HIPPOCAMPUS 17:1209–1229 (2007)

Network Dynamics of Hippocampal Cell-Assemblies Resemble Multiple Spatial Maps Within Single Tasks

Jadin Jackson[1] and A. David Redish[2*]

ABSTRACT: The firing of place cells in the rodent hippocampus is reliable enough to infer the rodent's position to a high accuracy; however, hippocampal firing also reflects the stages of complex tasks. Theories have suggested that these task-stage responses may reflect changes in reference frame related to task-related subgoals. If the hippocampus represents an environment in multiple ways depending on a task's demands, then switching between these cell assemblies should be detectable as a switch in spatial maps or reference frames. Place cells exhibit extreme temporal variability or "overdispersion," which Fenton et al. suggest reflects changes in active cell-assemblies. If reference-frame switching exists, investigating the relationship of the single cell variability described by Fenton and collegues to network level processes provides an entry point to understanding the relationship between cell-assembly-like mechanisms and an animal's behavior. We tested the cell-assembly explanation for overdispersion by recording hippocampal neural ensembles from rats running three tasks of varying spatial complexity: linear track (LT), cylinder-foraging (CF), and cylinder-goal (CG). Consistent with the reports by Fenton and colleagues, hippocampal place cells showed high variance in their firing rates across place field passes on the CF and CG tasks. The directional firing of hippocampal place cells on LT provided a test of the reference-frame hypothesis: ignoring direction produced overdispersion similar to the CF and CG tasks; taking direction into account produced a significant decrease in overdispersion. To directly examine the possibility of a network modulation of cell-assemblies, we clustered the firing patterns within each pixel and chained them together to construct whole-environment spatial firing maps. Maps were internally self-consistent, switching with mean rates of several hundred milliseconds. There were significant increases in map-switching rates following reward-related events on the LT and CG tasks, but not on the CF task. Our results link single cell variability with network-level processes and imply that hippocampal spatial representations are made up of multiple, continuous submaps, the selection of which depends on the animal's goals when reward is tied to the animal's spatial behavior.

KEY WORDS: place cell; neural ensemble; cell assembly; tetrode; overdispersion

INTRODUCTION

In 1949, Hebb proposed the concept of a "cell-assembly"—that within a self-recurrent network, subnetworks would form, which would be detectable by sets of related firing (Hebb, 1949). The hippocampus is noted for its self-recurrent connections (Lorento do Nó, 1933, 1934; Amaral and Witter, 1989) and thus forms a particularly useful candidate for the study of the cell-assembly hypothesis (McNaughton and Morris, 1987; Harris, 2005). The first-order behavioral correlates of rodent hippocampal pyramidal cells are their spatial firing correlates or *place fields* (O'Keefe and Dostrovsky, 1971; O'Keefe, 1976; Redish, 1999). The spatial correlates of hippocampal firing provide leverage to examine the dynamics of cell assemblies within the hippocampus allowing the distinction to be made between single cell response properties and phenomena that can only be observed at the population level.

While the firing of place cells is reliable enough spatially to infer the rodent's position to within 1 cm given only the current firing pattern in a hippocampal ensemble and the spatial tuning of each neuron (Wilson and McNaughton, 1993; Zhang et al., 1998; Brown et al., 1998), hippocampal cell firing also reflects the stages of complex tasks (Eichenbaum et al., 1987; Wiener et al., 1989; Sakurai, 1990, 1994; Cohen and Eichenbaum, 1993; Hampson et al., 1993; Deadwyler et al., 1996; Gothard et al., 1996b). For example, in linear environments, hippocampal firing depends on the direction of travel (McNaughton et al., 1983; O'Keefe and Recce, 1993; Markus et al., 1995; Gothard et al., 1996a; Redish et al., 2000) and in certain spatial alternation tasks, hippocampal firing depends on the future and past paths of the animal (Wood et al., 2000; Ferbinteanu and Shapiro, 2003; Bower et al., 2005; Ainge et al., 2005; Ferbinteanu et al., 2006). Theories have suggested that these nonspatial changes may reflect multiple submaps reflecting the different *reference frames* needed for goal planning (Touretzky and Redish, 1996; Redish and Touretzky, 1997; Redish, 1999; Touretzky and Muller, 2006) or rapid changes in attention to cues (Worden, 1992; McNaughton et al., 1994; Fenton et al., 1998; Zinyuk et al., 2000). If the hippocampus represents an environment in multiple ways depending on a task's demands, then switching between these reference-frames should be detectable as a switch in the active cell-assemblies.

Fenton and Muller (1998) found that place cells exhibit extreme temporal variability or *overdispersion* in that their firing patterns are much less reliable tempo-

[1] **Graduate Program in Neuroscience, University of Minnesota, Minnneapolis, Minnesota;** [2] **Department of Neuroscience, University of Minnesota, Minneapolis, Minnesota**
J. Jackson is curently at Swammerdam Institute of Life Sciences, University of Amsterdam.
Grant sponsor: NIH; Grant number: R01-MH06829; Grant sponsor: NSF IGERT; Grant number: 9870633.
*Correspondence to: A. David Redish, Department of Neuroscience, University of Minnesota, Minneapolis, Minnesota.
E-mail: redish@ahc.umn.edu
Accepted for publication 11 July 2007
DOI 10.1002/hipo.20359
Published online 31 August 2007 in Wiley InterScience (www.interscience.wiley.com).

rally than would be predicted by a spatial, inhomogeneous Poisson process based on the neuron's own spatial tuning curve. Fenton and colleagues suggested that the overdispersion they observed may be the result of the rat switching reference frames at a mean rate of 1 to 2 times per second (Lánský et al., 2001; Olypher et al., 2002). Lánský et al. (2001) characterized this overdispersion in terms of a doubly-stochastic Poisson process that switches between two mean spike emission rates at a mean interval of 1–2 s. Olypher et al. (2002) showed that the overdispersion decreases to approximately expected levels during navigation to a goal on a goal-directed task and suggested that the reference-frame switching may be due to changes in transient goals.

Later, Harris et al. (2003) reported an improvement in the prediction of a place cell's firing when the activity of other neurons in the ensemble was taken into account. The improvement in prediction resulting from these subgroups of neurons lasted for a short duration (~25 ms; Harris et al., 2003). It was argued that these neuronal dynamics are similar to Hebb's concept of cell assemblies and may be evidence of "internal cognitive processes" (Harris et al., 2003; Harris, 2005).

Taken together, these reports suggest that the extreme variability in the spatial responses of neurons may be due to internal dynamics of the hippocampus that result in the formation of cell assemblies. It is not known, however, whether the variability reported by Fenton and Muller (1998) is related to any ensemble-level modulation, and previous data suggest this may not be the case (Fenton and Muller, 1998). Furthermore, the effect of cell assembly-like dynamics on the spatial tuning of neurons and the relation of these dynamics to an animal's internal goals is not understood. We therefore hypothesized that if multiple maps are used by an animal during exploration, this should resemble ensemble-level modulation that may explain overdispersion. If so, switches between maps may be tied to task parameters, such as switching from goal seeking to foraging.

Using neural ensemble recording methods, we tested the cell-assembly hypothesis of Fenton et al. by recording hippocampal neural ensembles from rats running three different tasks of varying spatial complexity: a one-dimensional linear track, a two-dimensional cylinder-foraging task, and a two-dimensional cylinder-goal task. The directional firing of hippocampal place cells on linear tracks provides a baseline test of the reference-frame hypothesis: (1) taking direction into account produces two separate representational maps which should have less overdispersion than when rats run in two-dimensional tasks; (2) if overdispersion results from unexpected switching between internal maps, then ignoring directional differences on the linear track should produce overdispersion (i.e. serving as a positive control). The cylinder-foraging task provides a baseline since it is a commonly used experimental condition and the variability of place cell firing on this task has been previously documented (Fenton and Muller, 1998). The cylinder-goal task provides our experimental test for the goal-dependent reference-frame switching hypothesis since overdispersion is known to change in a behaviorally-dependent manner on this task, possibly indicating changes in the set of active reference frames (Olypher et al., 2002). Comparing these three tasks, we found that the directional firing on the linear track mimicked the variability observed in the foraging and goal tasks and that invoking cell-assembly concepts allowed the separation of multiple maps on all three tasks. These transitions between active maps depended on the behavioral state of the animal, switching just after food delivery on tasks where reward was tied to the animal's spatial behavior.

METHODS

Experimental Methods

Subjects

Male Brown-Norway Fisher-344 hybrid cross rats were housed individually in a specific pathogen-free (SPF) vivarium maintained on a synchronous day/night cycle. Animals were handled daily for 15 min for at least 1 week prior to beginning behavioral training. One day prior to commencement of behavioral training, animals were denied access to food in their home cage while water access remained ad libitum. Subsequently, animals received their full daily complement of food on the tasks based on their behavioral performance. All procedures were approved by the University of Minnesota IACUC and met all NIH guidelines for animal use in research.

Behavioral training

Food deprived rats were trained to run on a series of multiple tasks, including shuttling back and forth along a 137 cm by 15 cm linear track (LT), foraging for scattered pellets in an 92 cm-diameter cylindrical arena (CF), or navigating to a small, invisible goal for food reward in the same cylindrical arena (CG).

The linear track task was similar to the shuttling tasks studied by O'Keefe and Recce (1993) and others: 2–4 (depending on the animal and session), 45-mg food pellets (TestDiet, Richmond VA) were delivered via automated feeders (Med Associates, St. Albans VT) when the animal reached the end of the track. The linear track was placed in the same position in the room each day, thus making the location of reward-delivery as well as the location of reward-receipt constant across days in the room reference frame. Animals were required to alternate between track ends to receive food.

The cylinder foraging task (CF) was a variant of that studied by Muller et al. (1987) and others: food was delivered randomly into the cylinder at 10-s Poisson intervals. Forty-five milligram food pellets were delivered from one of three sites above the cylinder chosen randomly. The cylinder was placed in the same position in the room each day, but food-delivery was independent of the behavior of the animal. Because of large variations in pellet speed and ejection trajectory, the food distribution that reached the cylinder floor was highly uniform and random.

The cylinder-goal (CG) task was a variant of that studied by Rossier et al. (2000) and Olypher et al. (2002): in order to receive food, the rat had to enter a 7-cm diameter goal region. Once the rat entered the goal region, a tone was played and three pellets were delivered from the automated feeders. Like the cylinder-foraging (CF) task, the food scattered randomly upon dispensing. The goal was not rearmed until the rat had been outside a 14 cm surrounding region for 4 s. The cylinder was placed in the same position in the room (the same position as for the CF task). The goal was randomly placed within the cylinder on each session, but remained at a constant location within each session. Since the CF and CG tasks used the same arena, the arena was wiped down with 70% isopropyl alcohol before each CF and CG task to reduce carry-over of odor cues.

Naive rats were trained on each task individually until proficient. Training began with single 30–40 min sessions on a single task (task training order counter-balanced across rats) until an animal was proficient on that task: full coverage of arena on CF; at least 30 successful goal entries on CG; at least 50 trials on LT. On the CG task, the time between successful goal triggers, is also referred to as a *trial.* Once the training criterion for a task was met, animals were then trained on the next task to proficiency. This continued until animals had been trained on each task individually. This usually took about 1 week per task. The single-task training criteria above were intended to ensure robust performance on three-task sessions and do not apply to later training and recording since the three-task sessions were shorter in duration. Next, animals were familiarized with the three-task protocol for at least 4 days such that they encountered each ordering at least once: LT–CF–CG; LT–CG–CF; CF–CG–LT; CG–CF–LT. Thus, final-training and the post-implantation recording sessions consisted of 20-min exposure to each of the three tasks pseudorandomly ordered each day (goal location also varied pseudo-randomly each day) with a 5 min rest period before and after each task.

Surgery

Once a rat was running proficiently on all three tasks, it was implanted with a 14-tetrode microdrive (Kopf Neuro-Hyperdrive, David Kopf Instruments, Tujunga, CA; 12 tetrodes and 2 references) at (Bregma −3.8 mm A/P, 2.0–2.5 mm M/L, Paxinos and Watson, 1998). Rats were deeply anesthetized with an intraperitoneal injection of Nembutal (sodium pentobarbital, 40–50 mg/kg, Abbot Laboratories, North Chicago, IL), shaved on the scalp, and placed on a stereotax. A 0.5–2.0% isofluorane–oxygen mixture was then provided to maintain general anesthesia. About 0.1 mL Dualcillin (Phoenix Pharmaceutical Inc., Saint Joseph, MI) injection was administered to each hind limb. The scalp was disinfected first with alcohol then with Betadine (Purdue Frederick, Norwalk, CT). Skin and fascia were removed from the skull around the implantation site, and the wound was cauterized. Holes were drilled for the 8–9 jeweler's screws and 1-ground screw which were distributed around the implant to anchor it to the skull. Once the screws were in place, a craniotomy was opened above the target using a surgical trephine (Fine Science Tools, Foster City, CA), and the hyperdrive was lowered into place. Ground screws were soldered to a steel wire prior to implantation and were connected to a steel wire from the hyperdrive ground terminal using Amphenol pins during implantation. A Silastic (Dow Corning 3140) barrier filled the space between the hyperdrive bundle and the skull. Dental acrylic (Perm Reline and Repair Resin, The Hygenic Corp., Akron, OH) was used to fix the hyperdrive to the bone screws and seal the wound. After removal from the stereotax, 3 mL saline was administered subcutaneously. Some rats received regimens of 0.1 mL Baytril (2.27% enrofloxacin, Bayer Corp., Shawnee Mission, KS) injected subcutaneously each day following surgery for 3 days. Animals received 0.8 mL Children's Tylenol orally immediately upon waking and in their water supply (25 mL mixed in 0.275 L water) during recovery.

Recording

After surgery, electrodes were advanced into the pyramidal layer of the CA1 region of hippocampus over the course of approximately 1 week. The pyramidal layer was identified by the presence (and wave morphology) of strong high-frequency (100–200 Hz) ripples (Ylinen et al., 1995). Recordings were carried out in 10 ft × 10 ft room enclosed in copper screen. All electrophysiological and video tracker recordings were digitized and synchronously time-stamped by a 64-channel Cheetah Data Acquisition system (Neuralynx, Tucson AZ) and recorded to disk.

Extracellular action potentials were recorded at 32 kHz for a 1-ms window when the voltage crossed a threshold set by the experimenter on any of the four channels on a tetrode. A 1-ms window was taken for each action potential consisting of 32 samples per spike-waveform per channel. The signals were first amplified at the headstage with unity gain amplifiers, then passed through multistrand cables and a commutator before reaching variable gain amplifiers (1–50,000×). There, they were band pass filtered from 600 to 9,000 Hz for spike recordings using 48 channels of a Neuralynx 64 channel Cheetah system, or filtered from 1–475 Hz and sampled at 2 kHz for local field potential recordings (LFP) using 16 channels of the same Neuralynx Cheetah system. Binding of recording cables due to rotation of the rat was minimized by a torque-sensing, motorized 72-channel commutator (Neuralynx, Tucson, AZ; Dragonfly, Ridgeley, WV; AirFlyte, Bayonne, NJ) and on rare occasions corrected by the experimenter.

The positions of LEDs mounted on the animal's headstage were detected by a camera mounted in the center of the recording room's ceiling. The video frame data was sampled at 60 Hz and digitized and time-stamped by a Cheetah data acquisition system (Neuralynx, Tucson, AZ); for each video frame, pixels with luminance above an experimenter-defined threshold were recorded to disk. Real-time position data was accessed by in-house behavioral control software implemented in Matlab (The Mathworks, Natick, MA). This software used the serial ports to communicate with an experimental control box (constructed by

JCJ) to trigger food delivery (45 mg pellets: Research Diets, New Brunswick, NJ; food dispensers: Med-Associates, St. Albans, VT) and simultaneously signal the Cheetah recording system for a synchronous food delivery time-stamp (each feeder had a unique digital identification).

Position data were then preprocessed for post hoc analysis by extracting the center of mass of all pixels with suprathreshold luminance. Video interlacing effects were removed from the data through linear interpolation of odd and even position samples (two 30-Hz time-series) to produce two 60-Hz time series, which were then averaged to yield a single, stable 60-Hz time series.

Data Analysis

Place-fields

The tuning of a cell is the average or expected firing rate of a neuron measured over a given behavioral variable; in this case, this variable is the animal's spatial location. Spatial tuning curves, or place fields, were constructed by binning the task area into 11 pixel $\times$ 11 pixel (approximately 3 cm $\times$ 3 cm) bins and creating two 2-dimensional histograms: a histogram of the number of spikes emitted in each bin and a histogram of the number of video-tracker samples in each bin. The occupancy time for each bin was determined by dividing the number of positions samples per bin by the video sampling rate. The firing rate per bin was determined by dividing the spike count in each bin by the occupancy time in that bin.

SW and theta detection

To reduce the possibility of known network states, such as LIA and SW, contaminating the analysis, we removed all SW events and used only data from high-theta/low-delta periods. SW events and theta epochs were detected and defined as in Jackson et al. (2006).

SW events were extracted by down-sampling the LFP traces by a factor of 2 (using an anti-aliasing low-pass filter), and bandpass filtering from 100 to 250 Hz. Amplitude for each trace was found via Hilbert-transform and then averaged across traces. The distribution of log-transformed average amplitude was used to find samples more than 2.5σ from the mean power. Visual inspection of a subset of the data revealed ripple events synchronous across LFP channels. Threshold crossings shorter than 20 ms were removed, the remaining events were concatenated if less than 100 ms apart. Twenty millisecond was added to the beginning and end of each SW to capture the tails of the SW.

A similar method was used for detecting theta epochs. Theta times were extracted by down-sampling the LFP traces by a factor of 5 (using an anti-aliasing low-pass filter), and bandpass filtering from 6 to 10 Hz to obtain theta-band signals, and from 2 to 4 Hz to obtain delta-band signals. Amplitude for each trace's band was found via Hilbert-transform and then averaged across traces to obtain two averaged signals: an average theta-band amplitude and an average delta-band amplitude. The distribution of the log-transformed ratio (theta/delta) of average amplitudes was used to identify samples with a low power-ratio more than 1σ from the session mean, these were taken to be *non-theta* brain states. Visual inspection of a subset of the data revealed low theta amplitude epochs that clustered at locations of immobility (i.e. the linear track ends). These non-theta epochs were concatenated if less than 500 ms (the low frequency cutoff for the delta-band) apart, and events smaller than 100 ms (the high frequency cutoff for the theta-band) were removed. These high theta band power and low delta band power epochs tended to coincide with times when the animal was moving.

Measuring firing rate dispersion

In order to quantify the variability of single cell firing, a statistical measure is needed that allows the inclusion of neuronal response parameters. One possible model of the behavioral or stimulus dependence of a neuron's firing is that of an inhomogeneous Poisson point process where the intensity of the spike emission rate at any given time is dependent on the stimulus or behavioral value at that moment. This was the approach used by Fenton and Muller (1998) to examine the variability of hippocampal neurons as the animal passes through their place field. In order to compare the results of our experiment, we implemented the analysis described in their paper (Fenton and Muller, 1998).

First, a neuron's place field was taken as the largest contiguous body of bins with nonzero firing rate. Therefore, every bin in a place field shared a border with at least one other pixel in the place field. The center of a place field was taken as the 3×3 group of pixels with the highest mean firing rate.

An animal's pass through a place field was only considered if it satisfied all of the following conditions: (1) the pass went through the center of the place-field, (2) the pass lasted longer than 1 s, (3) the tracking of the animal was continuous throughout the pass with high theta rhythm and not interrupted by sharp-waves. The actual number of spikes S of a neuron during each pass through its place field that met the above conditions was compared with the expected number of spikes N given the spatial tuning specified by its place field. As described in Fenton and Muller (1998), for $N > 4$, the Poisson distribution can be approximated by a normal distribution with mean $\mu = N$ and variance $\sigma^2 = N$. Thus, the Z-transform of the number of spikes transmitted relative to N for each pass through a place field can be calculated as follows:

$$Z = \begin{cases} \dfrac{S - N - \frac{1}{2}}{\sqrt{N}} & \text{if } S \geq N, \\ \dfrac{S - N + \frac{1}{2}}{\sqrt{N}} & \text{if } S < N, \end{cases} \tag{1}$$

where S is the number of spikes actually emitted on each pass. The factor of 1/2 is a correction for the discrete distribution.

Pair-wise correlation analyses

Since our experiments generated large ensembles of simultaneously recorded neurons, we examined correlations in the dispersion *z*-score between pairs of simultaneously recorded neurons in our data. The overlap of place-fields for all neuron pairs in an ensemble was measured as the number of bins greater than zero firing rate in both neurons' tuning curves. The distribution of overlap values was divided such that pairs in the upper 75% of overlap scores for each task were chosen for analysis. The pass times for each neuron in a pair were then compared to find all passes that overlapped temporally by any amount. The dispersion z-score values for these matched passes were then added to the pool of pair-wise data. Correlations were then performed on this pooled data. Higher and lower overlap cutoffs (50% and 95%, when possible) were also assessed and qualitatively similar results were obtained. Data were compared with a control condition in which one neuron's *z*-score values were randomized across passes prior to matching the pass times for the pair.

Map splitting

If cell assemblies are subsets of coactive neurons bound together through repeated cofiring (Hebb, 1949), then their functional properties in the hippocampal network at a given location should be observable as distinctly uncorrelated (not necessarily anti-correlated) ensemble firing patterns on passes through the same location if different cell assemblies exist. We hypothesized that neurons are therefore switching between activity states depending on the activity states of other neurons in the ensemble (See H1, Fig. 1A). This is in contrast to switches in activity states independent of other neurons in the ensemble (See H0, Fig. 1A). To examine the spatial properties of such subensemble interactions, we clustered the firing patterns that occurred within each spatial bin of an environment to construct multiple whole-environment spatial firing maps. The process described below is depicted in Figure 1C. Sufficient numbers of simultaneously recorded neurons were required for all map-splitting and coherency analyses described below, therefore only the 10 sessions with greater than 25 simultaneously recorded neurons were used.

Clustering of firing rates within a spatial bin. Firing rates were calculated for each neuron by binning spikes into 10-ms bins and convolving the result with a 100-ms exponential decay function normalized within a 500-ms convolution window. For each 11 pixel $\times$ 11 pixel bin (3 cm $\times$ 3 cm bin), all ensemble firing patterns observed when the animal was in that bin were clustered using a *k*-means algorithm using correlation as the distance metric (i.e. $d = 1 - \rho$, where d is the distance between two samples and ρ is the correlation between those two samples). The correlation metric was used since it is bounded and robust to changes in whole ensemble excitability (i.e. fluctuations in ensemble mean firing rates) and therefore allows the grouping of correlated firing patterns.

Since previous research has suggested that the overdispersion phenomenon is best described using a two-state model (Lánský et al., 2001; Olypher et al., 2002), we set the *k*-means algorithm to output 2 clusters. Larger numbers of clusters were also tried; however, extracting more than 2 clusters fragmented the data too much for the analyses. Thus, there was insufficient data for dispersion analyses with more than 2 clusters. Figure 1 shows two patterns isolated from one 3 cm $\times$ 3 cm bin. Based on the clusters for each bin, we constructed two preliminary 2-dimensional firing rate maps (labeled Sub-Map 1 and Sub-Map 2 in Fig. 1C) by sorting each pixel's clusters into one or the other map by maximizing correlations of the cluster's mean firing pattern with the mean firing pattern of clusters in the neighboring pixels. Finally, the times when the animal's firing patterns were detected in each preliminary map were extracted and used to partition the behavior into either of the two representational states. (We use the term "state" in this paper to define the map, or stack of place fields, that the ensemble activity pattern best corresponds with. For example, saying the hippocampus is in "State 1" corresponds to the ensemble firing being more closely related to Map 1 than Map 2.) These times were referred to as the *map-switching times.* Ensemble switching times were then used to construct state-dependent place fields for each neuron. We refer to this process of deriving multiple spatial tuning curves from ensemble firing patterns as *map-splitting.*

Dispersion (leave-one-out). To test against the null hypothesis that neurons switch firing states independent of the rest of the ensemble (see H0, Fig. 1A), we examined the effect that switching maps would have on the dispersion of neuronal firing rates using a leave-one-out approach. For each neuron in an ensemble, the map-splitting and assembly analysis was performed on the ensemble excluding that neuron to obtain switching times for the ensemble. These ensemble switching times were applied to the left-out neuron to partition the position and spike data into separate states for the construction of spatial tuning-curves (i.e. place fields) for each state. The dispersion analysis was then run for the state times using these supposed state-dependent firing fields. Only neurons in ensembles that were large enough to perform the map-splitting analysis and only passes through a place field that were not interrupted by a map-switch were used. As an additional control for the partitioning of states, dispersion *z*-score distributions were calculated from spatial tuning-curves that were constructed by applying shuffled map-switching intervals from sessions drawn at random (including sessions from other animals) to each neuron's spike train. This was repeated 100 times resulting in dispersion values for two randomized maps each, totaling 200 *z*-score distributions for control dispersion variance estimates. The distribution of these variance estimates was approximately Gaussian. The distribution of variance estimates was used to calculate mean dispersion variance and 95% confidence inter-

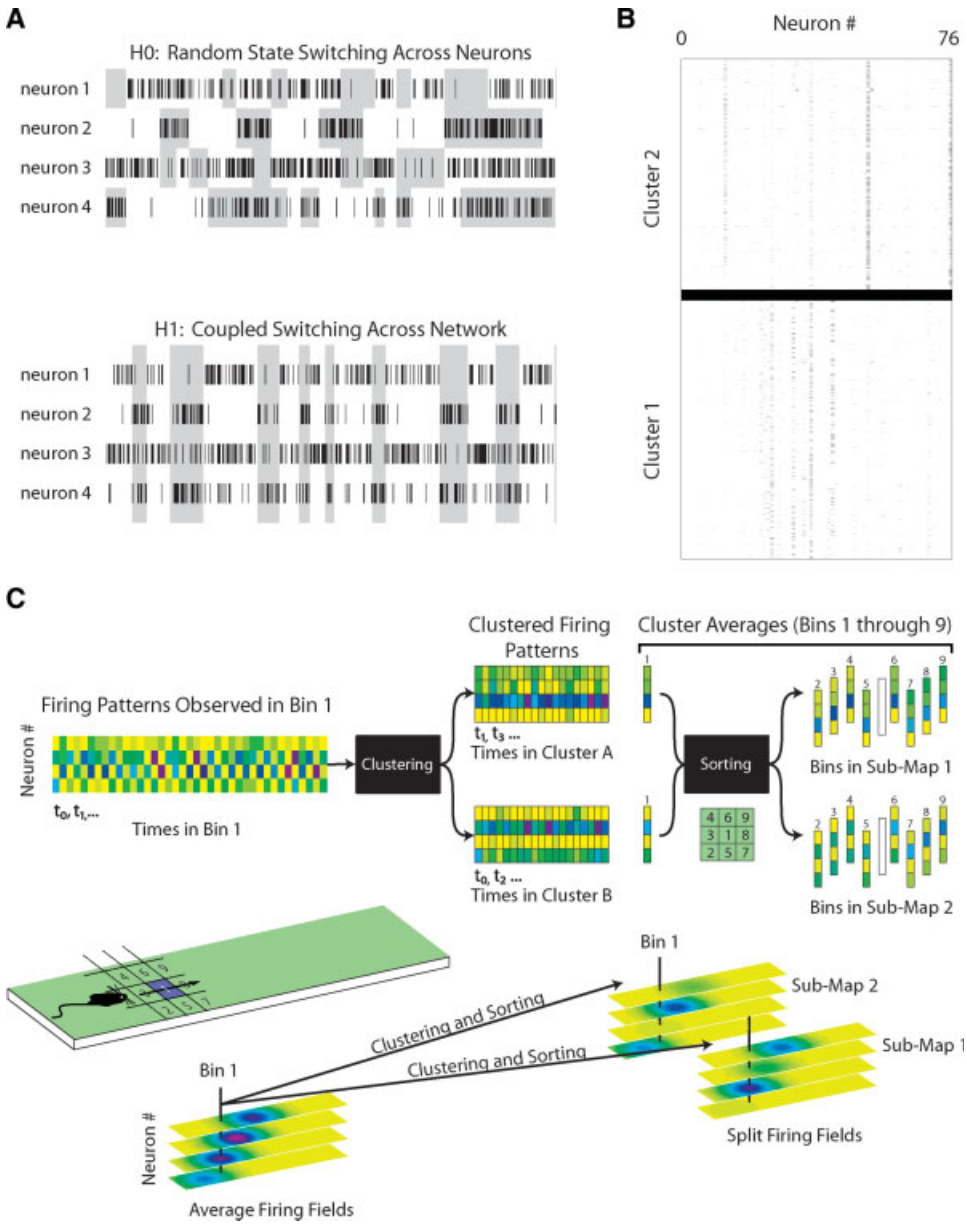

FIGURE 1. Splitting hypothetical maps. (A) Two hypotheses for state-dependent switching of neuronal firing rates. H_0 depicts random modulation of individual neuronal firing rates independent of modulation of other neurons in the ensemble. H_1 depicts random modulation of whole-ensemble firing rates. White and gray background depicts two states: neurons 1 and 3 fire maximally in the white state, and neurons 2 and 4 fire maximally in the gray state. (B) An example of actual firing rate vectors associated with each video tracker sample found within an 3 cm × 3 cm region were clustered using a *k*-means algorithm. Times have been sorted into one of two clusters. Times when no firing occurred within the ensemble are not shown. (C) Schematic of the splitting process. Average firing fields result from the customary method of averaging all firing patterns that occur in each spatial bin (e.g. Bin 1). To test whether multiple maps can be derived from hippocampal ensemble firing, the firing rate for each neuron (shading by firing rate) for each time detected in a bin of interest was clustered (Clustered Firing Patterns) then sorted into the map with the highest correlation for neighboring Bins (Bins in Sub-Map 1 and Bins in Sub-Map 2). Completing this process results in two sets of firing fields (i.e. maps; Sub-Map 1 and Sub-Map 2). The times associated with firing rates in each cluster were used to determine which map the animal was using at a given moment and to derive switching times between maps. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

vals under random splitting conditions with similar first-order state transition temporal statistics.

Map-switching measures. Whole-ensemble switching times were used for the rest of the analyses described below. Dwell time in each state was calculated as the product of the position sampling period and the number of position samples detected in each respective state, divided by the number of transitions into each state. Similarly, the switching rate was calculated as the number of transitions into each state divided by the prod-

uct of the position sampling period and the total number of position samples detected on the task. These values were averaged over the sessions available to this analysis.

The rate-difference ratio was used to quantify the differences in rate between place-fields in each map. The rate-difference ratio was calculated as

$$R = \frac{|\text{avg}(\text{PF}_{\text{Map}_1}) - \text{avg}(\text{PF}_{\text{Map}_2})|}{\max(\text{avg}(\text{PF}_{\text{Map}_1}), \text{avg}(\text{PF}_{\text{Map}_2}))}, \quad (2)$$

where $\text{avg}(\text{PF}_{\text{Map}_1})$ is the average firing rate of the place-field in Map_1, $\text{avg}(\text{PF}_{\text{Map}_2})$ is the average firing rate of the place-field in Map_2, max denotes the maximum of either the average rate in Map_1 or Map_2, and $|x|$ denotes the absolute value of x.

Coherency calculation. In order to test whether the maps derived by the map-switching analysis were consistent, stable features of the ensemble, we applied a coherency analysis previously shown to detect changes in reference frames (Redish et al., 2000) and aberrant network states (Jackson and Redish, 2003). The whole-ensemble switching times were used for this coherency calculation. *Coherency* measures the self-consistency of a representation within a neural ensemble (Jackson and Redish, 2003). Briefly, measuring coherency within a time window t over a behavioral variable (x, y, for place fields) entails calculating an activity packet (defined as the firing rate weighted tuning curves, Redish et al., 2000; Jackson and Redish, 2003), calculating an expected activity packet (using the expected firing rate rather than the actual firing rate, Jackson and Redish, 2003), and then comparing them using a comparison measure (Redish et al., 2000; Jackson and Redish, 2003; Johnson et al., 2005; Jackson, 2006; Johnson et al., in press). The actual and expected activity packets were compared using the I_{RMS} measure (Jackson and Redish, 2003),

$$I_{\text{RMS}}(t) = \frac{\sqrt{\int_x \int_y (A(x,y,t) - \hat{A}(x,y,t))^2 \, dx \, dy}}{\int_x \int_y \hat{A}(x,y,t) \, dx \, dy}. \quad (3)$$

I_{RMS} measures the difference or inconsistency between the two packets; $A(x,y,t)$ indicates the actual activity packet and $\hat{A}(x,y,t)$ indicates the expected activity packet:

$$A(x,y,t) = \frac{\sum_k \text{PF}_k(x,y) \cdot F_k(t)}{\sum_k \text{PF}_k(x,y)} \quad (4)$$

and

$$\hat{A}(x,y,t) = \frac{\sum_k \text{PF}_k(x,y) \cdot \text{PF}(\hat{x}(t), \hat{y}(t))}{\sum_k \text{PF}_k(x,y)}, \quad (5)$$

where $\hat{x}(t)$ and $\hat{y}(t)$ are the current x and y position of the animal at time t, $\text{PF}_k(x, y)$ is the tuning curve (or place field) of neuron k over the two-dimentional space x, y, $F_k(t)$ is the firing rate of neuron k at time t, and Σ_k is the sum over all neurons in the ensemble.

The I_{RMS} measure is sensitive to absolute differences in ensemble firing across the population. Other measures sensitive to relative differences or to similarities in ensemble firing across the population were also used and yielded qualitatively similar results (I_{STD}, I_{VAR}, for more details see Jackson, 2006). To calculate the *coherency ratio* $\text{C}_{\text{index}}^{1-2}$ to examine the maps derived from the map-splitting analysis, the data set was split in half by interleaving minutes. Thus, one half of the data set consisted of all *odd* one-minute blocks of data, and the other consisted of all *even* one-minute blocks of data. One half was used to derive the state-dependent place fields for the other given the switching times between maps and vice versa. This eliminates issues of tautology. The ensemble measure (i.e. I_{RMS}) of one half was then calculated given the place fields of the other half and vice versa. The coherency ratio was then defined as

$$\text{C}_{\text{index}}^{1-2}(t) = \frac{\textbf{Coherency}_\textbf{1}(t) - \textbf{Coherency}_\textbf{2}(t)}{\textbf{Coherency}_\textbf{1}(t) + \textbf{Coherency}_\textbf{2}(t)}, \quad (6)$$

where $\textbf{Coherency}_\textbf{1}$ and $\textbf{Coherency}_\textbf{2}$ are the coherency values for the two cluster-derived maps, respectively. $\textbf{Coherency}(t)$ was defined as the proportion of times in the same data set that the actual and expected activity packets matched as well or better than the sample of interest at time t:

$$\textbf{Coherency}(t) = 1 - \text{cdf}_{I_{\text{RMS}}}(I_{\text{RMS}}(t)) \quad (7)$$

where $\text{cdf}_{I_{\text{RMS}}}$ (I_{RMS} (t)) means that the I_{RMS} values from each half were concatenated and used to calculated the cumulative distributions cdf. (This does not result in a tautology since this is used only for a relative comparison. At most, this concatenation only lowers our sensitivity working against a significant result.)

This index $\text{C}_{\text{index}}^{1-2}$ of how coherent one state was with respect to the other will be above zero if the ensemble firing pattern is more similar to map 1 than to map 2, and less than zero if ensemble firing is more similar to map 2 than to map 1. The coherency of firing patterns occurring in one half of the session was calculated using place fields derived from ensemble firing given the switching times in the other half of that session, thereby circumventing any tautological issues. The coherency index was then aligned to switching times from one state to the other for each task for the half of the data that was not used to construct the place fields. Similar results were found when the tautology was used. Ensemble coherency values were calculated for each theta cycle, with each theta cycle identified by positive peaks in the filtered local field potential.

Behavioral analysis: peri-event time histograms. We constructed peri-event time histograms (PETH) to examine the relationship between task events and transitions between representational maps. For the PETH switching analysis, the time of switching was measured relative to the onset of key task events such as food delivery (on CF and LT) or to goal-entry (on CG). The number of switching times at each time lag were binned into 0.1 s bins for the 6 s before and after each event,

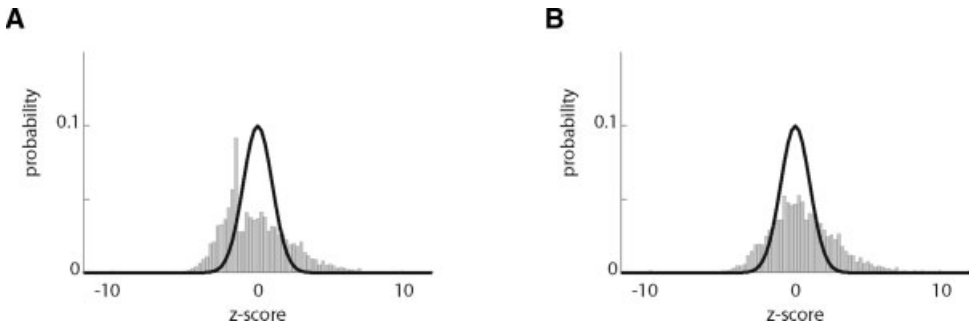

FIGURE 2. Overdispersion of hippocampal place cells on the Cylinder Foraging (CF) task. Figure shows histogram of *z*-scores for the number of spikes emitted on a pass through a place field given the expected number of spikes predicted by a Poisson point process model. (A) Place cells on the cylinder foraging (CF) task exhibited a comparable amount of excess variability to previous reports (e.g. see Fenton and Muller, 1998). Sharp peak at negative *z*-scores is due exclusively to no-fire passes. Black solid line represents expected variability of a Poisson point process model. (B) Data same as in A but excluding passes through a place field when no spikes were fired.

summed for all event times, and normalized by the number of position samples detected in each bin to yield the transition rate at each time lag from the event. The rates for each session were averaged and the standard error of the mean was calculated for each bin. The transition rate after an event was compared with the transition rate before the event using an unpaired *t*-test given the mean across sessions. The 95% confidence intervals corrected for multiple comparisons were found using a bootstrap. For each session, 50 randomly selected event times were drawn from a uniform distribution and PETH of switching times were created for each of these pseudosessions. The distribution of 50 runs $\times$ 120 bins followed a normal distribution very closely. The mean and standard deviation of this distribution were used to calculate the 95% confidence intervals corrected for multiple comparisons ($.025/N_{\text{times}}$, $.975/N_{\text{times}}$), where $N_{\text{times}} = 120$ (the number of bins per PETH).

RESULTS

Six hundred ninety-six spike-trains were recorded from 6 rats over 24 sessions in ensembles of up to 96 neurons/session (30 $\pm$ 31 neurons/session; mean $\pm$ SD) while rats ran three tasks. Neurons with session average firing rates of >2.0 or <0.04 spikes/s (i.e. firing less than 150 spikes in the entire recording session) were excluded from analysis. These restrictions effectively removed interneurons and cells not specifically activated by behavioral tasks leaving 600 neurons for further analysis. Rats ran 25 $\pm$ 12 (SD) trials/session on LT, 37 $\pm$ 15 (SD) trials/session on CG, and tended to cover the environment on CF.

Replication of Previous Single-Cell Variability Measurements: Overdispersion

While hippocampal neurons have highly reliable spatial responses, the temporal variability of these spatial responses is known to be quite large when compared to a noisy model based on Poisson statistics (Fenton and Muller, 1998). Consistent with the previous reports (Fenton and Muller, 1998; Lánský et al., 2001; Olypher et al., 2002), hippocampal place cells showed high variance in their firing rates across place field passes on the CF task ($\sigma^2 = 6.0$, $\mu = -0.02$, $N_{\text{passes}} = 3830$, $n_{\text{rats}} = 6$, χ^2 test, $p(\sigma^2 = 1.0) < 10^{-10}$; number of qualifying passes per place field = 10.37 $\pm$ 0.40 SEM). The mean of the dispersion *z*-scores was not significantly different from zero ($p(\mu = 0) = 0.59$, $t = -0.53$, $df = 3829$). Figure 2 shows the distribution of firing rate *z*-scores comparing the number of spikes emitted on each pass through a place field to the expected distribution of spikes given an inhomogeneous Poisson process model. The thin black line shows the expected distribution with zero-mean and unit variance. The sharp peak at negative *z*-scores was the result of passes through place fields when no spikes were fired. Thus, removing no-fire passes resulted in a smooth, unimodal distribution of similar width (Fig. 2B). The negative peak could not be explained by the inclusion of low-firing rate neurons since exclusion of neurons with less than 10 Hz in-field firing rates did not abolish this negative peak nor did it visibly alter the positive skew (data not shown).

It has been reported that this extreme temporal variability is reduced when animals are required to navigate to a goal. A similar overdispersion was seen on the CG task ($\sigma^2 = 6.7$, $\mu = 0.05$, $n_{\text{passes}} = 4419$, $p(\sigma^2 = 1) < 10^{-10}$; number of qualifying passes per place field = 11.25 $\pm$ 0.49 SEM). The mean of the dispersion z-scores was not significantly different from zero ($p(\mu = 0) = 0.19$, $t = 1.30$, $df = 4418$). Consistent with reports by Olypher et al. (2002), we observed a similar reduction in variance in the 5 s pregoal ($\sigma^2 = 5.1$, $N_{\text{passes}} = 896$, $n_{\text{rats}} = 6$, significantly less than overall dispersion, *F*-test, $p < 10^{-6}$, $F = 1.3$), there was a general reduction in the variance as the animals approached the goal: $\sigma^2_{4s} = 4.9$, $\sigma^2_{3s} = 5.1$, $\sigma^2_{2s} = 4.6$ (See Fig. 3). (All four conditions less than overall dispersion $p < 10^{-4}$, but the decrease as the animal approached the goal was not significant.)

Fenton and colleagues suggested that the overdispersion they observed may be the result of the rat switching goals or refer-

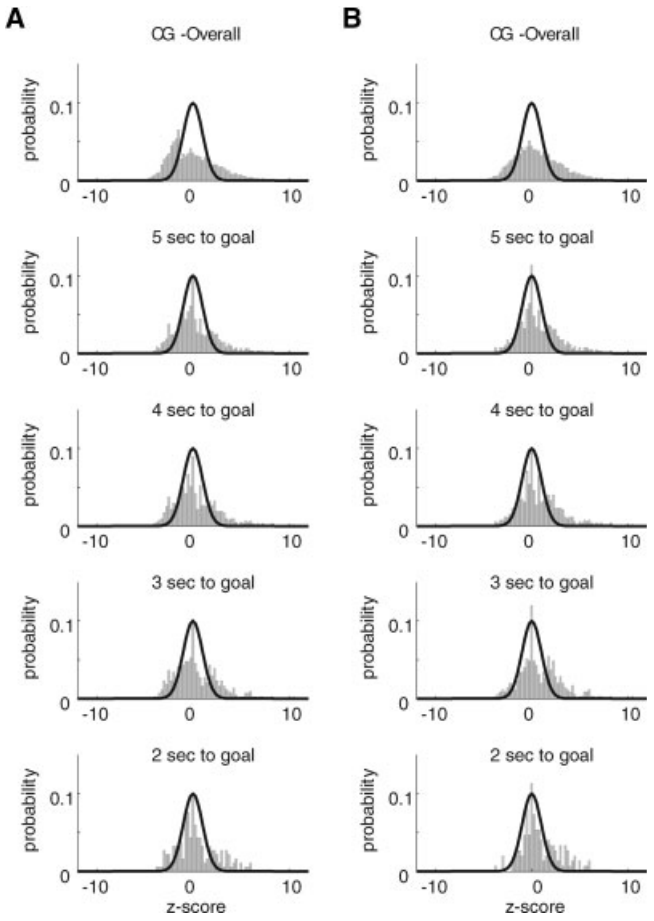

FIGURE 3. Overdispersion of hippocampal place cells on the Cylinder Goal (CG) task. (A) Overall variability on the CG task (σ^2 = 6.7, N_{passes} = 4419) and variability on approach to goal: 5 seconds prior to goal entry (σ^2_{5s} = 5.0, N_{passes} = 855); 4 seconds prior to goal entry (σ^2_{4s} = 4.8, N_{passes} = 648); 3 seconds prior to goal entry (σ^2_{3s} = 4.9, N_{passes} = 439); 2 seconds prior to goal entry (σ^2_{2s} = 4.3, N_{passes} = 185); Note that the primary change to the distribution of *z*-scores observed as the animal approaches the goal, is the reduction in the low-*z*-score peak associated with no-fire passes and an increase in the positive tail of the distribution. This indicates that the chance that at least some spikes are fired is increasing along with the overall firing rate. Black solid line represents expected variability of a Poisson point process model. (B) Data same as in A but excluding passes through a place field when no spikes were fired.

ence frames (Fenton and Muller, 1998; Lánský et al., 2001; Olypher et al., 2002). The directional firing of hippocampal place cells on linear tracks provides a baseline test of the reference-frame hypothesis (McNaughton et al., 1983, 1996; O'Keefe and Recce, 1993; Touretzky and Redish, 1996; Redish, 1999). Taking direction into account on the linear track produces two separate representational maps (Fig. 4), which should have less overdispersion than when rats run in 2 dimensional tasks. Likewise, we predicted that if the high variance observed in the CF and CG tasks is related to the switching of reference frames, then taking a task known to demonstrate two separate firing maps for each direction should generate overdispersion of similar magnitude to that seen on CF and CG.

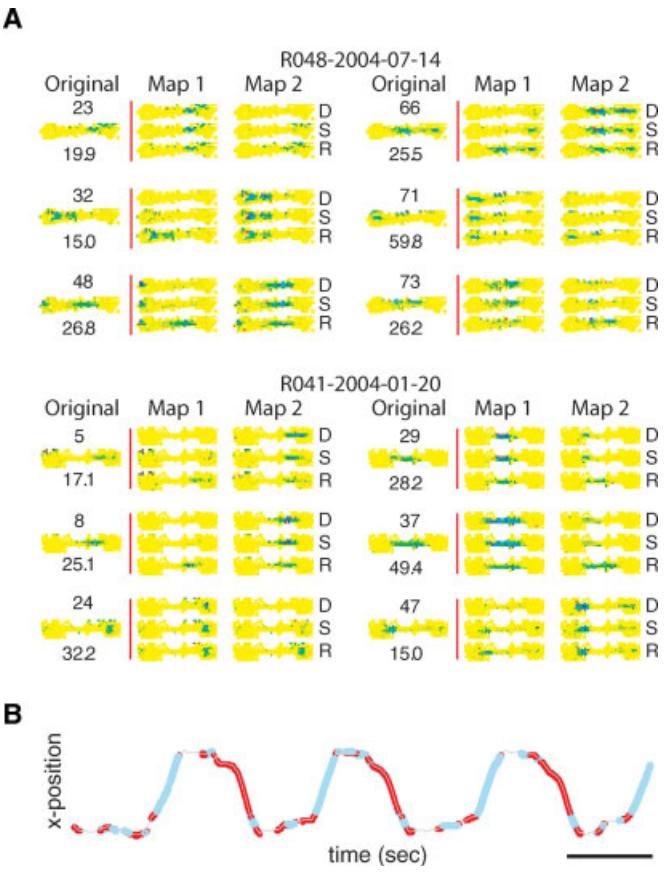

FIGURE 4. Map-splitting analysis yields directional specificity of maps (A) Original and split spatial firing fields of 6 sample neurons from each of two different rats. The firing fields shown to the right of the line were created by one of three methods: (D) splitting a session based on the animal's direction of travel, (S) splitting the session based on ensemble firing patterns using the map-splitting analysis (See methods), or (R) splitting the session based on randomized times. The assignment of Map 1 and Map 2 for S was based on which map (i.e. stack of place fields) correlated best with Map 1 vs. Map 2 derived by the directional split D. Randomly split maps used shuffled switching times, but were matched in a similar manner across tasks. Small number in the lower left-hand corner of each group is the firing rate associated with dark on the map. Not all neurons are represented, however all neurons shown for each session were recorded simultaneously. (B) An example of the positions where an animal was detected in each state by our map-splitting analysis. Red and green samples show times when hippocampal ensemble firing patterns were detected in one or the other map. Thin gray points show all positions sampled. Note that the red and green times segregate well to movement in one direction or the other. Gaps in red and green represent low-theta power epochs or sharp-wave ripples that were excluded from analysis; these generally correspond to resting periods at the track ends. Scale bar is 10 s, displacement along the x-direction is approximately 1 m. State splitting was accomplished by clustering firing rate vectors into two clusters using a k-means algorithm and sorting into one of two maps by maximizing correlation of cluster means with neighboring pixels based on cluster means. The times that the ensemble firing was clustered to one map or the other were used to partition the data set for construction of the place fields for each map. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

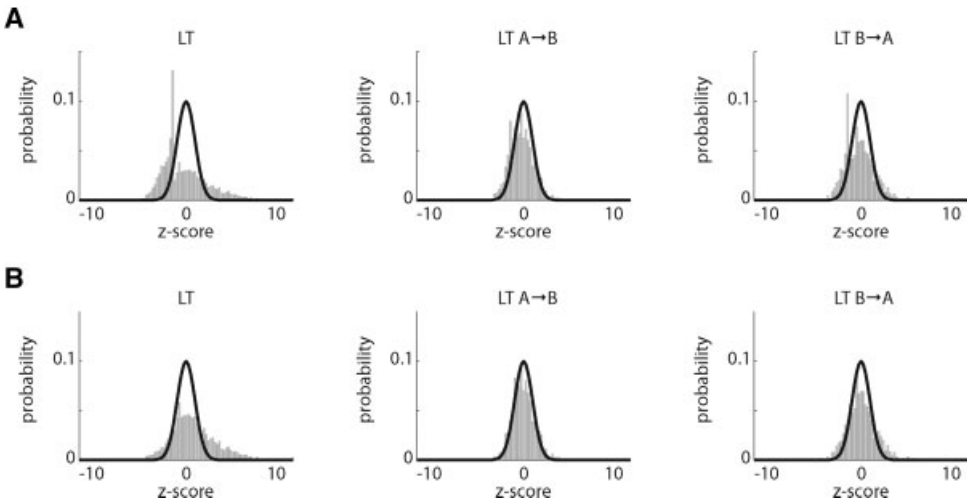

FIGURE 5. Overdispersion on the linear track (LT). (A) Ignoring directionality results in a highly variable distribution of firing rates ($\sigma^2 = 7.0$, $\mu = -0.14$, $N_{passes} = 3801$, $n_{rats} = 6$). However, splitting by direction (A-end to B-end versus B-end to A-end) resulted in a trend toward convergence of the actual firing rates with the Poisson-process based model (LT $A \rightarrow B$: $\sigma^2 = 2.1$, $N_{passes} = 922$; LT $B \rightarrow A$: $\sigma^2 = 2.8$, $N_{passes} = 1392$). Black solid line represents expected variability of a Poisson point process model. (B) Data same as in A but excluding no-fire passes.

Our predictions of the effects of switching representational maps based on the direction of travel on the neuronal response variability in linear track were confirmed. Ignoring directionality yielded a highly variable distribution of firing rates ($\sigma^2 = 7.1$, $N_{passes} = 3801$, $n_{rats} = 6$, $p(\sigma^2 = 1) < 10^{-10}$; number of qualifying passes per place field = 11.38 ± 0.59 SEM). The mean of the dispersion *z*-scores was significantly different from zero ($\mu = -0.14$, $p(\mu = 0) < 0.001$, $t = -3.36$, $df = 3800$). Importantly, splitting by direction resulted in a significant decrease in overdispersion (*F*-test, $p < 10^{-10}$), showing a strong trend toward convergence of the actual firing rates with the Poisson-process based model (See Fig. 5; LT $A \rightarrow B$: $\sigma^2 = 2.1$, $N_{passes} = 922$; LT $B \rightarrow A$: $\sigma^2 = 2.8$, $N_{passes} = 1392$; $n_{rats} = 6$.). Figure 5B shows how removing the few remaining no-fire passes from the directionally-split data results in a *z*-score distribution of near unit variance.

These results are summarized in Table 1.

Pair-wise Correlations

Since network-level switching of reference frames or the temporary transition between cell assemblies should result in correlated increases and decreases among some cell pairs with overlapping place-fields, we began our examination of the these hypotheses with a pair-wise analysis. For each task, the dispersion of cell pairs with the most overlap was compared on passes when the animal went through both place fields (the pairs with overlap in the top 75% among all pairs on a task, see Methods; LT: 16 or more pixels, approximately 144 cm^2 or more; CF: 20 or more pixels, approximately 180 cm^2 or more; CG: 22 or more pixels, approximately 198 cm^2 or more; higher and lower overlap cutoffs—50% and 95%, when possible—were also assessed and qualitatively similar results were obtained). On the cylinder foraging (CF) task, we found that there were indeed weak but highly significant positive correlations between dispersion *z*-score for neurons with overlapping place fields ($\rho = 0.047$, $P(\rho = 0) = 0.0048$, $N_{pairs} = 966$). This correlation was absent in randomized controls (randomizing the order of the z-scores for the mutual passes [passes that went through both neurons' place-fields] for one neuron in each pair, $\rho = 0.0092$, $P(\rho[randomized] = 0) = 0.59$) suggesting that sheer increases in the number of points were not resulting in a false appearance of coupling between cells. On the cylinder goal (CG) task, there was a similar level of correlation ($\rho = 0.049$, $P(\rho = 0) = 0.00095$, $N_{pairs} = 1105$) which was absent in the randomized controls ($\rho = 0.016$, $P(\rho = 0) = 0.27$). This correlation almost doubled when considering the five seconds prior to goal entry ($\rho = 0.095$, $P(\rho = 0) = 0.036$, $N_{pairs} = 323$). The linear track with its two-reference-frame nature was, however, not correlated significantly under any condition (LT non-directional: $\rho = -0.011$, $P(\rho = 0) = 0.48$, $N_{pairs} = 897$; LT

TABLE 1.

Task-Dependent Variability

	All passes		Exluding no-fire passes	
Task	σ^2	N_{passes}	σ^2	N_{passes}
LT	7.1	3,801	7.0	2,513
LT $A \rightarrow B$	2.1	922	1.9	822
LT $B \rightarrow A$	2.7	1,392	2.6	1,210
CF	6.1	3,830	5.9	2,993
CG	6.7	4,419	6.4	3,547
CG 5s	5.1	855	4.8	687
CG 4s	4.8	648	4.6	522
CG 3s	4.9	439	4.7	350
CG 2s	4.3	185	3.8	150

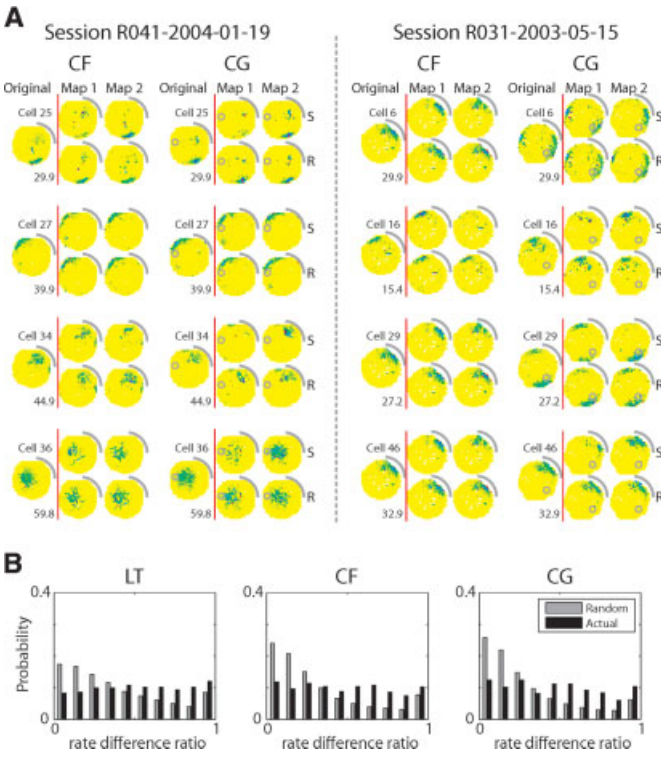

FIGURE 6. Two distinct maps are found on CF and CG. (A) Example place fields before and after map-splitting from two different sessions from two different rats. Original place-fields are shown to the left of the line for each group; split maps are shown on top row to the left of the line (S); and randomly split maps are shown on the bottom row (R). Firing rate vectors were clustered into two clusters using a *k*-means algorithm and sorted into one of two maps by maximizing correlations with neighboring pixels based on cluster means. Place fields were then constructed for times when the ensemble firing patterns were associated with each map. The assignment of Map 1 and Map 2 for CF and CG was based on which map (or stack of tuning curves) on CF correlated best with Map 1 vs Map 2 on CG. Randomly split maps used shuffled switching times, but were matched in a similar manner across tasks. Small number in the lower left-hand corner of each group is the firing rate associated with dark on the map. Gray thick arc represents cue-card location; gray circle represents goal location (note: goal location was different for every session). Not all neurons are represented, however all neurons shown for each session were recorded simultaneously. Note there are instances of switched maps from CF to CG (e.g. Cells 16 and 29 of Session R031-2003-05-15), instances of dropped place-fields between Map 1 and Map 2 (e.g. Cells 25 and 34 of Session R041-2004-01-19 and cell 16 Session R031-2003-05-15), instances where bimodal place fields are separated into two separate maps (e.g. Cells 6 and 29 of Session R031-2003-05-15), and even instances of rate remapping between maps where one field has a higher firing rate in one map than in the other (e.g. Cells 27 and 36 of Session R041-2004-01-19 and cell 46 Session R031-2003-05-15). Comparing CF and CG revealed slight differences in coding across maps (e.g. Cell 36 of Session R041-2004-01-19 and cells 16 and 29 Session R031-2003-05-15). (B) Map-splitting results in non-random rate-remapping across the ensemble. Distributions of the rate difference ratio between split maps is shown for each task for ensemble-based splitting ("Actual"; black bars) and for bootstrapped distributions of splitting based on a randomized switching times ("Random"; gray bars). Actual splitting on all tasks was significantly different from randomized splitting (LT: $p < 0.001$; CF: $p < 0.001$; CG: $p < 0.001$; sign-test on difference between random and actual rate ratios for all cells). [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

A → *B:* $\rho = -0.018$, $P(\rho = 0) = 0.60$, $N_{pairs} = 222$; LT *B* → *A*: $\rho = 0.037$, $P(\rho = 0) = 0.11$, $N_{pairs} = 359$). The randomized controls were also not significant (LT: $\rho = -0.018$, $P(\rho = 0) = 0.24$; LTAB: $\rho = -0.028$, $P(\rho = 0) = 0.42$; LTBA: $\rho = -0.015$, $P(\rho = 0) = 0.53$).

The significant correlation of *z*-score across cell pairs on CF and CG suggested that there was indeed a process coupled across cells, perhaps at the network level, that was influencing their variability. The fact that we did not see this on LT suggests that either the expected anti-correlation between place-cells that prefer opposite directions was swamping the expected correlation between cells that prefer the same directions or that hippocampal dynamics were fundamentally different on LT compared to tasks in the open arena.

Representational Maps

The cell-pair data suggested that the modulation of neuronal firing was at least partially explained by cell-assembly properties (i.e. including multiple cells, but not the entire ensemble). To directly examine the possibility of a network-level modulation of cell-assemblies (Lánský et al., 2001; Olypher et al., 2002; Harris et al., 2003; Harris, 2005), we clustered the firing patterns within each pixel and assembled them together to construct whole-environment spatial firing maps (see Methods; Fig. 1).

Because place cells on the linear track are directional (McNaughton et al., 1983; Muller et al., 1994; see Redish, 1999, for review), we expected this directionality to produce two "maps". Not surprisingly, on the linear track, the place fields of the maps derived by our map-splitting analysis closely matched those based on the animal's direction of movement (See Fig. 4). This same process was next applied to the CF and CG data.

Applying the map-splitting analysis to the CF and CG data resulted in similar splits on both tasks. Within the same task, there were clear instances where multiple-place-field neurons had one place field split between maps (e.g. Cells 6 and 29 of Session R031-2003-05-15; Fig. 6A), where place-fields were assigned to one map and not the other (e.g. Cells 25 and 34 of Session R041-2004-01-19 and cell 16 Session R031-2003-05-15; Fig. 6A), or where a place field in one map had a higher firing rate than in the other map (e.g. Cells 27 and 36 of Session R041-2004-01-19 and cell 46 Session R031-2003-05-15; Fig. 6A). Comparing CF and CG revealed slight differences in coding across maps (e.g. Cell 36 of Session R041-2004-01-19 and cells 16 and 29 Session R031-2003-05-15; Fig. 6A). The maps in Figure 6A were aligned so that ensemble Map 1 and ensemble Map 2 on CF were most correlated with Map 1 and Map 2 across the ensemble of CG, respectively.

The difference in spatial firing between maps (Fig. 6A) suggest a rate-remapping-like process similar to that described by Leutgeb et al. (2005), with a variable level of rate-modulation spanning 0–100%. To quantify the extent of these rate differences across maps, we compared the rate-difference ratio between maps for the random splitting and splitting based on ensemble switching times (Fig. 6B). As expected, random split-

TABLE 2.

Significant Changes in Coherency at Transitions

Task	Transition	P(before > after)	Transition	P(before < after)
LT	($S_1 \rightarrow S_2$)	3.2×10^{-5}	($S_2 \rightarrow S_1$)	1.4×10^{-8}
CF	($S_1 \rightarrow S_2$)	2.5×10^{-8}	($S_2 \rightarrow S_1$)	4.1×10^{-17}
CG	($S_1 \rightarrow S_2$)	3.3×10^{-7}	($S_2 \rightarrow S_1$)	9.4×10^{-16}

The probability that the coherency index $C^{1-2}_{index} =$ (**Coherency$_1$** − **Coherency$_2$**)/(**Coherency$_1$** + **Coherency$_2$**) of both states was not increased for transition ($S_2 \rightarrow S_1$) or decreased for transition ($S_1 \rightarrow S_2$) was calculated for each transition between states. Data shown in Figure 7.

ting often resulted in maps with very little rate difference with the majority of rate-difference ratios near zero ("random"; Fig. 6B). By contrast, splitting based on actual ensemble firing patterns often resulted in maps with a wide variety of rate differences ("actual"; Fig. 6B). These two distributions were significanly different for maps from each task (LT: $p < 0.001$; CF: $p < 0.001$; CG: $p < 0.001$; sign-test on difference between random and actual rate ratios for all cells).

Given the complexity of the map-switching analysis, it could be possible that these maps result from random assignments of transient ensemble firing patterns to "maps" that may not be consistent across the duration of a session. To be sure that these maps indeed represent coherent stable states of the ensemble, we tested the consistency of these states using a *coherency* measurement previously shown to detect network-state instabilities and reference-frame switches (Redish et al., 2000; Jackson and Redish, 2003). We defined our network states by the maps derived from a subset of a session's data and tested the switching dynamics of the rest of the data from that session. We then defined a coherency index that would be positive if the network was in a state that resembled Map 1 and negative if the network was in a state that resembled Map 2. There was a significant increase in the coherency index as the ensemble switched from state 2 to state 1, and a significant decrease in the coherency index when the ensemble switched from state 1 to state 2 (for statistics, see Table 2). This indicates that these states are stable and represent robust differences in the ensemble firing pattern, suggesting that the network did indeed encode at least two stable (yet separate) maps within the same environment (Fig. 7). Thus, the switching times found by our map-switching analysis indeed represented transitions in the network from one stable map to another. Similar results were found for other measures of network consistency (data not shown; see Jackson, 2006, for more details).

The temporal dynamics of state switching were next explored. Table 3 shows the transition statistics for the state

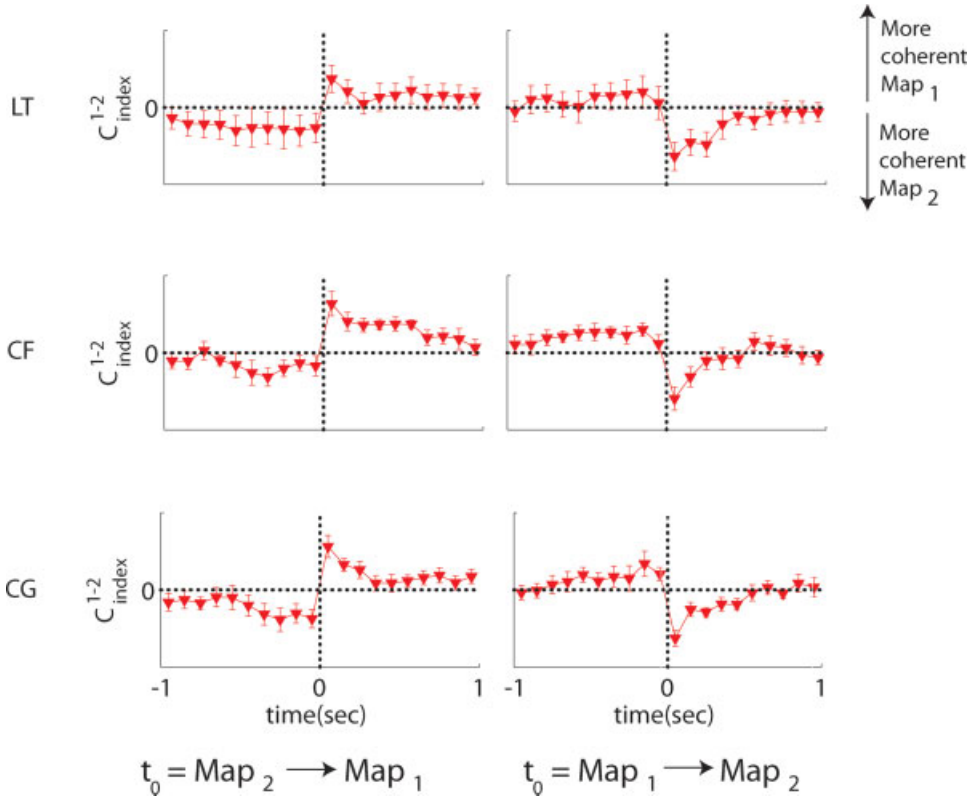

FIGURE 7. States are stable features of ensemble information processing. The coherency index C^{1-2}_{index} was measured for I_{RMS}, see Equations 3, 6. The other measures gave similar results. (left) Transitions into Map 1. (right) Transitions into Map 2. Note that each map remains coherent for some time before and after each transition. Error bars show standard error of the mean over sessions. Statistics are provided in Table 2. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

switching processes described above. The average switching rate $\mu^{\#}_{switch}$, defined as the reciprocal of the time between any switch between maps (i.e. the in-state dwell), was approximately 3.2 Hz; the average cycle rate $\mu^{\#\rightarrow\#}_{switch}$, defined as the reciprocal of time to get from one state back to the same state, was therefore approximately 1.6 Hz. The in-state dwell was approximately 380 ms; thus, these states often persisted through multiple theta cycles on average.

Map-Switching as an Explanation for Overdispersion

If the overdispersion described by Fenton and colleagues arises from effects occurring within each cell, independent of the other cells in the network, then map-switching times calculated from the other cells in the network would have no effect on the overdispersion of the cell in question. In contrast, if overdispersion does in fact reflect changing cell-assemblies, then map-switching determined from the other cells in the network should produce a reduction in the overdispersion of firing of the cell in question. This is similar to testing whether the cell-assembly-like effects described by Harris et al. (2003) can account for the single-cell variability described by Fenton and Muller (1998), while at the same-time testing the reference frame hypothesis for overdispersion (Lánský et al., 2001; Olypher et al., 2002).

We therefore tested the cell-assembly hypothesis using a leave-one-out approach. For each cell, map-switching times were calculated from all the cells except the cell in question, and then the overdispersion of the cell in question's firing was determined for each sub-map (see Methods). Only passes that were not interrupted by a map-switch were used. Because the cell whose overdispersion is being measured was not included in the detection of map-switching, any reduction in overdispersion must reflect dynamics of changing cell-assemblies. Thus, this is a direct test of the hypothesis that overdispersion reflects changing cell-assemblies.

Table 4 shows that there was much greater dispersion for the whole-task maps than for the task split maps. (Task original values are slightly different from the data in Table 1, because the data reported in Table 1 are from the subset of sessions with more than 25 neurons recorded simultaneously, see Methods.) To control for the possibility of this reduction in variance being a result of a simple inflation of explanatory power resulting from dividing the task into two states, the dispersion of each task was calculated given the randomized switching times. There was no comparable reduction in variance (Table 4).

TABLE 3.

State Transition Statistics

Task	$\mu^{\#\rightarrow\#}_{switch}$(Hz)	Dwell (s)
LT	1.5	0.35
CF	1.6	0.31
CG	1.6	0.33

TABLE 4.

Dispersion of Firing Patterns on Each Task Before and After Splitting by Representational Map: a Leave-One-Out Analysis

Task	Original	Map_1	Map_2	Randomized controls mean ± 95% CI
LT	6.7	2.9	2.4	5.9 ± (4.3, 7.5)
CF	5.2	2.5	4.6	4.7 ± (3.5, 6.1)
CG	5.8	3.5	2.8	4.7 ± (3.7, 5.8)

Behavioral Relevance

An interesting question that follows from this splitting of hippocampal representational states is whether these artificially derived states are of any relevance to the animal. We examined this by identifying the locations at which map-switching tended to occur and by testing whether the switching times were in any way related to the task requirements.

To examine the relevance of spatial position to map-switching, we generated rate-maps of the animal's position at map-switching times by dividing the number of map-switches in each spatial bin on the task by the time the animal spent in that bin (see Fig. 8). All three tasks were significantly different from random (ANOVA, measuring the effect of position on map-switch rates: LT, $F(1,15) = 18.98$, $p < 10^{-5}$; CF: $F(1,229) = 11.47$, $p < 10^{-5}$; CG: $F(1,229) = 8.37$, $p < 10^{-5}$). The CF and CG tasks were significantly different from each other (ANOVA, effect of task, $F = 24.63$, $p < 10^{-5}$, including an interaction between task and position, $F(1,227) = 1.45$, $p < 10^{-5}$).

To examine the relation of these temporal dynamics to the task behavioral parameters, a PETH was constructed from all state transition times derived from the ensemble state to yield the transition rate ratio leading up to each task event (viz., food delivery on LT or CF, and the qualifying tone on CG). These normalized PETHs are shown in Figure 9. There were significant increases in switching rates following reward-related events on the LT and CG tasks, but not on CF (LT: $P(\text{pre} = \text{post}) < 10^{-8}$; CG: $P(\text{pre} = \text{post}) < 0.003$; CF: $P(\text{pre} = \text{post}) = 0.87$; 2-sample unpaired t-tests comparing 6 s before and 6 s after food delivery). While reward in the LT and CG tasks was contingent on the location of the animal (thus providing the animal with goals that change with times of reward-delivery), the reward on the CF task was delivered independent of the behavior of the animal.

DISCUSSION

The data presented here, showing multiple self-consistent spatial firing maps across a hippocampal ensemble within a single task, suggest that the hippocampus can switch the reference

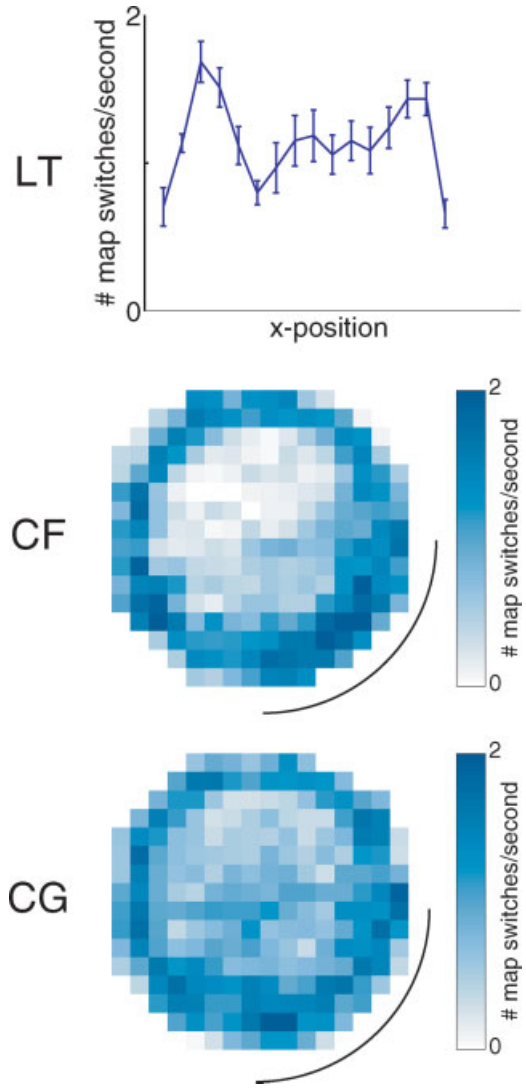

FIGURE 8. Spatial distribution of map switches in the three tasks. Number of map-switches were measured per second for each bin (total number of switches in that bin divided by the time spent in that bin). Because the linear track is effectively a linear task, the location of the map-transitions were collapsed across the width of the track to make position 1-dimensional. The location of the cue card is marked with the small arc in the CF and CG tasks. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

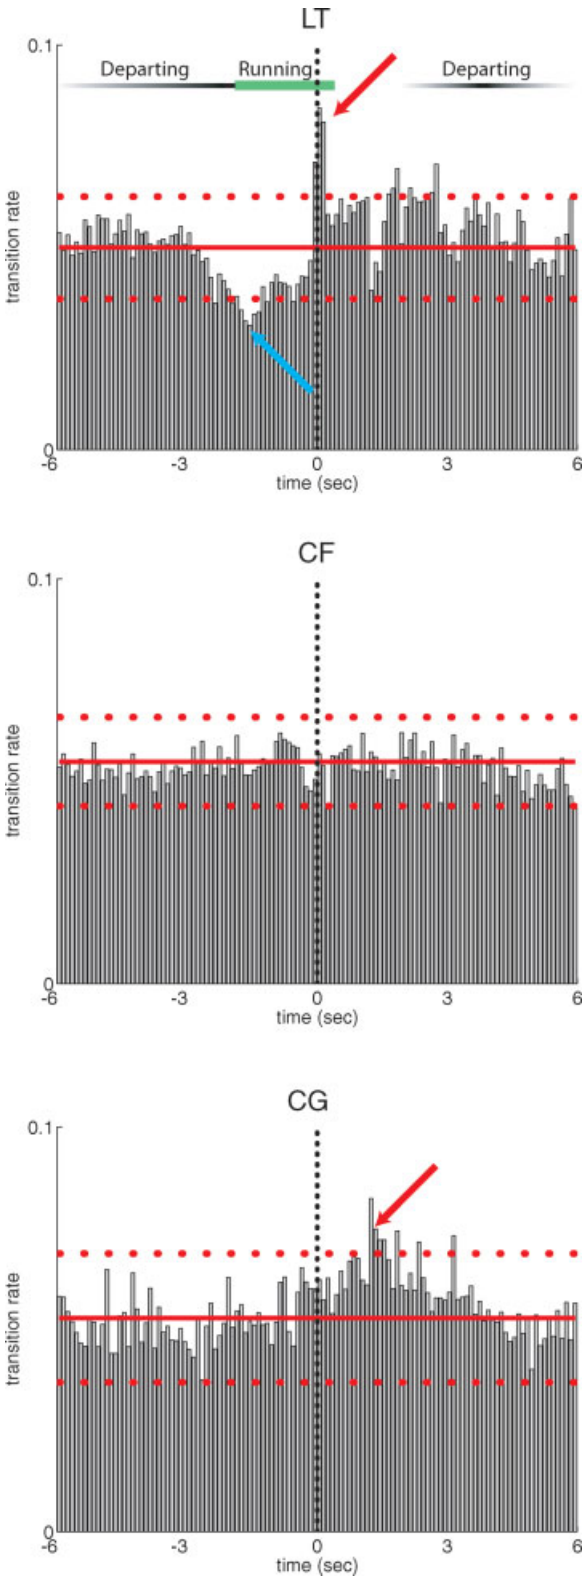

FIGURE 9

frame used to solve a task within a task depending on the task requirements. This is consistent with experimental data showing an internal switch of reference frames within task when this switch was required for proper task performance (Eichenbaum et al., 1987; Cohen and Eichenbaum, 1993; Markus et al., 1995; Gothard et al., 1996a,b; Redish, 1999; Redish et al., 2000; Zinyuk et al., 2000; Rosenzweig et al., 2003). Our results suggest a network-wide modulation occurring at behaviorally-relevant times, such as the delivery of food reward. However, given that these shifts occur also with a high frequency during noncontingent spatial tasks, it is possible that these shifts in network state correspond to internal cognitive

shifts in motivation or behavioral planning such as spatial target selection (e.g. switching from pellet foraging behavior to targeted navigation; Markus et al., 1995; Fenton et al., 1998; Olypher et al., 2002; Kentros et al., 2004).

A number of theorists have suggested that the hippocampal cognitive map may contain submaps or "maplets" (Worden, 1992; McNaughton et al., 1994; Samsonovich and McNaughton, 1997; Touretzky and Muller, 2006). The question of how unified the interaction between place cells is has been a contentious one over the years (McNaughton et al., 1994; Samsonovich and McNaughton, 1997; Redish et al., 1998; Redish, 1999; Harris et al., 2003; Harris, 2005; Touretzky and Muller, 2006) with data supporting both the idea that cells respond individually (Anderson and Jeffery, 2003; Tanila et al., 1997; Lee et al., 2004) and other data supporting changes occurring in the map as a whole (O'Keefe and Conway, 1978; Markus et al., 1995; Barnes et al., 1997; Leutgeb et al., 2005, 2007) as well as data suggesting partial remapping (Quirk et al., 1990; Skaggs and McNaughton, 1998; Knierim, 2002; Leutgeb et al., 2004; Vazdarjanova and Guzowski, 2004; Lee et al., 2004; Fuhs et al., 2005). Similarly, place cell responses to subgoals within a task (Eichenbaum et al., 1987; Wiener et al., 1989; Cohen and Eichenbaum, 1993; Hampson et al., 1993; Wood et al., 2000; Ferbinteanu and Shapiro, 2003) have been proposed to arise from goal-dependent submaps (Touretzky and Redish, 1996; Redish and Touretzky, 1997; Redish, 1999). The existence of multiple, coherent, goal-dependent maps within a single task provide direct evidence for the existence of submap fragments.

Reference Frame Switching

Since an ensemble of hippocampal place cells will generally tend to have a stable pattern of place-fields distributed uniformly throughout an environment, this is taken as the neural instantiation of a spatial map or reference frame (see, for review, O'Keefe and Nadel, 1978; Redish, 1999). Markus et al. (1995) found that within an open field foraging task, simply changing the task from foraging to directed running (tapping on 1 of 4 locations in sequence) for food caused place fields to partially remap. Thus, partial remapping can result from modification of task behavioral parameters and remapping with respect to goal-directed behavior could be taken to mean that an animal is referencing its position with respect to a particular spatial goal (Markus et al., 1995). In this sense, the maps derived from the CG task may be relevant to different behavioral states associated with the task (i.e. navigating to the goal and searching for pellets) similar to how previous studies distinguished between spatial reference frames based on behavioral requirements (Redish et al., 2000; Rosenzweig et al., 2003).

One further question is why CF would have multiple spatial reference frames with switching statistics similar to CG. Here the answer may lie in the qualitative description of the animal's behavior. On CF, the animal's behavior appears to alternate between pellet searching and less-directed wandering. It is possible that even these two behavioral states are differentiated in the CA1 code. This may be due to the animal's knowledge of the statistics of food delivery, which sometimes could result in long waiting times. However, it is important to note that one map in CF did not have dispersion z-scores different from randomized controls (see Table 4), this may be evidence that this second map may be composed of fragments of submaps perhaps relating to a variety of subgoals, such as individual pellet locations. Thus, fitting more than two maps to CF may lower the dispersion further, but to test this would require larger ensembles and longer recording times than were available for the current experiment. Alternatively, this high dispersion of one map may represent a less attentive state such as has been described in mice (Kentros et al., 2004).

Finally, why is the rate of switching between maps largest near the edges of the environment? It may be that the wall serves as a point of reference for the animal where the path integrator can be reliably reset. If this true, high contrast locations on the wall, such as the cue card, may be especially important and should demonstrate a higher likelihood of switching. In Figure 8, there appears to be a slight tendency toward a higher switching rate near the cue card on CF and perhaps more specifically at the cue-card edges on CG. The overall rate of switching on CG looks slightly more uniform than on CF. This may be due to increased goal-related switching (see Fig. 9); however, this would be averaged out over the surface of the arena due to random goal placement in each session. Thus, switching may be related both to goal-related and landmark-related resets. Another possible interpretation of the increased switching near the walls on CF is that switches may be related to when the animal changes from searching for more pellets to a more direct navigation toward pellets that have been delivered. Animals often favored the wall more during searching and crossed the center more during chasing. The pellets drop from above the rat and bounce after hitting the arena floor, making noise to which the animals clearly oriented. If there was a switch of maps that corresponded to the behavioral

FIGURE 9. State switching dynamics leading up to and following key task events. (*top*) State switching is suppressed during approach to the feeder on LT ($P(pre = post) < 10^{-8}$, 2 sample unpaired *t*-test comparing 6 s before and after food delivery). There is a significant drop in state switching (blue arrow) during running (green) with an abrupt increase at food delivery (red arrow; note animal is still in motion). Black gradient depicts departure-time distributions. (*middle*) There is no visible change in the transition rate between maps on CF ($P(pre = post) = 0.87$, 2 sample unpaired *t*-test comparing 6 s before and after food delivery). (*bottom*) There is a significant increase in the rate of state switching (red arrow) following the qualifying tone ($P(pre = post) = 0.0023$, 2 sample unpaired t-test comparing 6 s before and after food delivery). The solid red line is the mean of the boot-strapped distribution, the dotted red lines are the 95% confidence intervals corrected for multiple comparisons. The red arrow points to a significant increase in switching (this time approximately corresponds with the time that it takes the pellets to reach the arena floor). [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

switch from searching to chasing, this should be evident as an increase in switching rate just after pellet delivery (see Fig. 9). Our data does not support such an explanation of the higher switching at the wall since there is no clear relation between pellet delivery and map switching on CF.

Why are there no-fire Passes on Directionally Split LT Data?

While no-fire passes on cylindrical arena tasks would be expected for neurons with strong rate differences between maps, one would not expect there to be many times when a neuron would not fire during a pass through a directionally defined place field on the linear track. In the directionally filtered data, it is possible that the internal cognitive state of the animal near the ends of the track is different from our external expectations based on direction. The PETHs representing the map-switching on linear track with respect to reward delivery (which resulted in an audible click before the animal reached the food port), show a dramatic spike in map-switching just after food delivery while the animal is still running. Thus, the animal is switching maps while it is still moving in the same direction based on an *external* temporal event. No-fire passes in the directionally filtered data may then relate to a deviation of the animal's cognitive state from our expectation that movement direction is the only external feature governing the internal cognitive state.

Other Possible Sources of Network Switching

These results cannot be explained by phenomena such as phase precession and transitions in the behavioral state from theta to LIA. As presented in the Methods, these analyses were conducted only on periods of data with high-theta and low-delta power. As a further restriction, high-frequency events such as sharp-waves were removed. Thus, the network switching discussed here is not likely to be a result of gross fluctuations in the hippocampal processing state (i.e. from movement related theta rhythm to resting LIA states). Furthermore, the average dwell time was 380 ms, on the order of about three theta cycles. While faster and slower switches were observed, it is unlikely that the network switches reported here are simply due only to a switch from the end of each theta cycle to the beginning of the next. This time-scale is also much shorter than most continuous theta epochs in our data. Finally, the overdispersion results are averages over passes through place fields that lasted 1 s or more. Thus, these results should be robust to fluctuation on the order of a single theta cycle such as phase precession. Likewise, the overdispersion of the split maps used only passes where the network remained in one state throughout the pass (similar results were obtained after loosening this restriction, data not shown).

Comparison to Fenton and Colleagues

We replicated Fenton et al.'s findings of excessive variance in place cell discharge ("overdispersion") and its task dependence (Fenton and Muller, 1998; Lánský et al., 2001; Olypher et al., 2002). Consistent with expectations, we found that on a task with known reference-frame switching (the linear track), splitting by representational state (direction of travel) resulted in greatly reduced variability. In contrast to Fenton and Muller (1998), however, we found significant local interactions between neurons with overlapping place-fields. We also found that separating the ensemble firing patterns into two maps produced coherent and consistent patterns, thus implying that the overdispersion phenomenon was a consequence of a network-level process. Additionally, the variability of the firing patterns within these reference frames was greatly reduced. We found that transitions between these reference frames occurred at behaviorally relevant times, such as delivery of reward on spatially-contingent tasks (e.g. LT, CG).

Our results argue for a faster average switching rate with deeper modulation than suggested by Olypher et al. (2002). As mentioned earlier, the prominent, left-shifted peak in our data (see Fig. 2) is due to no-fire passes. Figure 6 depicts numerous state-remappings where a place field exists in one state and is absent in the other state. If an animal runs through this region in one state, spikes will be fired, but spikes will not be fired during passes in the other state. Markus et al. (1995) reported the fast remapping of spatial responses in an environment to changes in the task reward contingencies within same session. (Both Markus et al.'s and our animals had had extensive experience with the tasks before recording began.) It may be that training our animals on two behaviorally different tasks has produced hippocampal spatial representations that contrast these behaviorally different tasks ultimately resulting in deeper modulation between maps than in the data of Olypher et al. (2002). This would increase the incidence of no-fire passes on both tasks and contribute a large left-shifted component to the distribution of z-scores.

Controls: Are We Just Adding Parameters?

One possible interpretation for the reduced variance of the split times shown in Table 2, is that it is simply a phenomenon of explaining the variance by adding another parameter. This is unlikely for a number of reasons. First, comparing the dispersion of the split maps was significantly lower than the maps created by random switching (see Table 4). Second, if neuronal firing was independent of the rest of the ensemble, the leave-one-out analysis should have provided the same results as the random split, because the cell being tested was not included in the map-splitting process. Third, regardless of whether or not the smoothing that results from maximizing correlations with neighboring pixels within a map is responsible for the reduction in variance after splitting, there remains the spatial and temporal significance of the transition times between maps, which emphasize the salience of ensemble interactions that led to the reduced variance. For instance, the maps correlate well with the maps corresponding to the directions of travel on the linear track and tend to switch right after food-delivery (while the rat is running, just before the rat reaches the food-site).

This switch was unrelated to changes in local field potential activity (i.e., the rat remained in theta through the transition). Likewise, on CG, the transition between maps corresponds closely with the reward, with an increase in transition rate between maps after the tone indicating the animal reached the goal. Again, the rat remained in theta through the transition. Thus, the parameter of ensemble interactions is behaviorally relevant, and the reduction in variance after adding this parameter is more than would be expected by simply splitting the data arbitrarily into two states.

Representing Internal Goals

The increase of switching responses following reward-related cues (see Fig. 9) strongly suggests that these maps may be related to an animal's representation of internal goals with respect to the external world. A recent report by Hok et al. (2007), shows specific modulation of hippocampal firing during a waiting period on a goal-related task. They suggest that the hippocampus may code both spatial and goal-related information. Another example of goal-related modulation of hippocampal place firing was reported by Hollup et al. (2001), where a buildup of place fields was observed near the hidden platform location in an annular water maze. The tasks in both of these studies differ from our CG task in that their goal region was fixed across sessions. Nevertheless, the gross changes observed in hippocampal firing in these studies support the dynamic ensemble-level modulations observed in our data. Our data, however, extend these findings to tasks where no specific experimenter-defined goal is specified, as in the CF task. Together, these data suggest that goal-related modulation of hippocampal pyramidal cell firing can be reflected in the spatial responses of neurons across the hippocampal network at short time scales within a task.

Cell Assemblies

Cell assemblies are by definition neurons bound together by common inputs and reciprocal connections such that their spiking is more interrelated within an assembly than across assemblies. While CA1 itself lacks the recurrent connectivity of CA3, the inhibitory network within CA1 as well as the coupling of CA1 to CA3 and the trisynaptic-loop circuitry should provide the basis for the formation of cell-assemblies in CA1. If we apply the above definition to our hippocampal data, we see that there are multiple continuous cell assemblies. *Multiple* in the sense that distinct spatial tuning maps result from clustering the firing patterns observed on the task. *Continuous* in the sense that each map is extended in time and space such that the pattern of neurons active within one map gradually evolves as the animal moves through space until the hippocampus switches to the next map.

The evidence for these cell assemblies is in the pair-wise correlations. If neurons are part of the same cell assembly they could be bound together spatially and temporally, while neurons from a different cell assembly may overlap spatially and remain temporally independent or anticorrelated. This is consistent with recent reports (Harris et al., 2003; Harris, 2005; Lin et al., 2005). The ensemble activation patterns in our data overlapped to some extent in that there was a generalized similarity between the two maps extracted from CF and CG with variations in firing rate depending on which state the network was in and occasional remapping of neuronal responses to new preferred locations in the environment. This level of overlap is reminiscent of the partial and rate remapping phenomena observed across environments (Leutgeb et al., 2005; Anderson and Jeffery, 2003) and when changing tasks within the same environment (Markus et al., 1995).

While we can account for some of the variance in the overdispersion of place-cell firing by splitting the ensemble map into two spatial reference frames, there was still more variance than would be expected given a Poisson point process. Some of this variance may be due to the non-Poisson nature of hippocampal pyramidal cells (Ranck, 1973; Barbieri et al., 2001). Additional variability is likely to be generated from a number of sources such as speed modulation (McNaughton et al., 1983), changes in specialized, nonlocal firing patterns (Jensen and Lisman, 2000; Johnson and Redish, 2005; O'Neill et al., 2006; Jackson et al., 2006), as well as variations in plasticity as a function of experience and the regularity of spatial behavior (Mehta et al., 1997). There also remains the distinct possibility that other internal cognitive processes such as navigational planning and transient nonlocal encoding may also influence the temporal variability observed in hippocampal neurons (Jensen and Lisman, 1998, 2005; Koene et al., 2003; Johnson and Redish, 2006; Ferbinteanu et al., 2006).

In conclusion, we have presented data that link single cell variability in the hippocampus to network-level processes. This variability could be largely accounted for by the switching of network states made up of continuous assemblies of cells that formed multiple spatial maps of the environment. Taken together our results imply that hippocampal spatial representations of tasks are made up of multiple, continuous submaps, the selection of which depends on the animal's goals.

Acknowledgments

We thank Mallika Arudi, Deborah Bang, Dan Bernal, Chris Boldt, Giuseppe Cortese, Maniezheh Firouzi, Mandy Huber, Sarah Jutila, Morgan Little, and Kelsey Seeland for help with data collection. We thank Neil Schmitzer-Torbert, Adam Johnson, Andre Fenton, and Bruce McNaughton for helpful discussions and comments on early drafts of the manuscript.

REFERENCES

Ainge JA, van der Meer MAA, Wood ER. 2005. Disparity between sequence-dependent hippocampal activity and hippocampal lesion effects on a continuous t-maze task. Society for Neuroscience Abstracts Program No. 776.13. Online.

Amaral DG, Witter MP. 1989. The three-dimensional organization of the hippocampal formation: A review of anatomical data. Neuroscience 31:571–591.

Anderson MI, Jeffery KJ. 2003. Heterogeneous modulation of place cell firing by changes in context. J Neurosci 23:8827–8835.

Barbieri R, Quirk MC, Frank LM, Wilson MA, Brown EN. 2001. Construction and analysis of non-poisson stimulus-response models of neural spiking activity. J Neurosci Meth 105:25–37.

Barnes CA, Suster MS, Shen J, McNaughton BL. 1997. Multistability of cognitive maps in the hippocampus of old rats. Nature 388: 272–275.

Bower MR, Euston DR, McNaughton BL. 2005. Sequential-context-dependent hippocampal activity is not necessary to learn sequences with repeated elements. J Neurosci 25:1313–1323.

Brown EN, Frank LM, Tang D, Quirk MC, Wilson MA. 1998. A statistical paradigm for neural spike train decoding applied to position prediction from ensemble firing patterns of rat hippocampal place cells. J Neurosci 18:7411–7425.

Cohen NJ, Eichenbaum H. 1993. Memory, Amnesia, and the Hippocampal System. Cambridge, MA: MIT Press.

Deadwyler SA, Bunn T, Hampson RE. 1996. Hippocampal ensemble activity during spatial delayed-nonmatch-to-sample performance in rats. J Neurosci 16:354–372.

Eichenbaum H, Kuperstein M, Fagan A, Nagode J. 1987. Cue-sampling and goal-approach correlates of hippocampal unit activity in rats performing an odor-discrimination task. J Neurosci 7:716–732.

Ekstrom AD, Meltzer J, McNaughton BL, Barnes CA. 2001. NMDA receptor antagonism blocks experience-dependent expansion of hippocampal "place fields". Neuron 31:631–638.

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 USA 95:3182–3187.

Fenton AA, Wsierska M, Kaminsky Y, Bures J. 1998. Both here and there: Simultaneous expression of autonomous spatial memories in rats. Proc Natl Acad Sci USA 95:11493–11498.

Ferbinteanu J, Kennedy PJ, Shapiro ML. 2006. Episodic memory—From brain to mind. Hippocampus 16:704–715.

Ferbinteanu J, Shapiro ML. 2003. Prospective and retrospective memory coding in the hippocampus. Neuron 40:1227–1239.

Fuhs MC, VanRhoads SR, Casale AE, McNaughton B, Touretzky DS. 2005. Influence of path integration versus environmental orientation on place cell remapping between visually identical environments. J Neurophysiol 94:2603–2616.

Gothard KM, Skaggs WE, McNaughton BL. 1996a. Dynamics of mismatch correction in the hippocampal ensemble code for space: Interaction between path integration and environmental cues. J Neurosci 16:8027–8040.

Gothard KM, Skaggs WE, Moore KM, McNaughton BL. 1996b. Binding of hippocampal CA1 neural activity to multiple reference frames in a landmark-based navigation task. J Neurosci 16:823–835.

Hampson RE, Heyser CJ, Deadwyler SA. 1993. Hippocampal cell firing correlates of delayed-match-to-sample performance in the rat. Behav Neurosci 107:715–739.

Harris KD. 2005. Neural signatures of cell assembly organization. Nat Rev Neurosci 6:399–407.

Harris KD, Csicsvari J, Hirase H, Dragoi G, Buzsáki G. 2003. Organization of cell assemblies in the hippocampus. Nature 424:552–556.

Hebb DO. 1949. The Organization of Behavior. New York: Wiley. Reissued 2002 LEA.

Hok V, Lenck-Santini PP, Roux S, Save E, Muller RU, Poucet B. 2007. Goal-related activity in hippocampal place cells. J Neurosci 27:472–482.

Hollup SA, Molden S, Donnett JG, Moser MB, Moser EI. 2001. Accumulation of hippocampal place fields at the goal location in an annular watermaze task. J Neurosci 21:1635–1644.

Huxter J, Burgess N, O'Keefe J. 2003. Independent rate and temporal coding in hippocampal pyramidal cells. Nature 425:828–832.

Jackson J. 2006. Network Consistency and Hippocampal Dynamics: Using the properties of cell assemblies to probe the hippocampal representation of space. Ph.D. thesis, University of Minnesota.

Jackson JC, Johnson A, Redish AD. 2006. Hippocampal sharp waves and reactivation during awake states depend on repeated sequential experience. J Neurosci 26:12415–12426.

Jackson JC, Redish AD. 2003. Detecting dynamical changes within a simulated neural ensemble using a measure of representational quality. Netw Computat in Neural Syst 14:629–645.

Jensen O, Lisman JE. 1998. An oscillatory short-term memory buffer model can account for data on the Sternberg task. J Neurosci 18:10688–10699.

Jensen O, Lisman JE. 2000. Position reconstruction from an ensemble of hippocampal place cells: contribution of theta phase encoding. J Neurophysiol 83:2602–2609.

Jensen O, Lisman JE. 2005. Hippocampal sequence-encoding driven by a cortical multi-item working memory buffer. Trend Neurosci 28:67–72.

Johnson A, Jackson J, Redish AD. in press. Measuring distributed properties of neural representations beyond the decoding of local variables—Implications for cognition. In: Hölscher C, Munk MHJ, editors. Mechanisms of Information Processing in the Brain: Encoding of Information in Neural Populations and Networks. Cambridge UK: Cambridge University Press.

Johnson A, Redish AD. 2005. Observation of transient neural dynamics in the rodent hippocampus during behavior of a sequential decision task using predictive filter methods. Acta Neurobiol Exp 65:103. Poster 245.

Johnson A, Redish AD. 2006. Neural ensembles in ca3 transiently encode paths forward of the animal at a decision point: A possible mechanism for the consideration of alternatives. Soc Neurosci Abst.

Johnson A, Seeland KD, Redish AD. 2005. Reconstruction of the postsubiculum head direction signal from neural ensembles. Hippocampus 15:86–96.

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.

Knierim JJ. 2002. Dynamic interactions between local surface cues, distal landmarks, and intrinsic circuitry in hippocampal place cells. J Neurosci 22:6254–6264.

Koene RA, Gorchetchnikov A, Cannon RC, Hasselmo ME. 2003. Modeling goal-directed spatial navigation in the rat based on physiological data from the hippocampal formation. Neural Netw 16:577–584.

Lánský P, Fenton AA, Vaillant J. 2001. The overdispersion in activity of place cells. Neurocomputing 38–40:1393–1399.

Lee I, Yoganarasimha D, Rao G, Knierim JJ. 2004. Comparison of population coherence of place cells in hippocampal subfields of CA1 and CA3. Nature 430:456–459.

Leutgeb JK, Leutgeb S, Moser MB, Moser EI. 2007. Pattern separation in the dentate gyrus and CA3 of the hippocampus. Science 315:961–966.

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.

Leutgeb S, Leutgeb JK, Treves A, Moser MB, Moser EI. 2004. Distinct ensemble codes in hippocampal areas CA3 and CA1. Science 305:1295–1298.

Lin L, Osan R, Shoham S, Jin W, Zuo W, Tsien JZ. 2005. Identification of network-level coding units for real-time representation of episodic experiences in the hippocampus. PNAS 102:6125–6130.

Lorento do Nó R. 1933. Studies on the structure of the cerebral cortex. J Psychol Neurol 45:381–442.

Lorento do Nó R. 1934. Studies on the structure of the cerebral cortex II. Continuation of the study of the ammonic system. J Psychol Neurol 46:113–177.

Markus EJ, Qin Y, 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, Gerrard JL, Gothard K, Jung MW, Knierim JJ, Kudrimoti H, Qin Y, Skaggs WE, Suster M, Weaver KL. 1996. Deciphering the hippocampal polyglot: The hippocampus as a path integration system. J Exp Biol 199:173–186.

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.

McNaughton BL, Knierim JJ, Wilson MA. 1994. Vector encoding and the vestibular foundations of spatial cognition: Neurophysiological and computational mechanisms. In: Gazzaniga M, editor. The Cognitive Neurosciences. Cambridge MA: MIT Press. pp. 585–595.

McNaughton BL, Morris RGM. 1987. Hippocampal synaptic enhancement and information storage within a distributed memory system. Trend Neurosci 10:408–415.

Mehta MR, Barnes CA, McNaughton BL. 1997. Experience-dependent, asymmetric expansion of hippocampal place fields. Proc Natl Acad Sci USA 94:8918–8921.

Mehta MR, Quirk MC, Wilson MA. 2000. Experience-dependent asymmetric shape of hippocampal receptive fields. Neuron 25:707–715.

Muller RU, Bostock E, Taube JS, Kubie JL. 1994. On the directional firing properties of hippocampal place cells. J Neurosci 14:7235–7251.

Muller RU, Kubie JL, Ranck JB Jr. 1987. Spatial firing patterns of hippocampal complex-spike cells in a fixed environment. J Neurosci 7:1935–1950.

O'Keefe J. 1976. Place units in the hippocampus of the freely moving rat. Exp Neurol 51:78–109.

O'Keefe J, Conway DH. 1978. Hippocampal place units in the freely moving rat: Why they fire where they fire. Exp Brain Res 31:573–590.

O'Keefe J, Dostrovsky J. 1971. The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely moving rat. Brain Res 34:171–175.

O'Keefe J, Nadel L. 1978. The Hippocampus as a Cognitive Map. Oxford: Clarendon Press.

O'Keefe J, Recce M. 1993. Phase relationship between hippocampal place units and the EEG theta rhythm. Hippocampus 3:317–330.

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.

O'Neill J, Senior T, Csicsvari J. 2006. Place-selective firing of ca1 pyramidal cells during sharp wave/ripple network patterns in exploratory behavior. Neuron 49:143–155.

Paxinos G, Watson C. 1998. The Rat Brain in Stereotaxic Coordinates, 4th ed. New York: Academic Press.

Quirk GJ, Muller RU, Kubie JL. 1990. The firing of hippocampal place cells in the dark depends on the rat's recent experience. J Neurosci 10:2008–2017.

Ranck JB Jr. 1973. Studies on single neurons in dorsal hippocampus formation and septum in unrestrained rats. I. Behavioral correlates and firing repertoires. Exp Neurol 41:461–555.

Redish AD. 1999. Beyond the Cognitive Map: From Place Cells to Episodic Memory. Cambridge MA: MIT Press.

Redish AD, McNaughton BL, Barnes CA. 1998. Reconciling Barnes et al. (1997) and Tanila et al. (1997a, 1997b). Hippocampus 8:438–443.

Redish AD, Rosenzweig ES, Bohanick JD, McNaughton BL, Barnes CA. 2000. Dynamics of hippocampal ensemble realignment: Time vs. space. J Neurosci 20:9289–9309.

Redish AD, Touretzky DS. 1997. Cognitive maps beyond the hippocampus. Hippocampus 7:15–35.

Rosenzweig ES, Redish AD, McNaughton BL, Barnes CA. 2003. Hippocampal map realignment and spatial learning. Nature Neurosci 6:609–615.

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.

Sakurai Y. 1990. Hippocampal cells have behavioral correlates during the performance of an auditory working memory task in the rat. Behav Neurosci 104:253–263.

Sakurai Y. 1994. Involvement of auditory cortical and hippocampal neurons in auditory working memory and reference memory in the rat. J Neurosci 14:2606–2623.

Samsonovich AV, McNaughton BL. 1997. Path integration and cognitive mapping in a continuous attractor neural network model. J Neurosci 17:5900–5920.

Shen J, Barnes CA, McNaughton BL, Skaggs WE, Weaver KL. 1997. The effect of aging on experience-dependent plasticity of hippocampal place cells. J Neurosci 17:6769–6782.

Skaggs WE, McNaughton BL. 1998. Spatial firing properties of hippocampal CA1 populations in an environment containing two visually identical regions. J Neurosci 18:8455–8466.

Tanila H, Shapiro ML, Eichenbaum H. 1997. Discordance of spatial representation in ensembles of hippocampal place cells. Hippocampus 7:613–623.

Touretzky DS, Muller RU. 2006. Place field dissociation and multiple maps in hippocampus. Neurocomputing 69:1260–1263.

Touretzky DS, Redish AD. 1996. A theory of rodent navigation based on interacting representations of space. Hippocampus 6: 247–270.

Vazdarjanova A, Guzowski JF. 2004. Differences in hippocampal neuronal population responses to modifications of an environmental context: Evidence for distinct, yet complementary, functions of CA3 and CA1 ensembles. J Neurosci 24:6489–6496.

Wiener SI, Paul CA, Eichenbaum H. 1989. Spatial and behavioral correlates of hippocampal neuronal activity. J Neurosci 9:2737–2783.

Wilson MA, McNaughton BL. 1993. Dynamics of the hippocampal ensemble code for space. Science 261:1055–1058.

Wood ER, Dudchenko PA, Robitsek RJ, Eichenbaum H. 2000. Hippocampal neurons encode information about different types of memory episodes occurring in the same location. Neuron 27:623–633.

Worden R. 1992. Navigation by fragment fitting: A theory of hippocampal function. Hippocampus 2:165–87.

Ylinen A, Bragin A, Nadasdy Z, Jando G, Szabo I, Sik A, Buzsáki G. 1995. Sharp wave-associated high-frequency oscillation (200 Hz) in the intact hippocampus: Network and intracellular mechanisms. J Neurosci 15:30–46.

Zhang K, Ginzburg I, McNaughton BL, Sejnowski TJ. 1998. Interpreting neuronal population activity by reconstruction: Unified framework with application to hippocampal place cells. J Neurophysiol 79:1017–1044.

Zinyuk L, Kubik S, Kaminsky Y, Fenton AA, Bures J. 2000. Understanding hippocampal activity by using purposeful behavior: place navigation induces place cell discharge in both taskrelevant and task-irrelevant spatial reference frames. Proc Natl Acad Sci USA 97:3771–3776.

APPENDIX: KNOWN PHENOMENA WHICH CAN PRODUCE NON-POISSON CHANGES IN FIRING RATE

Place Field Expansion and Repetition

One possible mechanism that can produce a change in expected firing rate is experience-dependent modification of place-fields (Mehta et al., 1997, 2000; Shen et al., 1997; Ekstrom et al., 2001). To examine the effect of this phenomenon in our data, we measured the correlation between the firing rate *z*-score and the number of times the animal passed through a place field. On the CF task, there was no significant correlation between a neuron's *z*-score on a particular pass and the cumulative number of passes through that neuron's place field up to that pass ($\rho = 0.024$, $P(\rho = 0) = 0.13$). However, there was a significant positive correlation on the CG task ($\rho = 0.043$, $P(\rho = 0) = 0.0030$). On the LT task, the firing rate *z*-score was also correlated with the number of passes whether ignoring directionality or splitting by direction (LT: $\rho = 0.047$, $P(\rho = 0) = 0.0024$; LT $A \rightarrow B$: $\rho = 0.086$, $P(\rho = 0) = 0.0064$; LT $B \rightarrow A$: $\rho = 0.056$, $P(\rho = 0) = 0.029$).

The task difference seen above suggests that differences in the regularity of behavior may drive differences in overdispersion. To explore the contribution of cumulative stereotypy in the animal's spatial behavior to the changes in the variability of neuronal firing, we wanted to measure the development of path regularity in different regions of the environment. Therefore, we calculated the average amount of disorder observed in an animal's previous movements in a place-field for each subsequent path taken through that place-field (Jackson et al., 2006). We call this quantity *behavioral entropy*.

The *x, y*-position data was binned into identical bins as the tuning curves: 11 pixel × 11 pixel blocks (3 cm × 3 cm) and the transition probability from each bin into every other bin was updated as the animal transitioned from one bin to another as above. With each transition (e.g. moving from bin j to bin k), the Shannon entropy H_j of the transition probability from bin j to all other bins was calculated for the previously occupied bin:

$$H_j = \sum_i^N -p_{i,j} \log_2 p_{i,j} \tag{A1}$$

Thus, a cumulative record of the entropy of every location traversed in the animal's path throughout the session was compiled. For comparison with the dispersion *z*-score of each pass through a place fields, the average cumulative entropy of the path during that pass was used. Since the quantity H_j is related to the regularity of an animal's path through location j, we refer to H_j as the local behavioral entropy.

Since our video sampling is sufficiently high enough that an animal is unable to move more than one spatial bin in each of eight possible directions from its current spatial bin and since we only counted transitions from one bin to another bin, the maximum value for H_j corresponding to equal transitions to all neighboring bins is $H_j = 8(-1/8) \log_2(1/8) = 3$. Lower values of H_j correspond to more ordered local paths.

We compared the behavioral entropy of the task defined above with the firing rate *z*-scores. There were significant positive correlations between an animal's behavioral entropy and the magnitude of the dispersion observed on the linear track and the cylinder goal task (LT: $\rho = 0.078 \pm 0.020$, $P(\rho = 0) = 0.0022$; CG: $\rho = 0.081 \pm 0.029$, $P(\rho = 0) = 0.0085$). This correlation was even stronger on CG when considering the five seconds leading up to goal entry ($\rho = 0.19 \pm 0.05$, $P(\rho = 0) = 0.0015$). However, splitting by direction on LT yielded significant negative correlation in the $A \rightarrow B$ direction ($\rho = -0.12 \pm 0.051$, $P(\rho = 0) = 0.049$), and near zero correlation in the $B \rightarrow A$ direction ($\rho = -0.0069 \pm 0.059$, $P(\rho = 0) = 0.91$). More importantly, there was no significant correlation on the cylinder foraging task CF ($\rho = 0.013 \pm 0.045$, $P(\rho = 0) = 0.57$). These data suggest that changes in regularity of behavior do not explain the overdispersion seen on the CF task, implying that some other phenomenon is driving the overdispersion effect.

Velocity Dependence

The small but significant pair-wise correlations and correlations with number of passes and behavioral entropy would not

TABLE A1.

Velocity vs. Z-Score: Correlation and Regression Results

Task	ρ	$P(\rho = 0)$	Intercept	Slope (s/cm)	R^2	F	P-value
LT	0.13	6.3×10^{-16}	−0.60	0.024	0.016	62.8	3.0×10^{-15}
LT $A \rightarrow B$	0.11	1.4×10^{-3}	−0.47	0.008	0.009	8.1	4.6×10^{-3}
LT $B \rightarrow A$	0.10	4.3×10^{-4}	−0.54	0.013	0.011	15.0	1.1×10^{-4}
CF	0.14	3.9×10^{-18}	−0.87	0.051	0.020	77.2	$< 10^{-18}$
CG	0.12	6.8×10^{-15}	−0.78	0.048	0.015	65.5	7.8×10^{-16}
CG 5s	0.17	3.1×10^{-7}	−0.89	0.069	0.035	30.8	3.9×10^{-8}

fully explain the large variance observed in the over-dispersion phenomenon. However, there have been reports of speed and direction modulation of place cell activity (McNaughton et al., 1983; Markus et al., 1995; Huxter et al., 2003). Therefore, we examined the relationship between speed and the dispersion *z*-score. On all tasks, the correlations between *z*-score and velocity were positive and significant (see Table A1).

We tested this trend by performing a regression on each cell's dispersion as a function of velocity. The average slope was significantly above zero on CF and CG (CF: slope = 0.019 ± 0.0044, $P(\text{slope} = 0) = 0.000048$; CG: slope = 0.018 ± 0.0047, $P(\text{slope} = 0) = 0.00022$; slopes are mean ± SE). This relationship was also found on LT and LT $A \rightarrow B$, but was not significant on LT $B \rightarrow A$ (LT: slope = 0.010 ± 0.0049, $P(\text{slope} = 0) = 0.037$; LT $A \rightarrow B$: slope = 0.0077 ± 0.0028, $P(\text{slope} = 0) = 0.010$; LT $B \rightarrow A$: slope = 0.0036 ± 0.0033, $P(\text{slope} = 0) = 0.29$; slopes are mean ± SE). The intercepts were significantly negative on all task conditions except on LT $B \rightarrow A$ (negative, but not significant; data not shown).

Testing this trend across all neurons allows an estimate of the amount of variance in the *z*-score that can be explained by speed. On all tasks, there was an overall significant positive effect of velocity on the dispersion *z*-score for each pass (*LT*: $R^2 = 0.016$, $F = 62.8$, $p < 0.0000001$; LT $A \rightarrow B$: $R^2 = 0.009$, $F = 8.1$, $p < 0.005$; LT $B \rightarrow A$: $R^2 = 0.011$, $F = 15.0$, $p < 0.0002$; CF: $R^2 = 0.020$, $F = 77.2$, $p < 0.0000001$; CG: $R^2 = 0.015$, $F = 65.5$, $p < 0.0000001$; for more detail see Table A1). Thus, although speed has a significant effect on biasing the z-score toward higher firing rates at higher velocities, this effect is not sufficient to explain more than about 4% of the variance observed in our data as evidenced by R^2 values of 0.04 or less.