need to discretize R because R effectively is discretized already into a suitably low number of bins. (This is because by measuring responses as the number of spikes in 500 ms or less, these spike counts never exceeded 15–20 for hippocampal cells with their low rates.)

Information available in the responses of an ensemble of neurons

DECODING AND CROSS-VALIDATION PROCEDURE. In estimating the information carried by the responses of several cells, the analysis involved, first of all, constructing pseudosimultaneous population response vectors $\mathbf{r}$, occurring, as it were, in what were labeled as ''test'' trials [$\mathbf{r}$ is a vector with 1 element (or component) for each of the C cells considered]. Each response vector was compared with the mean population response vector to each stimulus, as derived from a different set of ''training'' or reference data, to estimate, by means of one of several decoding algorithms, as described later, the relative probabilities $P(s'|\mathbf{r})$ that the response $\mathbf{r}$ had been elicited by any one stimulus s' in the set. Summing over different test trial responses to the same stimulus s, we could extract the probability that by presenting stimulus s, the neuronal response would be interpreted as having been elicited by stimulus s' and from that the resulting measures of percent correct identification and of the information decoded from the responses.

Separating the test from the training data is called cross-validation and was performed in detail as follows, using the so called *jack-knife* technique. One of the available trials for each stimulus was used for testing, and the remaining trials for training. The resulting probabilities that s is decoded as s', however, were averaged over all choices of test trials, thus alleviating finite sampling problems, as described by Rolls, Treves and Tovee (1997b).

Algorithms for likelihood estimation

Several different decoding algorithms were used for estimating from the recorded response the likelihood of each stimulus, i.e., $P(s'|\mathbf{r})$. In the final analysis reported here, two are selected, namely the Euclidean distance and the dot product. The probability estimator (PE) algorithm, which tries to reconstruct the correct Bayesian probabilities from the data assuming a particular distribution of the neuronal responses such as Gaussian or Poisson (see Rolls et al. 1997b), was used but the results are not reported here because it was found that the sparse distribution of hippocampal cell responses fitted each of these distributions less well than in the case of inferior temporal cortex cells. The information and percent correct values obtained with the PE algorithm were, in any case, very similar (and usually slightly inferior) to those obtained with the Euclidean distance algorithm.

Both the algorithms that produced the results we report try to emulate the processing that could be performed by neurons receiving the output of the neuronal population recorded, thus extracting that portion of the information theoretically available that could be extracted with simple neurophysiologically plausible operations by receiving neurons. The DP (dot product) algorithm is simpler as it just computes the normalized dot products between the current firing vector $\mathbf{r}$ on a test trial and each of the mean firing rate response vectors in the training trials for each stimulus s'. (The normalized dot product is the dot or inner product of two vectors divided by the product of the length of each vector. The length of each vector is the square root of the sum of the squares.) Thus what is computed are the cosines of the angles of the test vector of cell rates with the mean response vector to each stimulus in turn. The highest dot product indicates the most likely stimulus that was presented, and this is taken as the best guess for the percentage correct measures. For the information measures, it is desirable to have a graded set of probabilities resulting from the decoding for which of the different stimuli was shown, and these were obtained from the dot products as follows. The S dot product values were cut at a threshold equal to their own mean plus SD, and the remaining nonzero ones were normalized to sum to 1. It is clear that in this case each operation could be performed by an elementary neuronal circuit (the dot product by a weighted sum of excitatory inputs, the thresholding by activity-dependent inhibitory subtraction, and the normalization by divisive inhibition).

The ED (Euclidean distance) algorithm calculates the stimulus likelihood as a decreasing function of the Euclidean distance between the mean response vector to each stimulus and the test vector. The specific function used was $\exp(-d^2/2\sigma^2)$, where $d = (|\mathbf{r}_s - \mathbf{r}|)$ and σ is the standard deviation of the responses calculated across all training trials and stimuli. The smaller this Euclidean distance is between the response vector of a test trial to a stimulus and the average response vector to a stimulus, the more likely it is that the stimulus on the test trial is the stimulus that produced that average response vector. Here response vector refers to the vector of firings of the set of cells in the ensemble. This measure is similar in principle to the biologically plausible dot product decoding considered before, in that both might be performed by a cell that received the test vector as a set of input firings and produced an output that depends on its synaptic weight vector, which represents the average response vector to a stimulus (see Rolls and Treves 1998). The slight additional complexity of the ED algorithm is that the lengths of both the mean response vectors and test vectors (which must be computed also by the DP algorithm for normalization) are used directly in combination with the dot product itself because $d^2 = \mathbf{r}_s \cdot \mathbf{r}_s - 2\mathbf{r} \cdot \mathbf{r}_s + \mathbf{r} \cdot \mathbf{r}$. The ED algorithm yields higher values for both percent correct and information, and thus it appears to minimize the loss in information due to the decoding step; we nevertheless report also values obtained with the DP algorithm to provide some indication of the extent to which the precise type of decoding used quantitatively affects the results.

Probability and frequency tables

Having estimated the relative probabilities that the test trial response had been elicited by any one stimulus, the stimulus which

turned out to be most likely, i.e., that which had the highest (estimated) probability, was defined to be the predicted stimulus, s^P. The fraction of times that the predicted stimulus s^P was the same as the actual stimulus, s, is directly a measure of the percent correct for a given data set. In parallel, the estimated relative probabilities (normalized to 1) were averaged over all test trials for all stimuli, to generate a table $P_N^R(s, s')$ describing the relative probability of each pair of actual stimulus s and posited stimulus s'. We also generated a second (frequency) table $P_N^F(s, s^P)$ from the fraction of times an actual stimulus s elicited a response that led to a predicted (most likely) stimulus s^P. The difference between the table P_N^R and the table P_N^F can be appreciated by noting that each vector comprising a pseudosimultaneous trial contributes to P_N^R a set of numbers (1 for each possible s') the sum of which is 1, whereas to P_N^F, it contributes a single 1 for s^P and zeroes for all other stimuli. Obviously each contribution was normalized by dividing, in both cases, by the total number N of (test) trials available (see Rolls et al. 1997b).

Information measures

Information values can be extracted from the joint probability tables $P(s, s')$ as from any other probability table $P(s, r)$. Again, when the probability table has to be estimated as the frequency table of a limited data sample, it becomes crucial to evaluate the effects of limited sampling on the information estimate. Because when considering the actual responses of several cells r turns into a multidimensional quantity (a vector, r), the minimum number of trials required to sample sufficiently the response space becomes very large [it grows exponentially with the dimensionality of that space, i.e., the number of cells considered (Treves and Panzeri 1995)]. This is what rules out, in the case of populations of more than very few cells, any attempt to evaluate directly the quantity $I(S, R)$ and forces us to resort to the (standard) procedure of deriving from the original frequency table of stimuli and responses an auxiliary table, of actual and potential stimuli, the latter being simply functions of the responses spanning a discrete set with a reduced number of elements (equal, that is, to the number of stimuli). In general, the information content of the auxiliary table will be less than that of the original table by an amount that depends on the severity of the manipulation performed. If the decoding is efficient, that is, if it extracts from the responses nearly all that the responses can tell about the stimulus, then, by construction, not much information is lost in the decoding step. The fact that a decoding operation may be a plausible part of the processing produced by the nervous system at some stage adds credibility to the procedure, which is in any case necessary, of estimating the information carried by several cells only after decoding their responses. Two types of auxiliary tables were derived here, called P_N^R and P_N^F see preceding text.

Information estimates corrected for the limited number of trials

The procedure introduced so far for estimating information values, both for single cells [from the probability table $P(s, r)$] and for ensembles of cells [from either the table $P^F(s, s^P)$ or $P^R(s, s')$] must be supplemented by a procedure that corrects the raw estimates for their limited sampling biases. In practice, in fact, because of the limited number of trials that can be collected, the various probability tables are not available, and one can at best approximate them with frequency tables, e.g., $P_N(s, r)$, computed on the basis of a (limited) number of trials N. If N is very large, the frequencies should get close to the underlying probabilities, but for any finite N there will be a discrepancy, which will result in an error in the estimated information gain. This error decreases as the number of trials for each stimulus increases. Because information quantities depend on probabilities not in a linear but in a greater than linear manner, the error deriving from this ''limited sampling'' does not cancel out on averaging many measurements; it is, instead, usually biased upward, resulting in an (average) overestimate of the information gain, as described by Tovee et al. (1993), Treves and Panzeri (1995), and Rolls and Treves (1998).

The net bias, or average error (usually an overestimating error), can be expressed analytically as a formal expansion in $1/N$, and the first few terms (in particular, the very first) of this expansion can be evaluated directly (Panzeri and Treves 1996) in a variety of situations. Simulation experiments have shown that the first term in the expansion is responsible for most of the discrepancy between the raw and correct information measures, whereas successive terms do not in fact correlate with the remainder of the discrepancy. This first term then can be subtracted from the raw estimates to produce corrected estimates. This procedure has been shown to improve significantly the reliability of information estimates based on limited data samples, as discussed also in an explicit comparison with an alternative procedure by Golomb et al. (1997). The procedure provides a good correction ($<5\%$ error) when the number of trials for each stimulus is larger than the number of bins R into which the neuronal responses are discretized [when decoding is used, this number of bins reduces to S (see preceding text; Golomb et al. 1997; Panzeri and Treves 1996)]. With respect to the correction based on another procedure, which we used in previous investigations (Tovee and Rolls 1995; Tovee et al. 1993), one should note that, while in several cases it yields information values that are close to those obtained with the present procedure, it cannot be applied to compute the stimulus-specific information $I(s, R)$. This is one case, therefore, in which it was essential to develop a novel procedure to correct for limited sampling.

With respect to the sensitivity to limited sampling of the different measures of the information carried by an ensemble of cells, it is worthwhile to note the following. In deriving P_N^F, each response is used to predict its stimulus. Although s^P spans only S values compared with the very large number of possible (multidimensional) rate responses, the auxiliary table is otherwise unregularized in that each trial of a limited total number produces a relatively large ''bump'' in $P_N^F(s, s^P)$. The result of this is that a raw estimate of $I(S, S^P)$ [which can be denoted as $I_N(S, S^P)$ to point out that it is obtained from a total of N trials] can still be very inaccurate, in particular, overestimated. The correction methods we use, on the other hand, are safely applicable when the subtracted term $[+I_N(S, S^P) - I(S, S^P)]$ is smaller than ~ 1 bit. With the present data, the subtracted term turns out to be safely small except in some cases when few cells are considered.

P_N^R, on the other hand, can be conceived of as being more ''regularized'' than P_N^F because each trial contributes not a relatively large bump to just to one bin s^P but smaller additions to several bins s'. The consequence is that the distortion in the information estimate due to limited sampling (small N) is smaller, and the subtraction of a suitable correction term $[I_N(S, S')] - I(S, S')$ is enough to produce accurate corrected estimates of information. The correction term to be used differs from that appropriate to correct $I(S, S^P)$ or $I(S, R)$, and we refer to Panzeri and Treves (1996) and to Rolls, Treves, and Tovee (1997b) for details. Here it is sufficient to note, again, that $I(S, S')$ (as best estimated with the present correction procedure) will in any case tend to a ''true'' value that, being based on a regularized probability distribution, is less than the value (unmeasurable except with few cells) attained by $I(S, R)$. The same applies to $I(S, S^P)$. In a previous paper (Rolls et al. 1997b), we have used $I(S, S')$ as a substitute for $I(S, S^P)$ when the latter could not be safely estimated due to limited sampling. In the present work, relatively more data are available, but we still prefer to report $I(S, S')$ values, because they are more smooth. In any case, we have calculated the information with both methods, and the values are very close for the decoding algorithms used in this paper. Residual limited sampling problems were en-

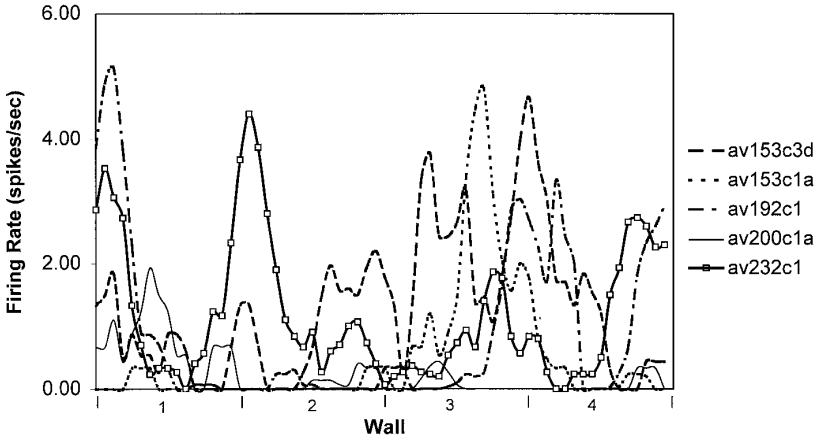

FIG. 2. Profiles of the response magnitudes (in spikes/s computed during 500-ms periods) of 5 different cells shown as a function of where on the walls the monkey was looking in the horizontal axis (extracted from data binned in just 64 horizontal bins).

countered in some cases in which ensembles of few cells with very low firing activity were considered (essentially because in those cases on most trials the ensemble firing vectors are equal to the null vector, and only the remaining trials effectively convey information), and to control for them, we used the procedure described next.

Calculation of the information directly from the responses of cells with sparse firing rate distributions

A new alternative procedure for information measurement contained in a population of cells developed for use with the low firing rates and sparse firing rate distributions characteristic of hippocampal cells is introduced next. With such distributions, it is possible to consider the responses of single neurons in even 100-ms periods as being binary without losing a significant amount of information, as we have quantitatively checked at the single cell level. We then binarize single cell responses by labeling them as 0 if no spike was emitted during the given time period and as 1 if at least 1 spike was emitted, and thus we represent the response by the whole ensemble as a binary response vector. As with C cells, there are only 2^C possible responses; it is feasible to collect as many trials as there are, in practice, different types of response for up to perhaps five cells (which thus would need 32 trials, but only if every possible response actually occurred; in practice responses where more than 1–2 cells fire are very rare, and the number of actual outputs is $<10-15$). The information analysis is reliable when the number of trials per stimulus is at least as large as the number of actual responses (see Panzeri and Treves 1996; Rolls and Treves 1998). With this type of analysis, there is no need for any decoding procedure, and the information can be calculated directly from *Eq. 2* with the finite sampling correction

TABLE 1. *Quantitative measures of the response characteristics of the different hippocampal neurons*

Cell	$I(S, R)$	I_{max}	I_{500}	a	a_{25}	a_{500}	Sp	Peak Rate	Site	Cell Number
av142	0.021	0.134	0.188	0.322	0.04	0.201	0.2	12	DENT	1
av148	0.049	0.535	0.154	0.053	0.004	0.035	0.0	9	CA1	2
av153c1	0.060	0.310	0.277	0.334	0.048	0.190	0.0	9	PSUB	3
av153c3	0.036	0.282	0.216	0.297	0.017	0.179	0.3	9	PHG	4
av180	0.128	0.264	0.153	0.296	0.046	0.230	0.2	12	PHG	5
av181	0.171	0.726	0.282	0.349	0.088	0.298	0.5	24	CA1	6
av191	0.170	0.830	0.488	0.388	0.113	0.346	0.0	24	PHG	7
av192c1	0.159	0.390	0.432	0.229	0.025	0.163	0.0	7.5	PSUB	8
av192c2	0.064	0.518	0.271	0.250	0.012	0.122	0.0	7.5	PSUB	9
av192c2cl2	0.054	0.236	0.361	0.208	0.003	0.049	0.0	7.5	PSUB	10
av197	0.162	0.579	0.312	0.323	0.080	0.338	0.6	18	PSUB	11
av200c1	0.056	0.351	0.162	0.124	0.012	0.135	0.0	4.5	PHG	12
av216	0.056	0.150	0.185	0.242	0.039	0.163	0.0	21	PHG	13
av221	0.144	0.754	0.348	0.310	0.001	0.018	0.2	19.5	PHG	14
av222	0.035	0.122	0.091	0.205	0.024	0.144	0.2	12	CA1	15
av229	0.180	1.057	0.589	0.068	0.083	0.366	0.0	13.5	CA1	16
av232	0.091	0.746	0.332	0.129	0.023	0.125	0.0	18	CA1	17
av273.1	0.241	1.007	0.749	0.399	0.174	0.406	0.0	30.75	PHG	18
av273.2	0.141	0.394	0.604	0.200	0.054	0.239	0.0	21	PHG	19
av296	0.133	0.276	0.376	0.131	0.035	0.176	0.1	18	PHG	20
az033c1	0.150	0.519	0.305	0.206	0.029	0.155	0.0	9	CA3	21
az033c2a	0.127	0.710	0.295	0.141	0.023	0.105	0.0	9.2	CA3	22
az033c2acl2	0.028	0.141	0.199	0.098	0.021	0.124	0.0	2.5	CA3	23
az034c1cl1	0.042	0.278	0.389	0.211	0.037	0.171	0.0	7.25	CA3	24
az102	0.012	0.098	0.108	0.072	0.009	0.090	0.0	1.0	CA1	25
az103	0.335	1.359	0.535	0.332	0.086	0.370	0.0	17.5	PSUB	26
Mean	0.109	0.491	0.323	0.227	0.043	0.189	0.1	13.2		

Note: a is calculated from 64 horizontal $\times$ 16 vertical spatial bins in which there were $\geq$1 s of data in 500-ms runs. $I(S, R)$ is the mutual information available in 100 ms about spatial view. I_{max} is the maximum information about any of the 16 stimuli or views available in 100 ms. I_{500} is the mutual information available in 500 ms about the spatial view. Sp, spontaneous rate; DENT, dentate gyrus; PHG, parahippocampal gyrus; PSUB, parasubiculum.

applied. The advantage of this procedure is that when it can be applied, it is more direct because it does not involve a decoding step and thus allows the relation between the number of cells and the information available to be specified accurately for small numbers of cells.

RESULTS

An example of the firing rate of a hippocampal pyramidal cell when the monkey was walking round the environment is shown in Fig. 1*A*. The inner set of four rectangular boxes show where the monkey looked on the four walls. (The top of each wall is furthest from the center.) The outer set of four boxes again represent the four walls, but in these, a spot indicates where the cell fired. It is clear that the cell has a spatial view field located mainly on wall 3. To show in more detail how the firing rate of the neuron varied as a function of where the monkey was looking, we show in Fig. 1*B* that for the same cell, the firing rate (gray scale calibration, *bottom*) projected onto a representation of the four walls of the room.

The profiles of the response magnitudes of a sample of the different cells plotted as a function of where on the walls the monkey was looking in the horizontal axis are shown in Fig. 2. The values for any horizontal point were averaged over the vertical values. (Further examples of similar plots for different cells are provided by Georges-François et al. 1998; Robertson et al. 1998; Rolls et al. 1997a; where full details of the response properties of hippocampal spatial view cells in primates are provided). From Fig. 2 it can be seen that each cell typically has its spatial view field centered at a different position relative to the other cells. It also can be seen that the view fields are typically quite extensive horizontally, with mean half-amplitude widths of 0.9 walls.

Information available in the responses of single neurons

A histogram showing the values of $I(S, R)$ (the average information in the responses of a cell about the stimulus set within 100-ms periods) for each of the 26 cells in the two monkeys for which sufficient data were collected for the analyses described here is provided in Fig. 3*A*. Most of the neurons had values for $I(S, R)$ in the range 0.05–0.2 bits, with the average across the population of neurons being 0.11 bits. All these neurons showed a significant differential response to the different walls in one-way analyses of variance. For a 500-ms period, the average information provided by these cells about which location the monkey was looking at was 0.32 bits. Thus there was a reasonable amount of information available in the firing rates of these neurons in a 100- and a 500-ms period about spatial location ''out there,'' even though the firing rates of the neurons were low, with a mean peak response to the most effective spatial location of 13.2 spikes/s (compared with a spontaneous rate of 0.1 spikes/s). Although $I(S, R)$ may not appear to be high, it should be remembered that this neuronal information measure is equivalent to the average of all the information contained in the responses to the individual stimuli (corrected for the minor differences in the number of trials). If many of the stimuli (walls) evoke a similar neuronal response, then the average information from the neuronal response about which stimulus was being looked at is low.

To understand the representation of individual stimuli by individual cells, the specific information $I(s)$ available in the neuronal response about each stimulus s in the set of stimuli S was calculated for each cell. If the neuron responded to one of the stimuli (1 of the half walls) and not to any other, then the maximum information contained when that effective stimulus was shown would be 4 bits, and of the other stimuli close to 0 bits (in fact, $\log_2(16/15) = 0.09$ bits). The maximum information values $I_{\max}$ of the different neurons about any one stimulus are shown in Fig. 3*B* and Table 1, again calculated for 100-ms periods of the neuronal response. The majority for 100 ms are in the range 0.2–0.8 bits. The mean value of $I_{\max}$ for the different cells was 0.49 bits for 100 ms. For a 500-ms period, the mean value of $I_{\max}$ was 1.2 bits (see Table 1).

Sparseness of the representation

The data above indicate that the encoding of information about spatial view in this population of neurons is not achieved by very finely tuned neurons, that is, the representation is likely to be achieved by distributed encoding. To quantify this, the measure a of the sparseness of the representation was calculated. The sparseness measure indicates the length of the tail of the distribution of neuronal responses to the stimuli such that low values indicate high selectivity

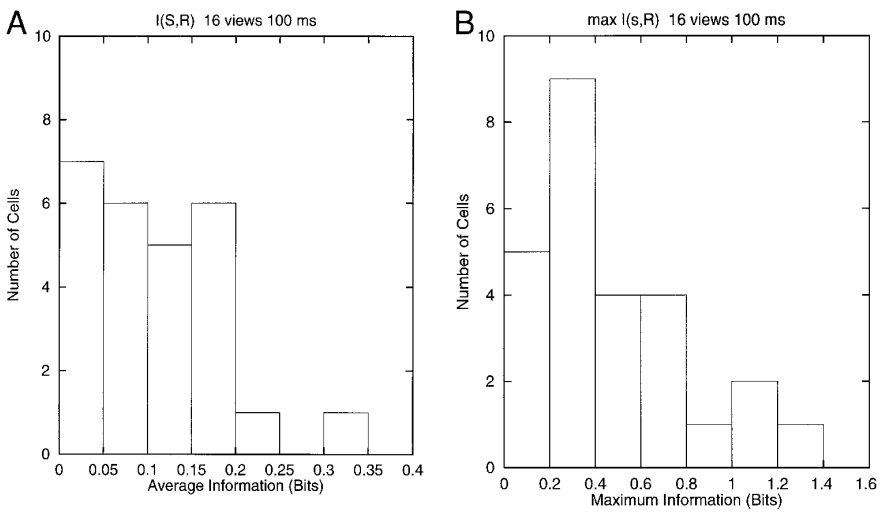

FIG. 3. *A*: histogram showing the values of $I(S, R)$ (the average information in the responses of a cell about the stimulus set in a 100-ms period) for each cell; here spatial views were binned into a total of 16 stimuli. *B*: maximum information values $I_{\max}$ of the different neurons about any 1 stimulus available in a 100-ms period.

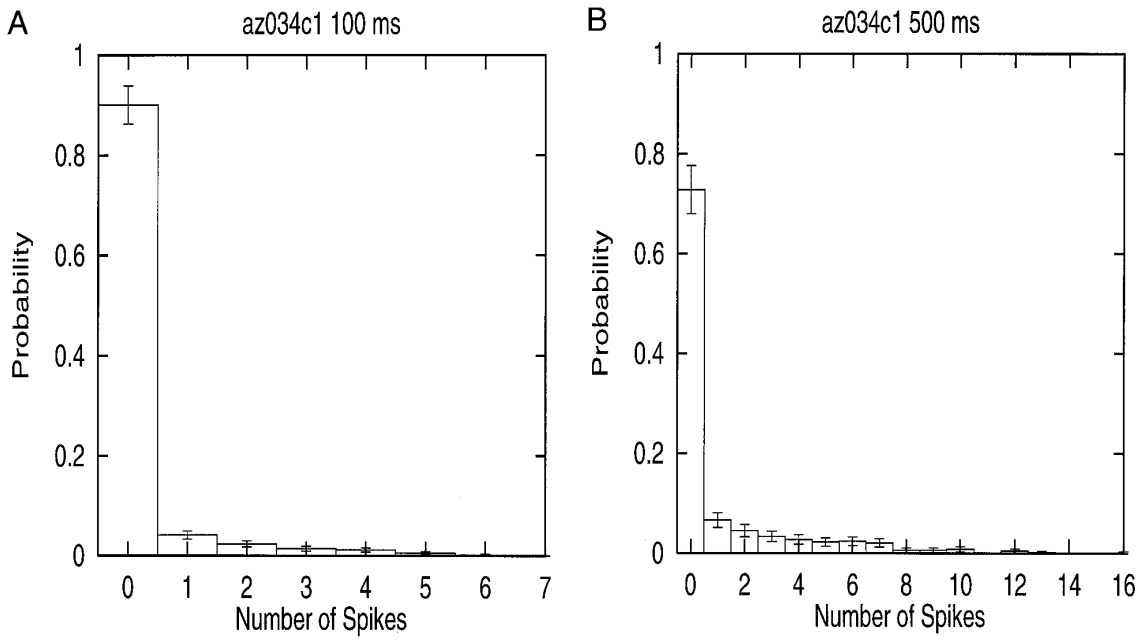

FIG. 4. *A*: distribution of the number of spikes emitted by a CA3 neuron in 100-ms periods. *B*: distribution of the number of spikes emitted by a CA3 neuron in 500-ms periods.

to one or a few of the stimuli in the set, and a value of 0.5 if the neurons had binary firing rates (e.g., firing or not) would indicate equal (firing) responses to half the stimuli and no response to the other half. Sparseness was calculated in the manner described in METHODS.

Table 1 shows the sparsenesses a of the neuronal responses (sampled in 500-ms periods) in the 64 horizontal $\times$ 16 vertical spatial locations into which the walls of the room were divided. It can be seen that none of the neurons had very low values for a, that is, the coding was relatively distributed. The mean value for a was 0.22, indicating somewhat distributed encoding.

Figure 4*A* shows the distribution of the number of spikes emitted by a typical neuron in 100-ms periods. This is of interest, because it shows the statistics of the spike arrivals that might be expected in other hippocampal neurons receiving from the recorded neuron, in a time that is already longer than or in the order of size of the synaptic and membrane time constants. For comparison, we show also the spike count distribution calculated during 500-ms time periods in Fig. 4*B* for the same CA3 cell. The averages across the population of 26 cells of the firing rate distributions are shown in Figs. 5 and 6. These distributions were calculated by normalizing the mean firing rate of each cell to 1, so that cells could be combined.

To obtain a quantitative measure to reflect this distribution, the sparseness of the neuronal responses was calculated from 25- and 500-ms time periods of the neuronal response [i.e., using the formula

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

where r_i is the firing rate in the ith 25-ms period of the

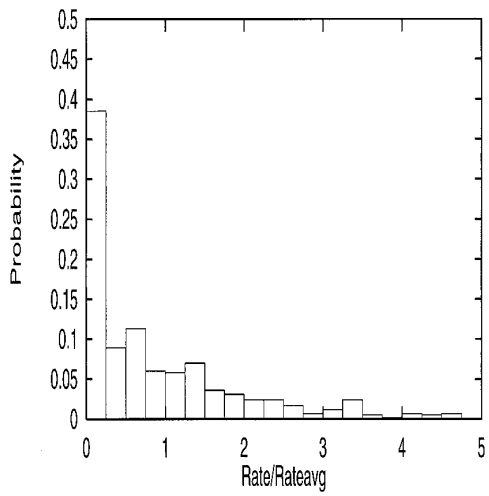

FIG. 5. Average across the population of cells of the number of spikes emitted during 100-ms periods. To compute this average distribution, the firing rate of each cell was normalized to 1.0 by dividing it by its mean.

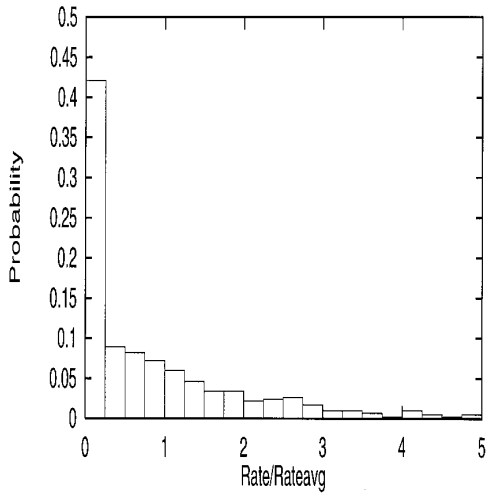

FIG. 6. Average across the population of cells of the number of spikes emitted during 500-ms periods. To compute this average distribution, the firing rate of each cell was normalized to 1.0 by dividing it by its mean.

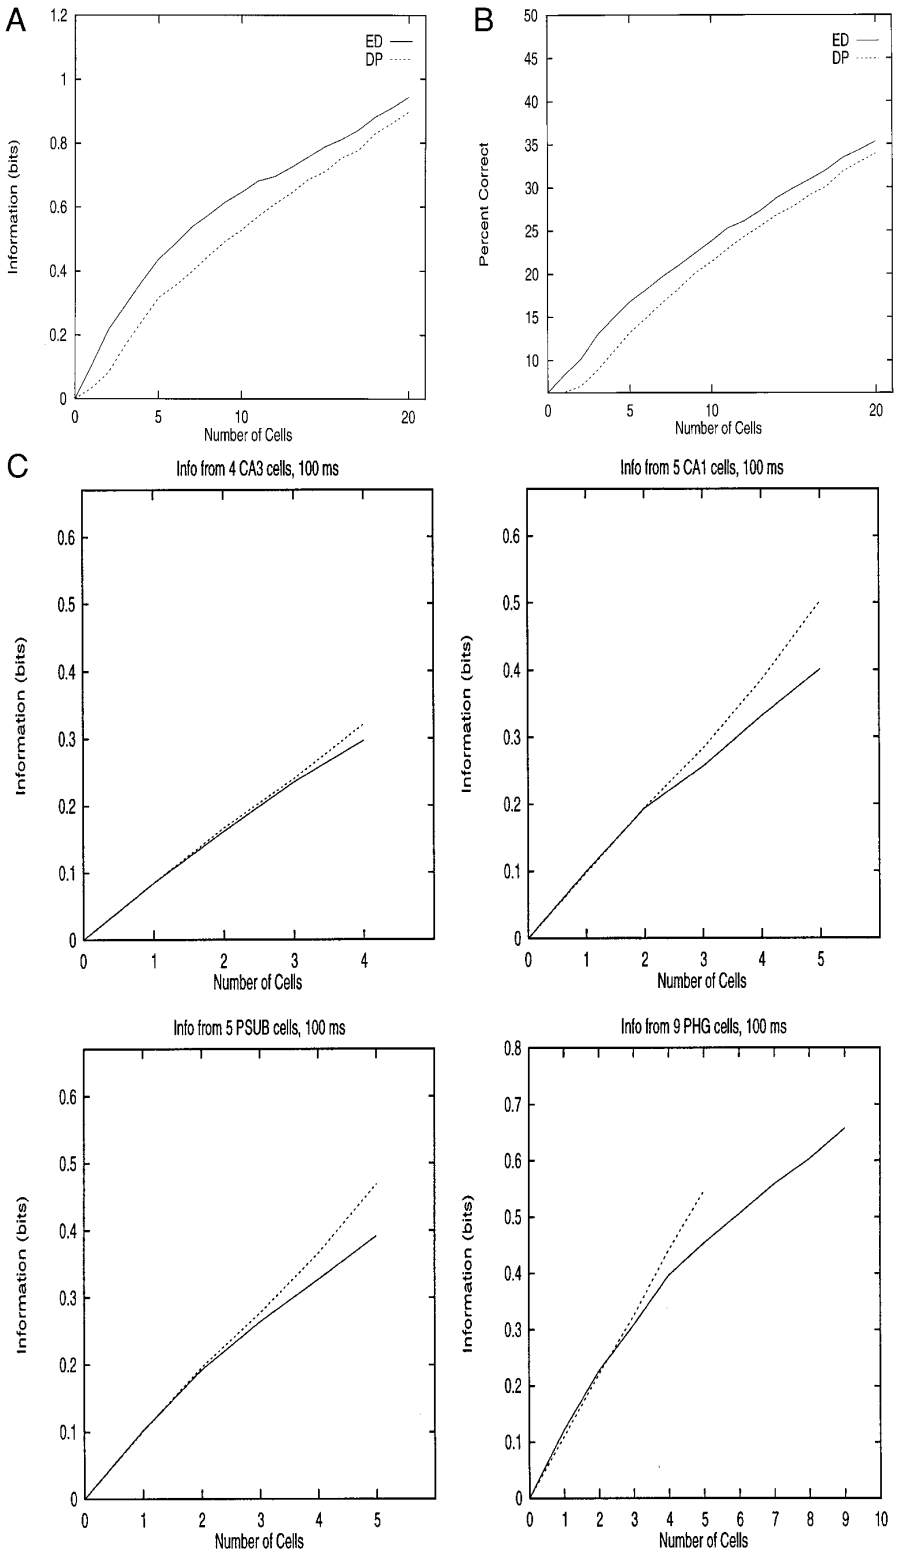

FIG. 7. *A*: values for the average information, $I(S, S')$, available in the responses of different numbers of these neurons on each trial, about which of the 16 stimuli (i.e., quarters of walls) is being looked at. ——, Euclidean distance decoding algorithm was used for estimating the relative probability of posited stimuli s'; – – –, dot product result. Twenty cells were recorded from the same (*av*) animal. *B*: percent correct predictions based on the same data used in *A*. *C*: values of $I(S, R)$ extracted by means of the direct procedure (– – –) are compared with the values of $I(S, S')$ calculated using the Euclidean distance decoding algorithm (——). Values were estimated from 100-ms periods. Values for the average information from cells in 4 different subregions are shown (CA3 and CA1, hippocampal pyramidal cell fields; PSUB, presubiculum; PHG, parahippocampal gyrus).

neuronal activity measured during >5 min while the monkey explored the environment]. Note that this is a measure of the sparseness of the spike count distribution, that is, of the number of time bins with 0, 1, 2 spikes, etc. irrespective of how the firing correlates with spatial view. The interest of this measure is that it summarizes the statistics of the spikes received by a neuron in short time periods down to as low as estimates of the time constant of the fast excitatory synapses, considering also dendritic delays (which, even combined, are not beyond 25 ms). It can be seen that the sparseness of the representation described in this way takes a small value. The mean value of a_{25} for the neurons in Table 1 was 0.04. This period was chosen because it is in the order of size of the time constants of the hippocampal synaptic events

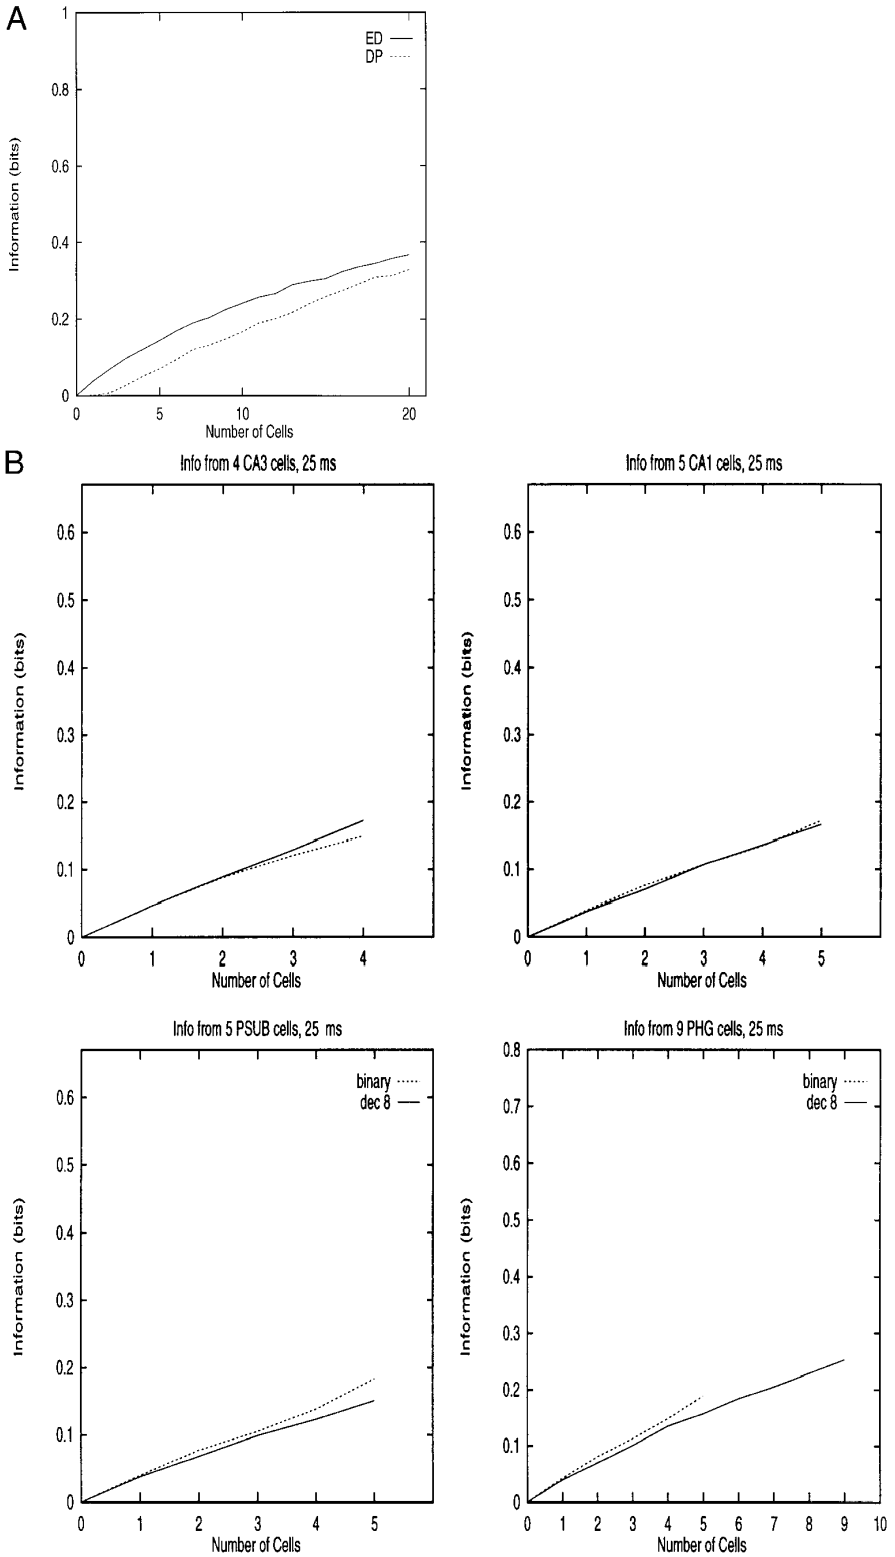

FIG. 8. *A*: results from the same set of cells analyzed as in Fig. 7*A* but with a shorter time for each trial (i.e., period within which the eyes are still and during which the spikes emitted by each cell are counted) of 25 ms. *B*: values of $I(S, R)$ for the same set of cells as in Fig. 7*C* extracted by means of the direct procedure (– – –) are compared with the values of $I(S, S')$ calculated using the Euclidean distance decoding algorithm (——). Values were estimated from 25-ms periods. Values for the average information from cells in 4 different subregions are shown.

which set the time scale of the operation of the system (see further Treves 1993). Essentially because over 25 ms these cells behave roughly as binary units, the result is that on average they fire (usually just 1 spike) during 4% of all 25-ms periods in the recording session. The results of a similar calculation for spike sparseness values based on 500-ms time periods are shown in Table 1 as a_{500}. The 500-ms period was chosen because it is of the order of time over which a new memory might be learned, and thus this is the sparseness value which will set the sparseness of the synaptic weight vector on a neuron that might be laid down for a single memory. This value of sparseness, denoted here as a_{500}, has a mean value of 0.19, close to the value a calculated from the mean responses to each spatial view. Note that this does

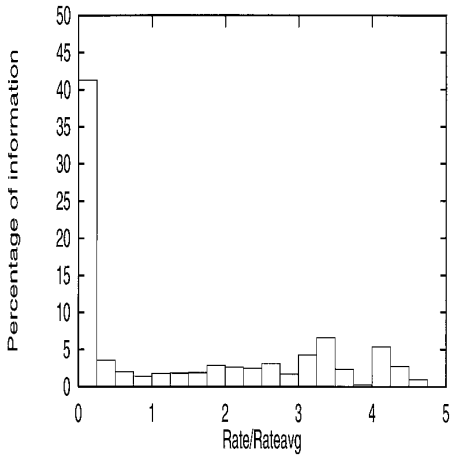

FIG. 9. Percentage of the information carried by the different levels of firing rate for this population of 26 cells. Average rate for each cell was normalized to 1 so that the cells could be combined. Firing rate measure on the abscissa is the firing rate expressed as a fraction of the mean rate for the cell.

not imply that on average these cells fire 19% of all 500-ms periods because over 500 ms they do not behave as binary units. Rather, the sparseness value obtained is a measure of the tail of the firing rate distribution computed over 500 ms, and it turns out not to matter very much whether it is calculated irrespective of spatial view, as in a_{500}, or from the mean responses to each view, as in a. Either way, this is roughly the value that might be expected to influence the number of memories stored in an associative network (see Rolls and Treves 1998).

Information available from an ensemble of these neurons

The values for the average information, $I(s, s')$, available in the responses of different numbers of these neurons on each trial, about which of the 16 stimuli (i.e., walls) is being looked at, are shown in Fig. 7*A* for 100-ms periods. The Euclidean distance decoding algorithm was used for estimating the relative probability of posited stimuli s'. The same data produced the percent correct predictions reported in Fig. 7*B*. It can be seen that the information rises approximately linearly with population size from its baseline level (which is zero for 0 cells) for the first four to five cells and after that increases less rapidly. The percent correct also rises approximately linearly with population size from its baseline (chance) level (which is $100/S = 6.25$ for the percent correct). The 20 cells were, of course, all from the same animal (av) and used 40 trials from every cell in almost all cases. (For 25-ms analysis, 40 trials were available for every cell for every stimulus; the number of trials was 2% less than ideally required for 100-ms time win-

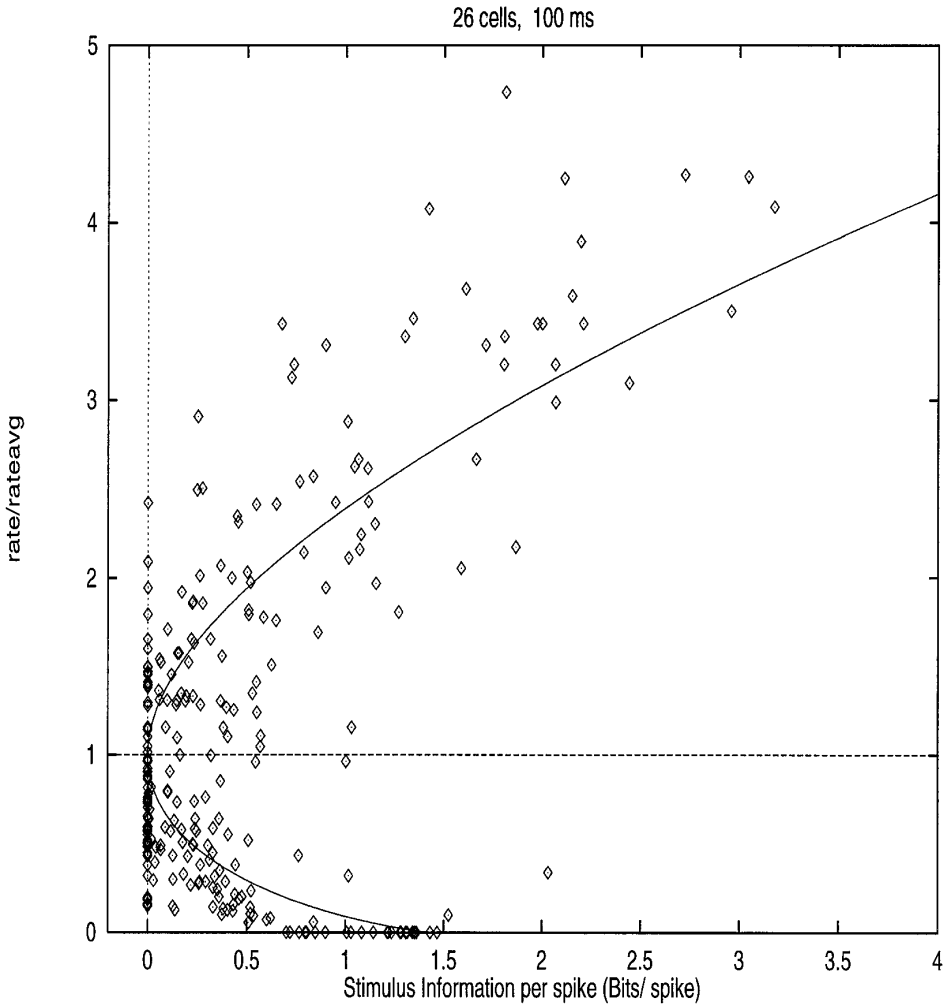

FIG. 10. Information $I(s, R)$ available in the response of the set of 26 hippocampal neurons about each of the stimuli in the set of 16 stimuli, each of which was a different part of space (abscissa), with the firing rate of the neuron to the corresponding stimulus plotted as a function of this on the ordinate. – – –, mean firing rate of the cell. Average rate of the cell on the ordinate was normalized to 1 so that the cells could be combined. Stimulus-specific information is divided by the mean number of spikes emitted by the cell on the abscissa and has the meaning of the information about a particular stimulus available in 1 mean interspike interval of the cell. ——, how the information per spike about a stimulus is related to the firing rate of a neuron to that stimulus in the limit of short time windows (see Rolls et al. 1997c).

dows, but we checked that the small number of random missing trials that needed replication had a negligible effect on the information analysis.)

To investigate to what extent the information does rise linearly, we applied the direct information measurement procedure possible with binary rate distributions and compared the results with the Euclidean distance decoding procedure in Fig. 7*C*. The analysis was performed separately for neurons in different parts of the hippocampal formation and separately for each animal. The comparison shows that the increase of information is closely linear with the number of cells when using the direct information measurement (which is possible for ≤ 5 cells) and that the measurement based on decoding (in this case ED decoding) underestimates the information for more than about four cells. The underestimation probably is related to the sparseness of the firing rate distributions of hippocampal neurons, which makes the decoding step lose some of the information. The conclusion is that the somewhat less than linear increase in information apparent in Fig. 7*A* for more than about four cells is probably just due to the inefficiency of the decoding procedure when applied to these low firing rate neurons. (The fact that the increase of information with the direct information measurement method may appear to be close to supralinear in the number of cells is probably that with 5 cells the finite sampling correction for the limited number of trials is operating at its limit, given that the number of trials per stimulus was 40 and the dimensionality of the binary response space is 32.)

The results from the same set of cells analyzed with dot product decoding also are shown in Fig. 7, *A* and *B*. The reason that the information is zero and the percent correct is at chance with one cell for DP decoding is, obviously, that then the dot product of the test trial vector of cell responses with any of the average response vectors to the stimuli is essentially meaningless.

The multiple cell information analysis for the same set of cells analyzed as in Fig. 7 but with a shorter time for each trial (i.e., period within which the eyes are still, and the number of spikes is measured) of 25 ms are shown in Fig. 8.

It is possible to show how much of the information is carried by the different levels of firing rate, given the mean firing rates elicited by each stimulus and the corresponding $I(s, R)$ values that have been the subject of this paper. The result is shown in Fig. 9 for 100 ms, averaged over the 26 cells. It is of considerable interest that much of the information was available from the firing rates that were below the mean (normalized to 1 in Fig. 9), related to the fact that low firing rates were very common. The mode of this distribution is between 0.0 and 0.25 with respect to the mean firing rate across all stimuli. This is linked to the fact that information is a relative measure. This results in some information at very low rates relative to the mean rate. Given the high probability of very low rates for hippocampal cells, the total information conveyed by low rates is thus high, as shown in Fig. 9. It is also shown in Fig. 9 that there is a dip in the information available in those rates that are near the mean rate for each cell. This is related to the fact that stimuli that evoke a firing rate response close to the mean across all stimuli carry little information. This point is made more explicitly in Fig. 10, which shows the information available about each stimulus in relation to the firing rate response to the same stimulus. Figure 10 shows that firing rates above, or below, the mean convey information.

The sites at which these 26 cells were recorded are shown in Fig. 11. Ten were in the hippocampal pyramidal cell fields CA3 or CA1. They were probably hippocampal pyramidal cells, as shown by the large amplitude action potentials, very low spontaneous firing rates in this type of experiment (mean 0.2 spikes/s), and relatively low peak firing rates (mean 10.5 spikes/s) (cf. Feigenbaum and Rolls 1992). Sixteen were in the overlying cortical areas or paracortical areas, including the parahippocampal gyrus, which connect the hippocampus to other cortical areas. The mean spontaneous firing rates of these cortical neurons in this type of experiment was 0.1 spikes/s, and the peak firing rates had a mean of 15.1 spikes/s.

DISCUSSION

The neurophysiological results described here show that the information about where (on the walls of the room) the monkey was looking increases approximately linearly with the number of cells in the ensemble. This shows that the information conveyed by a hippocampal neuron is roughly independent of that carried by other hippocampal neurons. Put another way, the number of stimuli, in this case locations in space, that can be encoded by a population of neurons in this part of the brain increases approximately exponentially as the number of cells in the sample increases. That is, the log of the number of stimuli increases approximately linearly as the number of cells in the sample is increased. This is in contrast to a local encoding scheme (of ''grandmother'' cells), in which it is the number of stimuli encoded that increases linearly with the number of cells in the sample. The conclusion is that one of the attractive potential properties of distributed encoding, that the number of stimuli that can be encoded increases exponentially with the number of cells in the representation, is expressed by this population of hippocampal neurons. A mechanism that has been suggested to contribute to this is the pattern separation (or orthogonalization) performed by the dentate granule cells operating as a competitive network and by the mossy fiber projection to the CA3 cells (Rolls 1989; Rolls and Treves 1998; Treves and Rolls 1992).

That an exponentially increasing capacity with an increase in cell number is a potential property of a distributed representation can be seen clearly from the following example. Consider the number of stimuli that can be encoded by a population of C neurons without noise. If local encoding is used (i.e., a single neuron specifies the stimulus, that is grandmother cell encoding), and the representation is binary (e.g., the neuron is either active or not), then C different stimuli can be encoded. (One different neuron is on for each stimulus.) If distributed encoding is allowed, then 2^C different stimuli can be encoded. (2^C is the number of different combinations of C binary variables.) The fundamental question addressed in this paper is the extent to which the hippocampal system can use the potential advantage of distributed representations to encode a very large (exponentially large) number of different stimuli in a population of neurons. The potential advantage only will be usefully realized to the extent that each member of the population of neurons has different responses to each stimulus in a set of

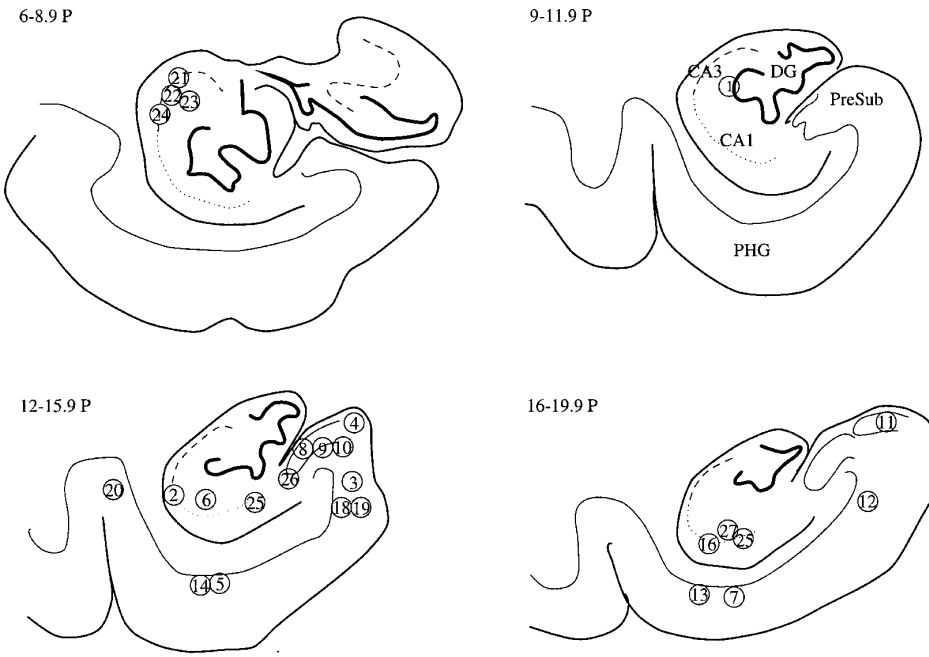

FIG. 11. Hippocampal and parahippocampal sites at which different spatial view cells were recorded. Cells are numbered, and cross-refer to Table 1.

stimuli (with, e.g., different combinations of neurons firing to each stimulus) and to the extent that the responses of a neuron on a given trial are not too noisy. That is, the standard deviation of the responses of a neuron to the same stimulus on different trials must not be too great, and the responses to different stimuli must be reliably different to each other. Evidence on this issue only can be obtained by examining the response properties of real neurons in the brain. The results described in this paper show the extent to which these conditions are met, that is that the neurons do have sufficiently different view field centers (see Fig. 2), and the firing of each neuron is sufficiently reliable and independent (Figs. 7 and 8).

The results described here also show that a reasonable amount of information about spatial location is provided by primate hippocampal neurons. For example, the average information provided by these cells about which location the monkey was looking at was 0.32 bits, and the mean across cells of the maximum information conveyed about which location was being looked at was 1.20 bits, measured in a period of 0.5 s. In a study performed in rats, the information from an ensemble of hippocampal place cells about the rat's location has been estimated as on average ~0.3 bits in a period of 0.5 s (Treves et al. 1996).

Two different algorithms were used to estimate which of the average response vectors (1 for each stimulus) most closely matched the vector of cell responses being produced by a test stimulus. The ED algorithm was found to be more powerful and appropriate given the low firing rates of hippocampal neurons than decoding methods based on Gaussian or Poisson firing rate distributions. In addition, it was found that with another neurally plausible algorithm (the DP algorithm) that calculates which average response vector the neuronal response vector was closest to by performing a normalized dot product (equivalent to measuring the angle between the test and the average response vector), the same generic results were obtained with similar percent correct and only a 15–20% reduction in information compared with the more efficient (ED) algorithm. This is an indication that the brain could use the exponentially increasing capacity for encoding stimuli as the number of neurons in the population increases. The details of the decoding that may be used by actual neurons do matter but in a quantitatively minor way (both the ED and DP algorithm require an estimate of the Euclidean ''length'' of the firing rate vector, an operation that could be performed by feedforward inhibition, but then use this quantity in slightly different ways). For example, in an autoassociative memory [which we believe may be implemented in the hippocampus (see Rolls 1989; Treves and Rolls 1994)], which computes effectively the dot product on each neuron between the input vector and the synaptic weight vector, most of the information available would in fact be extracted (see Rolls and Treves 1990, 1998; Treves and Rolls 1991).

The new procedure for information measurement contained in a population of cells developed for use with the low firing rates and sparse firing rate distributions described here, which calculates the information directly from binarized neuronal response vectors, confirmed with precision that the information available did increase linearly with the number of cells in an ensemble. The algorithm can be used when there are as many trials as there are actual response vectors, in practice up to about five cells. It is very helpful from a methodological point of view, because it allows independent confirmation of the operation of the decoding procedures used in the other algorithms.

One of the important points made here is that because the representational capacity of a set of neurons increases exponentially, neurons in the next brain region would each need to sample the activity of only a reasonable number (e.g., a few hundred) of what might be a much larger cell population and yet still obtain information about which of many stimuli (e.g., locations in space) had been seen. This would be useful for recall of information from the hippocampus via backprojection pathways to the neocortex (see Treves and Rolls 1994).

Comparison of the results shown in Figs. 7 and 8 (with 100- and 25-ms periods, respectively) highlight the value of having large numbers of neurons of the type described here. They make it clear that part of the value is that information can be made available very rapidly about which stimulus is present if the responses of a population of neurons, rather than just a single neuron, are considered. Moreover, the fact that the representation provided by each neuron is apparently independent to that provided by other neurons means that the information is available very rapidly from whichever subset of neurons is taken. This rapid availability of information from a ''population'' of neurons is one factor that contributes to the very rapid processing of information in the brain, for even in a short time much information is available from the population, allowing the information from one cortical area to be extracted very rapidly by the next (see further Rolls 1994; Rolls and Tovee 1994; Rolls and Treves 1998).

A point that certainly merits further investigation is the effect of generating pseudosimultaneous trials (as performed here), rather than recording simultaneously from large populations of cells (Wilson and McNaughton 1993). Particularly in exploring fine points such as the presence of trial-to-trial correlations in the responses, it is helpful to have some evidence about simultaneously recorded cells in the primate hippocampus, to check, for example that the simultaneously recorded cells do convey independent information, consistent with the linear increase in information with the number of cells in the ensemble described here. In fact, we do have preliminary evidence that this is indeed the case in the primate hippocampus, in that six of the cells described here were recorded as three pairs and with this simultaneous recording still conveyed information that was largely independent. In particular, with simultaneous recording the information increased linearly with the number of cells as found for the nonsimultaneous recordings. Further, the information values obtained from the three cell pairs when they were analyzed as simultaneously recorded were on average 4% more than the values when they were treated as not being simultaneously recorded. (The treatment for nonsimultaneous analysis involved simply randomly shuffling the order of the trials for each stimulus.) The redundancy (see Rolls and Treves 1998) was on average 0% for the simultaneously recorded analysis and 4% for the nonsimultaneously recorded analysis. The result indicates that the cells do carry almost independent information. We are continuing with simultaneous recordings and will provide a full report on simultaneously recorded cells in the primate hippocampus in future. However, we note further evidence that the conclusion described here is reasonable, in that in analyses of cells recorded simultaneously in the rat, the information provided by different hippocampal cells is also independent, given that shuffling the rat data to produce nonsimultaneously recorded virtual trials makes little difference to the information analyses (Treves et al. 1996).

These experiments also showed that the representation provided by these hippocampal neurons, is very sparse, with $a_{25} = 0.04$. Twenty-five milliseconds is the order of the time scale of the time constants of synaptic transmission. If only 0 or 1 spike was produced by a neuron in this time period, then we could treat the neuronal system as a network of binary neurons rather than as one with graded firing rates. The probability distribution of different numbers of spikes shown in Figs. 5 and 6 indicate that the neurons very rarely produce more than one spike in 100 ms (and do so with the very low probability of 0.01 in a time period of 25 ms), so that on the time scale of operation of these neurons in the hippocampus, it may be appropriate to consider them as binary variables. Now, to maximize the number of memories stored in an autoassociative attractor neural network such as that which could be implemented by the hippocampal CA3 neurons, it can help to have sparse and binary representations (Rolls 1989, 1995; Rolls et al. 1997d; Treves and Rolls 1991, 1994). However, we note that the time scale of the operation of the synaptic modification involved in learning may be considerably longer, on the order of 100 ms or more, partly because of effects such as the relatively slow unbinding of glutamate from the *N*-methyl-D-aspartate receptor. Another factor lengthening the time scale may be the behavior, in that the animal may process the data for times on the order of ≥ 0.5 s, for example, by looking for ≥ 0.5 s at a location in space where an object is present. This may mean that the actual sparseness of the firing relevant to the synaptic representation laid down for a memory may be more like the value a_{500}, which was on average 0.19. Indeed, the storage capacity in depending on the synaptic matrix and not the instantaneous firing rates of the neurons is likely to reflect this value of the sparseness or, even more, the value of a shown in Table 1, which is 0.22 for these spatial neurons considered alone. Effectively, we interpret the attractors as set up by the learning that might occur over the order of ≥ 0.5 s, so that what is important is the sparseness a rather than the sparseness of the neuronal spikes arriving over 25-ms periods. The network as a whole, when operating as an attractor network as has been suggested for CA3, then would be working in the ''low firing rate regime'' and with rather sparse representations (Rolls and Treves 1998).

The representation in the hippocampus may be more sparse than that in the temporal visual cortical areas where values of 0.6 are common (Rolls and Tovee 1995). This may allow more information to be represented in the pattern of firing of temporal cortical visual neurons than in hippocampal neurons. It has been suggested that this difference in the type of coding is that the more distributed encoding in the visual cortex allows much information to be represented about what is being seen and that the more sparse binary encoding in the hippocampus allows many memories to be stored at the cost of less information per memory than would be possible with a more distributed representation. Indeed, the amount of information present in a hippocampal memory now can be estimated. If each CA3 spatial neuron represents on average 0.3 bits of information about spatial location in 500 ms, if (conservatively) 5% of hippocampal CA3 neurons represented spatial information, and if the neurons are tested in a sufficiently large spatial world with the neurons coding nonredundantly (see further Rolls et al. 1997b), then the information about spatial location in any one hippocampal memory in 1,000,000 CA3 neurons might be as high as $5\% * 1{,}000{,}000 * 0.3$ bits = 15,000 bits. If a neuron downstream had access to the outputs of, say, 1,000 of these CA3 cells, it would ''see'' then 15 bits of spatial information, which is still a considerable amount (they allow precise discriminations to be made between $2^{15} \approx 30{,}000$ locations).

Many spatial view (or ''space'' or ''view'') neurons have

been found in this series of experiments in the locomoting monkey (for a description of 40 spatial view cells, see Rolls et al. 1997a). No place cells have been found that responded based on where the monkey was, as contrasted with where it was looking in the environment. These cells in the primate hippocampus are thus unlike place cells found in the rat (Muller et al. 1991; O'Keefe 1979). Primates, with their highly developed visual and eye movement control systems, can explore and remember information about what is present at places in the environment without having to visit those places. Such view cells in primates thus would be useful as part of a memory system in that they would provide a representation of a part of space that would not depend on exactly where the monkey was and that could be associated with items that might be present in those spatial locations. An example of the use of such a representation in monkeys might be in enabling a monkey to remember where it had seen ripe fruit; in humans, it might help in remembering where they had seen a person or where they had left keys. The representations of space provided by hippocampal view-responsive neurons thus may be useful in forming memories of spatial environments (for example of where an object such as ripe fruit has been seen).

The representation of space in the rat hippocampus, which is of the place where the rat is, 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 thus may tend to reflect the place where the rat is. Although the representation of space in rats therefore may be in some ways analogous to the representation of space in the primate hippocampus, the difference does have implications for theories of hippocampal function. In rats, the presence of place cells has led to theories that the rat hippocampus is a spatial cognitive map and can perform spatial computations to implement navigation through spatial environments (Burgess et al. 1994; O'Keefe 1991; O'Keefe and Nadel 1978). The details of such navigational theories could not apply in any direct way to what is found in the primate hippocampus. Instead, what is applicable to both the primate and rat hippocampal recordings is that hippocampal neurons contain a representation of space (for the rat, primarily where the rat is, and for the primate, primarily of positions ''out there'' in space) that is a suitable representation for the spatial information encoded in an episodic memory system. In primates, this would enable one to remember for example where an object was seen. In rats, it might enable memories to be formed of where particular objects (e.g., defined e.g., by olfactory, tactile, and taste inputs) were found. Thus, at least in primates and possibly also in rats, the neuronal representation of space in the primate hippocampus may be appropriate for forming memories of events (which usually have a spatial component). Such memories would be useful for spatial navigation, for which according to the present hypothesis the hippocampus would hit the memory component but not the spatial computation component.

Finally, the spatial representation found would be ideal for association within the hippocampus to a representation of an object to implement an episodic memory. We have suggested that such an episodic memory could be laid down in the hippocampus using the neuronal network process of association, implemented by the recurrent collateral axons of the CA3 neurons (see Rolls 1989, 1995; Rolls and Treves 1998; Treves and Rolls 1994).

This research was supported by Medical Research Council Grant PG8513579 and the Human Frontier Science Program. P. Georges-François is supported by European Community Marie Curie Research Training Grant ERBFMBICT961279.

Present address of A. Treves: S.I.S.S.A.-Programme in Neuroscience, Via Beirut 2-4, 34013 Trieste, Italy.

Address reprint requests to E. T. Rolls.

Received 15 July 1997; accepted in final form 19 December 1997.

REFERENCES

ANGELI, S. J., MURRAY, E. A., AND MISHKIN, M. Hippocampectomized monkeys can remember one place but not two. *Neuropsychologia* 31: 1021–1030, 1993.

BURGESS, N., RECCE, M., AND O'KEEFE, J. A model of hippocampal function. *Neural Networks* 7: 1065–1081, 1994.

FEIGENBAUM, J. D. AND ROLLS, E. T. Allocentric and egocentric spatial information processing in the hippocampal formation of the behaving primate. *Psychobiology* 19: 21–40, 1991.

FOSTER, T. C., CASTRO, C., A., AND MCNAUGHTON, B. L. Spatial selectivity of rat hippocampal neurons: dependence on preparedness for movement. *Science* 244: 1580–1582, 1989.

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

GEORGES-FRANÇOIS, P., ROLLS, E. T., AND ROBERTSON, R. G. Spatial view cells in the primate hippocampus: allocentric view not head direction or eye position. *J. Neurosci.* 1998 In press.

GOLOMB, D., HERTZ, J. A., PANZERI, S., RICHMOND, B. J., AND TREVES, A. How well can we estimate the information carried in neuronal responses from limited samples? *Neural Comput.* 9: 649–665, 1997.

MULLER, R. U., KUBIE, J. L., BOSTOCK, E. M., TAUBE, J. S., AND QUIRK, G. J. Spatial firing correlates of neurons in the hippocampal formation of freely moving rats. In: *Brain and Space*, edited by J. Paillard. Oxford: Oxford, 1991, p. 296–333.

O'KEEFE, J. A review of the hippocampal place cells. *Prog. Neurobiol.* 13: 419–439, 1979.

O'KEEFE, J. The hippocampal cognitive map and navigational strategies. In: *Brain and Space*, edited by J. Paillard. Oxford, UK; Oxford Univ. Press, 1991, p. 273–295.

O'KEEFE, J. AND NADEL, L. *The Hippocampus as a Cognitive Map*. Oxford, UK: Clarendon Press, 1978.

O'KEEFE, J. AND SPEAKMAN, A. Single unit activity in the rat hippocampus during a spatial memory task. *Exp. Brain Res.* 68: 1–27, 1987.

ONO, T., TAMURA, R., NISHIJO, H., AND NAKAMURA, K. Neural mechanisms of recognition and memory in the limbic system. In: *Brain Mechanisms of Perception and Memory: From Neuron to Behavior*, edited by T. Ono, L. R. Squire, M. E. Raichle, D. I. Perrett, and M. Fukuda. New York: Oxford Univ. Press, 1993, p. 330–355.

OPTICAN, L. M. AND RICHMOND, B. J. Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. III. Information theoretic analysis. *J. Neurophysiol.* 57: 162–178, 1987.

PANZERI, S. AND TREVES, A. Analytical estimates of limited sampling biases in different information measures. *Network* 7: 87–107, 1996.

PARKINSON, J. K., MURRAY, E. A., AND MISHKIN, M. A selective mnemonic role for the hippocampus in monkeys: memory for the location of objects. *J. Neurosci.* 8: 4059–4167, 1988.

RICHMOND, B. J. AND OPTICAN, L. Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. II. Quantification of response waveform. *J. Neurophysiol.* 57: 147–161, 1987.

ROBERTSON, R. G., ROLLS, E. T., AND GEORGES-FRANÇOIS, P. Spatial view cells in the primate hippocampus: effects of removal of view details. *J. Neurophysiol.* 79: 1145–1156, 1998.

ROLLS, E. T. Functions of neuronal networks in the hippocampus and neocortex in memory. In: *Neural Models of Plasticity: Experimental and Theoretical Approaches*, edited by J. H. Byrne and W. O. Berry. San Diego: Academic Press, 1989, p. 240–265.

ROLLS, E. T. Brain mechanisms for invariant visual recognition and learning. *Behav. Processes* 33: 113–138, 1994.

ROLLS, E. T. A model of the operation of the hippocampus and entorhinal cortex in memory. *Int. J. Neural Syst.* 6, *Suppl.*: 51–70, 1995.

ROLLS, E. T. A theory of hippocampal function in memory. *Hippocampus* 6: 601–620, 1996.

ROLLS, E. T., CRITCHLEY, H. D., AND TREVES, A. The representation of olfactory information in the primate orbitofrontal cortex. *J. Neurophysiol.* 75: 1982–1996, 1996.

ROLLS, E. T., MIYASHITA, Y., CAHUSAC, P.M.B., KESNER, R. P., NIKI, H., FEIGENBAUM, J., AND BACH, L. Hippocampal neurons in the monkey with activity related to the place in which a stimulus is shown. *J. Neurosci.* 9: 1835–1845, 1989.

ROLLS, E. T. AND O'MARA, S. M. View-responsive neurons in the primate hippocampal complex. *Hippocampus* 5: 409–424, 1995.

ROLLS, E. T., ROBERTSON, R. G., AND GEORGES-FRANÇOIS, P. Spatial view cells in the primate hippocampus. *Eur. J. Neurosci.* 9: 1789–1794, 1997a.

ROLLS, E. T. AND TOVEE, M. J. Processing speed in the cerebral cortex and the neurophysiology of visual masking. *Proc. R. Soc. Lond. B Biol. Sci.* 257: 9–15, 1994.

ROLLS, E. T. AND TOVEE, M. J. Sparseness of the neuronal representation of stimuli in the primate temporal visual cortex. *J. Neurophysiol.* 73: 713–726, 1995.

ROLLS, E. T. AND TREVES, A. The relative advantages of sparse versus distributed encoding for associative neuronal networks in the brain. *Network* 1: 407–421, 1990.

ROLLS, E. T. AND TREVES, A. *Neural Networks and Brain Function*. Oxford: Oxford, 1998.

ROLLS, E. T., TREVES, A., AND TOVEE, M. J. The representational capacity of the distributed encoding of information provided by populations of neurons in the primate temporal visual cortex. *Exp. Brain Res.* 114: 149–162, 1997b.

ROLLS, E. T., TREVES, A., TOVEE, M., AND PANZERI, S. Information in the neuronal representation of individual stimuli in the primate temporal visual cortex. *J. Comput. Neurosci.* 4: 309–333, 1997c.

ROLLS, E. T., TREVES, A., FOSTER, D., AND PEREZ-VICENTE, C. Simulation studies of the CA3 hippocampal subfield modelled as an attractor neural network. *Neural Networks* 10: 1559–1569, 1977d.

SCOVILLE, W. B. AND MILNER, B. Loss of recent memory after bilateral hippocampal lesions. *J. Neurol. Neurosurg. Psychiatry* 20: 11–21, 1957.

SMITH, M. L. AND MILNER, B. The role of the right hippocampus in the recall of spatial location. *Neuropsychologia* 19: 781–793, 1981.

SQUIRE, L. R. AND KNOWLTON, B. J. Memory, hippocampus, and brain systems. In: *The Cognitive Neurosciences*, edited by M. Gazzaniga. Cambridge, MA: MIT Press, 1994, p. 825–837.

TOVEE, M. J. AND ROLLS, E. T. Information encoding in short firing rate epochs by single neurons in the primate temporal visual cortex. *Visual Cognit.* 2: 35–58, 1995.

TOVEE, M. J., ROLLS, E. T., TREVES, A., AND BELLIS, R. P. Information encoding and the responses of single neurons in the primate temporal visual cortex. *J. Neurophysiol.* 70: 640–654, 1993.

TREVES, A. Graded-response neurons and information encodings in autoassociative memories. *Phys. Rev. A* 42: 2418–2430, 1990.

TREVES, A. Mean-field analysis of neuronal spike dynamics. *Network* 4: 259–284, 1993.

TREVES, A. AND PANZERI, S. The upward bias in measures of information derived from limited data samples. *Neural Comput.* 7: 399–407, 1995.

TREVES, A. AND ROLLS, E. T. What determines the capacity of autoassociative memories in the brain? *Network* 2: 371–397, 1991.

TREVES, A. AND ROLLS, E. T. Computational constraints suggest the need for two distinct input systems to the hippocampal CA3 network. *Hippocampus* 2: 189–199, 1992.

TREVES, A. AND ROLLS, E. T. A computational analysis of the role of the hippocampus in memory. *Hippocampus* 4: 374–391, 1994.

TREVES, A., SKAGGS, W. E., AND BARNES, C. A. How much of the hippocampus can be explained by functional constraints? *Hippocampus* 6: 666–674, 1996.

WILSON, M. A. AND MCNAUGHTON, B. L. Dynamics of the hippocampal ensemble code for space. *Science* 261: 1055–1058, 1993.