

Attention-Like Modulation of Hippocampus Place Cell Discharge

André A. Fenton,^{1,3,4} William W. Lytton,^{1,2,3} Jeremy M. Barry,¹ Pierre-Pascal Lenck-Santini,^{1,5} Larissa E. Zinyuk,⁴ Štěpán Kubík,⁴ Jan Bureš,⁴ Bruno Poucet,⁵ Robert U. Müller,^{1,3} and Andrey V. Olypher^{1,4,6}

Departments of ¹Physiology and Pharmacology and ²Neurology and ³The Robert F. Furchgott Center for Neural and Behavioral Science, State University of New York, Downstate Medical Center, New York, New York 11203, ⁴Institute of Physiology, Academy of Sciences of the Czech Republic, CZ-142 20 Prague 4, Czech Republic, ⁵Laboratory of Neurobiology and Cognition, Centre National de la Recherche Scientifique and Université de Provence, 13331 Marseille cedex 3, France, and ⁶Department of Biology, Emory University, Atlanta, Georgia 30322

Hippocampus place cell discharge is an important model system for understanding cognition, but evidence is missing that the place code is under the kind of dynamic attentional control characterized in primates as selective activation of one neural representation and suppression of another, competing representation. We investigated the apparent noise (“overdispersion”) in the CA1 place code, hypothesizing that overdispersion results from discharge fluctuations as spatial attention alternates between distal cues and local/self-motion cues. The hypothesis predicts that: (1) preferential use of distal cues will decrease overdispersion; (2) global, attention-like states can be decoded from ensemble discharge such that both the discharge rates and the spatial firing patterns of individual cells will be distinct in the two states; (3) identifying attention-like states improves reconstructions of the rat’s path from ensemble discharge. These predictions were confirmed, implying that a covert, dynamic attention-like process modulates discharge on a ~ 1 s time scale. We conclude the hippocampus place code is a dynamic representation of the spatial information in the immediate focus of attention.

Introduction

The ensemble firing of rodent hippocampal place cells signals the animal’s current location and aspects of its locomotor trajectory (Wilson and McNaughton, 1993). For these reasons, the hippocampus is believed to form a cognitive map of space that reflects environmental memories (Lever et al., 2002; Leutgeb et al., 2004; J. K. Leutgeb et al., 2005; Wills et al., 2005) and is engaged in navigation (Hollup et al., 2001; Lenck-Santini et al., 2002; Kubie et al., 2007). Although the map (i.e., the ensemble place code) can be drastically modified when the animal is in a different environment or context (O’Keefe and Nadel, 1978; Müller and Kubie, 1987; Müller et al., 1991; Markus et al., 1995; S. Leutgeb et al., 2005) it is apparently stable for long intervals when the rat performs a single task within a single environment. Thus, although the cognitive map seems to be a valuable model for human cognition, there is an important discrepancy: In contrast to the

presumed map stability, human cognition is dynamic and information active in mind changes from moment to moment, with the focus of attention (Spiers and Maguire, 2006, 2007). In fact, a major advance in understanding the impact of attention on neural function was the discovery that when an extrastriate visual neuron responds to attended receptive field stimuli, responses to unattended stimuli are concurrently suppressed (Moran and Desimone, 1985; Spitzer et al., 1988).

The hippocampal neural code for location is an ensemble place code (Wilson and McNaughton, 1993; Fenton et al., 2008) and although dynamical ensemble responses have been described (Gothard et al., 1996; Redish et al., 2000; Jackson and Redish, 2007; Johnson and Redish, 2007) they have not been associated with the attention-like modulation present in higher visual areas of the primate. Thus, a fundamental question is whether the hippocampal map statically reflects spatial information or whether it dynamically reflects changes of information that are the immediate foci of attention.

Our basic contention is that place cell activity can be modulated on an ~ 1 s time scale by a behaviorally silent attention-like process (Olypher et al., 2002). A clue in this direction is that place cell discharge may be unreliable, often far faster or slower than predicted by the average location-specific rate (Fenton and Müller, 1998; Jackson and Redish, 2007) (Fig. 1A). The statistics of this unreliability (“overdispersion”, Fig. 1) fit the idea that place cell discharge switches between two states as the rat’s attention switches between a representation of position relative to distal cues and one relative to local and self-motion cues (Olypher et al., 2002) (Fig. 1B). This dynamic-attention hypothesis makes three predictions. (1) Preferential use of one cue class should decrease overdispersion. (2) Ensemble discharge as a function of time is

Received Nov. 10, 2009; revised Feb. 2, 2010; accepted Feb. 8, 2010.

This work was supported by National Science Foundation Grant IOS-0725001 to A.A.F., National Institutes of Health Grants MH082417 to A.A.F., NS045612 to W.W.L., and NS20686 to R.U.M., Grant AV0250110509 to the Institute of Physiology, Academy of Sciences of the Czech Republic, and an award to S.K. from the foundation Nadání Josefa, Marie a Zdeňky Hlávkových. We are grateful to Dr. A. D. Redish for valuable discussions and comments on an early draft of this manuscript.

Correspondence should be addressed to André A. Fenton, Department of Physiology & Pharmacology, The Robert F. Furchgott Center for Neural and Behavioral Science, State University of New York, Downstate Medical Center, 450 Clarkson Ave., Brooklyn, NY. E-mail: afenton@downstate.edu.

A. A. Fenton’s present address: Center for Neural Science, New York University, New York, NY 10003.

P. P. Lenck-Santini’s present address: Dartmouth Medical School, Lebanon, NH 03755.

L. E. Zinyuk’s present address: Bristol University, Department of Physiology and Pharmacology, School of Medical Sciences, MRC Centre for Synaptic Plasticity, Bristol BS8 1TD, UK.

DOI:10.1523/JNEUROSCI.5576-09.2010

Copyright © 2010 the authors 0270-6474/10/304613-13\$15.00/0

separable into two states in which the spatial firing patterns and rates of individual cells differ substantially. (3) State-specific spike trains predict current location better than observed spike trains. We provide evidence for each prediction and conclude that place cell activity is modulated by a behaviorally-silent attention-like process that spontaneously switches between distinct ensemble place codes. We conclude that the hippocampal code dynamically represents spatial information in the immediate focus of attention.

Some of these data have been reported in abstract form (Fenton et al., 2000a; Lytton et al., 2007).

Materials and Methods

Electrophysiology. Hippocampal place cells were recorded from the CA1 subfield using standard techniques that have been described (Fenton et al., 2000b, 2008). All procedures received institutional approval and conformed with National Institutes of Health and European Union guidelines for the care and protection of animals. Briefly, adult male Long-Evans rats (350–450 g) were implanted with 8 or 32 wire electrodes made from 25 μm Formvar-insulated Nichrome wires. The wires were configured as tetrodes for the recordings of large hippocampal ensembles. The electrodes were mounted in a microdrive for implantation.

The rat was deeply anesthetized for electrode implantation (Nembutal >50 mg/kg, i.p.). The animal was then mounted in a stereotaxic frame, a mid-line incision was made in the scalp, and a trephine hole was made in the skull above the dorsal hippocampus (posterior 3.8 mm and lateral 2.5 mm to bregma). The electrodes were lowered to ~1 mm above the CA1 pyramidal cell layer (ventral 2.5 mm to bregma). The microdrive assembly was fixed to the skull by 4–6 bone screws and dental cement.

One to 2 weeks after surgery, screening for single unit activity began. The microdrive was connected to a buffering preamplifier and the signals were passed along a counterbalanced light-weight cable to main amplifiers. The signals were filtered (300 Hz–6 kHz), amplified to the order of a volt, digitized (30 kHz) and stored for standard waveform discrimination offline. The electrodes were advanced in small (<25 μm) steps until stable single unit activity could be detected. During recordings, action potential activity was recorded in parallel with the position extracted from the television signal of an overhead camera.

Behavioral conditions. To test the prediction that forcing rats to attend to one of two stimulus subsets will reduce overdispersion we performed a data-mining study. We measured the discharge variance of place cells in groups of hungry rats trained in four different circumstances (see Table 1; Fig. 2A). The task for groups 1 (forage/stable) and 2 (forage/variable) was to find and eat food pellets randomly scattered into cylindrical apparatuses. There were two major differences between these groups: (1) The visual environment. The cylinder for forage/stable rats was gray and was polarized by a single white cue card that occupied 100° of arc on the

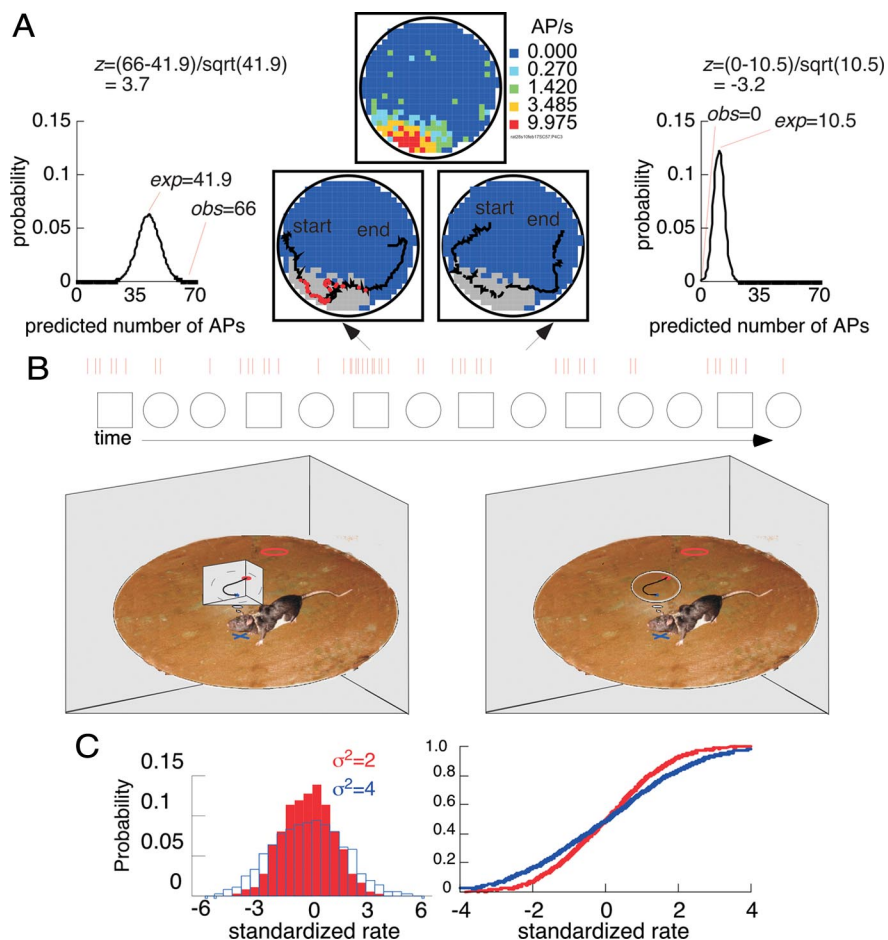


Figure 1. Overdispersion and the dynamic-attention hypothesis. **A**, Example of the unreliability of location-specific place cell discharge. The color-coded firing rate map for a place cell is shown (upper map) along with two 10 s paths that passed through the firing field (lower maps). The cell fired 66 action potentials (APs) along the pass depicted on the left and zero APs along the pass on the right. Given the firing rate map, the likelihood of the two observations can be quantified by calculating a standardized firing rate for each pass (z). This calculation is depicted graphically using the histograms that correspond to each pass. On the assumption that the firing rate map describes the average rate in a location, integrating the rates along the path generates an expected number of APs during the pass. The expectation was 41.9 APs for the pass on the left and 10.5 APs for the pass on the right. The corresponding Poisson probability distribution of the number of expected APs for each pass is also shown. The standardized rate for a single pass is the normalized difference between the observed (obs) and expected (exp) activity, a standard normal deviate (z). Overdispersion is characterized by the variance of the distribution of standardized firing rates for a large number of passes. **B**, Diagram depiction of the dynamic-attention hypothesis that accounts for the place cell example in **A**. When the rat is at a location it chooses to attend to either distal cues (left) or local cues (right) to locate itself. The location-specific firing rate of each cell is modulated by the attentional state. In the raster of the diagrammed cell, the cell fires more when the rat attends to distal cues (squares) and less when the rat attends to local cues (circles). According to this model, despite being at the same locations, the rat was in the attentional state associated with higher firing rate during the pass on the left and during the pass on the right the rat was in a different attentional state, one associated with lower firing rates. **C**, Example of how differences in place cell overdispersion will be compared. Two distributions of 1000 randomly-generated numbers with the averages equal to zero are depicted as histograms (left) and cumulative probability functions (right). The distributions only differ in that the variance of one distribution is 2 and the other is 4. It is easy to see that the distribution with the lower variance has the steeper cumulative probability function. The distributions intersect at probability = 50% because their medians are the same. The dynamic-attention hypothesis predicts that training rats to focus attention on one cue subset will reduce the variance of the standardized firing rates.

wall. The cylinder for forage/variable rats was uniformly black and polarizing information was provided by multiple stimuli visible above the top rim of the cylinder. (2) Rotational state. The cylinder was always stationary for forage/stable rats whereas for forage/variable rats the cylinder and floor were rotated at 1 rpm during alternating recording sessions but were stationary during interleaved sessions. Only results from stationary sessions were analyzed.

Groups 3 (navigate/stable) and 4 (navigate/variable) were trained in the place preference task (Bures et al., 1997b; Rossier et al., 2000). In this task, rats are required to go to and stay in a small circular region for >1 s to release a pellet that can subsequently be found and eaten. The visual

Table 1. Description of the behavioral conditions, data used for data-mining, and summary of overdispersion results

Group	Cylinder	Rotation	Number of passes with $\text{exp} \geq 5$	Number of rats (cells)	Overdispersion (σ^2)
Forage/stable (Save et al., 2000)	Gray	Stable	1329	10 (37)	5.5
Forage/variable (Bures et al., 1997a; Zinyuk et al., 2000)	Black	Stable/rotating	1147	7 (25)	5.2
Navigate/stable (Lenck-Santini et al., 2002)	Gray	Stable	2245	4 (42)	3.5
Navigate/variable (Zinyuk et al., 2000; Fenton et al., 2002)	Black	Stable/rotating	2474	4 (61)	2.1

The citation indicates the publication in which the recording experiment was reported. exp is the expected number of spikes along a pass.

environment and lack of cylinder rotation were the same for navigate/stable as for forage/stable rats. Similarly, the visual environment and cylinder rotation were the same for navigate/variable as for forage/variable rats. Again, only data from stationary sessions were used.

The four groups were selected to separate the effects of navigation and attention on overdispersion. The influence of attention was estimated by comparing overdispersion in the navigate/stable and navigate/variable groups. Although data from the navigate groups were collected in the same conditions, rats in the navigate/stable group could always identify the navigation goal using either local or distal cues, whereas rats in the navigate/variable were conditioned by arena rotation to preferentially use stationary distal cues to locate the goal. Since conditions for rats in the navigate/stable and navigate/variable groups differed according to intermittent rotation we also characterized overdispersion in forage/stable and the forage/variable groups that were never trained to navigate to a specific goal location.

Quantifying overdispersion. Place cells were selected for study if their spatial firing patterns were location-specific (coherence >0.25 ; spatial information >0.5 bits/AP). To calculate overdispersion we used the time-averaged firing rate map of a cell to identify episodes for analysis; the selection was done in two different ways.

In the first method, the entire session was divided into 5 s intervals. For each interval we calculated the expected number of spikes, exp , as:

$$\text{exp} = \sum r_i t_i, \quad (1)$$

where r_i is the time-averaged rate at location i , and t_i is the time spent in location i during the pass. Only intervals during which $\text{exp} \geq 5.0$ AP were used to calculate overdispersion since the overall firing rate of place cells is ~ 1.0 AP/s.

For each selected 5 s interval, we then calculated z , the normalized standard deviation of obs, the observed number of spikes as:

$$z = \frac{(\text{obs} - \text{exp})}{\sqrt{\text{exp}}}. \quad (2)$$

z measures the deviation of observed discharge from expected in standard deviation units. Overdispersion in turn is the variance of the z distribution for a set of passes. The outcome of this calculation was found to be indistinguishable from the somewhat different method previously used (Fenton and Muller, 1998).

In the second method, we compared overdispersion during two kinds of searching in the navigating rats. When rats were observed to approach the pellet release zone, we took the 5 s before arrival to be goal-directed searching. In contrast, when a pellet was successfully released to scatter into the cylinder, we took the 5 s after departure from the goal zone to be undirected searching. If an approach or departure trajectory did not go through the firing field it was discarded since $\text{exp} = 0$ and z was undefined. We again restricted analysis to passes for which $\text{exp} \geq 5$ since such passes virtually always went through the middle of the firing field and since 1.0 AP/s is a reasonable estimate of the overall discharge rate for place cells. Using other rate thresholds from >0 AP/s to 5 AP/s altered the numbers of passes selected for study but did not change the conclusions.

Decoding two location-specific states. To test the second and third predictions of the dynamic-attention hypothesis, we used a heuristic algorithm to organize the spike trains of ensembles of simultaneously recorded place cells into time series that switched between two mutually exclusive states. The activity patterns of three large place cell ensembles were recorded from different rats that had only ever been trained in a fixed environment to forage for scattered food in a 76 cm diameter gray

cylinder with a white cue card ($n = 53$, 23 cells) or a 76 cm square white box with a 45 cm wide black and white striped card on one wall ($n = 54$ cells). These three ensembles were studied with the NEURON Neural Query System (Lytton, 2006) using the cDNA microarray C-Clustering Library (de Hoon et al., 2004). The first four steps of the algorithm are depicted in Figure 4A (see below). First, ensemble activity during each time step was defined as an “activity vector” of spike counts whose dimension was equal to the number of cells in the ensemble and whose elements were the number of spikes fired by a cell during a time step. Step 2 associated each activity vector with the rat’s location during the time step; if the rat moved, its position was taken as the average position weighted by the time spent in each square region. Step 3 collected all the activity vectors that occurred at a location, thereby establishing the distribution of activity vectors for each location in the environment. Step 4 partitioned the activity vector distribution at each location into 2 categories using a k -medians algorithm based on the Pearson correlation between pairs of vectors. Specifically, the partitioning at each location maximized vector correlations within a category and minimized vector correlations between the categories. This produced two groups of activity vectors at each location such that the correlation of vectors within a group was maximized and the correlation of vectors between the groups was minimized. To test the separation algorithm, activity vectors were shuffled and pseudorandomly assigned to one of two categories at each location, preserving the number of vectors in each category obtained with the k -medians partitioning procedure. Using the k -means procedure, as in Jackson and Redish (Jackson and Redish, 2007) produced indistinguishable results from the k -medians procedure. Examples of activity vector partitioning are given in Figure 4B (see below). The quality of group separation at each location was quantified as the ratio of average distance between vector pairs within states and the distance between the pair of average group vectors according to:

SQ = 1

$$= \frac{\frac{2}{N(N-1)} \sum_{1 \leq i < j \leq N} |x_i - x_j| + \frac{2}{M(M-1)} \sum_{1 \leq i < j \leq M} |w_i - w_j|}{\left| \frac{1}{N} \sum_{i=1}^N x_i - \frac{1}{M} \sum_{i=1}^M w_i \right|}, \quad (3)$$

where $\{x_i\}_{i=1}^N$ are the N vectors of one state and $\{w_i\}_{i=1}^M$ are the M vectors of the other state, and $||$ is the Euclidean distance. The term on the right was subtracted from 1 so that SQ = 1 corresponds to ideal separation.

Decoding two global states. Once the activity vectors were partitioned into two groups at each location, they were sorted into two global states as illustrated in Figure 5A (see below). Starting at the location with the best group separation, one group was assigned to state A and the other group to state B. Next, the groups in the nearest neighbor with the best separation were assigned to either state A or state B depending on the correlations between the average group vectors in the two locations. Specifically, the assignment was done so that groups of the same state in adjacent locations had large intergroup correlations and groups from different states had small intergroup correlations. This process was carried out iteratively to cover all locations in the environment. Finally, once the two global states were established the individual activity vectors for each state were put into time order to create two state-specific time series of discharge activity. To test if state separation had systematic effects, a randomized control pair of state-specific time series was generated by shuffling the state identities of each vector group at each location, while preserving the vector grouping at each location.

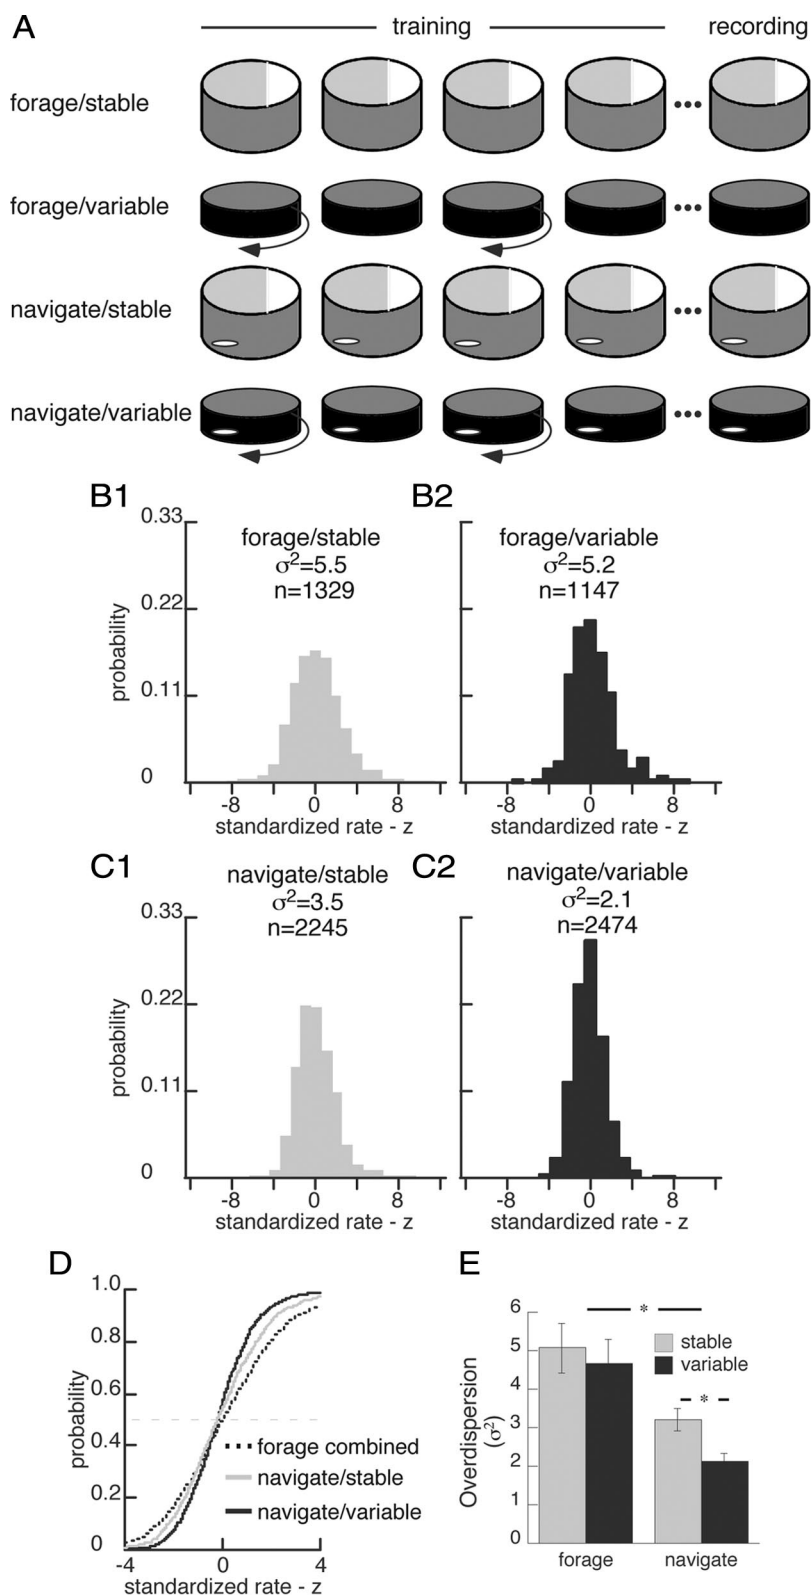


Figure 2. Prediction 1. Navigation training reduced overdispersion and navigation training to use the distal cue subset reduced overdispersion even more. **A**, Schematic description of the behavioral conditions of the four groups of rats during the training and recording phases of the experiment. Group 1: Rats in the forage/stable group foraged for sporadically scattered food in a gray cylinder that was surrounded by curtains. A white card on the cylinder wall provided a polarizing landmark. The recording conditions were identical to the training conditions. Group 2: Rats in the forage/variable group foraged for scattered food in a black cylinder with a view of the surrounding room. The cylinder was stable on approximately half the training trials and it rotated at 1 rpm on the other trials. Only recordings during the stable trials were analyzed. Group 3: Rats in the navigate/stable group were also trained in the gray cylinder with a white card and surrounding curtains. The conditions were essentially the same as for the

To estimate the importance of selecting specific temporal and spatial resolutions for determining the initial activity vectors, the entire process was repeated for different time steps ranging from 400 to 1600 ms and for square regions ranging from 2.5 to 20 cm on a side. In each case, a location was kept in the analysis only if included at least 4 activity vectors. Numerical evaluation of the results (see Results) indicate that the process is not sensitive to the temporal and spatial resolution ranges, suggesting that the procedure is robust (see Fig. 7B).

State-specific comparisons. State-specific activity vectors were calculated for every 20 s interval during a recording session by selecting all the same-state activity vectors (based on a shorter time step) that occurred during the 20 s interval. The similarity for each pair of 20 s state-specific activity vectors was computed using Pearson's correlation and then displayed as a color-coded matrix (see Fig. 5B).

The time series of states was used to generate two state-specific firing rate maps for each cell. Only time steps assigned to a particular state were used to generate the corresponding state-specific map. For each cell, properties of the two state-specific firing rate maps were compared with the properties of the maps gener-

←

forage/stable rats except these rats performed a place preference task in which the food was only scattered whenever the rat visited an unmarked goal zone (white oval). Group 4: Rats in the navigate/variable group performed the same place preference navigation in the same physical conditions as the forage/variable group. On alternate trials the arena was either stable or rotating, but the navigation goal was always at a fixed location in the room. Although only recordings during the stable trials were analyzed, the rotating sessions discouraged the use of local cues for identifying the goal location because only distal cues could identify the stationary goal location. According to the dynamic attention hypothesis, overdispersion in the stable sessions should be less in the navigate/variable group than in the navigate/stable group even though the overt behavior was the same. This is because rats in the navigate/variable group were conditioned to attend to distal cues and to ignore local cues whereas rats in the navigate/stable group could always use both cue sets to locate the place preference goal. The environments also differed for the navigate/variable and navigate/stable groups, so the effect of the environment itself on overdispersion was evaluated by comparing overdispersion in the forage/stable and forage/variable groups. **B, C**, Overdispersion in the four groups was characterized as the variance (σ^2) of the standardized firing rate along passes through a set of firing fields (n values given). Overdispersion was similar in the two forage groups and less in the navigate groups. **D**, Cumulative probability plots of the standardized firing rate distributions illustrate overdispersion was greatest in the forage groups, moderate in the navigate/stable group, and least in the navigate/variable group. **E**, Overdispersion was also characterized by analyzing the firing variance in the set of passes through an individual cell's firing field. Despite having different numbers of passes, each cell contributed one estimate of overdispersion. This analysis confirmed overdispersion was largest and similar in the two forage groups, moderate in the navigate/stable group and least in the navigate/variable group. (* $p < 0.05$).

ated by the control partitioning in which activity vectors were randomly assigned to each partition.

The state-specific spatial firing rate patterns of an individual cell were compared by calculating the pixel-to-pixel spatial correlation of the two firing rate maps. The correlation is reported, but for statistical comparisons the correlation coefficient (r) was transformed to standard units (z):

$$z = \frac{1}{2} \ln \left(\frac{r+1}{r-1} \right) \quad (4)$$

Correlations between the state-specific firing rate maps and the correlations between the control shuffle-separated firing rate maps were compared by paired t tests.

Quantitative comparisons of the state-specific location-independent firing rate levels were made by calculating the distinction between the overall firing rates in the two states. The rate distinction was defined as:

$$\frac{(\text{rateA} - \text{rateB})}{\max(\text{rateA}, \text{rateB})} \quad (5)$$

where rateA and rateB were the location-independent firing rates in states A and B, respectively. The state-specific rate distinctions were compared with paired t tests to the control rate distinctions created by shuffling activity vectors.

The influence of an individual cell's activity on the state separation was eliminated by using a "leave-one-out" algorithm. A new time series of states was computed using ensemble activity vectors of dimension $n - 1$ that excluded the cell. The time series was then used to define the state-specific firing of the single cell for the quantitative comparisons of state-specific spatial firing patterns and state-specific position-independent firing rates.

Predicting current location. The ensemble activity vector at each time step was used to reconstruct the rat's path using a simple template-matching method (Wilson and McNaughton, 1993; Fenton and Muller, 1998). At each location the firing rate of each cell in the ensemble was used to construct a location-specific template firing rate vector. The predicted current location was the position that maximized the projection of the current firing rate vector onto one of the location-specific template vectors. If there was no activity during a time step no prediction was made. This method was chosen because it explicitly tests how well location-specific firing rate itself predicts the current location (Wilson and McNaughton, 1993; Fenton and Muller, 1998; Fenton et al., 2008). Note that the reconstruction at each time step is independent of the activity and prediction at other time steps. The average distance between the reconstructed and observed locations (in pixels) was used to evaluate the accuracy of the prediction.

Current location was predicted by three ensemble activity time series. The first was the raw time series. This was simply the observed time series, without regard to global states. Two predictions from the global state-specific time series were also evaluated. For both predictions, the state-specific template firing rate vectors at each location were computed from the state-specific time series of place cell activity. Then, for predicting current location, the state at each time step was used to select between the two state-specific maps of template activity vectors. In one set of predictions, the 'correct' state-specific map was always selected. In the other set of predictions, the 'incorrect' state-specific map was always selected. The dynamic-attention hypothesis predicts that using the correct state map should improve reconstruction from the raw, state-independent spike trains. If specific knowledge of the current state itself was responsible for the improvement, then using the incorrect state map should worsen the reconstruction.

Reconstruction was performed at four temporal resolutions (400, 800, 1200, and 1600 ms time steps) and four spatial resolutions (2.5, 5, 10, and 20 cm pixels). The corresponding reconstruction from the raw data was used to compute a standardized reconstruction error. The standardized error was used to compare the accuracy of reconstruction at these different resolutions of time and space:

$$\frac{(\text{state} - \text{raw})}{(\text{rate} + \text{raw})} \quad (6)$$

where state is a state-specific prediction error, and raw is the raw, state-independent prediction error.

Results

Overdispersion in place cell firing

First, we characterized overdispersion in a database of place cell recordings from rats trained in four different circumstances (Table 1). Detailed group descriptions are provided in Methods. Briefly, rats in the navigate/variable group selectively used distal stimuli instead of local and self-motion stimuli for navigating to a goal. They did so because the goal has a fixed relationship to distal cues whereas the apparatus floor sometimes rotated relative to the goal, rendering local and self-motion cues unreliable. In contrast, rats in the navigate/stable group could navigate to the goal using any combination of distal, local and self-motion stimuli because the floor never rotated. There was no constant goal for rats in the forage/variable and forage/stable control groups but the physical conditions were identical to those for the corresponding navigate group. To characterize firing during every pass through each cell's firing field we calculated a standardized rate (z) as the deviation of the observed firing from an expectation derived from the session average (Fenton and Muller, 1998). The variability of firing was then expressed as the variance of the firing deviations for all passes. For each group of rats, the average standardized rate for the observed passes through the firing fields could not be distinguished from zero. In contrast, the variance for each group was reliably >1 . In agreement with earlier observations (Fenton and Muller, 1998; Jackson and Redish, 2007) place cell discharge variance is much greater than would be generated by an inhomogeneous Poisson process and is therefore said to show overdispersion (Olypher et al., 2002).

Preferential use of stimuli reduces overdispersion

Next, we tested the prediction of the dynamic-attention hypothesis that preferential use of one stimulus subset should reduce overdispersion. Several comparisons of standardized firing variance among the four groups are of interest (Fig. 2). First, the variance for forage/stable rats was similar to that for forage/variable rats ($F_{(1328,1146)} = 1.06$; $p = 0.16$), indicating that neither the visual environment nor intermittent rotation affected overdispersion. Moreover, the variance of both forage groups was numerically equal to the variances seen in earlier work (Fenton and Muller, 1998; Jackson and Redish, 2007), bolstering the conclusion that overdispersion is a robust phenomenon.

Critically, the firing variance was considerably lower for navigate/stable than for forage/stable rats ($F_{(1329,2245)} = 1.575$; $p \approx 0$). There was an even larger variance decrease when forage/variable rats were compared with navigate/variable rats ($F_{(1147,2474)} = 2.420$; $p \approx 0$). Thus, requiring rats to solve a spatial problem causes striking reductions of temporal firing variance. This variance reduction may be a form of the firing field stabilization that occurs when animals are required to attend more closely to environmental conditions (Zinyuk et al., 2000; Kentros et al., 2004; Muzzio et al., 2009). The reduction of variance by training rats to navigate to a hidden-goal is independent of the apparatus floor stability; navigation reduced overdispersion when the floor was stationary so that local cues could be useful but also when the floor rotated on half the trials, which trains rats to not use local cues because they are useless for finding the hidden goal (Fenton and Bures, 2003).

We also found that the discharge variance is lower for navigate/variable rats than navigate/stable rats (Fig. 2C; $F_{(2245,2474)} = 1.607$; $p \approx 0$). This is unexpected from the firing variance equality of forage/variable and forage/stable rats; it means that when rats

are required to solve a spatial problem, intermittent rotation has consequences that are not observed during simple pellet retrieval. In our view, this implies that rats continue to attend to local cues during pellet chasing even if they can be unstable on alternate trials, but largely cease to use potentially unstable cues if they might be misleading in getting to a goal. The overall consequence is that the activity of a cell along a path during navigation would more closely resemble the activity expected if only distal cues affected discharge by reducing an attentional source of place cell firing overdispersion. The results of these analyses confirm the first prediction of the dynamic-attention hypothesis that preferentially using one stimulus subset reduces overdispersion.

In addition to computing discharge variance for all passes by each group of rats, we also began by obtaining the variance for individual cells and then took the average across cells within each group. This alternative procedure did not affect any of the conclusions; as seen in Figure 2E, discharge variance is equal for forage/stable and forage/variable rats, is greater for each forage group than for the corresponding navigation group and is lower for navigate/variable than navigate/stable rats.

In a separate analysis, we asked if discharge variance is lower during ‘approaches’ when rats are walking toward the hidden goal than during ‘departures’ when rats leave the hidden goal after successfully releasing a pellet. This was motivated by the possibility that rats might attend more closely to the environment when navigating toward a specific place than at the beginning of a bout of foraging. We saw no difference for navigate/stable rats (approaches: $\sigma^2 = 3.7$; departures: $\sigma^2 = 3.9$; $F_{(357,949)} = 1.05$; $p = 0.26$) but for navigate/variable rats the discharge variance during departures ($\sigma^2 = 1.9$) was just detectably greater than during approaches ($\sigma^2 = 1.7$; $F_{(685,672)} = 1.14$; $p < 0.05$).

The different outcome for navigate/variable and navigate/stable rats is interesting because both groups run faster and in straighter paths during approaches than departures (Fig. 3A) (Rossier et al., 2000). While goal approaches and departures are only a subset of the passes through the firing field, making their comparison have lower statistical power than the group comparisons in Table 1 and Figure 2, the statistical power is nonetheless substantial. We therefore contend that these behavioral differences are not an adequate explanation of the lower overdispersion during goal approaches in the navigate/variable group. We also sought other potential behavioral origins of overdispersion. For instance, we investigated running speed which is known to cause firing rate variations (McNaughton et al., 1983; Czurkó et al., 1999). We confirmed that firing rate increases with running speed. Thus, the correlation between speed and standardized fir-

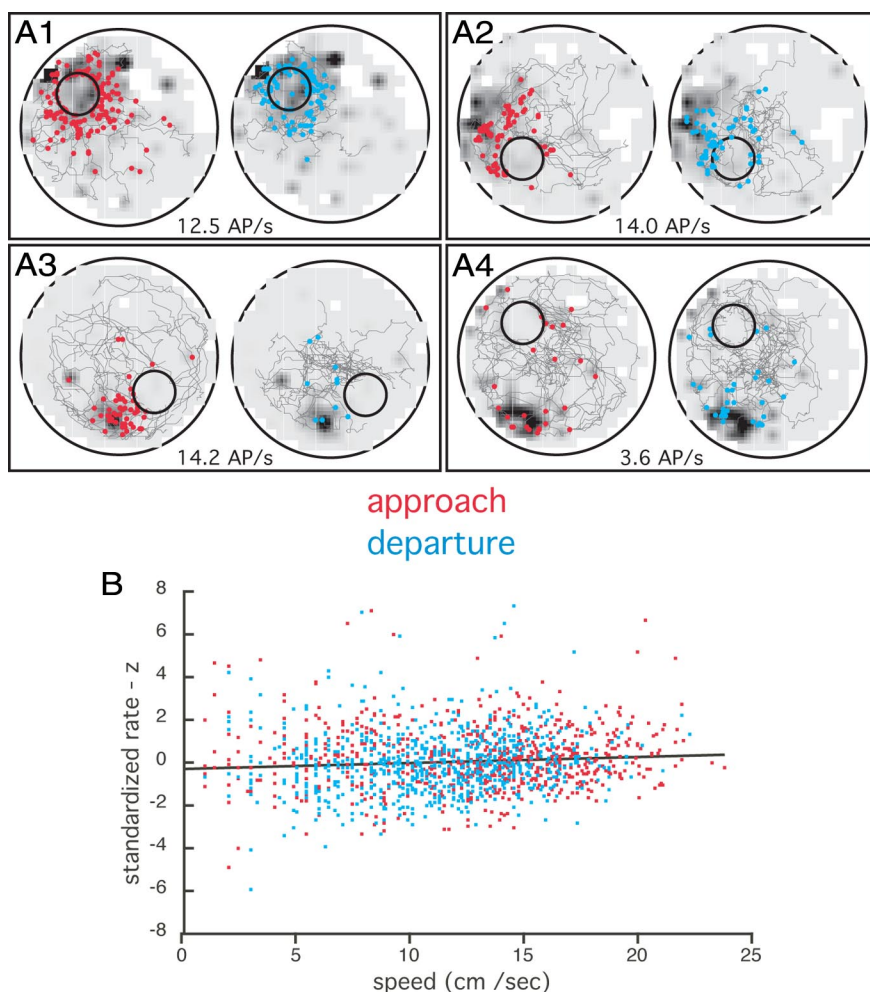


Figure 3. Differences on approaches to (red) and departures (blue) from the goal in the navigate/variable group and the relation of standardized firing to speed. **A1–A4**, The 5 s approaches to (left) and departures from (right) the goal in one session are superimposed on the gray-scaled session-averaged firing rate maps of 4 place cells. The gray-scale from light gray to black represents the range of firing rates from 0 AP/s to the peak (given below each map). White indicates the pixel was not visited. Gray-levels were assigned by linear interpolation of the rates between the centers of adjacent pixels in the averaged firing rate distributions. The track of the rat is superimposed on the rate map. Dots mark action potentials along the paths. Approaches are generally longer and straighter and cover more of the arena than departures. Nonetheless, firing is confined to a similar region on the approaches and departures. **B**, The speed on each approach and departure is plotted against the standardized firing rate for the behavioral episode. Although the speeds of approaches and departures overlap substantially, speed on approaches tend to be faster. The regression line ($y = 0.15x - 0.32$) on all the data indicates that speed is positively correlated with discharge but accounts for $< 1\%$ of the variance ($r = 0.09$).

ing rate variations (z) during passes was significant (Fig. 3B) but the effect is small; the modulation of standardized firing rate by speed accounted for $< 1\%$ ($r^2 = 0.008$) of the variance, making it unlikely that the substantial group overdispersion differences are due to speed, in agreement with an earlier analysis (Fenton and Muller, 1998). Moreover, the prior work detected no significant contribution to overdispersion from time in the firing field, length of the field path, tortuosity (crookedness) of the path or the rat’s direction through the field (Fenton and Muller, 1998). Similarly, we previously compared overdispersion in forage/variable rats (called ‘foragers’) and navigate/variable rats (called ‘navigators’) during the identical foraging task; the only difference is that the navigators had previously been trained to selectively use distal cues over local cues to find the reward location (Olypher et al., 2002). Overdispersion was high in the foragers ($\sigma^2 = 4.9$) but low in the navigators ($\sigma^2 = 2.3$) despite identical physical conditions and overt behaviors. We conclude that lower overdispers-

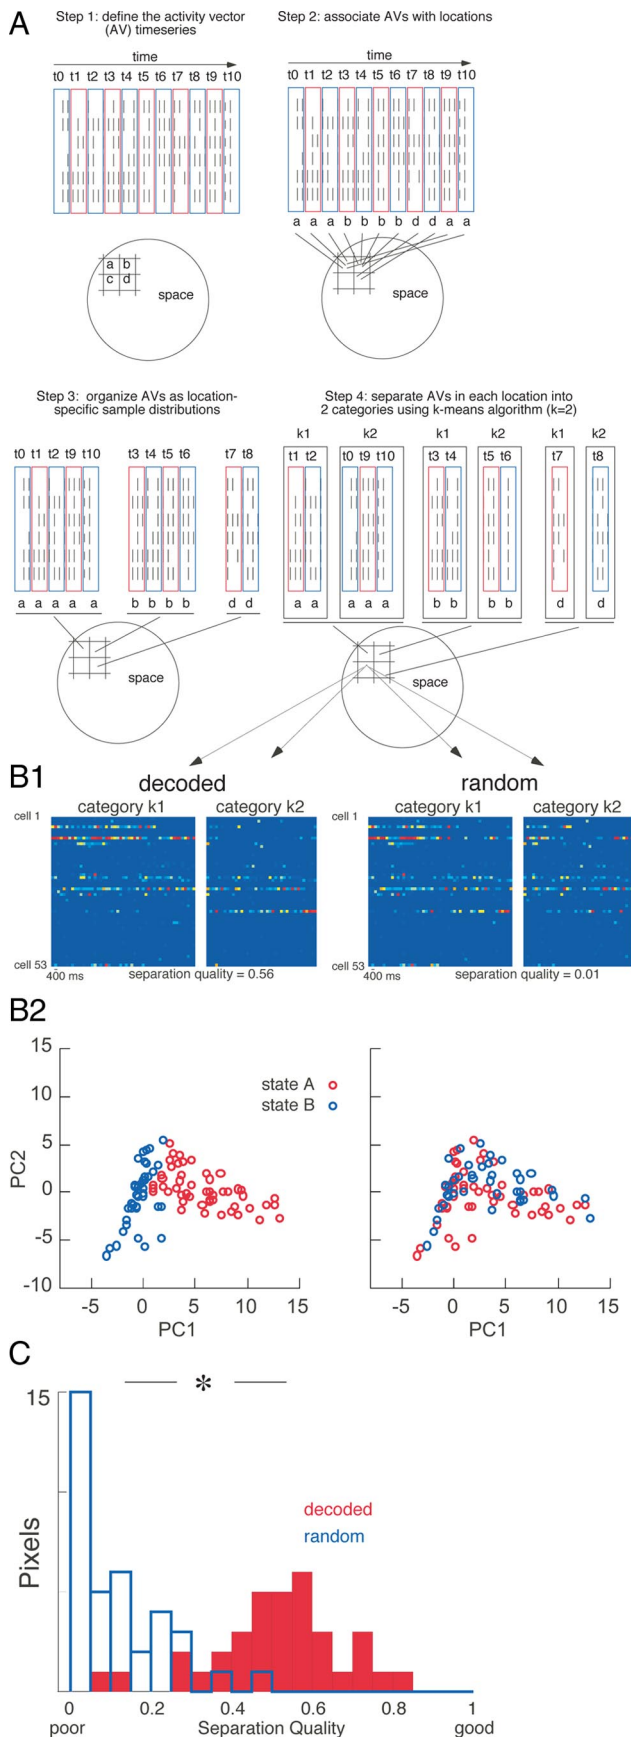


Figure 4. Partitioning the local activity of a 53-place cell ensemble into two categories. **A**, Summary of the algorithm. Step 1: Generate an activity vector (AV) for each time step t_n . Step 2: Assign each AV to the average location in the time step. Step 3: Gather AVs for each location. Step 4: Partition the AVs at a location into two categories k_1 and k_2 . **B**, Activity vectors partitioned in two ways for a 10 cm square pixel. **B1**, The vectors on the left were well-separated (separation quality = 0.56), close to the average yielded by the partitioning algorithm. The set of vectors on the right were randomly assigned to the two categories, yielding poor separation (separation quality = 0.01) that was typical for the random assignment. The random assignment was done so as to preserve the number of vectors in each category produced by the partitioning algorithm. The blue-to-red color code corresponds to 0 to 8 APs per 400 ms. **B2**, The well-separated and poorly-separated categorization was confirmed by principal component analysis. Plotting the first two principal components (PCA1, PCA2) for each AV illustrates that the AVs form two clusters after the partitioning but not after the random assignment. **C**, The AVs were well separated in far more pixels by the partitioning algorithm than by random assignment (* $t_{59} = 10.0$; $p \approx 0$).

sion in the navigate groups arise because of differences in attention-like task demands and conditioning.

Functionally-distinct ensemble states decoded from place cell ensemble discharge

The second and third predictions of the dynamic attention hypothesis propose that overdispersion is a consequence of coherent, ensemble-level switching among multiple attention-like states. Following Jackson and Redish, a CA1 place cell ‘state’ is a temporally delimited ensemble activity profile characterized by a particular set of location-specific spatial firing patterns at any place in an environment (Jackson and Redish, 2007). There are two corollaries to this view. First, the precise across-cell pattern for a state is itself location-specific and may recur any time the rat returns to a certain place. Second, there may be more than one state so that the across-cell pattern at each place in the environment may be different at different times. To test predictions of the dynamic-attention hypothesis we examined the across-cell discharge of three large CA1 place cell ensembles ($n = 54, 53, 23$ cells) from rats trained to forage for scattered food so as to ensure a high degree of overdispersion. The 53 and 23 cell ensembles were recorded in a 76 cm diameter gray cylinder with a white cue card; the 54-cell ensemble was recorded in a 76 cm square white box with a 45 cm black and white striped card on one wall.

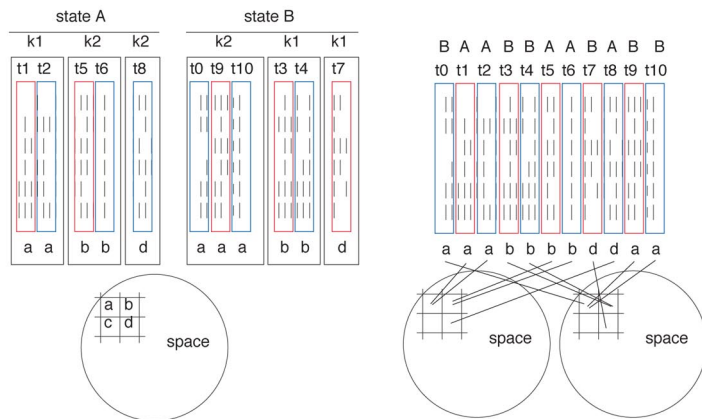
The second prediction of the dynamic-attention hypothesis is that ensemble discharge can be decoded into two states such that the overall firing rate and the environment-wide spatial firing pattern of individual cells in one state differ substantially and reliably from their rate and pattern in the other state. To test this possibility we used the algorithm summarized in Figure 4A to separate activity vectors into two categories, k_1 and k_2 , for each 2.5 cm square pixel at 400 ms temporal resolution; the example in Figure 4, B and C, is for the 53-cell ensemble. Although we did not bin and smooth the temporal activity profiles as in Jackson and Redish and we used a k -medians rather than k -means procedure, the partitioning algorithm was essentially the same (Jackson and Redish, 2007). The two activity vector sets for a single pixel of average separation quality are shown on the left of Figure 4B for the partitioning algorithm. On the right of Figure 4B are the same two sets separated by randomly assigning each activity vector to one of the two categories. It is clear from the example of Figure 4B1 that partitioning did a much better job of separating the activity vectors than did random assignment. Figure 4B2 confirms this impression using principal component analysis, which illustrates that the activity vectors were separated into two numerically distinct categories by the partitioning algorithm but not by the random assignment. This result is shown in Figure 4C for all well-sampled pixels in the environment. We conclude that

←


Step 4. Partition the AVs at a location into two categories k_1 and k_2 . **B**, Activity vectors partitioned in two ways for a 10 cm square pixel. **B1**, The vectors on the left were well-separated (separation quality = 0.56), close to the average yielded by the partitioning algorithm. The set of vectors on the right were randomly assigned to the two categories, yielding poor separation (separation quality = 0.01) that was typical for the random assignment. The random assignment was done so as to preserve the number of vectors in each category produced by the partitioning algorithm. The blue-to-red color code corresponds to 0 to 8 APs per 400 ms. **B2**, The well-separated and poorly-separated categorization was confirmed by principal component analysis. Plotting the first two principal components (PCA1, PCA2) for each AV illustrates that the AVs form two clusters after the partitioning but not after the random assignment. **C**, The AVs were well separated in far more pixels by the partitioning algorithm than by random assignment (* $t_{59} = 10.0$; $p \approx 0$).

A

Step 1: associate one of the two AV groups in each location with state A and the other group with state B by comparing state-specific AV similarity across neighbouring locations. Step 2: reconstitute the AV timeseries \square to define the timeseries of states A and B



B



C

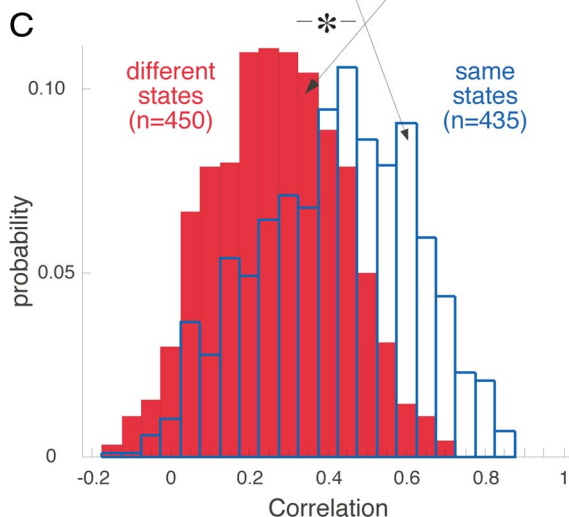


Figure 5. Distinct global ensemble activity states during foraging. **A**, Summary of the algorithm for assigning location-specific categories k1 and k2 across neighboring pixels to global states A and B. Step 1: activity vector correlations between category k1 and

different place cell ensemble activity patterns at a given location may form distinct categories.

Next, we asked if the separated categories at a given location can be related to the separated categories at all other locations to form two global (i.e., location-independent), mutually exclusive states. Starting this decoding process at the pixel with the best separation quality, we arbitrarily designated activity vectors in category k1 as belonging to state A and vectors in category k2 as belonging to state B. We then computed 4 correlations, namely k1-k1, k1-k2, k2-k1, k2-k2 between vector categories in the first pixel and those in a neighboring pixel. If the average of the k1-k1 and k2-k2 correlations was greater than the average of the k1-k2 and k2-k1 correlations, k1 in the second pixel was assigned to state A and k2 to state B; otherwise the assignments were reversed. The two global states were formed by iteratively carrying out this assignment process for each pixel in the environment; the algorithm is summarized in Figure 5A. We computed two state-specific, location-independent ensemble activity vectors for each 20 s episode by grouping the activity from all 400 ms time steps decoded to be in the same state. We then compared each 20 s activity vector to every other 20 s vector by generating a correlation matrix in which the 20 s vectors were sorted by time order within each state (Fig. 5B). From the color-coded representation of the correlation matrix it is apparent that activity within a state was more similar than activity between states (Fig. 5C). Thus, in standard conditions, it was possible to decode two global, location-independent activity states from place cell ensemble activity.

We then used this decoding to determine if the firing properties of individual place cells differ between the states. In Figure 6A we show conventional (raw), state-specific and shuffled firing rate maps for 7 cells of the 53-cell ensemble; similar maps for the entire ensemble are in supplemental Figure 1 (available at www.jneurosci.

← category k2 are computed for a base pixel and a neighboring pixel. The assignment to state A or B is determined as stated in the text. Step 2: regenerate the rat's trajectory tagging each data sample with one global state or the other. **B**, A correlation matrix comparing all pairs of 20 s ensemble activity vectors from a 10 min, 53-cell ensemble recording. The 20 s intervals are organized chronologically. **C**, The correlations of ensemble activity from moments in the same states (A-A, B-B) were greater than the correlations from moments when the ensemble was judged to be in different states (A-B, B-A). (* $t_{884} = 16.9$; $p \approx 0$).

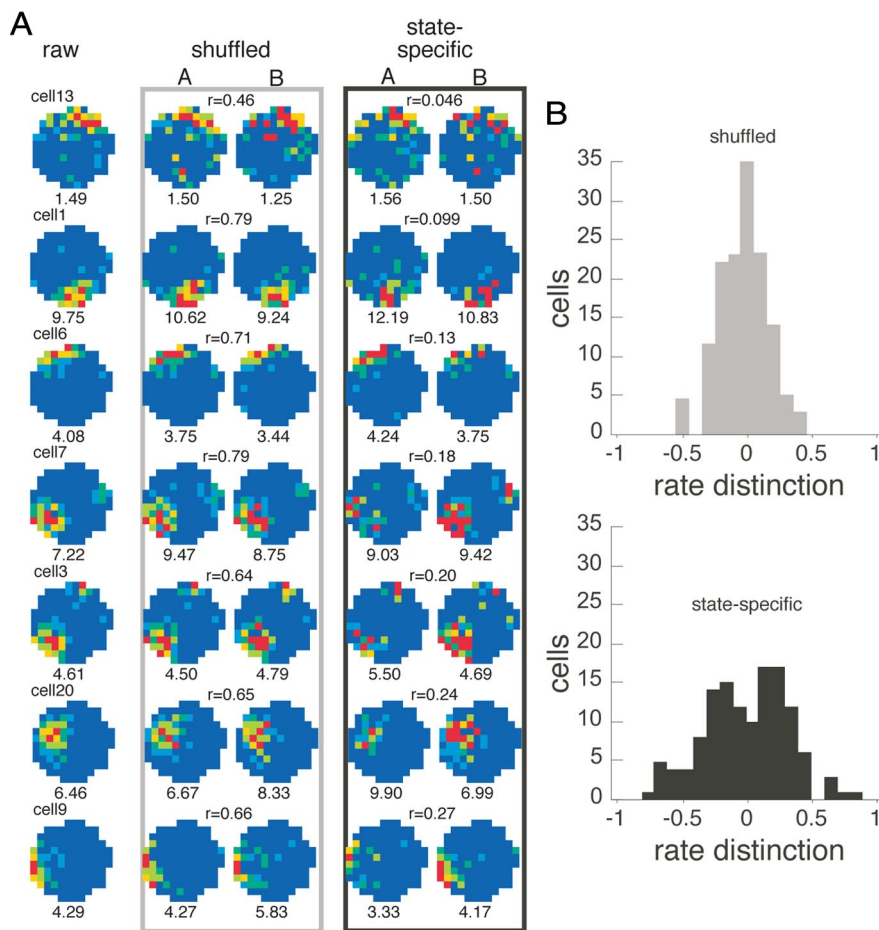


Figure 6. Prediction 2. Decoding of activity vectors produces distinct properties for individual place cells in each of the two global states. **A**, Firing rate maps from 7 cells of the 53-cell ensemble. For each cell the standard (raw) firing rate map is depicted along with the state-specific firing rate maps and the shuffled state maps, which help assess the likelihood that the decoding can occur by chance. The color code is the same for each set of 5 maps. The median firing rate in the red (peak rate) category is given below each map. The correlation between each shuffled pair or state-specific pair of firing rate arrays is given above the corresponding pairs of maps. The cells are sorted in ascending order of the state-specific correlation. **B**, Histograms of the signed rate distinction after random (top) and state-specific assignment (right). Note that the randomly generated histogram is unimodal and narrow whereas the histogram made with state-specific assignment is much broader and has a clear bimodal character. The implication is that there are approximately equal numbers of cells that discharge more rapidly in each of the states.

org as supplemental material) online. Note that the shuffling procedure preserved the observed individual activity vectors but pseudo-randomly assigned the vectors to one of the two location-specific partitioned categories. For the shuffled rate maps, the location-specific partitioned categories were randomly assigned to the global states before composing the maps. The shuffled states produced rate maps that are very similar to each other and that look like degraded versions of the conventional map whereas by inspection, the state-specific maps show greater differences. The results were in line with these observations when the data from the three ensembles were analyzed together ($n = 130$ cells). The pixel-by-pixel correlation between the shuffled rate maps were large and all were significant (minimum $r = 0.28$; average $r = 0.68 \pm 0.01$) whereas the similarities were substantially smaller between the state-specific firing rate maps (minimum $r = -0.02$; average $r = 0.37 \pm 0.02$; paired $t_{129} = 20.9$; $p \approx 0$). Indeed, 22% (28/130) of the correlations between state-specific firing rate maps were not significant, suggesting the maps were unrelated, and thus the difference was unlikely to be the result of rate remapping (S. Leutgeb et al., 2005). In summary, the state-specific firing

rate maps of individual cells have less resemblance to each other and may even be independent by this criterion.

We also performed an alternative shuffling procedure that destroyed the individual observed activity vectors. During a particular visit to a pixel the activity of each cell was randomly chosen (without replacement) from the set of all visits to the pixel. The separation quality (0.43 ± 0.03 in the 53-cell ensemble) remained the same as the separation quality without shuffling (0.47 ± 0.04 ; $t_{72} = 0.83$, $p = 0.41$) but this created pairs of similar firing rate maps. For the 130-cell dataset, the distribution of spatial correlations ($r = 0.67 \pm 0.023$) was indistinguishable from that of the firing rate map pairs that were created after shuffling that preserved the observed activity vectors ($r = 0.68 \pm 0.01$; paired $t_{129} = 0.1$, $p = 0.92$). Thus whether the ensemble data were shuffled in a way that preserved or destroyed the observed activity vectors the result was the same. Since both these control procedures failed to produce distinct state-specific spatial firing rate distributions, while in contrast, the decoding algorithm generated distinct state-specific spatial firing rate distributions (Fig. 6), we conclude that two functionally-distinct global activity states can be decoded from ensemble recordings during foraging for scattered food.

The dynamic-attention hypothesis predicts that the two location-independent, state-specific firing rates may differ for individual cells. In agreement, rate distinctions between the shuffled-state maps (0.14 ± 0.013) were much smaller than rate distinctions between the two global states (0.30 ± 0.028 ; $t_{129} = 8.2$; $p \approx 0$). Such state-specific rate distinctions during continuous recording sessions in constant conditions

were at least as large as the rate variations observed by moving rats between colored or differently shaped enclosures (S. Leutgeb et al., 2005). The distribution of rate differences was also consistent with the dynamic-attention hypothesis in that approximately half the 130 cells had a higher rate in each state (Fig. 6B). Indeed, this bimodality of firing rate fluctuations within a single ensemble provides further evidence against the possibility that a group difference in some behavioral variable accounts for the group differences in overdispersion that were identified in the data-mining study (Table 1; Fig. 2).

Similar results to those in Figure 6 were obtained from a leave-one-out analysis in which, for each cell in an n -cell ensemble, the separation into two activity states was determined by the activity of the $n - 1$ remaining cells after excluding the cell (Jackson and Redish, 2007). The implication is that the distinct firing rates and spatial firing patterns of a single cell were identified because the cell is part of a network in which two global activity states could be decoded from CA1 ensemble place cell activity, consistent with the dynamic-attention hypothesis.

We also estimated the persistence of the two states by calculating the dwell time before switching within the three ensemble

recordings. Within the total of 49 min of data, based on 400 ms resolution, the shortest interval used for decoding, the average dwell time in a state was 913 ± 13 ms, with the longest interval 6.4 s. According to the runs test ($z = 9.2$; $p \approx 0$), the dwell time was longer than random switching between the two states. These data indicate that CA1 place cell ensemble activity is modulated into distinct global activity states on a time scale of approximately one second.

Using decoded states improves prediction of current location

The third prediction of the dynamic-attention hypothesis is that separating correctly decoded state-specific spike trains will predict current location better than the raw state-independent spike trains we observed. We analyzed ensemble spike trains to reconstruct the rat's path using an algorithm that chooses the best match between the current ensemble activity vector and the average ensemble activity at each location (Wilson and McNaughton, 1993; Fenton and Muller, 1998; Fenton et al., 2008). The three ensemble recordings during open-field foraging were studied at four temporal (400, 800, 1200, 1600 ms) and four spatial (2.5, 5, 10, 20 cm/pixel) resolutions. Depending on the particular combination of temporal and spatial resolutions that were used, the average reconstruction error ranged between 0.4 and 5.9 pixels. This accuracy was substantially better than chance because the average error for each reconstruction was smaller than randomly selecting the rat's location at each time step by at least a 50% margin (average reconstruction errors between 1.4 and 13.4 pixels) (Fenton et al., 2008).

The reconstruction accuracy of the rat's path was improved when the reconstruction used the decoding of the ensemble state to decide which set of firing rate maps was more appropriate for locating the rat (Fig. 7). Figure 7A illustrates the reconstruction error along three 10 s episodes from the 53-cell ensemble recording. The reconstructions were made at 400 ms temporal resolution and 2.5 cm spatial resolution in these examples. Differences between the observed location of the rat and the location predicted by the raw ensemble spike trains (Fig. 7A, column 1) appear larger than the differences between the observations and the predictions from the state-decoded ensemble spike trains (Fig. 7A, column 2). We considered whether this improvement might be trivial, for example, a consequence of the mere fact that an additional parameter ('state') was used for the reconstruction. We reasoned that if the decoded state information was useful, then the reconstruction should be worse, if at each time step we deliberately chose to match the current activity vector with the wrong state-specific map. On the other hand, the wrong choice of state should also improve reconstruction if the improvement was trivial. Repeating the state-specific calculations with the opposite ('wrong') state classification at each time step produced worse reconstruction as illustrated in column 3 of Figure 7A.

Figure 7B shows color-coded representations of the average reconstruction error at the 16 combinations of spatial and temporal resolution that we tried for the 53-cell ensemble. Smaller errors are indicated by hotter colors, which are most common in the reconstructions that used the correct state decoding (Fig. 7B, middle). This presentation suggests that the accuracy of reconstructing the rat's location using the decoded states was as good or better than predictions using either the wrong state or the raw, state-independent spike trains for each combination of spatial and temporal resolution. These figures also illustrate that although the reconstruction error tended to decrease at lower temporal resolutions and higher spatial resolutions, the error is not a simple function of the two resolutions. Consequently for statis-

tical comparisons, the average reconstruction error that used the state decoding at each combination of the temporal and spatial resolution was standardized using the corresponding reconstruction error from the raw spike trains (see Fig. 7C legend).

All standardized reconstruction errors using decoded spike trains from the three ensembles were smaller than 0, which is a qualitative indication that the decoded state information consistently improved reconstruction compared with the observed spike trains that ignore state (Fig. 7C). Repeating the state-specific calculations with the opposite ('wrong') state classification at each time step produced worse reconstructions than using the 'correct' state decoding (Fig. 7C; $t_{47} = 9.4$; $p \approx 0$). Furthermore, these reconstructions were also worse than the state-independent predictions ($t_{47} = 2.2$; $p = 0.03$). Together, these findings confirm the third prediction of the dynamic-attention hypothesis that correct knowledge of the ensemble state will improve predictions of current location.

Discussion

Our results confirm three predictions of the dynamic-attention hypothesis. First, overdispersion is reduced when rats were trained to use one of two stimulus subsets (Fig. 2). Second, the collective activity of large place cell ensembles is separable into two global (environment-spanning) place codes ('maps') (Fig. 5). Individual place cells have distinct spatial firing patterns and rates (Fig. 6) in the two place codes, which alternated with an average period of approximately a second. Third, the accuracy of predicting the rat's moment-to-moment position is improved when the switching between place codes is taken into account (Fig. 7). It is our contention that these switches correspond to transitions between distinct attentional states, although it cannot be excluded that these attention-like switches reflect some other rapid alternation of the information being processed by the hippocampus, or some other internally-generated activity patterns (Johnson and Redish, 2007; Pastalkova et al., 2008) associated with cognition (Johnson et al., 2009).

Behavioral control of overdispersion

Our argument that place cell activity reflects a form of dynamic attention is based on the notion that the overdispersion of firing by single cells during navigation to a goal is a consequence of a rat's ability to use either of two in-register cue subsets to locate itself and to navigate (Fenton et al., 1998). Prior work showed that place preference training could reduce overdispersion (Olypher et al., 2002; Jackson and Redish, 2007). Nevertheless, because either the introduction of navigation or changes in attention could explain the reduction, we set out to disambiguate the roles of navigation and attention. Our key analysis compared overdispersion in recordings from the navigate/stable and navigate/variable groups. During recordings, rats in both groups navigated to a fixed goal on a stationary arena, but by intermittent arena rotation the navigate/variable rats had learned that the local cue set could not reliably guide them to the hidden goal. We reasoned that rendering one subset of cues irrelevant for locating a goal forced the hippocampal navigational system to attend preferentially to the remaining subset with the observable outcome that the noisiness of discharge would be reduced. We found that overdispersion decreased in both groups compared with their respective control groups (forage/stable and forage/variable). The control rats never received navigation training, but had been exposed to the identical physical conditions and foraging behavior as the corresponding navigation group. Since overdispersion was the same in the control groups, the data indicate that naviga-

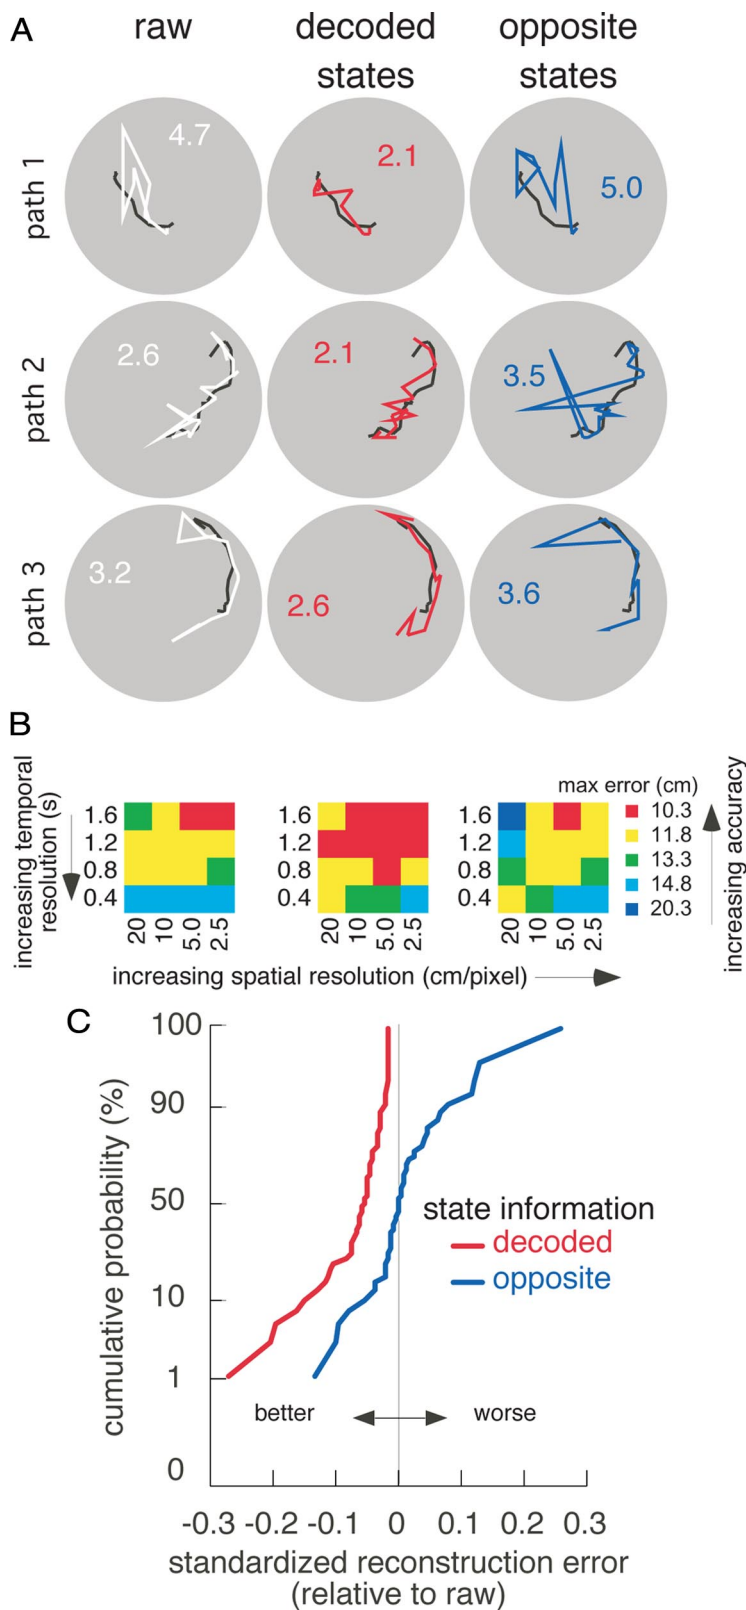


Figure 7. Prediction 3. Decoding global ensemble states improves reconstruction of the rat's location. At each time step, an estimate of the decoded state of the current ensemble activity vector was either ignored (raw) or used to predict the current location. If the decoded state was used, the reconstruction was made by finding the best match between the vector and the state-specific template vectors that were constructed from the time series of state-specific activity vectors. The state-specific template vectors were either always chosen to match the state of the current ensemble vector (decoded), or be opposite to that state (opposite). **A**, Three 10 s paths that were reconstructed at 400 ms and 2.5 cm resolution. The rat's path (black) and the reconstructed path using the decoded ensemble state (red) were typically more similar than the estimates from the raw data (white) or estimates based on the incorrect ensemble state (blue). The numbers give the average prediction error in pixels for the

tion itself reduces overdispersion, confirming prior work. The main finding provides the strongest evidence to date for the dynamic attention hypothesis; overdispersion was reduced even more in the navigate/variable group than the navigate/stable group. This reduction is definitive evidence that overdispersion is influenced by a covert, internal variable, which we suggest is attention. More detailed comparisons of the present results to prior work are available in a supplemental Discussion online.

This effect of reduced overdispersion is closely in line with the classic definition of attention: "Everyone knows what attention is. It is the taking possession by the mind in clear and vivid form, of one out of what seem several simultaneously possible objects or trains of thought . . . It implies withdrawal from some things to deal effectively with others . . ." (James, 1890). In the absence of rotational training, we think that the hippocampus switches at a rate of ~1 Hz between two co-equal representations of space so that when one is active the other is suppressed. After rotational training, the switching frequency and therefore overdispersion decreases because the system dwells longer in the state associated with distal cues. It is this increased dwell that is the neural analog of paying more attention to the distal cues.

Two other behavioral methods of reducing overdispersion are currently known, each of which suggests that additional factors besides cue switching may be crucial in understanding how attention-like processes affect hippocampal activity. First, we saw that simply requiring rats to perform a hidden goal task reduces overdispersion

path. **B**, Summary of 16 reconstruction attempts with the 53-cell ensemble. Current position was reconstructed at 400, 800, 1200, and 1600 ms temporal resolutions and at 2.5, 5, 10, and 20 cm spatial resolutions. The average error from each of the 16 combinations of the temporal and spatial resolutions is represented as a color-coded matrix. The maximum error in each color category is given; hotter colors indicate better accuracy. Reconstruction accuracy using the decoded spike trains appears equivalent or better for each attempt. **C**, Summary of 48 reconstruction attempts. Location was reconstructed for three recordings (ensemble sizes = 54, 53, and 23 place cells) at the four temporal and spatial resolutions. The average error for each set of reconstructions was normalized by the average reconstruction error using the observed, state-independent data: $\frac{(\text{state} - \text{observed})}{(\text{state} + \text{observed})}$. Reconstructions using the decoded state (red) were always better than reconstructions from the raw state-independent data. Not only was this improvement significant ($t_{47} = 9.4; p \approx 0$), but it was a result of accurate state decoding, because reconstructions using the incorrect state (blue) were often worse than the raw, state-independent predictions (paired $t_{47} = 2.2; p = 0.03$).

from $\sigma^2 = 5.5$ to $\sigma^2 = 3.5$. This decrease might be due to slower switching between two equally likely cue subsets or to a tendency to dwell longer using the preferred subset, but yet other connections between firing variance and attention are plausible. Similarly, the relatively low overdispersion ($\sigma^2 = 2.1$) of place cell firing when rats run in one direction on a linear track (Jackson and Redish, 2007) is not immediately ascribable to switching between cue subsets but is attributed to the predominant use of a single map, which again suggests that attention fluctuates less during performance of a task with fixed goal locations.

Attention has also been implicated in the stabilization of firing fields in mice in which location-specific firing is apparently less reliable than in rats. Thus, increasing the complexity of behavioral demands on the mice from unreinforced running to pellet foraging to turning off aversive sensory stimulation by going to a hidden goal progressively increases the reproducibility of spatial firing on a time scale of hours, suggesting the stability was due to increased attention, although a contribution of increased arousal could not be excluded (Kentros et al., 2004). A subsequent study also demonstrated increased firing field reproducibility by cue-selective training of mice to use one of two stimulus subsets to obtain reward, which controlled for arousal, although potentially stabilizing effects of learning itself could not be excluded (Muzzio et al., 2009). Since neither differences in arousal nor learning account for the decreases in overdispersion we observed as a consequence of prior cue-selective training and going directly to a specific goal, the prior work with mice and the present results from rats support our contention that hippocampal activity is modulated by attention. The effects on firing field stabilization across hours may be due to constraining the hippocampus to spend a larger fraction of its time in a particular state if the stabilization is viewed as a form of variance reduction. Furthermore, partitioning the hippocampal population into transiently active subsets may bias cells in the active subset to discharge synchronously and desynchronize from other subsets that represent alternate states. This temporal structure could facilitate synaptic plasticity within the active subset and promote segregation between desynchronized subsets to stabilize LTP-dependent memories (Pastalkova et al., 2006; Serrano et al., 2008) and maintain firing field stability. Providing a temporal structure that can promote plasticity would complement the stabilizing action of a prolonged dopaminergic signal that may correlate with forms of attention (Li et al., 2003; Kentros et al., 2004; Lavin et al., 2005).

Taking these examples into account, it is likely that the procedure of partitioning population activity into two states is only a first approximation; additional states may become evident with larger recording ensembles, longer recording times and more complex behavioral manipulations.

Partitioning of place cell firing into two mutually exclusive global states

Overdispersion characterizes the putative attention-like switching of individual place cell discharge between control by distal and local sources of spatial information; the role of self-motion is uncertain. We also found, at the ensemble level, that the activity of a place cell sample can be sorted into two alternating time series (states) under the assumption that all cells in the sample switch synchronously between distal and local control. In our partitioning method, switching between the two states may occur at the end of each 400 ms interval but finer grained time slices could be used if the number of cells was larger. We began by separating time spent in each pixel into two clusters of similar firing rate population vectors and then iteratively matched the

vector pairs in neighboring pixels to form two environment-spanning, mutually exclusive place codes.

The ability to form two mutually exclusive, internally consistent global place codes satisfies the second prediction of the dynamic attention hypothesis and furthermore, reproduces a very similar treatment of ensemble place cell data recorded on a linear track, during pellet chasing in a cylinder and during goal seeking in the same cylinder (Jackson and Redish, 2007). The algorithms used by Jackson and Redish were done at higher temporal resolution, used a smoothing process we omitted and was based on k -mean rather than k -median clustering of population firing vectors, but the outcomes are extremely similar. Of special note are the findings that overdispersion in both studies is numerically equal during pellet foraging, that it is reduced equally in the hidden goal task and that the characteristic switching interval is calculated to be <1 s in both cases although our value of 0.91 s is somewhat greater than the value of 0.38 s seen by Jackson and Redish (2007).

Improved position reconstruction using two state-specific place codes

The dynamic-attention hypothesis proposes that there is alternation between distinct hippocampal place codes, one in which cells are activated by distal stimuli and one in which cells are activated by local sources of information. If this is true but the rat's head position is nevertheless computed according to the position in which the grand average population vector best matches the population vector for a given short interval (Wilson and McNaughton, 1993), the average position reconstruction error should be greater than if the currently active place code is used for the match. The finding that reconstruction accuracy is indeed improved by matching to the currently active place code satisfies the third prediction of the dynamic attention hypothesis and implies that partitioning has revealed a signal source previously expected from theoretical considerations (Olypher et al., 2002) but not seen empirically. These two place codes may correspond to the two hippocampal states that can be defined by the features of ongoing gamma oscillations (Senior et al., 2008; Colgin et al., 2009). We emphasize that for individual cells, the two firing rate maps resemble each other since each place cell fires in approximately the same region regardless of which place code is used (S. Leutgeb et al., 2005); this must be the case since in a small apparatus firing is generally confined to a single firing field. Thus, the reconstruction error introduced by using the "wrong" place code to identify the current position cannot be very large in stationary environments.

Summary

The attention-like modulation of place cell firing reported here supports the view that the hippocampus dynamically signals responses to a selected, time-varying subset of available features and is not a static representation of location. In this way, the navigational system of the rat is shown to more closely resemble human spatial cognition which is characterized by rapid, flexible switching of attention among overlapping or coincident frames, only one of which can be active at a given moment. The choice of whether to activate the, room, building or city frame is done according to a person's immediate needs.

References

- Bures J, Fenton AA, Kaminsky Y, Zinyuk L (1997a) Place cells and place navigation. *Proc Natl Acad Sci U S A* 94:343–350.
- Bures J, Fenton AA, Kaminsky Y, Rossier J, Sacchetti B, Zinyuk L (1997b)

- Dissociation of exteroceptive and idiothetic orientation cues: effect on hippocampal place cells and place navigation. *Philos Trans R Soc Lond B Biol Sci* 352:1515–1524.
- Colgin LL, Denninger T, Fyhn M, Hafting T, Bonnevie T, Jensen O, Moser MB, Moser EI (2009) Frequency of gamma oscillations routes flow of information in the hippocampus. *Nature* 462:353–357.
- Czurkó A, Hirase H, Csicsvari J, Buzsáki G (1999) Sustained activation of hippocampal pyramidal cells by 'space clamping' in a running wheel. *Eur J Neurosci* 11:344–352.
- de Hoon MJ, Imoto S, Nolan J, Miyano S (2004) Open source clustering software. *Bioinformatics* 20:1453–1454.
- Fenton AA, Bures J (2003) Navigation in the moving world. In: *The neurobiology of spatial behaviour* (Jeffery K, ed), pp. 240–258. Oxford: Oxford UP.
- Fenton AA, Muller RU (1998) Place cell discharge is extremely variable during individual passes of the rat through the firing field. *Proc Natl Acad Sci U S A* 95:3182–3187.
- Fenton AA, Wesierska M, Kaminsky Y, Bures J (1998) Both here and there: simultaneous expression of autonomous spatial memories in rats. *Proc Natl Acad Sci U S A* 95:11493–11498.
- Fenton AA, Zinyuk L, Bures J (2000a) Place cell discharge along search and goal-directed trajectories. *Eur J Neurosci [Suppl]* 43.09.
- Fenton AA, Csizmadia G, Muller RU (2000b) Conjoint control of hippocampal place cell firing by two visual stimuli. I. The effects of moving the stimuli on firing field positions. *J Gen Physiol* 116:191–209.
- Fenton AA, Bures J, Cimadevilla JM, Olypher AV, Wesierska M, Zinyuk L, eds (2002) *Place cell activity during overtly purposeful behavior (in dissociated reference frames)*. Boston: Kluwer Academic.
- Fenton AA, Kao HY, Neymotin SA, Olypher A, Vayntrub Y, Lytton WW, Ludvig N (2008) Unmasking the CA1 ensemble place code by exposures to small and large environments: more place cells and multiple, irregularly-arranged, and expanded place fields in the larger space. *J Neurosci* 28:11250–11262.
- Gothard KM, Skaggs WE, McNaughton BL (1996) Dynamics of mismatch correction in the hippocampal ensemble code for space: interaction between path integration and environmental cues. *J Neurosci* 16:8027–8040.
- Hollup SA, Molden S, Donnett JG, Moser MB, Moser EI (2001) Accumulation of hippocampal place fields at the goal location in an annular water-maze task. *J Neurosci* 21:1635–1644.
- Jackson J, Redish AD (2007) Network dynamics of hippocampal cell-assemblies resemble multiple spatial maps within single tasks. *Hippocampus* 17:1209–1229.
- James W (1890) *The principles of psychology*. New York: Henry Holt.
- Johnson A, Redish AD (2007) Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *J Neurosci* 27:12176–12189.
- Johnson A, Fenton AA, Kentros C, Redish AD (2009) Looking for cognition in the structure within the noise. *Trends Cogn Sci* 2:55–64.
- Kentros CG, Agnihotri NT, Streater S, Hawkins RD, Kandel ER (2004) Increased attention to spatial context increases both place field stability and spatial memory. *Neuron* 42:283–295.
- Kubie JL, Fenton A, Novikov N, Touretzky D, Muller RU (2007) Changes in goal selection induced by cue conflicts are in register with predictions from changes in place cell field locations. *Behav Neurosci* 121:751–763.
- Lavin A, Nogueira L, Lapish CC, Wightman RM, Phillips PE, Seamans JK (2005) Mesocortical dopamine neurons operate in distinct temporal domains using multimodal signaling. *J Neurosci* 25:5013–5023.
- Lenck-Santini PP, Muller RU, Save E, Poucet B (2002) Relationships between place cell firing fields and navigational decisions by rats. *J Neurosci* 22:9035–9047.
- Leutgeb JK, Leutgeb S, Treves A, Meyer R, Barnes CA, McNaughton BL, Moser MB, Moser EI (2005) Progressive transformation of hippocampal neuronal representations in "morphed" environments. *Neuron* 48:345–358.
- Leutgeb S, Leutgeb JK, Treves A, Moser MB, Moser EI (2004) Distinct ensemble codes in hippocampal areas CA3 and CA1. *Science* 305:1295–1298.
- Leutgeb S, Leutgeb JK, Barnes CA, Moser EI, McNaughton BL, Moser MB (2005) Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science* 309:619–623.
- Lever C, Wills T, Cacucci F, Burgess N, O'Keefe J (2002) Long-term plasticity in hippocampal place-cell representation of environmental geometry. *Nature* 416:90–94.
- Li S, Cullen WK, Anwyl R, Rowan MJ (2003) Dopamine-dependent facilitation of LTP induction in hippocampal CA1 by exposure to spatial novelty. *Nat Neurosci* 6:526–531.
- Lytton WW (2006) Neural Query System: Data-mining from within the NEURON simulator. *Neuroinformatics* 4:163–176.
- Lytton WW, Barry JM, Olypher AV, Muller RU, Fenton AA (2007) Hippocampal place-cell-ensemble firing-rate modulation suggests attention-like switching during foraging in a fixed environment. *Soc Neurosci Abstr* 33:90.93.
- Markus EJ, Qin YL, Leonard B, Skaggs WE, McNaughton BL, Barnes CA (1995) Interactions between location and task affect the spatial and directional firing of hippocampal neurons. *J Neurosci* 15:7079–7094.
- McNaughton BL, Barnes CA, O'Keefe J (1983) The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. *Exp Brain Res* 52:41–49.
- Moran J, Desimone R (1985) Selective attention gates visual processing in the extrastriate cortex. *Science* 229:782–784.
- Muller RU, Kubie JL (1987) The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *J Neurosci* 7:1951–1968.
- Muller RU, Kubie JL, Bostopck EM, Taube JS, Quirk GJ (1991) Spatial firing correlates of neurons in the hippocampal formation of freely-moving rats. In: *Brain and space* (Paillard J, ed), pp 296–333. Oxford: Oxford UP.
- Muzzio IA, Levita L, Kulkarni J, Monaco J, Kentros C, Stead M, Abbott LF, Kandel ER (2009) Attention enhances the retrieval and stability of visuospatial and olfactory representations in the dorsal hippocampus. *PLoS Biol* 7:e1000140.
- O'Keefe J, Nadel L (1978) *The hippocampus as a cognitive map*. Oxford: Clarendon.
- Olypher AV, Lánský P, Fenton AA (2002) Properties of the extra-positional signal in hippocampal place cell discharge derived from the overdispersion in location-specific firing. *Neuroscience* 111:553–566.
- Pastalkova E, Serrano P, Pinkhasova D, Wallace E, Fenton AA, Sacktor TC (2006) Storage of spatial information by the maintenance mechanism of LTP. *Science* 313:1141–1144.
- Pastalkova E, Itskov V, Amarasingham A, Buzsáki G (2008) Internally generated cell assembly sequences in the rat hippocampus. *Science* 321:1322–1327.
- Redish AD, Rosenzweig ES, Bohanick JD, McNaughton BL, Barnes CA (2000) Dynamics of hippocampal ensemble activity realignment: time versus space. *J Neurosci* 20:9298–9309.
- Rossier J, Kaminsky Y, Schenk F, Bures J (2000) The place preference task: a new tool for studying the relation between behavior and place cell activity in rats. *Behav Neurosci* 114:273–284.
- Save E, Nerad L, Poucet B (2000) Contribution of multiple sensory information to place field stability in hippocampal place cells. *Hippocampus* 10:64–76.
- Senior TJ, Huxter JR, Allen K, O'Neill J, Csicsvari J (2008) Gamma oscillatory firing reveals distinct populations of pyramidal cells in the CA1 region of the hippocampus. *J Neurosci* 28:2274–2286.
- Serrano P, Friedman EL, Kenney J, Taubenfeld SM, Zimmerman JM, Hanna J, Alberini C, Kelley AE, Maren S, Rudy JW, Yin JC, Sacktor TC, Fenton AA (2008) PKMzeta maintains spatial, instrumental, and classically conditioned long-term memories. *PLoS Biol* 6:2698–2706.
- Spiers HJ, Maguire EA (2006) Thoughts, behaviour, and brain dynamics during navigation in the real world. *Neuroimage* 31:1826–1840.
- Spiers HJ, Maguire EA (2007) Decoding human brain activity during real-world experiences. *Trends Cogn Sci* 11:356–365.
- Spitzer H, Desimone R, Moran J (1988) Increased attention enhances both behavioral and neuronal performance. *Science* 240:338–340.
- Wills TJ, Lever C, Cacucci F, Burgess N, O'Keefe J (2005) Attractor dynamics in the hippocampal representation of the local environment. *Science* 308:873–876.
- Wilson MA, McNaughton BL (1993) Dynamics of the hippocampal ensemble code for space. *Science* 261:1055–1058.
- Zinyuk L, Kubie S, Kaminsky Y, Fenton AA, Bures J (2000) Understanding hippocampal activity by using purposeful behavior: place navigation induces place cell discharge in both task-relevant and task-irrelevant spatial reference frames. *Proc Natl Acad Sci U S A* 97:3771–3776.