Neuron, Vol. 48, 345–358, October 20, 2005, Copyright ©2005 by Elsevier Inc. DOI 10.1016/j.neuron.2005.09.007

Progressive Transformation of Hippocampal Neuronal Representations in "Morphed" Environments

Jill K. Leutgeb,[1] Stefan Leutgeb,[1] Alessandro Treves,[1,2] Retsina Meyer,[1,3] Carol A. Barnes,[1,3] Bruce L. McNaughton,[1,3] May-Britt Moser,[1] and Edvard I. Moser[1,*]

[1]Center for the Biology of Memory
Norwegian University of Science and Technology
NO-7489 Trondheim
Norway

[2]Cognitive Neuroscience Sector
International School for Advanced Studies
Trieste I-34014
Italy

[3]Neural Systems, Memory and Aging
University of Arizona
Tucson, Arizona 85724

Summary

Hippocampal neural codes for different, familiar environments are thought to reflect distinct attractor states, possibly implemented in the recurrent CA3 network. A defining property of an attractor network is its ability to undergo sharp and coherent transitions between pre-established (learned) representations when the inputs to the network are changed. To determine whether hippocampal neuronal ensembles exhibit such discontinuities, we recorded in CA3 and CA1 when a familiar square recording enclosure was morphed in quantifiable steps into a familiar circular enclosure while leaving other inputs constant. We observed a gradual noncoherent progression from the initial to the final network state. In CA3, the transformation was accompanied by significant hysteresis, resulting in more similar end states than when only square and circle were presented. These observations suggest that hippocampal cell assemblies are capable of incremental plastic deformation, with incongruous information being incorporated into pre-existing representations.

Introduction

How can the brain reactivate representations of the external world in a constantly changing physical environment? It is commonly thought that this ability reflects the operation of attractor-based neural networks (Hopfield, 1982; Amit, 1989). An attractor network has one or several preferred positions or volumes in state space, such that arbitrary external inputs prompt neuronal activity to evolve dynamically and to approach one of these attractors, usually the one most closely correlated with the inputs. Attractor networks may support a number of cerebral functions with discrete values such as sensory pattern recognition (Bartlett and Sejnowski, 1998; Fdez Galan et al., 2004; Rotshtein et al., 2004), categorical perception (Wyttenbach et al., 1996; Freedman et al., 2001), and execution of movement trajectories (Lukashin et al., 1996).

Attractor boundaries are thought to be dynamic, reflecting the continuous modification of connection strengths within and between cell ensembles as new information is encoded into the network. Thus, attractor networks with modifiable synaptic weights are especially powerful models for storage and retrieval of memory at the neural population level (Rolls and Treves, 1998). Several factors indicate that the hippocampus contains a recurrent autoassociative attractor-based neural network for episodic and spatial memories (Marr, 1971; McNaughton and Morris, 1987; Treves and Rolls, 1992, 1994; McClelland and Goddard, 1996; Hasselmo et al., 1995). By means of the extensive associative connections of the CA3 subfield, the hippocampus is potentially able to store a very large number of arbitrary representations of discrete input patterns (Treves and Rolls, 1991, 1994) and/or a large number of spatial representations organized into two-dimensional continuous attractors, or charts (Samsonovich and McNaughton, 1997; Battaglia and Treves, 1998). Consistent with these models, recent studies of hippocampal place cells have shown that hippocampal neurons organize into assemblies whose synchrony exceeds that of their common input (Harris et al., 2003), but the possible role of attractor-like dynamics in hippocampal ensembles is still poorly understood.

If distinct attractors exist, then attractor network models predict a sharp and coherent transition as the network is probed with inputs that vary progressively along a continuum from one pre-established basin of attraction to another. In these models, synaptic plasticity is typically switched off when retrieval dynamics are tested. Small changes in input, compared to a familiar condition, are then expected to cause pattern completion, resulting in output representations that are more similar than their inputs, whereas larger changes should more frequently cause pattern separation, yielding outputs that are more dissimilar than their inputs (McClelland and Goddard, 1996; Rolls and Treves, 1998). Indirect evidence suggests that hippocampal assemblies perform both of these processes (Guzowski et al., 2004). Pattern completion is expressed as the persistence of hippocampal spatial representations after removal or modification of significant subsets of the original input (O'Keefe and Conway, 1978; Muller and Kubie, 1987; Quirk et al., 1990; Nakazawa et al., 2003). Pattern separation is expressed as the occasional tendency of hippocampal representations to show complete "remapping" in the presence of minor changes in the physical input patterns (Muller and Kubie, 1987; Bostock et al., 1991; Markus et al., 1995). However, although such nonlinearities would be expected in an attractor network, they could also merely mirror inhomogeneities in the influence of sensory inputs, with some stimuli contributing more to neuronal firing than others (Gothard et al., 1996; Hetherington and Shapiro, 1997; Hollup et al., 2001). Completion and separation might be assessed during systematic transformation

*Correspondence: edvard.moser@cbm.ntnu.no

of the same input stimuli to resolve this question. One recent study using such an approach has shown that sharp transitions between hippocampal network states can be induced by morphing the shapes of two distinct recording enclosures (Wills et al., 2005). Wills and colleagues used a particular protocol to obtain rapid discrimination, and testing was performed by presenting intermediate shapes in a scrambled order, presumably countering the cumulative effect of gradual synaptic changes. In the present study, we asked whether representations remain discontinuous under conditions more similar to those encountered during encoding and retrieval of episodic and spatial memories in natural behavior when changes in inputs are gradual and plasticity is not switched off.

Results

To examine the dynamics of neural ensembles during progressive transitions between pre-existing network states, we recorded hippocampal ensemble activity when input patterns were changed incrementally between two well-learned extremes, in the context of a single experiment (Figure 1). Seven rats were trained daily to search for randomly scattered food in square and circular versions of the same enclosure presented in random order. Pretraining continued until the two conditions reliably activated different subsets of neurons in the associative CA3 network. Ensemble activity was then recorded from 136 pyramidal cells in CA3 and 252 pyramidal cells in CA1 while the square box was morphed into a circle, or vice versa, through a series of five intermediate shapes. The sequence repeated in alternating directions on subsequent days (220 CA3 and 305 CA1 cells).

Remapping during Pretraining in the Square and Circular Shapes

All rats were pretrained for 16–19 days in the square and circle (Figure 1B). At the end of training, neuronal activity was recorded simultaneously from CA3 and CA1. Different subsets of neurons were active in CA3, in which attractor-based networks of the hippocampus are thought to be located (Figure 2A). Stable remapping was observed also in CA1 (Figure 2B), although here, the remapping was only partial, as in previous studies using similar training regimes (Lever et al., 2002; Leutgeb et al., 2004).

The difference in the subset of active neurons in the two extreme shapes was first compared with a measure of rate similarity, or "overlap" (Leutgeb et al., 2004) (see Experimental Procedures). For each cell, the firing rate in the less-active shape was divided by the rate in the more-active shape. The ratios were corrected for overlap expected by chance so that identical firing rates would result in a score of 1 and independent rate distributions in a mean score of 0 (see Experimental Procedures). Cumulative distribution functions were then computed in order to emphasize possible differences in the frequency distributions without having to bin the data (Figure S1; see also Experimental Procedures). In CA3, the representations for the two shapes were nearly but not completely independent with an overlap score of 0.10 ± 0.03 for the average rates (mean ± SEM)

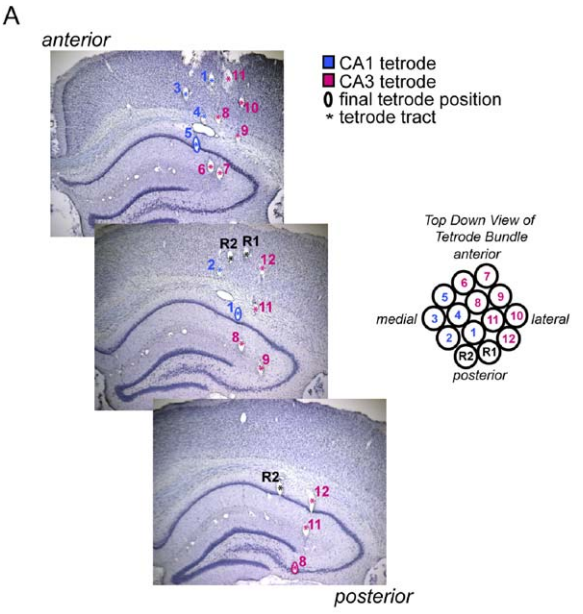

Figure 1. Recording Locations and Experimental Design

(A) Top-down view of a 14-tetrode bundle implanted at AP 3.8 and ML 3.0 in rat 10683 (right). Tetrodes R1 and R2 are reference tetrodes. The tracks are shown in three different coronal sections of the implanted brain (left). The brain was cut at an angle of about 10° compared to the electrode tracks to visualize short segments (asterisk; CA1 in blue, CA3 in red). The tip of each electrode was found by comparison with adjacent sections, and three examples (Tetrode 1, 5, and 8) are outlined (0).

(B) Randomly ordered sequences of a total of eight squares and circles were presented for 10 min each during pretraining. Pretraining lasted 16–19 days.

(C) Morphing experiments started either with the square or the circle. The initial shape (shape 1) was then gradually transformed to the opposite shape (shape 7) through a series of intermediate shapes (2–6). At the end of the sequence, the initial shape was presented a second time (1′). The morph sequence was repeated in the opposite direction on the second day. Together, the forward and backward sequence complete a morph set.

(Figure 2C). Repeated recordings in the same enclosure yielded overlap scores of 0.52 ± 0.03. In CA1, the overlap between the square and the circle was 0.26 ± 0.03, which was smaller than between trials with identical shapes (0.40 ± 0.02, $t = 3.43$, d.f. = 375, $p < 0.001$) but much larger than in CA3, as previously reported (Leutgeb et al., 2004).

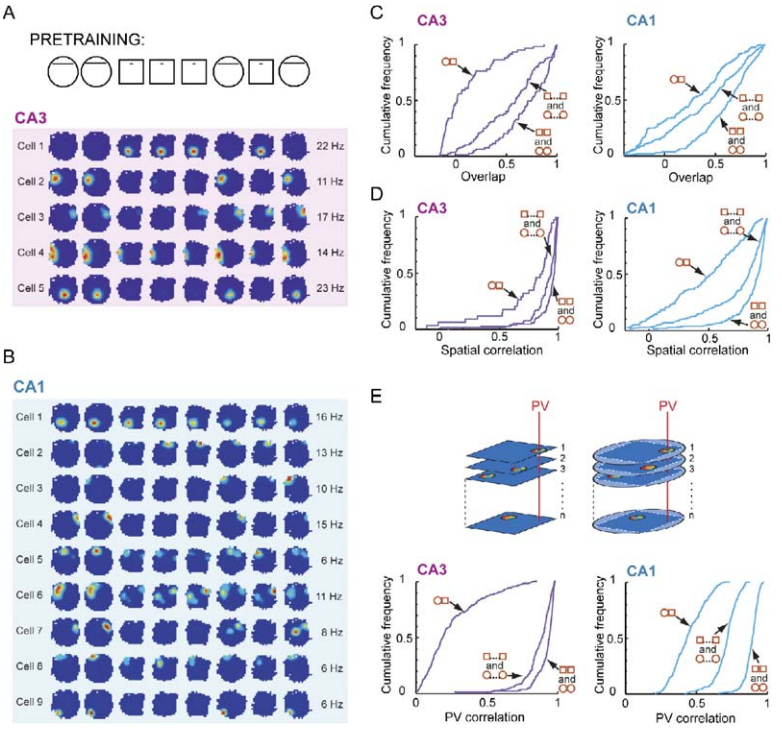

Figure 2. Pattern Separation in CA3 and CA1 at the End of Training in the Familiar Shapes

(A and B) Color-coded rate maps for place cells recorded in square and circular versions of the morph box with one representative tetrode in CA3 (A) and one in CA1 (B) (rat 10708). For each cell (i.e., each row), the rate maps were scaled to the maximum firing rate within the entire pretraining sequence. Peak rates are indicated to the right. The color scale is from blue (silent) to red (peak rate). Pixels not sampled are white. Silent cells are not shown ([A] n = 4; [B] n=1). When squares and circles were presented in a random order during pretraining, most CA3 cells were on in one shape and nearly off in the opposite shape.

(C) Cumulative distribution functions (see Figure S1 and Experimental Procedures) for overlap between firing rates in CA3 (left) and CA1 (right) in opposite box shapes (red circle and red square) and in identical box shapes (pairs of red squares or circles). For the identical box shapes, overlap scores are plotted separately for consecutive trials (without stippled lines) and the trials with longest separation (with stippled lines indicating three to six intervening shapes). Note the increased frequency of low overlap scores in CA3 compared to CA1 when box shapes are different; this difference indicates a more pronounced orthogonalization in CA3 (Leutgeb et al., 2004).

(D) Cumulative distribution functions for spatial correlations between pairs of rate maps. Symbols as in (C). The similarity of the distributions in CA3 suggests that in each cell's inactive state, residual spikes occur in the area corresponding to the place field in the active state (Leutgeb et al., 2005).

(E) Change in population activity between square and circle. The rates of all recorded CA3 or CA1 cells were stacked into 256 composite population vectors (PVs), one for each of the 16 × 16 pixels (25 cm^2) that was shared between the morph square and morph circle, and the correlation between the population vectors was computed at each corresponding location (top). Bottom panels show cumulative distribution functions for population vector correlations between square and circle in CA3 (left) and CA1 (right). Symbols as in (C). Note the pronounced change in CA3 population activity between the two box shapes.

Although the subset of active neurons changed almost completely in CA3, most cells that were active in one shape discharged at least a few spikes also in the other. To determine whether these spikes occurred at the same location in both configurations, we computed the spatial correlations for the common parts of the rate maps in the square and the circle for each cell (see Experimental Procedures). In CA3, the spatial correlation between square and circle was only marginally lower than between repeated trials in the same shape (0.74 ± 0.05 versus 0.89 ± 0.02; t = 3.90, d.f. = 118, p < 0.001) (Figure 2D and Figure S2), suggesting that most residual spikes occurred in the area corresponding to the firing field in the alternative shape. In CA1, the spatial correlation between rate maps in the square and the circle was lower (0.46 ± 0.03 for different shapes versus 0.73 ± 0.02 for repeated trials in the same shape, t = 8.51, d.f. = 356, p < 0.001). The limited spatial remapping in CA3 agrees with recent data (Leutgeb et al., 2005) indicating that hippocampal place cell ensembles can be altered in two fundamentally different ways: one in which firing rates change substantially while firing locations remain constant (rate remapping) and one in which rate and location vary independently (global remapping); in the present case, replacing the enclosure induced primarily rate remapping.

Finally, we estimated the change in hippocampal population output for each spatial bin that was shared between the square and the circle with a measure that is sensitive to location as well as magnitude of firing in the sample. Composite population vectors were constructed for each location in each box by stacking the rate maps of all recorded cells to obtain the firing rates of the cell population in each spatial bin (Figure 2E). In CA3, population vectors from corresponding locations in the square and circular box were only weakly correlated at the end of pretraining (0.26 ± 0.01; mean correlation ± SEM), indicating strong pattern separation between the two shapes of the box (Figure 2E). Recordings from identical boxes remained highly correlated (0.87 ± 0.01). A similar but less extensive difference was observed in CA1 (different shapes: 0.44 ± 0.01; same shape: 0.70 ± 0.01, t = 29.95, d.f. = 510, p < 0.001). These analyses suggest that at each corresponding location of the square and the circle, the population output was radically changed between the two configurations, especially in CA3, but also significantly in CA1.

Progressive Transformation in CA3 Ensembles during Morphing

Once stable differences between the representations for the two box shapes had been verified, one shape

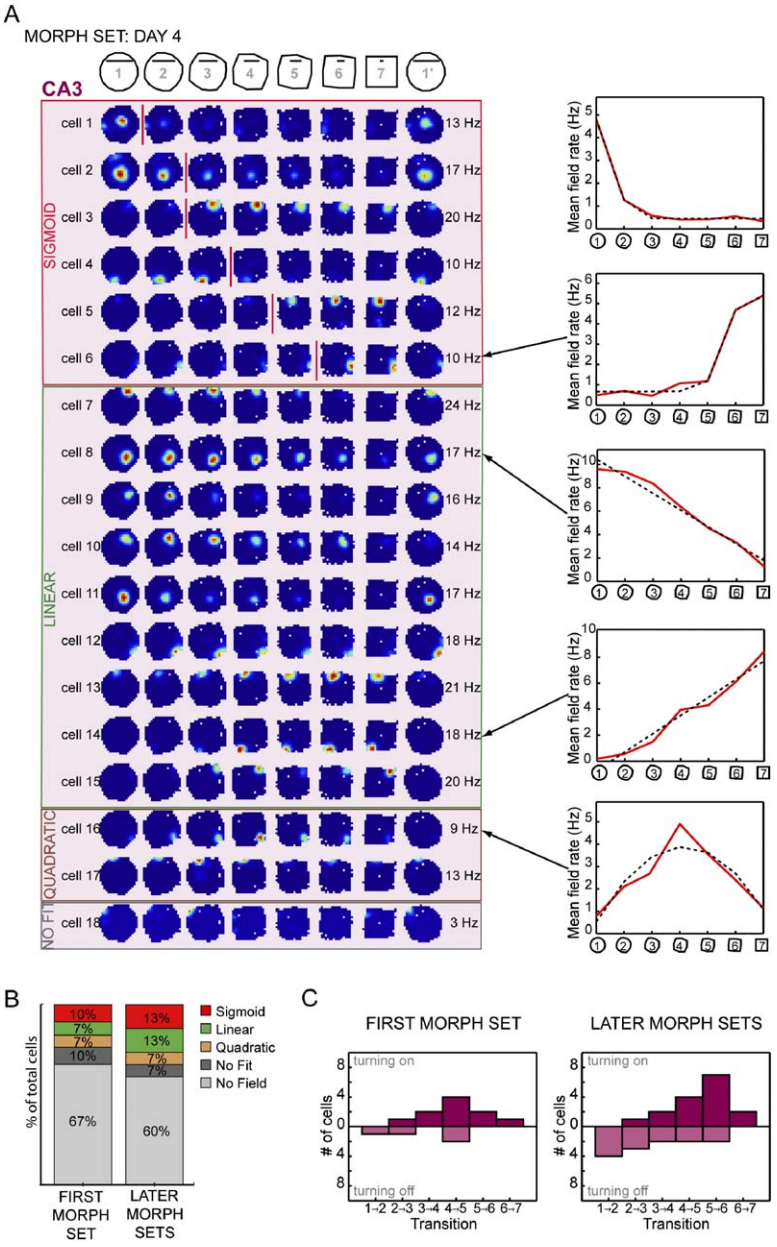

Figure 3. Gradually and Abruptly Changing Subpopulations Coexist in a Simultaneously Recorded CA3 Cell Population

(A) Rate maps for all active cells in a morphing trial for rat 10841 (n = 38 CA3 cells; scaled as in Figure 2). Silent cells are not shown. The mean rate for all pixels within the place field was calculated for each of the seven shapes (1 and 1′, circle; 7, square; 2–6, intermediate shapes), and the rates were then fitted to sigmoid, linear, or quadratic functions. In the example, while morphing a circle to a square, the fields of six neurons showed significant sigmoid curve fits with transition points after shapes 1, 2, 3, 4, and 5 (indicated by red lines). Nine fields showed more explained variance with linear fits, and two with quadratic fits. One field was not fitted and 20 cells had low rates without identified fields. Examples of fitted fields are shown on the right (mean rate in red; fitted curve in black).

(B) Distribution of CA3 cells with linear and nonlinear responses during morphing. Of all CA3 cells with fields (133 of 356), similar proportions had either sigmoid or linear fits, and smaller proportions had either a quadratic or no significant fit. Consistent with the decreased hysteresis in later morph sets, we observed a small increase in sigmoid and linear cells at the expense of unfitted and silent cells.

(C) The transition points of the small subset of CA3 cells with sigmoid curve fits in (B) were distributed across the entire morphing sequence (left: first morph set, i.e., first two days; right: later morph sets). Top histogram: cells that turned on (e.g., cell 5 in A). Bottom: cells that turned off (e.g., cell 1 in A).

was transformed to the other via five intermediate shapes (Figure 1C). Contrary to expectations based on attractor network models, the gradual transformation of the environment was accompanied by an equally gradual change of the ensemble representation in CA3 (Figure 3). The dissimilarity in ensemble activity built up linearly, with equally large changes in overlap values and population vector correlations between successive shapes in the morph sequence (Figure 4).

To determine whether the gradual change in average measures of representational similarity masked a range of sudden transitions, we examined the dynamics of individual neurons. Abruptly shifting neurons were defined as neurons with a significant sigmoidal curve fit in which the explained variance was larger than with a linear or quadratic curve fit. The explained variance was measured with F values, which are, by definition, corrected for differences in degrees of freedom (d.f. = 2 for linear fits, d.f. = 4 for sigmoid fits). Of the 356 CA3 neurons recorded during all morph experiments, 133 had an identified field. Only 42 of these neurons showed rate changes that were optimally fit by a sigmoidal curve (Figures 3A and 3B). In these neurons, the transition from high to low rates or vice versa was distributed across the entire morphing sequence, with a slight overrepresentation of cells that turned on in the later shapes (Figure 3C). The remaining 91 fields either failed to exhibit significant rate changes (n = 28; p values for all curve fits > 0.05), or the explained variance was larger with a linear curve fit (n = 39) or a quadratic fit (n = 24). There was no significant correlation between the degree of remapping in CA3 (as measured by population vector correlations) and the proportion of cells showing a sigmoidal shift (n = 17 sessions, r = −0.01, n.s.). The proportion of cells with sigmoid and linear fits increased slightly as the morph sequence was

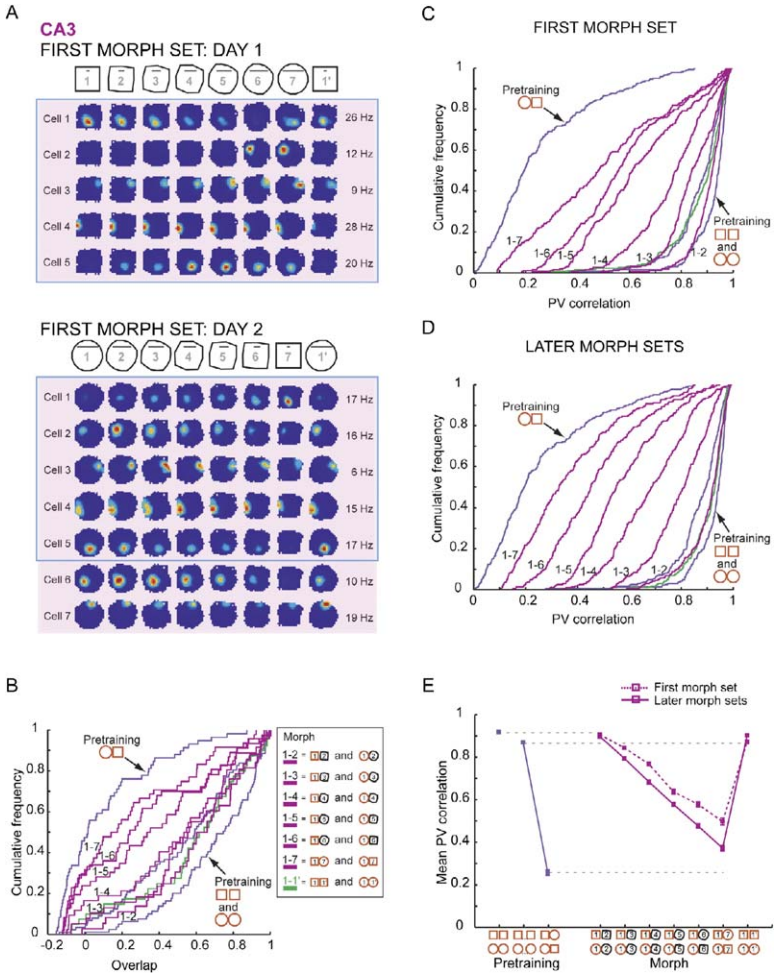

Figure 4. Hysteresis in CA3 during the Morph Sequence

(A) Rate maps for all place cells recorded with one representative tetrode in CA3 during the first morph set (rat 10708). Silent cells not shown (n = 6 for days 1 and 2). Cells outlined by blue boxes were also recorded during pretraining (Figure 2A). Note persistence of firing in several of the cells. For example, cell 2, which was selective for the circle during pretraining, did not turn on until the shape before the circle when the morphing sequence was started with the square (day 1). However, when the sequence started with a circle on day 2, pronounced firing persisted for four intermediate shapes.

(B) Cumulative distribution functions for overlap of CA3 firing rates between the initial shape and the various configurations of the morph box (red lines, 1–2 to 1–7; green line, 1–1′; pairs of small boxes indicate which shapes were compared; end shapes in red and intermediate shapes in black). Purple traces indicate overlap between end shapes before morphing started (as in Figure 2C). When comparing the same shapes after an intervening morphing sequence (1 versus 7), the overlap with the initial representation was larger than when the transition was direct. Overlap scores were restored to control levels when the initial shape was repeated after the morph sequence (1 versus 1′).

(C) Cumulative distribution functions for correlations between population vectors in the initial shape and each successive shape through the morph sequence (colors and symbols as in [B]). Note the significant bias toward higher correlations (hysteresis) as the opposite shape was approached.

(D) Cumulative distribution functions for later morph sets. Note the reduction of hysteresis with repeated exposure to the intermediate shapes.

(E) Mean CA3 population vector correlation during pretraining, early morph sets, and later morph sets. The square-circle (1–7) population vector correlation during morph sequences (red) was higher than the corresponding square-circle correlation after direct transitions in pretraining (purple).

repeated (Figures 3B and 3C). Taken together, these analyses strongly point to a progressive transformation of the active CA3 ensemble, with cells gradually fading in or out, shifting at different points in the sequence, or firing selectively in the intermediate shapes. There was no coherent all-or-none transition to the "opposite" representation during any recording session.

The gradual transformation of the CA3 representation was accompanied by significant hysteresis, i.e., a lag in the change of the rate distribution (Figure 4). This was apparent when comparing activity in the end shapes of the morph sequence (square versus circle) with activity in the same shapes during pretraining (either the first square and the last circle, or the first circle and the last square). Overlap scores increased from 0.10 ± 0.03 during pretraining to 0.24 ± 0.05 during morphing ($t = 2.44$, d.f. = 105, $p < 0.05$) (Figure 4B). Hysteresis was also apparent in the population vector correlations (Figures 4C and 4E). During the first morph sequence in each direction, the mean correlation with the initial set of population vectors decreased in approximately equal steps, from 0.91 ± 0.004 (shape 1 versus 2 in Figure 1C) to 0.50 ± 0.02 (shape 1 versus 7, i.e., square and circle). The end value (0.50 ± 0.02) was significantly larger than the correlation between the same shapes (square and circle) during pretraining (0.26 ± 0.01; $t = 11.63$, d.f. = 510, $p < 0.001$). When the first shape in the sequence was reintroduced after the morphing sequence, the original rate distribution was reinstated (Figures 4B and 4C). Repetition of the morph sequence reduced the hysteresis (Figures 4D and 4E). Overlap values decreased from 0.24 ± 0.05 to 0.19 ± 0.03; population vector correlations decreased from 0.50 ± 0.02 to 0.37 ± 0.01. The correlations were still elevated compared to those obtained during direct transitions ($t = 6.43$, d.f. = 510, $p < 0.001$). As in the pretraining stage, there was no change in the location of residual firing fields in CA3 (Figure S2). Together, these data suggest that experience in the intermediate shapes attenuates the differences between the previously established end-point representations in CA3.

Progressive Transformation in CA1 Ensembles during Morphing

CA1 cells were recorded simultaneously with the CA3 cells. The transition from a "square" to a "circle"

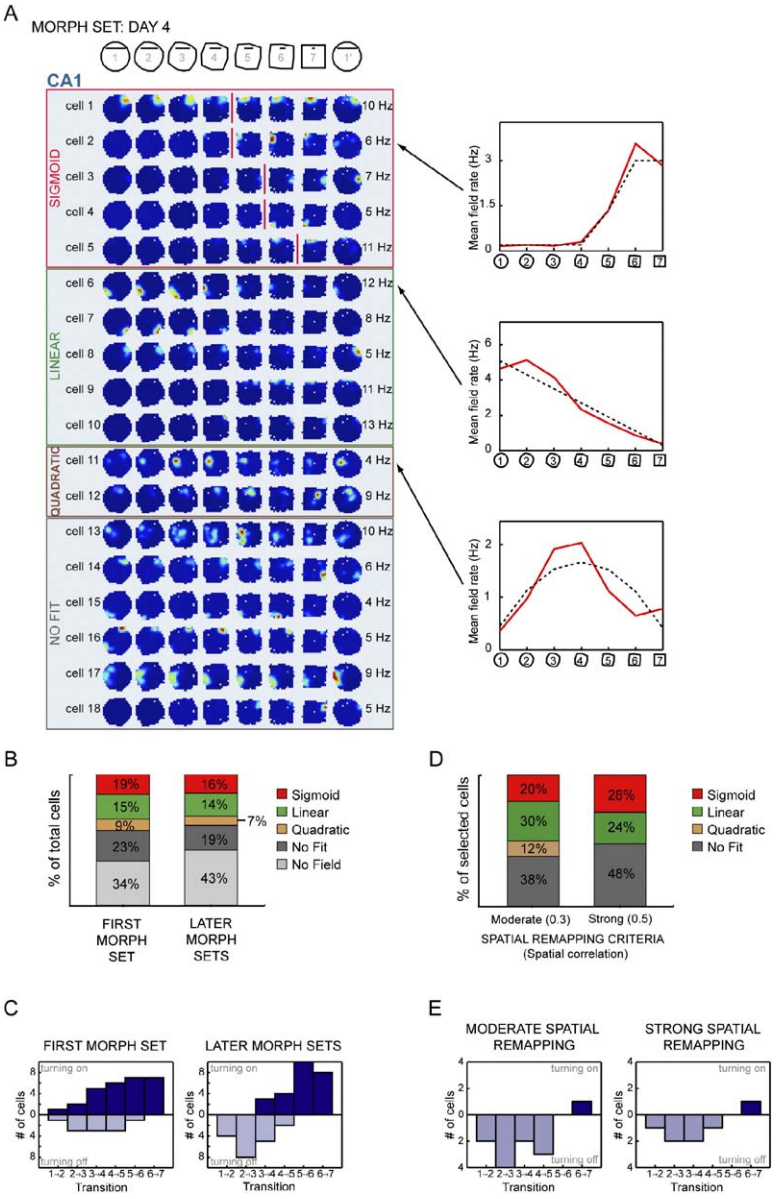

Figure 5. Gradually and Abruptly Changing Subpopulations also Coexist in Simultaneously Recorded CA1 Cells

(A) Rate maps for all active cells in a morphing trial for rat 10841 (n = 25 CA1 cells; simultaneously recorded with the CA3 cells in Figure 3). Silent cells are not shown. While the circle morphed to a square, the fields of five neurons showed significant sigmoid curve fits with transition points after shapes 4, 5, and 6 (indicated by red lines). Five additional fields showed linear fits, and two showed quadratic fits. Examples of fitted fields are shown on the right with the mean field rate in red and the fitted curve in black. Six fields were not fitted and seven cells had low rates without identified fields.

(B) Distribution of cells with linear and nonlinear responses during morphing. Of all CA1 cells with fields (341 of 557), similar proportions had linear, sigmoid, or no significant fits, and a smaller proportion had quadratic fits.

(C) Transition points of CA1 cells with sigmoid curve fits were not coherent. Top histogram: cells that turned on. Bottom: cells that turned off. During later morph sets, cells that turned off tended to have earlier transitions, whereas those that turned on had later transitions.

(D) Distribution of linear and nonlinear responses in the subset of selected CA1 cells with moderate or strong spatial remapping between square and circle. Moderate spatial remapping was defined as a difference of at least 0.3 in the spatial correlation between trials in the same shape and in opposite shapes. Strong remapping was defined as a difference of at least 0.5 (Wills et al., 2005).

(E) Transition points for the subsets of cells with sigmoid curve fits in (D).

representation was as smooth in CA1 as in CA3 (Figure 5). The decrease in population vector correlations, overlap, and individual spatial correlations was approximately linear (Figure 6). Regression analyses for individual cells showed that only 83 of 341 (24.3%) active cells were best fit with a sigmoid function (Figures 5A and 5B). The remaining 258 cells were linear (97 cells) or quadratic (43 cells), or none of the fits were significant (118 cells; these cells were usually those with no change in firing rate). Transition points for the subset of neurons with a sigmoid curve fit were widely scattered, perhaps with a weak overrepresentation among the later shapes during the first morph set (Figure 5C). When the morph sequence was repeated on subsequent days, the transition points for cells that turned off moved toward earlier shapes, whereas transitions for cells that turned on moved toward later shapes (Figure 5C).

The strong hysteresis in CA3 was not mirrored in the CA1 recordings. Although representations in CA3 generally lagged behind when the shapes were transformed incrementally, CA1 representations in the square and the circle were equally different during morphing and pretraining (Figure 6). The overlap of firing rates between square and circle was comparable (morphing: 0.18 ± 0.03; pretraining: 0.26 ± 0.03; t = 1.85, n.s.), whereas the population vector correlations were even lower during morphing than during pretraining (0.38 ± 0.01 versus 0.44 ± 0.01; t = 5.81, d.f. = 510, p < 0.001). There was also significant divergence of the firing locations, as expressed in a reduction of spatial correlations compared to trials with identical shapes, but the reduction was lower during morphing than pretraining (spatial correlations of 0.57 ± 0.03 versus 0.46 ± 0.03, respectively; t = 2.45, d.f. = 254, p < 0.05) (Figure S2). There was no further development with repetitions of the morph sequence (Figure 6D).

Although the morph procedure did not by itself lead to prolonged hysteresis in the CA1 representation, there

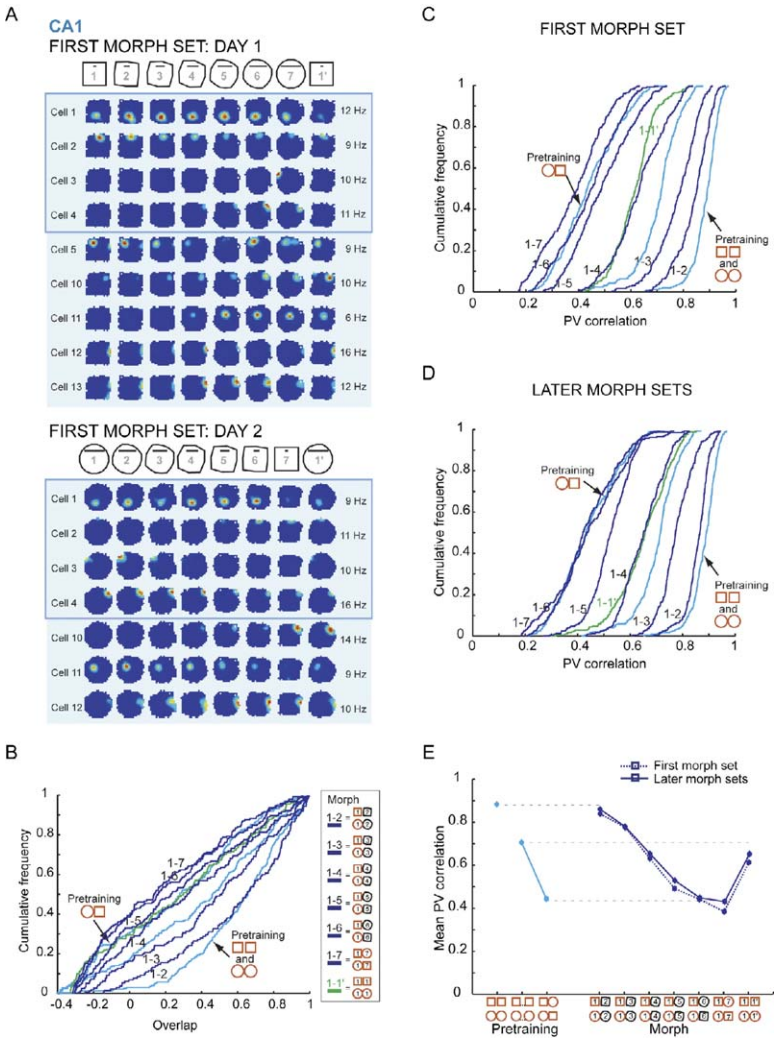

Figure 6. Carry-Over Effects in CA1

(A) Color-coded rate maps showing all fields recorded with one representative tetrode in CA1 during the first morph set (rat 10708). The CA1 fields were recorded simultaneously with the CA3 fields in Figures 2 and 4. Cells in blue boxes were recorded on both days as well as during pretraining (Figure 2B). Silent cells not shown (day 1, n = 1; day 2, n = 2). Although some cells (e.g., cell 4) exhibited hysteresis, many (e.g., cell 10 and cell 11) showed pronounced carry-over effects after reversal to the initial shape at the end of the morph sequence (compare 7 with 1′).

(B) Cumulative distribution functions showing overlap scores in CA1 between the initial shape and each successive shape. Note a lower degree of rate similarity during morphing (dark blue traces) than during direct transitions (pretraining, light blue traces) between the end shapes. Two curves are shown for direct transitions between identical shapes: one for the consecutive trials (right) and one for the pairs of trials with longest temporal separation (middle). Overlap scores did not return to initial levels when the first shape was presented again at the end of the morph sequence (1 versus 1′; green trace).

(C) Cumulative distribution functions for the correlation coefficients of population vectors during the first morph set (pretraining in light blue, morphing in dark blue). Population vector correlations within the morph sequence started at approximately the same-shape level (1 versus 2) and reached, during later parts of the morph sequence (1–6 to 1–7), the level of dissimilarity of the pretraining condition.

(D) Cumulative distribution functions for later morph sets. No differences were apparent between early and late sets, and representations of the end shapes remained as orthogonalized during morphing as during direct transitions in pretraining.

(E) Mean CA1 population vector correlation during pretraining, early morph sets, and late morph sets.

were significant carry-over effects between consecutive trials. During pretraining, overlap values were higher between consecutive trials with the same shape than between the first and last trial with the same shape (0.61 ± 0.02 versus 0.40 ± 0.02; t = 6.87, d.f. = 457, p < 0.001). The increased overlap between adjacent trials was relatively insensitive to the geometric shape of the boxes used on those trials. For example, when the initial shape was reintroduced at the end of the morph sequence, the CA1 representation failed to revert fully back to its original state but instead remained more similar to the preceding opposite shape (population vector correlation for shape 7 versus 1′: 0.68 ± 0.01; see for example, cell 10 in Figure 6A). No carry-over effects between opposite shapes were seen in simultaneous recordings from CA3. These data suggest that CA1 cells, unlike CA3 cells, transfer nonspatial information from previous trials in addition to responding to changes in the spatial environment.

Extent of Spatial Remapping

The continuous nature of the transition in our study differs strongly from the abrupt and coherent shift observed in CA1 with a different training procedure (Wills et al., 2005). In the latter study, only cells with strong spatial remapping were included in the analyses. We asked whether cells that displayed strong spatial remapping were more likely to show a steep transition in our data. The Wills criterion was a difference of at least 0.5 in the spatial correlation between trials in the same shape and in opposite shapes (square and circle). This value cannot be compared directly to the present sample because of differences in place field calculation and smoothing, but we examined the sharpness of individual transitions with a range of thresholds, two of which are reported here (spatial correlation differences of 0.5 and 0.3, respectively). With the most stringent criterion (0.5), only 5 CA3 cells and 25 CA1 cells passed (1.4% and 4.5%, respectively). Only 9 of these cells (2 in CA3 and 7 in CA1) were best fit with a sigmoid function (see Figure 5D for the CA1 cells). With the more relaxed criterion (0.3), 10 CA3 cells and 60 CA1 cells passed, of which only 16 cells (4 in CA3 and 12 in CA1) were best fit with a sigmoid function (Figure 5D). With both criteria, transition points for the selected cells with sigmoid curve fits were as distributed as for the whole population (Figure 5E). Thus, the transition between the end representations

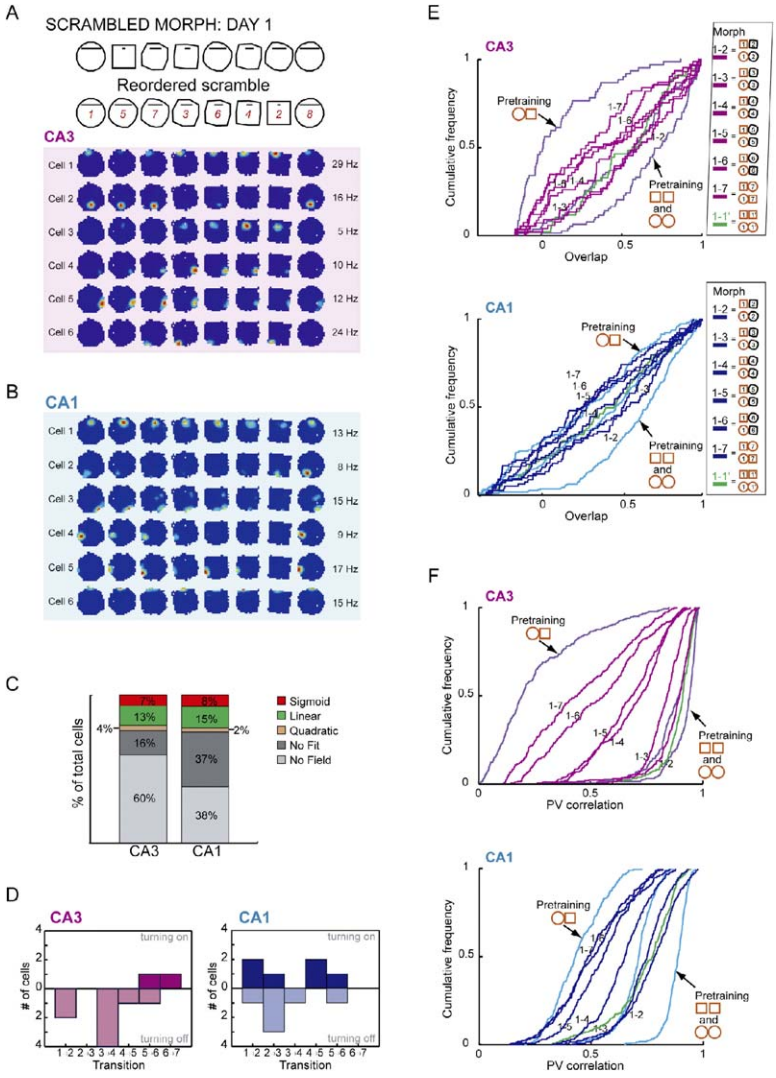

Figure 7. Progressive Changes in Ensemble Representations in Scrambled Morph Sequences

(A and B) Rate maps for all active cells on two representative tetrodes in CA3 (A) and CA1 (B) of rat 10841. Silent cells not shown ([A] n = 7; [B] n = 2). Maps are presented according to difference in shape rather than chronological order. Red numbers indicate order of presentation.

(C) Distribution of cells with linear and nonlinear responses in CA3 and CA1.

(D) Transition points of CA3 and CA1 cells with sigmoid curve fits. Top histogram: cells that turned on. Bottom: cells that turned off.

(E) Cumulative distribution functions for overlap scores in CA3 (pretraining in purple, scrambled morph in red) and CA1 (pretraining in light blue, scrambled morph in dark blue).

(F) Cumulative distribution functions for population vector correlations in CA3 and CA1 during scrambling. The first shape tended to bias the representation of later shapes, particularly in CA3.

was as gradual and noncoherent in the subset of strongly remapping cells as in the rest of the cell sample.

Scrambled Morph Sequences

The progressive transformation of network activity in our experiments could reflect the incremental nature of the test paradigm. The sequential presentation of the intermediate shapes caused significant hysteresis in the hippocampal representation, which may have prevented a coherent switch to the representation associated with the alternate end shape. To examine whether the transition would be more discontinuous under noncumulative test conditions, we finally presented the intermediate shapes in scrambled order.

The transition from a "square" to a "circle" representation was gradual even with a scrambled presentation sequence (Figures 7A and 7B). Regression analyses for individual cells showed that only 22 of the 298 active CA3 and CA1 cells were best fit with a sigmoid function (Figure 7C). Transition points in this subgroup were distributed, as in the incremental morph paradigm (Figure 7D). The remaining cells were linear (42 cells), quadratic (9 cells), or without a significant fit (81 cells).

Scrambling of the presentation sequence did not prevent hysteresis (Figures 7E and 7F). The overlap of CA3 firing rates in the square and the circle was higher during scrambled morphing than when transitions were direct (0.32 ± 0.04 and 0.10 ± 0.03, respectively; t = 3.82, d.f. = 112, $p < 0.001$). When the shapes were ordered according to similarity with the initial shape, the mean population vector correlation in CA3 decreased progressively from 0.88 ± 0.01 (shape 1 versus 2) to 0.46 ± 0.01 (shape 1 versus 7, i.e., square and circle). The correlation of opposite shapes was significantly stronger than during previous direct transitions (t = 10.49, d.f. = 510, $p < 0.001$). In CA1, the overlap of firing rates in the square compared to the circle decreased to a similar extent within the scrambled sequence (0.30 ± 0.04) as during pretraining or during sequential presentation (t = 0.76 and t = 1.11, respectively, n.s.). The mean population vector correlation also decreased but to a smaller extent than with sequential presentations (Figure 7F) (0.77 ± 0.01 for shape 1 versus 2; 0.51 ± 0.01 for shape 1 versus 7). Taken together, these data suggest that sequence randomization is not sufficient to prevent the impact of the initially activated cell population on

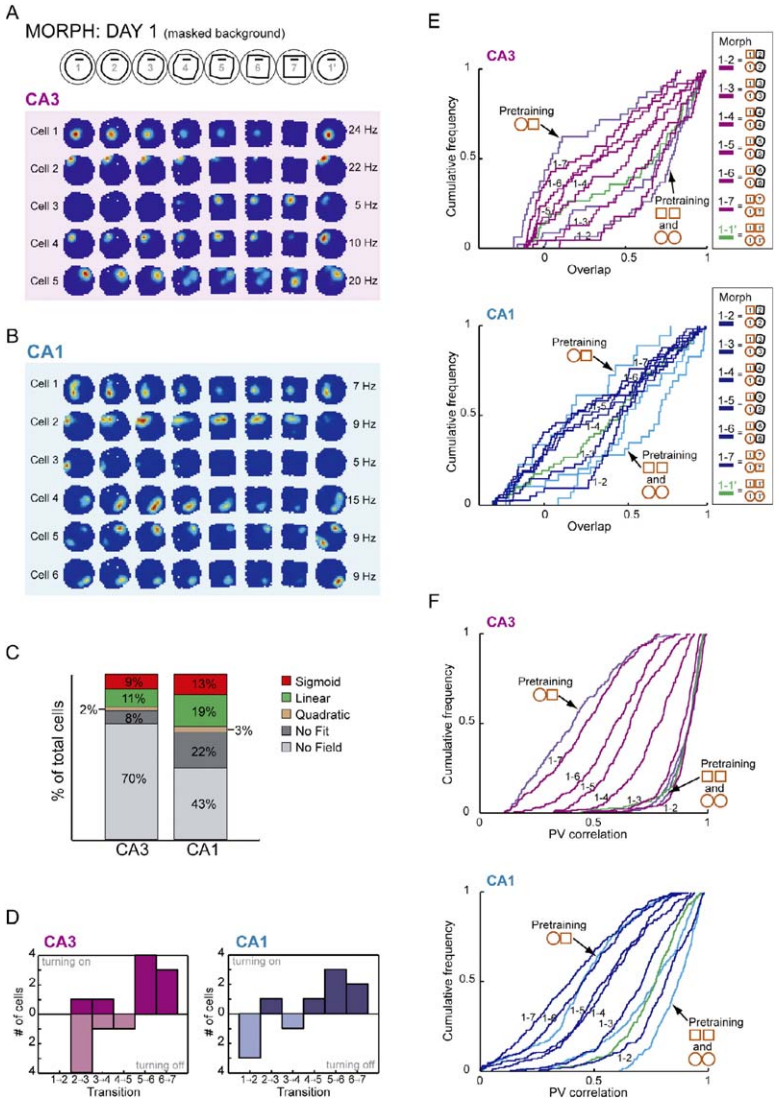

Figure 8. Progressive Transformation of the Hippocampal Representation When the Background Cues Are Masked during the Morphing

(A and B) Rate maps for all active cells on two representative tetrodes in CA3 (A) and CA1 (B), respectively (rat 11121). Silent cells not shown ([A] n = 8; [B] n = 5).

(C) Distributions of cells with linear and nonlinear responses in CA3 and CA1.

(D) Transition points of CA3 and CA1 cells with sigmoid curve fits. Top histogram: cells that turned on. Bottom: cells that turned off.

(E) Cumulative distribution functions for overlap scores in CA3 and CA1 during morphing without background cues available (color schemes as defined in previous figures). Scores from experiments with direct transitions (i.e., pretraining) in the same rats are indicated for comparison.

(F) Cumulative distribution functions for population vector correlations in CA3 and CA1 during morphing without background cues available.

activity during subsequent shapes. The data also show that network states change progressively between similar but distinct shapes regardless of the order.

Masking the Background Cues

One possible limitation of our experiments is that training and testing were conducted without hiding the background cues. The combination of changing box shapes and constant distal background cues may have reduced population coherence in the same way that differential rotation of proximal and distal cues causes disparate responses in CA1 (Shapiro et al., 1997; Fenton et al., 2000; Knierim, 2002; Lee et al., 2004). Thus, we repeated the experiment with curtains around the morph box in order to mask the distal cues such that the only remaining salient cues were the shape of the box and a constant polarizing cue card.

As in the previous experiments, the transition from the "square" to the "circular" representation was continuous and gradual (Figures 8A and 8B). Only 26 of the 100 active cells were best fit by a sigmoid function (Figure 8C). Their transition points were as distributed as in experiments with background cues visible (Figure 8D). The mean population vector correlation in CA3 decreased progressively from 0.91 ± 0.004 (shape 1 versus 2) to 0.46 ± 0.01 (shape 1 versus 7, i.e., square and circle) (Figure 8F). A similar gradual decrease in the similarity of the representations was apparent in CA1 (Figures 8E and 8F).

The introduction of curtains did not increase the number of cells that passed the threshold for remapping (spatial correlation difference of 0.5 or 0.3). With the stronger criterion, 2 of 166 CA3 cells and 9 of 88 CA1 cells showed spatial remapping. With the weaker criterion, the numbers were 6 of 166 and 15 of 88, respectively. There was no increase in the proportion of cells with sigmoid curve fits in these selected cells (strong criterion: 0 CA3 cells and 1 CA1 cell; weak criterion: 1 CA3 cell and 4 CA1 cells). Thus, taken together, these data suggest that representations changed progressively regardless of whether background cues were visible or not.

Discussion

We used a morphing procedure to determine whether hippocampal ensemble activity falls into discontinuous

states when hippocampal input changes incrementally and plasticity is allowed to take place. Before morphing, pyramidal cells exhibited several-fold changes in firing intensity, leading to a nearly complete replacement of the active subset of place cells between square and circle, especially in CA3. When the environments were morphed, the ensemble activity changed smoothly between the well-learned end representations. Changes in the environment appeared to be assimilated at the same time as the original representation gradually faded out. Our results suggest that hippocampal cell assemblies can effectively exist along continua, with novel stable or quasistable intermediate states emerging when incoming information is incongruent with previous representations.

The progressive modification of the hippocampal ensemble representation could reflect several types of network dynamics. One possibility is that place cells are controlled by selected subsets of stimuli in the environment and that the ensemble representation reflects the sum of such sensory inputs to the hippocampus at the time of recording. This would be consistent with the observation that individual stimuli or combinations of stimuli can exert strong control over the firing of place cells (Gothard et al., 1996; Hetherington and Shapiro, 1997; Shapiro et al., 1997; Hollup et al., 2001), but that possibility is, at least in its simplest form, ruled out by the strong hysteresis that accompanied the progression of firing in CA3 throughout the morph sequence. In each intermediate shape, ensemble activity in CA3 depended on the direction of morphing and the history of testing as much as the particular inputs that came into the hippocampus at the time of recording. Carry-over effects were observed also in CA1, but here, activity was transferred regardless of differences in shape. Because the overall sensory input was identical in the two morphing directions, it is unlikely that specific stimuli or stimulus combinations were alone responsible for the ensemble activity in CA3 and CA1 in each shape. This conclusion is reinforced by studies showing continued activation of place cells when animals are moved away from the firing field during intervening light sleep (Jarosiewicz and Skaggs, 2004).

The hysteresis of the hippocampal ensemble activity across successive morph shapes in CA3 thus suggests that place cells also express information stored intrinsically in the network. Although the results contradict the notion of discrete global attractors, they do not rule out local attractor states of a more continuous nature. The two-dimensional nature of the pre-established spatial representations, in fact, implies that the many locations within a single environment are represented by different subpopulations of active place cells. This multiplicity of active subpopulations may blur the distinction between representations of distinct shapes and allow for small changes to the environment to be accommodated while remaining within the current two-dimensional attractor manifold. The continuous transition of the hippocampal representation may also reflect reshaping of the attractor manifold. In theoretical studies, networks generally operate in two regimes: a learning phase, in which the memory states are established and stored, and a recall phase, in which further plasticity is turned off. In the living brain, instead, learning and recall can be intermixed and new information may be accommodated at the same time as the circuit performs pattern completion (Paulsen and Moser, 1998; Hasselmo et al., 2002; Kunec et al., 2005). New information may be added to the existing representations, as suggested by the emergence of cells with maximal activity in the intermediate shapes. The incorporation of new information could be interleaved with retrieval of stored information (Hasselmo et al., 2002; Harris et al., 2003; Kunec et al., 2005), or old and new information could be merged into an undifferentiated intermediate representation.

The induction of a discontinuous transition in CA1 in the Wills study suggests that global nonlinear dynamics can be evoked in the hippocampus under different training or testing conditions. Several factors might contribute to the different outcomes of the studies. First, our morphing procedure was incremental, whereas Wills and colleagues used a semiscrambled sequence. Pre-established representations may be more prone to plastic modification when the environment is changed gradually (M. Tsodyks, personal communication). With a scrambled sequence, network activity may fall back into preformed states instead of accommodating the changing input. The progressive transformation of network activity during the scrambling trials in the present study speaks against a critical role for test sequence, but the scrambling was performed after incremental morphing and, thus, might be biased by the possible development of new representations in the morphing phase. Moreover, when rats were placed in intermediate shapes in the Wills study, the similarity with the end shapes became gradually stronger during the course of the trial. This is opposite of what would be expected if plasticity acted to incorporate new features into a continuously interpolated representation between the two extremes.

A second difference was related to the nature of the input stimuli. In both experiments, morphing was based on a transformation of the shape of the local environment. Yet the differences in training and test procedures may have provided hidden contextual cues on top of the spatial ones. In the present study, the sequence of squares and circles was random during training, whereas the morphing was incremental. In the Wills study, the rats experienced squares and circles on alternating trials during training and continued to be tested with square-like and circle-like shapes on a partly alternating basis also during the probe test (see Supporting Figure 1E in Wills et al. [2005]). Such contextual differences may have contributed to the formation of different types of ensemble representation. The shape-independent transfer of CA1 firing patterns across successive trials (e.g., between shape 7 and 1′ in the morph sequence) supports the notion that CA1 cells are responsive to contextual information as well as geometric configuration.

A third difference concerns the protocol used to induce remapping. In the Wills study, remapping was obtained almost instantly by a procedure in which strong differences in input on the initial trials produced a clear spatial remapping that in four of six rats, was preserved during subsequent training. Remapping was observed in CA1 already on the first day with square and circle training, and the orthogonalization was apparently

almost complete (Figure 1 in Wills et al. [2005]). Similar fast-and-complete remapping has been reported with other training regimes based on color reversal or odor substitution (Bostock et al., 1991; Kentros et al., 1998; Anderson and Jeffery, 2003). These forms of remapping differ strongly from the slow and partial remapping observed in CA1 here and in previous studies with primarily geometrical reconfigurations (O'Keefe and Burgess, 1996; Lever et al., 2002; Leutgeb et al., 2004). The latter type of remapping in CA1 is accompanied by strong but exclusive rate remapping in CA3 (Leutgeb et al., 2005). These contrasting types of remapping raise the possibility that strong orthogonalization of firing locations may be prerequisite for inducing bistability in hippocampal ensemble dynamics.

The dissimilar response patterns in the Wills study and the present study also introduce the possibility that different brain circuits are involved in fast and slow remapping. If remapping depends on attractor networks in the hippocampus, these are likely to be implemented in the dentate gyrus and CA3, which possess the expansion recoding and the recurrent architecture required to perform pattern separation and completion (McNaughton and Nadel, 1989; Rolls and Treves, 1998). However, nonlinear response patterns in CA1 may also reflect output from attractor networks in the neocortex (Hafting et al., 2005), which may influence place cells through direct connections from the entorhinal cortex (Brun et al., 2002). Bistable neuronal responses have been reported during morphing of categorical representations in prefrontal cortex (Freedman et al., 2001) as well as in fusiform gyrus and anterior temporal lobe (Rotshtein et al., 2004), suggesting that inputs could be orthogonalized into perceptual categories before they reach the hippocampal formation. Different forms of remapping in the hippocampus may differ in the extent to which the effect originates in upstream neocortical areas. The apparent lack of remapping in superficial layers of the entorhinal cortex (Quirk et al., 1992; M. Fyhn, T. Hafting, A. Treves, M.-B. Moser, and E.I. Moser, 2004, Soc. Neuroci., abstract) speaks against a neocortical origin for any type of hippocampal remapping, but the conditions and the extent of remapping in this area are far from established. The persistence of hippocampus-dependent spatial memories after fast remapping in CA1 (Jeffery et al., 2003) suggests that some forms of orthogonalization in this subfield may be neocortical in origin, perhaps reflecting an early perceptual process or a resetting of the rat's path integrator in response to salient differences in cues and context.

The main finding of our study is that hippocampal neuronal ensembles can progress gradually and noncoherently between pre-established network states that differ strongly in terms of which cells are active. Transitional representations do not negate the existence of attractor networks, but they suggest that hippocampal cell assemblies can incrementally incorporate new, incongruous information into well-learned spatial representations. Several previous studies have shown that firing properties of individual place cells can be significantly modified in the absence of general remapping in the hippocampal cell population under conditions in which only a part of the recording environment is reconfigured. When a barrier is introduced into an otherwise free space (Muller and Kubie, 1987) or when a salient part of the environment is dislocated or replaced (O'Keefe and Burgess, 1996; Fenton et al., 2000; Fyhn et al., 2002; Paz-Villagran et al., 2004), proximal place fields may appear, disappear, or move without substantial changes in cells with more remote firing locations. Thus in the normal operation of the system, pattern completion and pattern separation are accompanied by refinements of existing representations. Instantaneous switches to different network states (e.g., Knierim et al., 1998) may occur with more global types of remapping that also include the spatial domain, but within a certain dynamical range, hippocampal ensembles appear to accommodate existing schemes to new information rather than fluctuating arbitrarily between pre-established states. This property is likely to be fundamentally important for the ability to encode and retrieve consecutive inputs as parts of continuous and distinguishable sequences or episodes.

Experimental Procedures

Subjects

Seven male Long Evans rats were housed and food deprived to 80%–90% of free-feeding weight as described previously (Leutgeb et al., 2005). Testing occurred in the dark phase.

Surgery and Electrode Preparation

Rats were anesthetized i.p. with Equithesin (1 ml/250 g) and implanted with a "hyperdrive" (Kopf) with fourteen independently movable tetrodes, prepared as reported previously (Merrill and Ainsworth, 1972; Leutgeb et al., 2005). These tetrodes were implanted above the right hippocampus (AP 3.8, ML 3.0, DV 1.5). Six to seven of the tetrodes were later lowered to CA3c, and 5–6 tetrodes were lowered to CA1. The two additional tetrodes were used to record a reference signal from the corpus callosum and an EEG signal from the stratum lacunosum-moleculare.

Pretraining with End Shapes

Rats were trained in different shapes of the same enclosure (the morph box) in the same room (Figure 1B). The morph box consisted of 16 gray serially connected aluminium walls (width: 19 cm and 1 cm joint) that were shaped either as a square (80 cm × 80 cm; 50 cm high) or as a 16-sided polygon (50 cm radius, 50 cm high, 16 sides separated by angles of 157.5°; for simplicity referred to as a "circle"). The morph box rested on a 1 m high table and was always located at the same place in the room. A polarizing white cue card was placed in the morph box at a constant location. The internal cue card was varied in length depending on the shape of the morph box (5 rats; square: 24 cm high × 10 cm wide; "circle:" 24 cm high × 70 cm wide); or the cue card remained constant (2 rats, 50 cm high × 20 cm wide). In the two rats with the constant cue card, distal background cues were masked by black curtains encircling the recording environment (180 cm diameter). A flowerpot pedestal, where the animal slept and rested on a towel, was placed between the test box and the experimenter, outside the curtains when these were used.

Starting 3–4 days after surgery, rats were trained to run in the morph box by throwing crumbs of chocolate biscuits into the box. Each rat always received the same food reward and was always trained with the recording cables connected. The rats were trained from the first day on to forage for up to eight 10-min trials during which the morph box was presented in either its square or circular shape in a random order. The floor of the morph box was switched between each shape and cleaned with water between each trial. The rat was allowed to rest for 5 min on the pedestal between trials, and training sessions began and ended with a 15 min sleep trial. Training continued for 16–19 days after surgery before the morph sequence was introduced. Spike activity was recorded during random sequences of square and circular shapes for 1–3 days before the box was morphed.

At the end of the day, after the animal had rested in its home cage after training, the tetrodes were moved either toward CA3c or to the central part of CA1 while screening for hippocampal spike activity. During screening and recording, the hyperdrive was connected to a multichannel, impedance matching, unity-gain headstage. The output of the headstage was conducted via a lightweight multiwire tether and commutator to a data acquisition system containing 64 digitally programmable amplifiers. Unit activity was amplified by a factor of 3000–5000 and band-pass filtered at 600 Hz to 6 kHz. Spike waveforms above a threshold of 55 μV were timestamped and digitized at 32 kHz for 1 msec. Light emitting diodes on the headstage were tracked at 50 Hz. EEG from stratum lacunosum-moleculare was recorded continuously in the 1–450 Hz band.

Tests in Morphed Environments

Rats were prepared for recording in the morph sequence when clearly distinguishable units were obtained in hippocampal subfields CA1 and CA3 and remapping was observed in area CA1 for 1–2 sampled tetrodes. On the day of recording, electrodes were not moved at all or only <20 μm. Spikes were recorded simultaneously from subfields CA3 and CA1 while the rat foraged in the morph box. The box was gradually changed from square to circular or from circular to square through five intermediate shapes. The shape of the morph box was maintained by inserting the flexible wall into a groove traced into the Plexiglas floor (one floor for each shape). This arrangement minimized the time required to change shape while maximizing the similarity of the environments in all respects but shape. We fixed the midpoint morph AB as an irregular octagon of angles alternating between 112.5° and 157.5° (Figure 1C), so that four of its angles would be equal to all shape B angles, and four would be close to the square angles of shape A. Two intermediate morphs, shapes AAAB and AAB, were defined as octagons with alternating angles of 97.5° and 172.5°, and of 105° and 165°, respectively, which we obtained from the linear equation $(90 + q \times 22.5)°$ alternated with $(180 - q \times 22.5)°$ with q = 1/3 and q = 2/3. Two additional intermediate morphs, shapes ABB and ABBB, were defined as 16 gons with repeating sets of four consecutive angles defined by the linear formula: $157.5°$, $(180 - r \times 22.5)°$, $(112.5 + r \times 45)°$, $(180 - r \times 22.5)°$, again with r = 1/3 and r = 2/3. Thus, ABB had repeating angles of 157.5°, 172.5°, 127.5°, and 172.5°, whereas ABBB had angles of 157.5°, 165°, 142.5°, and 165°. When the background cues were not masked with curtains, the white cue card also varied in width depending on the shape of the morph box, with its width increasing by 10 cm with each additional intermediate shape (i.e., square A: 24 cm high × 10 cm wide; intermediate AAAB: 24 cm high × 20 cm wide; intermediate AAB: 24 cm high × 30 cm wide, and so forth).

After reaching the opposite shape, the initial shape in the morph sequence was repeated. The morphing procedure was also performed on the subsequent day but in the reverse direction and was repeated on subsequent days in alternating directions. Finally in the same rats, the sequence of intermediate shapes was presented in a random, scrambled order (Figure 7A).

Data Analysis

Spike Sorting and Place Fields

Procedures for spike sorting, cell classification, and construction of rate maps ("place fields") were as reported previously (Leutgeb et al., 2005). Firing rate distributions and location of firing were determined for trial pairs in the morph sequence and during pretraining. Care was taken to compare trials with an equal number of intervening tests at the various stages of testing. For example, the number of intervening trials between the square and circle was five for the morph sequence. For pretraining, either the first square and the last circle or the first circle and the last square were compared. With this procedure, the mean number of intervening shapes was 4.5.

Distribution of Rate and Firing Locations

Cumulative distribution functions were used to visualize changes in the distributions of firing rate differences, spatial correlations, and population vector correlations (Figure S1). For every real number x, the cumulative distribution functions for a variable X is given by $f(X \leq x)$ in which f is the frequency of cases at or below x. The definition implies that cumulative distribution functions always increase continuously and monotonically from left (in which the cumulative frequency is 0) to right (in which the cumulative frequency is 1).

Overlap

The "overlap" between the active cell populations in two shapes was calculated by dividing, for each cell, the mean firing rate in the less-active shape by the mean firing rate in the more-active shape. Cells with average firing rates of more than 0.25 Hz in at least one shape were considered active, and the ratios for active cells were averaged for the cell population. These values were then compared to the values expected from two independent distributions, estimated by randomly pairing the firing rate of a cell in shape A to the firing rate of another cell in shape B. The permutation procedure was repeated 10,000 times. The expected overlap was then used to correct the measured overlap by dividing the expected overlap minus the measured overlap by the expected overlap minus 1. The resulting overlap scores are equal to 1 for the comparison of a cell population with itself and equal to 0 for comparisons between independent cell populations.

Spatial Correlations

Spatial firing patterns on different trials in the same box were compared with a spatial correlation procedure. Each rate map was smoothed and binned into matrices of 5 × 5 cm pixels, and the rates of firing in pixels at the same location were correlated for each cell. Pixels visited less than 150 ms in either trial and pixels not common to the square and circular maps (outside a square of 16 × 16 bins or 80 × 80 cm^2) were excluded.

Population Vector

For the entire population of cells recorded in each condition, rate vectors were constructed by arranging all place fields for all animals and sessions in an x-y-z stack, in which x and y represent the two spatial dimensions (binned into 16 × 16 pixels of 25 cm^2 each), and z represents the cell-identity index. The distribution of mean rates along the z axis for a given x-y location (bin) represents the composite population vector for that location. Only bins that were shared between all configurations (a square of 16 × 16 bins or 80 × 80 cm^2) were considered. Cells from CA3 and CA1 were analyzed separately.

Curve Fits

In-field firing rates were estimated by first calculating an average rate map for all 7 morph shapes and then finding the largest number of continuous pixels exceeding 20% of the maximum rate. The average firing within these pixels was calculated for each shape, and a cell was considered to have a field if its average in-field rate was higher than 2 Hz in at least one of the morph shapes. Three different curves were fit to the in-field firing rates of active cells: (1) a linear regression; (2) a quadratic regression with a constant, a linear, and a negative quadratic term; and (3) a sigmoid function. Fits with p values < 0.05 were considered significant, and each cell was assigned to the category with the highest explained variance. The explained variance was measured with the F values, which are, by definition, corrected for the different degrees of freedom (d.f. = 2 for linear fits, 3 for quadratic fits, and 4 for sigmoid fits).

Histology

Brains were cut coronally at 30 μm and stained with cresyl violet. Each section through the relevant part of the hippocampus was collected for analysis. All tetrodes of the 14-tetrode bundle were identified, and the tip of each electrode was found by comparison with adjacent sections. The recordings from a tetrode were included in the data analysis if its deepest position was in the CA3 or CA1 pyramidal cell layer. The electrodes had not been moved after the recordings.

Supplemental Data

The Supplemental Data for this article can be found online at http://www.neuron.org/cgi/content/full/48/2/345/DC1/.

Acknowledgments

We thank J. Lisman, R.G.M. Morris, O. Paulsen, and M.P. Witter for discussion and M. Fyhn, I.M.F. Hammer, K. Haugen, K. Jenssen, R. Skjerpeng, and H. Waade for practical or technical assistance. The work was supported by a Centre of Excellence grant from the Norwegian Research Council.

Received: June 30, 2005
Revised: August 22, 2005
Accepted: September 2, 2005
Published: October 19, 2005

References

Amit, D.J. (1989). Modeling Brain Function: The World of Attractor Networks (New York: Cambridge University Press).

Anderson, M.I., and Jeffery, K.J. (2003). Heterogeneous modulation of place cell firing by changes in context. J. Neurosci. *23*, 8827–8835.

Bartlett, M.S., and Sejnowski, T.J. (1998). Learning viewpoint-invariant face representations from visual experience in an attractor network. Network *9*, 399–417.

Battaglia, F.P., and Treves, A. (1998). Attractor neural networks storing multiple space representations: a model for hippocampal place fields. Phys. Rev. E *58*, 7738–7753.

Bostock, E., Muller, R.U., and Kubie, J.L. (1991). Experience-dependent modifications of hippocampal place cell firing. Hippocampus *1*, 193–205.

Brun, V.H., Otnass, M.K., Molden, S., Steffenach, H.A., Witter, M.P., Moser, M.-B., and Moser, E.I. (2002). Place cells and place recognition maintained by direct entorhinal-hippocampal circuitry. Science *296*, 2243–2246.

Fdez Galan, R., Sachse, S., Galizia, C.G., and Herz, A.V. (2004). Odor-driven attractor dynamics in the antennal lobe allow for simple and rapid olfactory pattern classification. Neural Comput. *16*, 999–1012.

Fenton, A.A., Csizmadia, G., and Muller, R.U. (2000). Conjoint control of hippocampal place cell firing by two visual stimuli. I. The effects of moving the stimuli on firing field positions. J. Gen. Physiol. *116*, 191–209.

Freedman, D.J., Riesenhuber, M., Poggio, T., and Miller, E.K. (2001). Categorical representation of visual stimuli in the primate prefrontal cortex. Science *291*, 312–316.

Fyhn, M., Molden, S., Hollup, S., Moser, M.B., and Moser, E.I. (2002). Hippocampal neurons responding to first-time dislocation of a target object. Neuron *35*, 555–566.

Gothard, K.M., Skaggs, W.E., Moore, K.M., and McNaughton, B.L. (1996). Binding of hippocampal CA1 neural activity to multiple reference frames in a landmark-based navigation task. J. Neurosci. *16*, 823–835.

Guzowski, J.F., Knierim, J.J., and Moser, E.I. (2004). Ensemble dynamics of hippocampal regions CA3 and CA1. Neuron *44*, 581–584.

Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., and Moser, E.I. (2005). Microstructure of a spatial map in the entorhinal cortex. Nature *436*, 801–806.

Harris, K.D., Csicsvari, J., Hirase, H., Dragoi, G., and Buzsaki, G. (2003). Organization of cell assemblies in the hippocampus. Nature *424*, 552–556.

Hasselmo, M.E., Schnell, E., and Barkai, E. (1995). Dynamics of learning and recall at excitatory recurrent synapses and cholinergic modulation in rat hippocampal region CA3. J. Neurosci. *15*, 5249–5262.

Hasselmo, M.E., Bodelon, C., and Wyble, B.P. (2002). A proposed function for hippocampal theta rhythm: separate phases of encoding and retrieval enhance reversal of prior learning. Neural Comput. *14*, 793–817.

Hetherington, P.A., and Shapiro, M.L. (1997). Hippocampal place fields are altered by the removal of single visual cues in a distance-dependent manner. Behav. Neurosci. *111*, 20–34.

Hollup, S.A., Molden, S., Donnett, J.G., Moser, M.-B., and Moser, E.I. (2001). Accumulation of hippocampal place fields at the goal location in an annular watermaze task. J. Neurosci. *21*, 1635–1644.

Hopfield, J.J. (1982). Neural networks and physical systems with emergent collective computational abilities. Proc. Natl. Acad. Sci. USA *79*, 2554–2558.

Jarosiewicz, B., and Skaggs, W.E. (2004). Hippocampal place cells are not controlled by visual input during the small irregular activity state in the rat. J. Neurosci. *24*, 5070–5077.

Jeffery, K.J., Gilbert, A., Burton, S., and Strudwick, A. (2003). Preserved performance in a hippocampal-dependent spatial task despite complete place cell remapping. Hippocampus *13*, 175–189.

Kentros, C., Hargreaves, E., Hawkins, R.D., Kandel, E.R., Shapiro, M., and Muller, R.U. (1998). Abolition of long-term stability of new hippocampal place cell maps by NMDA receptor blockade. Science *280*, 2121–2126.

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

Knierim, J.J., Kudrimoti, H.S., and McNaughton, B.L. (1998). Interactions between idiothetic cues and external landmarks in the control of place cells and head direction cells. J. Neurophysiol. *79*, 425–466.

Kunec, S., Hasselmo, M., and Kopell, N. (2005). Encoding and retrieval in the CA3 region of the hippocampus: a model of theta phase separation. J. Neurophysiol. *94*, 70–82.

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

Leutgeb, S., Leutgeb, J.K., Treves, A., Moser, M.-B., and Moser, E.I. (2004). Distinct ensemble codes in hippocampal areas CA3 and CA1. Science *305*, 1295–1298.

Leutgeb, S., Leutgeb, J.K., Barnes, C.A., Moser, E.I., McNaughton, B.L., and Moser, M.-B. (2005). Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. Science *309*, 619–623.

Lever, C., Wills, T., Cacucci, F., Burgess, N., and O'Keefe, J. (2002). Long-term plasticity in hippocampal place-cell representation of environmental geometry. Nature *416*, 90–94.

Lukashin, A.V., Amirikian, B.R., Mozhaev, V.L., Wilcox, G.L., and Georgopoulos, A.P. (1996). Modeling motor cortical operations by an attractor network of stochastic neurons. Biol. Cybern. *74*, 255–261.

Marr, D. (1971). Simple memory: a theory for archicortex. Philos. Trans. R. Soc. Lond. B Biol. Sci. *262*, 23–81.

Markus, E.J., Qin, Y.L., Leonard, B., Skaggs, W.E., McNaughton, B.L., and Barnes, C.A. (1995). Interactions between location and task affect the spatial and directional firing of hippocampal neurons. J. Neurosci. *15*, 7079–7094.

McClelland, J.L., and Goddard, N.H. (1996). Considerations arising from a complementary learning systems perspective on hippocampus and neocortex. Hippocampus *6*, 654–665.

McNaughton, B.L., and Morris, R.G.M. (1987). Hippocampal synaptic enhancement and information storage within a distributed memory system. Trends Neurosci. *10*, 408–415.

McNaughton, B.L., and Nadel, L. (1989). Hebb-Marr networks and the neurobiological representation of action in space. In Neuroscience and Connectionist Theory. M.A. Gluck, and D.E. Rumelhart, eds. (Hillsdale, NJ: Lawrence Erlbaum Associates), pp. 1–63.

Merrill, E.G., and Ainsworth, A. (1972). Glass-coated platinum-plated tungsten microelectrodes. Med. Biol. Eng. *10*, 662–672.

Muller, R.U., and Kubie, J.L. (1987). The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. J. Neurosci. *7*, 1951–1968.

Nakazawa, K., Sun, L.D., Quirk, M.C., Rondi-Reig, L., Wilson, M.A., and Tonegawa, S. (2003). Hippocampal CA3 NMDA receptors are crucial for memory acquisition of one-time experience. Neuron *38*, 305–315.

O'Keefe, J., and Conway, D.H. (1978). Hippocampal place units in the freely moving rat: why they fire where they fire. Exp. Brain Res. *31*, 573–590.

O'Keefe, J., and Burgess, N. (1996). Geometric determinants of the place fields of hippocampal neurons. Nature *381*, 425–428.

Paulsen, O., and Moser, E.I. (1998). A model of hippocampal memory encoding and retrieval: GABAergic control of synaptic plasticity. Trends Neurosci. *21*, 273–278.

Paz-Villagran, V., Save, E., and Poucet, B. (2004). Independent coding of connected environments by place cells. Eur. J. Neurosci. *20*, 1379–1390.

Quirk, G.J., Muller, R.U., and Kubie, J.L. (1990). The firing of hippocampal place cells in the dark depends on the rat's recent experience. J. Neurosci. *10*, 2008–2017.

Quirk, G.J., Muller, R.U., Kubie, J.L., and Ranck, J.B., Jr. (1992). The positional firing properties of medial entorhinal neurons: description and comparison with hippocampal place cells. J. Neurosci. *12*, 1945–1963.

Rolls, E.T., and Treves, A. (1998). Neural Networks and Brain Function (Oxford, UK: Oxford University Press).

Rotshtein, P., Henson, R.N., Treves, A., Driver, J., and Dolan, R.J. (2004). Morphing Marilyn into Maggie dissociates physical and identity face representations in the brain. Nat. Neurosci. *8*, 107–113.

Samsonovich, A., and McNaughton, B.L. (1997). Path integration and cognitive mapping in a continuous attractor neural network model. J. Neurosci. *17*, 5900–5920.

Shapiro, M.L., Tanila, H., and Eichenbaum, H. (1997). Cues that hippocampal place cells encode: dynamic and hierarchical representation of local and distal stimuli. Hippocampus *7*, 624–642.

Treves, A., and Rolls, E.T. (1991). What determines the capacity of autoassociative memories in the brain? Network *2*, 371–397.

Treves, A., and Rolls, E.T. (1992). Computational constraints suggest the need for two distinct input systems to the hippocampal CA3 network. Hippocampus *2*, 189–199.

Treves, A., and Rolls, E.T. (1994). Computational analysis of the role of the hippocampus in memory. Hippocampus *4*, 374–391.

Wills, T.J., Lever, C., Cacucci, F., Burgess, N., and O'Keefe, J. (2005). Attractor dynamics in the hippocampal representation of the local environment. Science *308*, 873–876.

Wyttenbach, R.A., May, M.L., and Hoy, R.R. (1996). Categorical perception of sound frequency by crickets. Science *273*, 1542–1544.