Neuron
Article

CA3 Retrieves Coherent Representations from Degraded Input: Direct Evidence for CA3 Pattern Completion and Dentate Gyrus Pattern Separation

Joshua P. Neunuebel[1,3,4] and James J. Knierim[1,2,3,*]
[1]Krieger Mind/Brain Institute, Johns Hopkins University, Baltimore, MD 21218, USA
[2]Solomon H. Snyder Department of Neuroscience, Johns Hopkins University School of Medicine, Baltimore, MD 21205, USA
[3]Department of Neurobiology and Anatomy, University of Texas Medical School at Houston, Houston, TX 77030, USA
[4]Janelia Farm Research Campus, Howard Hughes Medical Institute, Ashburn, VA 20147, USA
*Correspondence: jknierim@jhu.edu
http://dx.doi.org/10.1016/j.neuron.2013.11.017

SUMMARY

Theories of associative memory suggest that successful memory storage and recall depend on a balance between two complementary processes: pattern separation (to minimize interference) and pattern completion (to retrieve a memory when presented with partial or degraded input cues). Putative attractor circuitry in the hippocampal CA3 region is thought to be the final arbiter between these two processes. Here we present direct, quantitative evidence that CA3 produces an output pattern closer to the originally stored representation than its degraded input patterns from the dentate gyrus (DG). We simultaneously recorded activity from CA3 and DG of behaving rats when local and global reference frames were placed in conflict. CA3 showed a coherent population response to the conflict (pattern completion), even though its DG inputs were severely disrupted (pattern separation). The results thus confirm the hallmark predictions of a longstanding computational model of hippocampal memory processing.

INTRODUCTION

The hippocampus is crucial for spatial, contextual, and episodic memory (Eichenbaum, 2004; O'Keefe and Nadel, 1978; Squire et al., 2004), but the precise computations performed by the hippocampus in support of these functions are unknown. It is thought that the hippocampus integrates external sensory information from the lateral entorhinal cortex (LEC) with self-motion-based spatial information from the medial entorhinal cortex (MEC) to create context-specific representations necessary for the recall of individual events (Knierim et al., 2006; Manns and Eichenbaum, 2006; Suzuki et al., 1997). A longstanding computational theory suggests that, to maximize the storage of information with minimal interference, associative networks such as the hippocampus perform two competing, yet complementary, processes (Guzowski et al., 2004; Hasselmo and Wyble, 1997; McClelland and Goddard, 1996; McNaughton and Morris, 1987; McNaughton and Nadel, 1990; O'Reilly and McClelland, 1994; Rolls and Treves, 1998). Pattern separation refers to the ability of the network to reduce the overlap between similar input patterns before they are stored in order to reduce the probability of interference in memory recall. Pattern completion refers to the ability of the network to retrieve stored output patterns when presented with partial or degraded input patterns. In many models of the hippocampus, the dentate gyrus (DG) region is regarded as a preprocessing stage that performs pattern separation on entorhinal cortex inputs. In contrast, the extensive network of recurrent collaterals in CA3 may produce attractor dynamics that result in pattern completion (or generalization) when input representations are similar to stored memories (attractor basins) or pattern separation when input representations are more distinct (Guzowski et al., 2004; Rolls and Treves, 1998).

Previous studies have provided evidence consistent with the hypothesized roles of the DG and CA3 in pattern separation and pattern completion (for reviews, see Santoro, 2013; Yassa and Stark, 2011). However, a rigorous test of these functions requires measuring both the input and output representations of the brain structures, to test explicitly whether the outputs are more similar (pattern completion) or less similar (pattern separation) than the inputs. Previous investigations of CA3 and DG lacked critical information about the inputs, making it uncertain whether the putative pattern separation or completion was inherent to the region under investigation or merely a reflection of processing that already occurred upstream (e.g., Gold and Kesner, 2005; Gilbert et al., 2001; Rolls and Kesner, 2006; Kesner et al., 2000; McHugh et al., 2007; Nakazawa et al., 2002; Lee et al., 2004). This uncertainty is magnified by the paucity of published studies on the nature of DG neural representations in freely moving animals. Other studies that investigated both input and output patterns worked under experimental conditions in which CA3 reflected pattern separation, precluding a test of its hypothesized pattern completion functions (Bakker et al., 2008; Leutgeb et al., 2007).

An experimental protocol utilizing local-global reference frame conflicts has been shown to result in CA3 neural responses that resemble pattern completion. Lee et al. (2004) showed that the population of CA3 cells responded to the local-global conflict more coherently than did the population of CA1 cells. In the

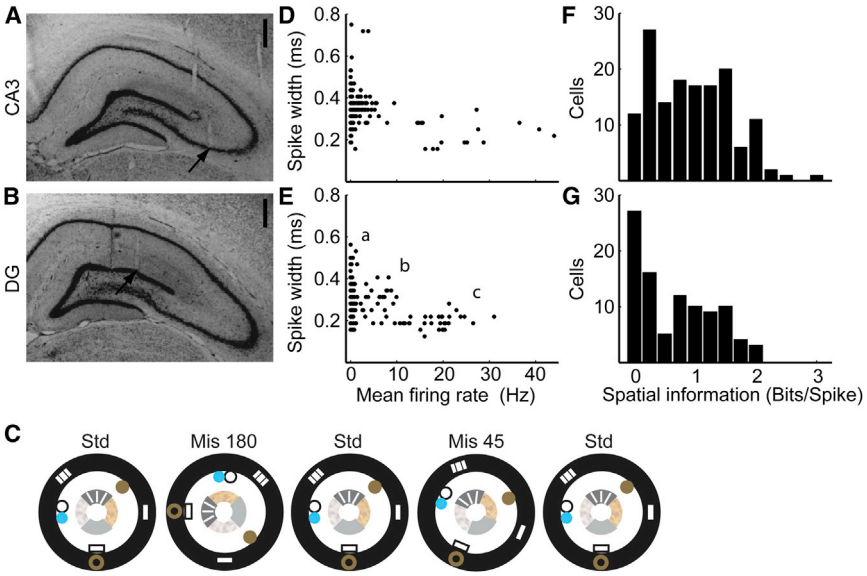

Figure 1. Basic Properties of CA3 and DG Neural Firing and Experimental Procedures

(A and B) Recording location examples show tetrodes targeting CA3 (A) and DG (B). Because the transverse axis of the hippocampus is angled relative to the midline, the DG tetrodes targeted sites medial and posterior to the regions sampled by CA3 tetrodes. Scale bar equals 500 μm, and arrows indicate the end of the tetrode tracks. Other tracks are visible that ended in adjacent sections.

(C) One day of the experimental protocol consisted of three Std sessions interleaved with two cue-mismatch sessions. The mismatch angles depicted are 180° and 45°.

(D and E) Putative cell types from CA3 (D) and DG (E) were differentiated by the mean firing rates (Hz; abscissa) and spike widths (ms; ordinate) of all well-isolated cells recorded in the first Std session of the day. For CA3 cells, two distinct groups were observed (putative principal cells with a mean firing rate <10 Hz and putative interneurons with a mean firing rate ≥10 Hz). Three groups of cells were apparent in DG: (a) <2 Hz, (b) 2–10 Hz, and (c) >10 Hz.

(F and G) The distribution of spatial information scores (Skaggs et al., 1996) from CA3 (F) was significantly higher than for DG (G) (Mann-Whitney U test, $Z = 3.1$, $p < 0.03$). See Figure S6 for the information score distribution of the DG neurons that fired <2 Hz.

present study, we recorded single-unit activity simultaneously from CA3 and DG in the same protocol to directly test whether the DG input patterns to CA3 were degraded in the cue-conflict environment (as predicted from pattern separation hypotheses) and whether the CA3 representation of the altered environment was more similar to the familiar environment, compared to its DG inputs. The results show that CA3 produced an output pattern closer to the originally stored representation than the degraded input patterns from the DG, providing direct, quantitative, neurophysiological evidence for pattern completion of severely degraded inputs in the DG-CA3 circuit.

RESULTS

Single-unit activity was recorded from the DG and CA3 of freely moving rats using multitetrode arrays (Figures 1A and 1B). The CA3 tetrodes were localized to the pyramidal cell layer, primarily in the CA3a and CA3b regions. The DG tetrodes were localized in (or just external to) the granule cell layer (n = 33) or in the hilus (n = 23); none of the tetrode tips encroached upon the CA3c layer. As discussed at length in a paper that presented data from the same animals as those reported here (Neunuebel and Knierim, 2012), it is impossible to distinguish from histology alone whether any individual, extracellularly recorded unit is a mature granule cell; an adult-born, immature granule cell; or a cell in the hilus. The cells recorded here had properties that are consistent with previous publications of DG neural activity (Gothard et al., 2001; Jung and McNaughton, 1993; Leutgeb et al., 2007; Neunuebel and Knierim, 2012). Approximately half of the cells in our sample fired in multiple, irregularly spaced subfields as rats ran in a large, open field after the main experiment each day (Figures S1 and S2 available online). This pattern of activity was ascribed to granule cells by Leutgeb et al. (2007). (Subsequent work has suggested that such patterns may be preferentially associated with newborn granule cells or hilar cells [Alme et al., 2010; Neunuebel and Knierim, 2012].) Other cells that were included in the analysis either fired at a low rate (<1 Hz) with little spatial specificity or had a single place field in the open field. These cells were considered by Neunuebel and Knierim (2012) as likely candidates for mature granule cells. Because it is unknown whether the different firing profiles corresponded to different morphological cell types, we did not segregate different subtypes of DG cells in the primary analyses (but we confirmed the main results by analyzing different subclasses in the Supplemental Information).

Rats ran clockwise (CW) around a track centered in a black-curtained, circular environment (Figure 1C). Four local cues tiled the surface of the track, and six global cues were placed on or near the curtains (Knierim, 2002). The rats experienced a standard configuration of these cue sets for an average of 16 days before the experiments started. Recording sessions consisted of three standard (Std) sessions separated by two mismatch (Mis) sessions, in which graded changes in sensory input were produced by rotating the global cues CW and the local cues counterclockwise (CCW) by the same amount, for net cue mismatches of 45°, 90°, 135°, and 180°. Over the course of 4 days, the rats experienced each mismatch amount twice. During the experiment, 399 CA3 units and 341 DG units were recorded during baseline (BL) sessions in which the rat sat quietly or slept in a towel-lined dish. These BL sessions were used to judge stability of the recordings within a day. Subsets of these cells were active during any given behavioral session and were analyzed quantitatively.

During behavior on the track, the CA3 cells could be classified into putative principal cells and putative interneurons on the basis of firing rates and spike widths (Figure 1D), with putative principal cells having lower firing rates and wider spikes and putative interneurons having higher firing rates and narrower spikes. The DG units, on the other hand, segregated into three clusters of

points (Figure 1E). Preliminary analysis of the data showed that some DG units fired at very low rates on the track, but the locations of these spikes were consistent across individual laps. The DG mossy fiber synapse onto CA3 is very large, and it is conceivable that even these low-rate cells might exert a significant influence on the CA3 network response. Therefore, to include these very low-rate cells and to exclude the high-firing-rate putative interneurons, we restricted our analyses to cells with $\geq$20 spikes in a given session and a mean firing rate <10 Hz, respectively. We applied these criteria to both CA3 and DG to remain consistent between regions and to remain consistent with the criteria of our previous report on the entorhinal inputs to the DG and CA3 (Neunuebel et al., 2013), thus allowing direct comparison to those results. Approximately 37% of CA3 (146) and 28% of DG (96) units met these inclusion criteria during the first Std session of the day (although many additional cells met the inclusion criteria in later sessions of the day; Figures S1–S3). Figures 1F and 1G show that both CA3 and DG had a similar range of spatial information scores, although the median score of the CA3 neurons was greater than the DG neurons (CA3: median 0.9, IQR 0.4–1.4; DG: median 0.6, IQR 0.1–1.2; Mann-Whitney U test, $z = 3.1$, $p < 0.03$).

Similar to previously published reports of hippocampal subregions CA3 and CA1 (Lee et al., 2004) and their entorhinal inputs (Neunuebel et al., 2013), individual cells responded differently to the double rotation manipulation. Classification of cells into different response types leads to arbitrary distinctions in many cases, and, therefore, we did not perform statistical or quantitative analyses on these categories. We present them here to provide a useful description of the types of single-unit responses that underlie the quantitative population analyses presented below. The firing fields of some cells rotated CCW with the local cues or CW with the global cues. This rotation was determined by correlating the Mis rate map with the Std 1 rate map at each of 72 rotational increments (each increment was 5°) and using the location of the peak correlation as an indicator of cue control. The peak correlation was required to exceed a threshold set at 0.6, which was a level that most reliably captured experienced observers' evaluations of the similarity between two rate maps. Figure 2 shows examples of CA3 cells that were classified as CCW (cells 1–5) and CW (cells 6b and 7). Cell 1 had a firing field near the 4 o'clock position on the track in both Std sessions (Std 1 and Std 2). In the 180° cue-Mis session, the firing field rotated CCW, which indicated local cue control. The maximum correlation occurred when the rate map was rotated 275° degrees CW (i.e., 85° CCW) and surpassed the 0.6 threshold (green line). Cells 2–5 are other examples of local-cue-controlled cells. A smaller number of CA3 place fields were controlled by the global cues, rotating their firing fields CW (cells 6b and 7). In agreement with Lee et al. (2004), the local-cue-dominated firing fields (n = 101) far outnumbered the global-cue-dominated fields (n = 40).

A number of cells (n = 60) had place fields that met inclusion criteria in both the Std and Mis sessions, but their responses were considered ambiguous (i.e., the peak rotational correlation was <0.6). Many of these cells had fields that became more diffuse during the Mis session (cells 8 and 9). For other cells, the activity criteria were reached in only one of the two sessions. Cells 10 and 11 are examples of strong fields (n = 64) that developed during the Mis session despite firing only a few spikes in the Std session. Cells 6a and 12 are examples of fields (n = 64) that were present in the Std session but were silent in the Mis session. Cell 6 is classified as disappear during the first Mis session (labeled as 6a) and as CW for the second Mis session (labeled 6b), showing that, as in prior reports (Knierim, 2002; Lee et al., 2004), the same cell could respond differently to the manipulation in different sessions. The appearance or disappearance of the fields was not an artifact of recording instability, since the tetrode cluster patterns were similar during BL sessions recorded before and after the behavioral sessions, indicating that the same cells were present throughout the experiment.

Examples of DG responses to the double rotation are shown in Figure 3. Similar to CA3, DG cells could also be classified as CCW, CW, ambiguous, appear, or disappear. Many of the DG cells fired in single or multiple locations on the track (Jung and McNaughton, 1993; Leutgeb et al., 2007; Neunuebel and Knierim, 2012) and fired consistently across the Std sessions. For example, cell 1 had two small fields (at 3 o'clock and 8 o'clock) in both Std sessions. During the 135° Mis session, one field apparently rotated to the 10 o'clock position and stretched in length while the other field rotated to 6 o'clock and the size remained similar. Because the peak correlation was just below the threshold, the cell's response was considered ambiguous, a classification consistent with the changes in field size that indicated that the response was not simply a rotation of the firing fields. Cells 2 and 3 are other examples of cells that met activity criteria in both sessions and could not be described as simple rotations (n = 49). Other cells had place fields that rotated either CW or CCW (cells 4–7), according to the >0.6 rotational correlation threshold. Unlike CA3, there was not a large difference between the number of rate maps that rotated CCW (n = 36) or CW (n = 23). Similar to CA3, some of the cells were classified as appear (n = 49; e.g., cells 8–10) or disappear (n = 29; e.g., cells 11–12b).

To analyze differences at the neural population level between CA3 and DG in response to the double rotation manipulation, we created spatial correlation matrices from the population firing rate vectors at each location on the track (Figure S4) (Gothard et al., 2001; Lee et al., 2004; Neunuebel et al., 2013; Yoganarasimha et al., 2006). The mean firing rate of every cell in the sample (normalized to its peak rate) was calculated for each 1° bin of the circular track to create 360 firing rate vectors. The firing rate vectors of a Std session (Std 1) were correlated with the firing rate vectors from either the next Mis session (Mis) or the next Std session (Std 2). The Std 1-versus-Std 2 correlation matrices for CA3 produced a band of high correlation on the main diagonal, showing that most CA3 cells fired at a similar location in both Std sessions (Figure 4, column 1). In every Mis session, CA3 maintained a band of highly correlated activity (Std 1 versus Mis; Figure 4, column 2). This band shifted downward from the main diagonal (dashed line), indicating that the CA3 representation was controlled coherently by the local cues (see below for analyses demonstrating that the high-correlation bands match precisely the angles of rotations of the local cues). However, the correlation structure degraded with increasing mismatch amounts, indicating that increasing cue mismatches caused increasing changes to the CA3 representation (Lee et al., 2004).

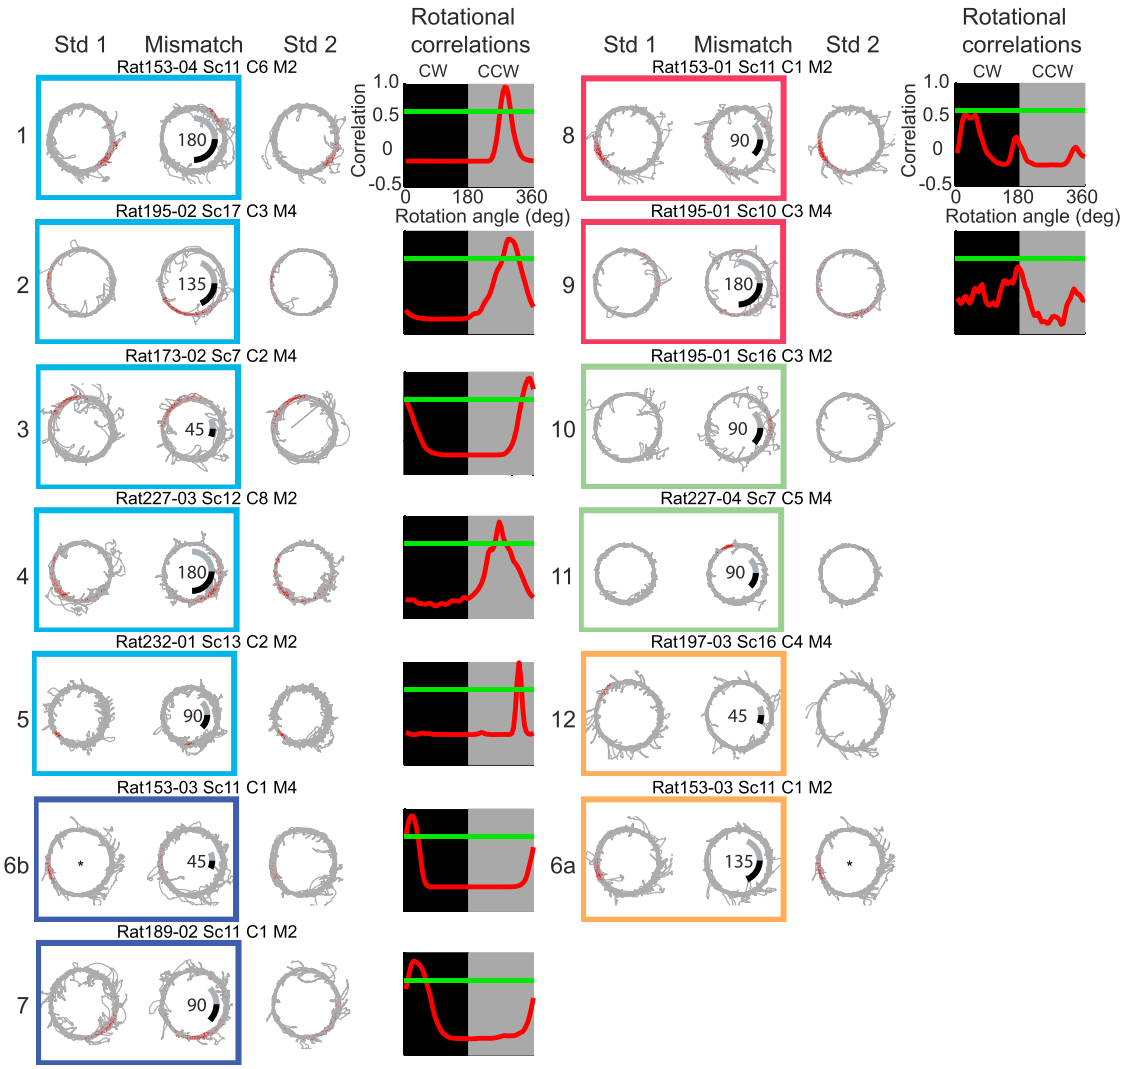

Figure 2. CA3 Cellular Responses

Example spike (red points) and trajectory (gray line) plots of CA3 cells. Values in the center of Mis sessions indicate the total mismatch angle. The gray and black lines show the amount of the local and global cue rotations, respectively. Boxes enclosing rate maps from Std1 and Mis sessions indicate cells categorized as CW (navy blue), CCW (cyan), appear (green), disappear (orange), or ambiguous (maroon). A plot of the rotation correlation analysis between the Std1 and Mis sessions (red line) is shown to the right of each set of rate maps. Peak correlations above 0.6 (green line) located in the black or gray box indicated that the fields rotated CW or CCW, respectively. Asterisks indicate that the Std sessions for cell were the same sessions.

The critical question for the present study is (1) whether the coherence in the CA3 response was a reflection of an active pattern completion (or pattern generalization or error correction) computation performed by CA3 circuitry on degraded or corrupted input patterns or (2) whether the coherence was merely a passive reflection of input patterns that were already coherent prior to CA3. To answer this question, we analyzed the input representations from the cells recorded from the DG in the same animals. In the Std 1-versus-Std 2 correlation matrix (Figure 4, column 3), the DG showed a band of high correlation at the main diagonal, consistent with a reproducible pattern between standard cue configuration sessions (although the correlation matrices showed a noisier overall distribution compared to CA3; see below). In the critical Std 1-versus-Mis matrices, the DG showed only weak evidence of a high-correlation band (Figure 4, column 4), consistent with the computational models that propose a pattern separation function for the DG. These results demonstrate that the CA3 representation remains stable in the presence of a severely degraded input from DG, implying that CA3 can retrieve a previously stored pattern based on that degraded input.

The Std 1-versus-Std 2 matrices for the DG (Figure 4, column 3) showed less coherence than their CA3 counterparts (Figure 4, column 1) (i.e., in addition to the high correlations along the main diagonal of the DG Std 1-versus-Std 2 matrices, there were additional pixels with high correlations away from the main diagonal). This finding raises the question of whether the DG supported coherent spatial representations even in the Std sessions. The

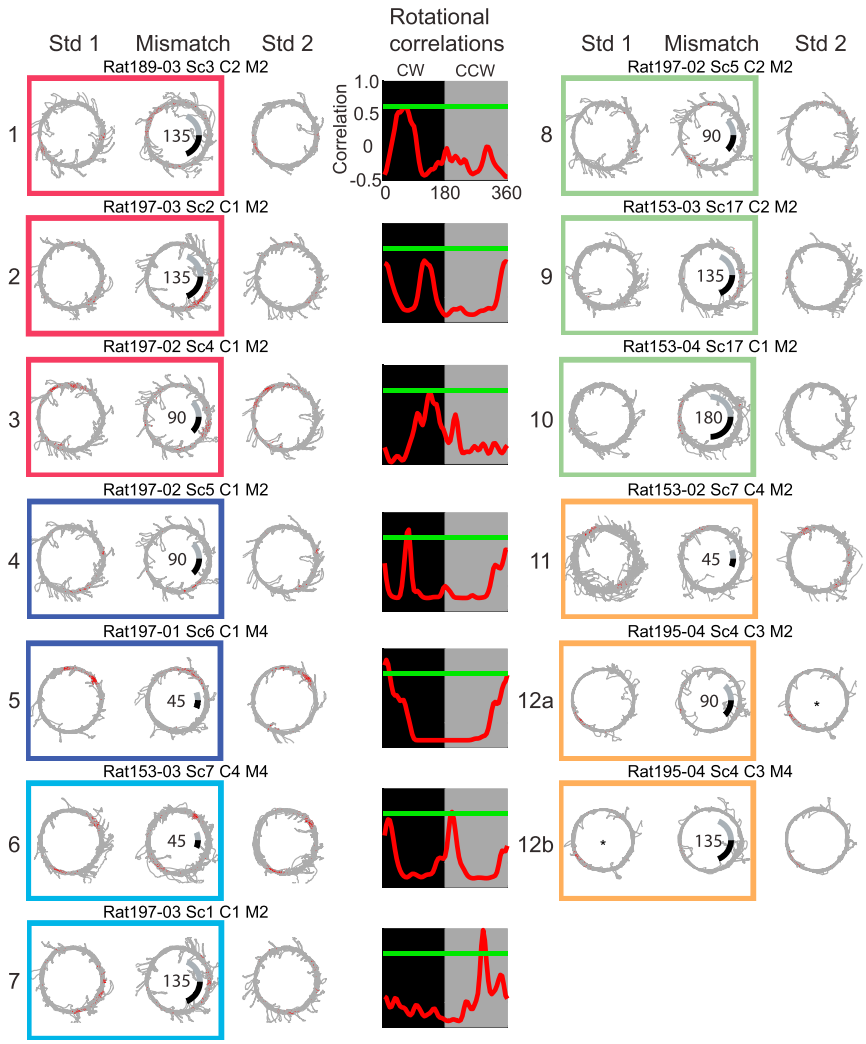

Figure 3. DG Cellular Responses

The figure format is identical to the CA3 cellular responses seen in Figure 2. See also Figures S1 and S2.

matrices used all of the putative principal cells with mean firing rates <10 Hz. As shown in Figure 1D, the DG population shows three clusters of cells in the spike-width-versus-firing-rate scatterplot: (a) a group of cells with very low rates (<2 Hz) and a wide distribution of spike widths; (b) a group of cells with moderate firing rates (2–10 Hz) and medium spike widths; and (c) a group of cells with high firing rates (>10 Hz) and narrow spike widths. The last group consisted of presumed interneurons that were dropped from the analysis. It is possible that the middle class of cells was not as spatially modulated as the low-rate cells, and inspection of the rate maps confirmed this suspicion for most (but not all) of these cells (Figures S1 and S2). We thus generated the correlation matrices based only on the low-rate cluster of cells (group a). These matrices showed a cleaner coherence band in the Std 1-versus-Std 2 matrices compared to the larger data set (Figure 4, column 5), demonstrating that the low-rate cells had highly reproducible firing patterns in the Std sessions. Moreover, the Std 1-versus-Mis matrices still showed a strong lack of spatial coherence, especially in the Mis sessions >45° (Figure 4, column 6). When we restricted analysis further to include only the cells that meet minimum spatial information criteria, the correlation matrix patterns were preserved (Figure S5). Because it is not known which of these cell classes are the principal neurons that project to CA3 (and therefore constitute the input pattern to that region), and to be conservative, we continued to analyze the combined data from all cells with firing rates <10 Hz (i.e., groups a and b in Figure 1E; columns 3 and 4 of Figure 4).

To statistically analyze the population responses, the mean correlations of pixels in each of the 360 diagonals of the correlation matrix were calculated and plotted in polar coordinates (Figure 5). The stability of the CA3 representations was evidenced by the peak correlation that occurred near 0° for all Std 1-versus-Std 2 correlation matrices (gray polar plots). For all Mis sessions (red plots), the peak correlations shifted CCW by approximately the same amount that the local cues were rotated. For DG, the peak correlations for the Std 1-versus-Std 2 comparison occurred near 0° (gray plots), indicating that the DG spatial representations were stable (see Figure S6 for the very low-rate [<2 Hz] DG cells included in Figure 4, columns 5 and 6, and Figure S5 for the polar plots of the cells that met the minimum spatial information criteria). The correlations between the Std and 45° Mis sessions maintained a peak centered near 0° (blue plots), but the distributions became much more circular with the larger mismatch angles, indicating a lack of coherence. The maximum correlations for DG corresponded to either CW (45° and 90° mismatch) or CCW (135° and 180° mismatch) rotations, although a small peak was also evident at the local-cue-predicted angle for the 90° mismatch.

The degree of coherence between the representations of the Std session and the Mis session is reflected in the sharpness and unimodality of the polar plots. To compare the different brain regions, we calculated the mean vector for each Std-Mis polar plot (i.e., the mean of all vectors originating at the origin of the plot and ending at each data point on the plot) and used bootstrapping statistical methods (Efron and Tibshirani, 1991) to compare the length of the mean vector between the hippocampal subregions (see Experimental Procedures). Collapsed across the four mismatch angles, CA3 mean vectors were significantly greater than DG mean vectors ($p < 0.001$) (Figure 5B). When comparing individual mismatch angles, the CA3 mean

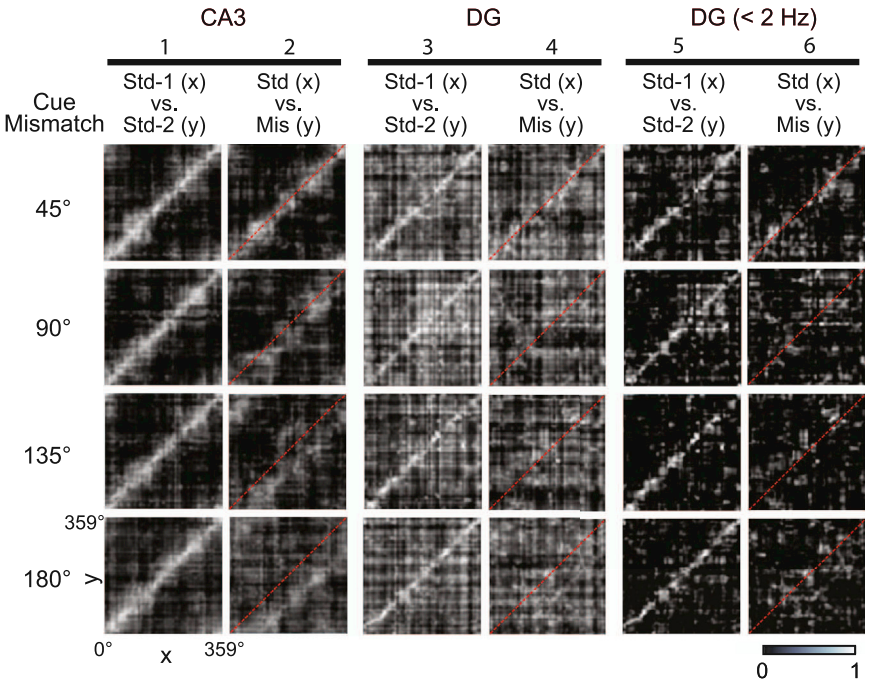

Figure 4. Population Responses to Cue-Mismatch Manipulations

Spatial correlation matrices were produced by correlating the normalized firing rate vectors for a Std session with those of the following Mis or Std session (Figure S4). CA3 representations maintained coherence in all Mis sessions (column 2), indicated by the bands of high correlation (white) shifting below the identity line (dashed line), despite the decorrelated DG representations found in the input (columns 4 and 6). See also Figures S4, S5, and S6.

vectors were significantly greater than DG for the 45° ($p < 0.002$), 90° ($p < 0.001$), and 135° ($p < 0.001$) mismatch angles. The mean vectors were not significantly different for the 180° mismatch ($p = 0.166$), although inspection of the polar plots shows a much narrower tuning curve for CA3 than DG (see also Figures S5 and S6).

To determine if the patterns observed in the population were mirrored in the firing properties of single cells, we examined the coherence of control that the local and global cues had on individual units (Figure 6). Cue control was determined by conducting a rotational analysis (see Experimental Procedures) on the subset of cells that met activity criteria in consecutive Std and Mis sessions. The mean vector length for CA3 cells was significant for all mismatch angles (Rayleigh test, $p < 0.001$), indicating significant clustering of the cell responses. Furthermore, the direction of the mean vector in all Mis sessions corresponded to a local cue rotation. DG responses were more variable, and only the mean vectors of the smallest two mismatch angles were significant (Rayleigh test; 45° and 90°; $p < 0.04$). The angle of the mean vector corresponded to the direction of a local cue rotation for three of the four mismatch angles (45°, 135°, and 180°).

The preceding analyses pooled data that were recorded across many sessions and rats. Because this pooling may have combined heterogeneous patterns of responses across data sets, it is important to know whether these results hold up at the level of individual data sets with simultaneously recorded neurons (Lee et al., 2004). Due to the sparse firing in DG, there were limited data sets with large ensembles of active cells in both CA3 and DG. Thus, we examined data sets with ≥ 2 simultaneously recorded cells from CA3 or ≥ 2 simultaneously recorded cells from DG, in which all cells met activity criteria in both the Std and the Mis sessions (examples of simultaneous recordings for each area and mismatch angle are shown in Figure 7A). On average, there were more cells recorded simultaneously in CA3 than in DG (Figure 7B). Because the size of the mean vector is dependent on the number of cells in the ensemble, we were unable to perform a simple comparison of the average lengths of the mean vectors between the regions. To illustrate the problem, Figure 7C shows the results of a simulation in which the mean vectors of ensembles of increasing size were calculated for angles randomly distributed around a circle. As the ensemble size increased, the magnitude of the mean vector decreased, even though the data were randomly distributed in all cases. Thus, a positive result may occur by chance due to unequal average ensemble sizes. To circumvent this problem, rather than comparing the magnitude of the mean vector itself, we compared the proportion of data sets in each region that were significantly clustered (Figure 7D). Significant clustering of a data set was determined by comparing the mean vector length of a sample with the randomized data produced for a sample of identical size. If the vector length from the data was >95% of the vectors from the randomized data (Figure 7C), the vector length was considered significant. CA3 had a larger proportion of significantly clustered data sets than its DG input (CA3 [11/40], DG [1/29]; $\chi^2(1) = 6.77$, $p < 0.01$). When we restricted the analysis to the subset of data in which ≥ 2 cells were recorded simultaneously from CA3 together with ≥ 2 cells recorded simultaneously from DG, the proportions were almost identical, although the smaller number of data sets reduced the statistical significance to a trend (CA3[5/19], DG[1/23]; $\chi^2(1) = 3.05$, $p = 0.08$). These results provide strong evidence that simultaneously recorded CA3 cells respond more cohesively than the input from DG.

DISCUSSION

One of the key goals of systems and cognitive neuroscience is to understand the transformations of neural representations and the rules governing these transformations, as information is processed from one stage of a circuit to another stage. This goal is aided by theoretical and computational studies that make explicit predictions about the different processing stages. One of the best known and oldest computational theories of neural

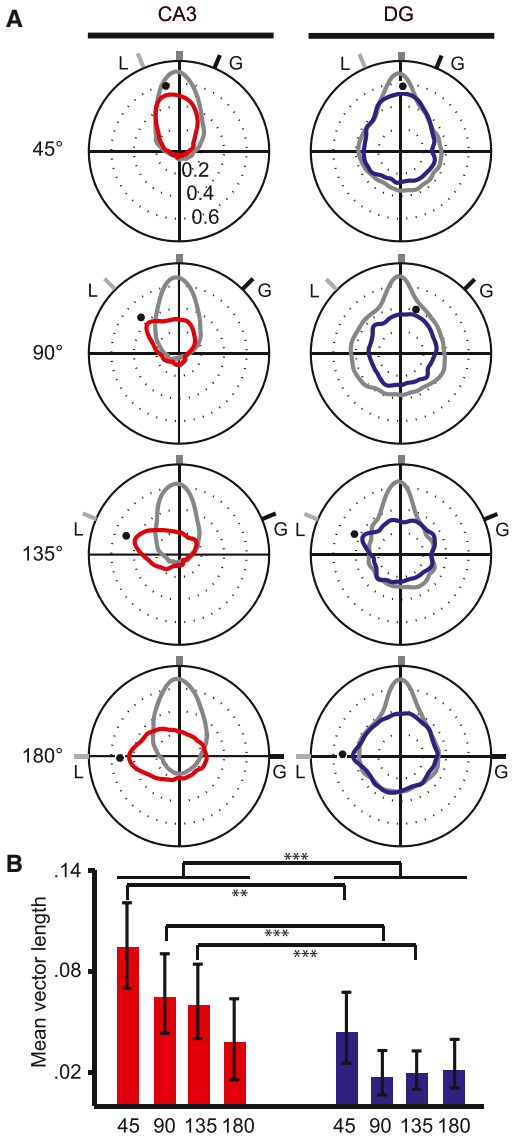

Figure 5. Quantifying Input and Output Representations

(A) Polar plots were created from the spatial correlation matrices to represent the population activity between Std 1-versus-Std 2 (gray) and Std-versus-Mis (color) sessions. Each polar plot was created by calculating the average correlation along each diagonal of the corresponding correlation matrix to convert the 2D matrix into a 1D polar plot. The gray and black tick marks labeled "L" and "G" indicate the rotation angles of the local and global cue sets, respectively. The black dots indicate the angle at which the population correlations for the Std-Mis comparisons were maximum. For CA3, the maximum correlations closely followed the rotation of the local cues.

(B) Mean vectors were calculated to quantify the coherence of the representations between sessions. Error bars represent the 95% confidence interval calculated with a bootstrap analysis. Collapsed across mismatch angles, CA3 had significantly larger mean vectors than its DG input. With respect to individual mismatch angles, the CA3 mean vectors were significantly larger than the DG mean vectors for the 45°, 90°, and 135° angles. ***$p < 0.001$; **$p < 0.002$.

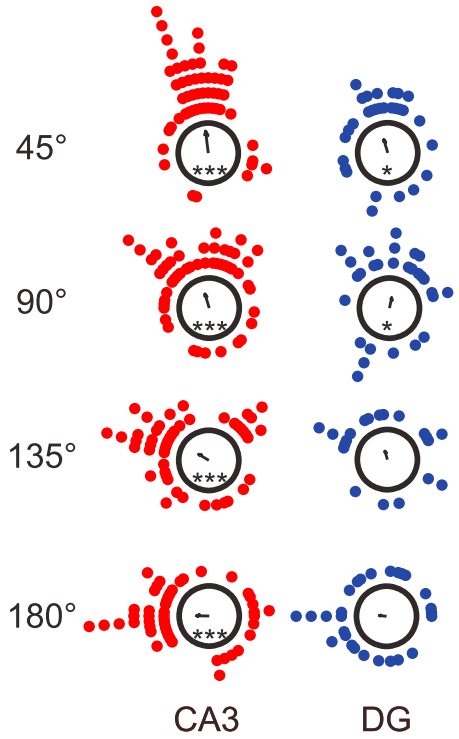

Figure 6. Analysis of Individual Cell Rotation Amounts

Each dot indicates the amount that the cell's spatial firing pattern rotated between the Std and Mis sessions (CA3, left; DG, right). The arrows at the centers of the polar plots denote the mean vector. The mean vector length for CA3 cells was significant for all mismatch angles, whereas the mean vector length for the DG was variable (Rayleigh test; CA3, all angles, $p < 0.001$; DG, 45° and 90°, $p < 0.04$). CA3 followed the local cues for all mismatch angles and DG followed the local cues for three out of the four mismatch angles (45°, 135°, and 180°). ***$p < 0.001$; *$p < 0.04$.

information processing postulates that the DG region of the hippocampus performs a pattern separation process on its input from the EC, whereas attractor circuitry in the downstream CA3 region can perform pattern separation or pattern completion based on the relative strengths of the embedded attractors and the exact nature of the external inputs from the EC and DG regions (Marr, 1971; McNaughton and Morris, 1987; McNaughton and Nadel, 1990; Rolls and Treves, 1998). Numerous studies have shown evidence consistent with a pattern separation function in both DG and CA3 and a pattern completion function in CA3 (e.g., Gilbert et al., 2001; Rolls and Kesner, 2006; Kesner et al., 2000; McHugh et al., 2007; Nakazawa et al., 2002; Lee et al., 2004; for review, see Yassa and Stark, 2011). However, most of the evidence in these studies was indirect, as the studies typically measured the output of the region without measuring the input representations, or they used behavioral tasks to try to assess the underlying neural representations. Because pattern separation and pattern completion, by definition, require knowledge of the transformation of an input representation to an output representation (Guzowski et al., 2004; McClelland and Goddard, 1996; O'Reilly and McClelland, 1994; Santoro, 2013), it is impossible to know whether the output reflects an operation intrinsic to a particular brain region without knowledge of the properties of the inputs to that region. In conjunction with previously published data (Neunuebel et al., 2013) on the EC representations under the

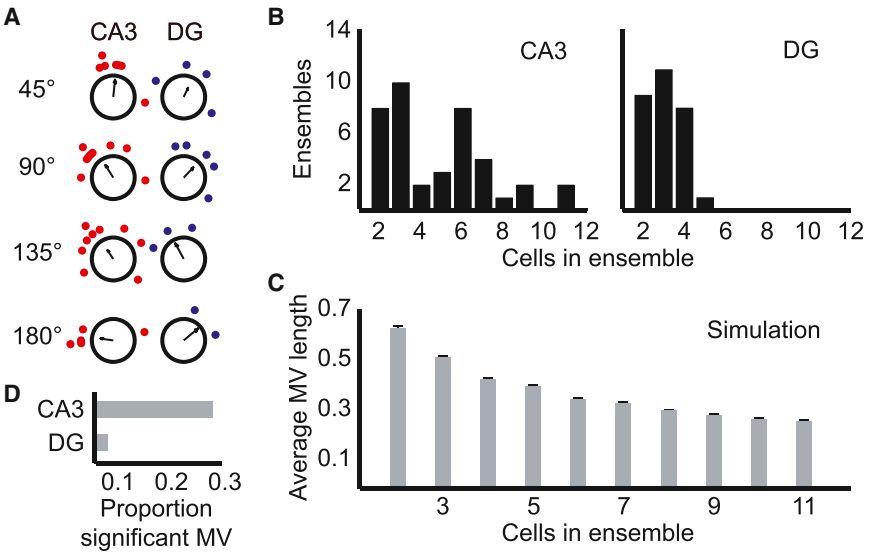

Figure 7. Ensemble Coherence

(A) Examples of simultaneous recordings from each region. Data within a circular plot were recorded simultaneously, but different plots come from different data sets.

(B) Histograms show the number of cells in each simultaneously recorded data set for the four regions. On average, CA3 had more cells per data set than DG.

(C) Simulations showed that small sample sizes could artificially increase the size of the mean vector. Data points (i.e., rotation angles) were randomly selected with replacement from a uniform distribution of orientations (1–360°) to calculate the expected value of the mean vector based on the samples coming from a random distribution. Simulations were run 1,000 times for each of 10 sample sizes (n = 2–11 samples), and the average length of all 1,000 mean vectors (average MV length) was plotted as a function of sample size. The mean vector length was largest for ensembles with two cells and decreased nonlinearly as the sample size increased. To account for the effect of sample size on the real data, the mean vector length of an ensemble was considered significant at $\alpha = 0.05$ when it was greater than 950 of the 1,000 mean vector lengths from the simulated data run with an equal number of cells in the sample. Error bars show SEM.

(D) The proportion of significant mean vector lengths for each region. CA3 showed more significant clustering than its DG input (CA3 [11/40], DG [1/29]; χ^2 (1) = 6.77, $p < 0.01$).

same manipulation, the present study provides direct evidence that the DG performs a pattern separation operation on its EC inputs and that CA3 performs a pattern completion operation on its DG and EC inputs.

Neunuebel et al. (2013) demonstrated that cells from the superficial layers of MEC (the layers that project to the hippocampus) were predominantly controlled by the global cues. Cells from the superficial layers of LEC, in contrast, showed a weak spatial representation, as expected, but there was a detectable signal at the population level that was controlled by the local cues. The LEC and the MEC are the primary inputs to the DG; therefore, the strong loss of coherence of the DG cells in the current experiment, given the coherent response of the MEC input, is a strong indication of a pattern separation function of the DG. In contrast, the data from CA3 provide the strongest neurophysiological evidence to date of a pattern completion (error correction) function in this region. The CA3 response was controlled more strongly by the local cues on the track, in agreement with a prior study using the same protocol (Lee et al., 2004). Because the MEC cells were controlled predominately by the global cues, consistent with their relationship to the global-cue-dominated head direction cell system (Hargreaves et al., 2007; Sargolini et al., 2006; Yoganarasimha et al., 2006; Zugaro et al., 2001), the MEC input could not have simply driven the local-cue-dominated CA3 response (although a weak subset of MEC cells that were local cue driven might have contributed). The LEC response, although controlled by local cues, was spatially very weak, and thus seemingly incapable of solely driving the strong spatial response patterns of CA3. The DG response was inconsistent and clearly less correlated than the CA3 response. Thus, these data provide conclusive evidence that the CA3 representations of the standard and altered environments were more correlated with each other than any of the input representations were correlated with each other, fulfilling the classic, computational definition of pattern completion.

This work bears some resemblance to the study of Gothard et al. (2001), who recorded DG and CA1 on a track that was parametrically changed in length by sliding a start box along the track on each trial. As the rat ran on the track, both the DG and CA1 representations switched abruptly from a reference frame defined by the start box to a global reference frame defined by the room. There was no evidence, however, of pattern separation occurring in that experiment. The experiment that most closely resembles the current work is a study by Leutgeb et al. (2007), who recorded from the DG and CA3 regions during manipulations in which they gradually morphed the geometry of a recording enclosure from a square to a circle. Both the Leutgeb study and the present study investigated how DG and CA3 representations recorded from the same animals changed as the result of parametric changes to the environment. Leutgeb et al. (2007) showed that DG neurons gradually changed their firing fields in response to the increasing changes in the geometry. Importantly, the authors also recorded from three units in the molecular layer of the DG that were grid cells; because the molecular layer is the site of synaptic connection between MEC and DG, the authors inferred that these units were probably axons of MEC grid cells. The firing fields of these grid cells did not change appreciably in response to the morphing manipulation. Thus, these data showed clear evidence that the DG representation changed more than its putative MEC input. It is not known how this manipulation affected the LEC inputs. Nonetheless, as it is unlikely that the LEC provides a strong spatial signal to the DG, these data show pattern separation in DG (although not of the classic "expansion recoding" type; Marr, 1969; McNaughton and Nadel, 1990).

Interestingly, the pattern separation in the Leutgeb et al. (2007) study appeared fairly linear, as the population vector changed gradually with increasing amounts of change to the environment. This result is in sharp contrast to the nonlinear effects in the present data, in which the DG population remained coherent in the 45° Mis session and then abruptly became highly noncoherent in the larger Mis sessions (especially evident in Figures S5 and S6). In the Leutgeb study, the CA3 population responses also changed gradually with the increasing morphing of the geometry, suggesting that this apparent pattern separation in CA3 was primarily driven by the upstream pattern separation performed by the DG, and not a reflection of any active computational processing of CA3 itself. Although Leutgeb et al. (2007) suggested tentatively that the slightly greater correlation between CA3 representations of small changes to the environment geometry, compared to DG representations, may reflect a pattern completion process, it is just as likely that these greater correlations were the result of CA3 receiving highly correlated patterns of activity directly from the MEC, rather than pattern completing the slightly altered DG representations. The present study, on the other hand, in conjunction with the EC data from Neunuebel et al. (2013), provides clear and convincing evidence that CA3 does indeed perform a pattern completion function (as defined computationally) on degraded input from DG and EC. Thus, the present study is a critical complement to the Leutgeb study, as well as other studies from those investigators (Leutgeb et al., 2004, 2005). Under certain conditions, CA3 representations are completely orthogonal in different environments, which may reflect a pattern separation function imposed by the DG as well as attractor dynamics in CA3. However, under other conditions, CA3 performs the long-hypothesized pattern completion and error correction functions long attributed to its recurrent collateral circuitry, conclusively demonstrated in the present study.

An important question arising from these results is why the pattern-completed CA3 representation follows the weak local cue signal provided by the LEC input, rather than the strong global cue signal provided by the MEC input. When grid cells were discovered in MEC (Hafting et al., 2005) and poor spatial selectivity was shown in LEC (Hargreaves et al., 2005), it was assumed by many investigators that grid cells were the primary spatial drive onto place cells (e.g., Hafting et al., 2005; McNaughton et al., 2006; Monaco and Abbott, 2011; Savelli and Knierim, 2010; Solstad et al., 2006; but see Burgess et al., 2007; Kropff and Treves, 2008; O'Keefe and Burgess, 2005). More recent data have shown that this simple model is at best incomplete, as (1) major disruption of grid cells can leave a significant amount of spatial tuning intact in place cells (Brandon et al., 2011); (2) during development, adult-like place field firing patterns appear before well-formed grid-cell firing patterns (Langston et al., 2010; Wills et al., 2010); (3) cue-card manipulations in a cylindrical environment can cause discordant responses of grid cells and place cells (Song et al., 2012, Soc. Neurosci. conference); and (4) inactivation of hippocampal place cells can cause the loss of gridness in MEC cells (Bonnevie et al., 2013). Neunuebel et al. (2013) proposed one model of how the LEC inputs may cause the CA3 attractor to form at the local-cue-predicted location based on a speculation that the rats pay attention first to the local cues when they are placed on the track. Another possible explanation is related to the different subtypes of cells in the DG. Immature, adult-born granule cells are hyperexcitable and hyperplastic (Ge et al., 2007), and these cells may dominate the responses of the DG (Alme et al., 2010; Neunuebel and Knierim, 2012). A recent study has shown that immature granule cells receive input preferentially from the LEC than from the MEC (Vivar et al., 2012). If these cells are the dominant drive onto the CA3 population, at least when the rat initially enters an environment, then they may override the input from the MEC cells. Simulations show that only a small bias input is required to cause an attractor bump to form at a particular location (Zhang, 1996), and it is possible that the bias caused by LEC-driven, highly active immature neurons may drive the CA3 response by seeding the recurrent collateral circuitry of CA3 to form the CA3 activity bump at locations corresponding to the local cues. Different physiologically defined DG cell types have different spatial firing profiles (Neunuebel and Knierim, 2012), but it is not known how these profiles map onto specific, morphologically defined cell types (e.g., developmentally born granule cells, adult-born mature granule cells, adult-born immature granule cells, hilar cells, interneurons, etc.). Understanding this mapping will be necessary to further understand the computations of the DG and how they influence the downstream, putative attractor circuitry in CA3.

EXPERIMENTAL PROCEDURES

Subjects and Surgery

Seven male, Long-Evans rats (Charles River Laboratories) were individually housed with ad libitum access to food and water during a 12 hr light/dark circadian cycle (lights off at noon). When rats were ~5–6 months old and had been habituated for ~14 days, a custom-built recording drive that contained 18 independently moveable tetrodes and 2 references was surgically implanted over the right hemisphere. The drives were positioned such that the most anterior-lateral tetrodes (n = 5) targeted CA3a and the most posterior-medial tetrodes (n = 13) targeted the DG. To optimize drive placement, recordings were performed during surgery to find the lateral edge of CA3, which served as a landmark for the mediolateral placement of the drives; the most lateral tetrode ranged from 3.2–4.9 mm lateral to bregma and 3.2–4.4 mm posterior to bregma. The Institutional Animal Care and Use Committees at John Hopkins University and the University of Texas Health Science Center at Houston approved surgical protocols, which were performed under aseptic conditions and complied with standards from the National Institutes of Health.

Training and Recording

Prior to surgery, rats were familiarized daily to human contact and sleeping in a small dish (~25.4 cm) located on a pedestal (each lasting 30 min/day over a 2 week period). Rats recovered from the surgical procedure for 5–7 days, and then their body weight was reduced to 80%–90% of the free-feeding weight. After a daily session of advancing tetrodes, rats were trained in a cue-controlled environment to run laps around a circular track (outside and inside diameters of the track were 76 cm and 56 cm, respectively) for an average of 16 days. The track, which was centered in a black-curtained enclosure with six salient cues located at the periphery, was divided into four 90° segments that were textured with different materials (Knierim, 2002). During the initial training sessions, chocolate sprinkles were dispersed around the track and rats gradually learned to continuously navigate CW for the reward. A cardboard panel was placed in front of any rats attempting to move CCW. As behavior progressively improved, the reward was eventually reduced to one to two random locations per lap.

The double rotation experiments were conducted for 4 days. Two BL sessions (BL1 lasting 1 hr and BL2 lasting 30 min), which were separated by

2 hr as the rat was returned to its home cage, were recorded prior to the start of the experiment. The BL sessions consisted of periods when the rat slept or was resting quietly in its holding dish. During behavior, rats ran five track sessions. Track sessions consisted of three Std sessions (Std: local and global cue relationship remained constant) interleaved with two Mis sessions (Mis: local and global cues were rotated by equal increments, but in opposite directions, producing mismatch angles of 45°, 90°, 135°, or 180°). For example, a 180° mismatch represents a 90° CCW local cue rotation plus a 90° CW global cue rotation. Mismatch angles were chosen in pseudorandom order such that each angle was chosen once during the first 2 days of recording and once again during the second 2 days. After the fifth session of the day, the cells were recorded as the rats foraged in an open field (135 × 135 cm) surrounded by white wooden walls (30 cm high) in the room that housed the recording equipment and where the BL sessions were recorded. All experiments concluded with a 30 min BL session.

Electrophysiological Recordings

A Cheetah Data Acquisition System (Neuralynx, Bozeman) concurrently obtained up to 72 channels (18 tetrodes) of single-unit data and 21 channels of local field potential (LFP) activity. Neural signals were detected simultaneously on four fine microwire electrodes (gold-plated nichrome [12 μm] or unplated platinum-iridium [17 μm]) that were wound together to form a tetrode. The signals were amplified 1,000–5,000 times and filtered between 0.6 and6 KHz (for units) or 1 and 300 Hz (for LFP). The spike waveforms above a threshold of 30–70 μV were sampled for 1 ms at 32 kHz, whereas LFPs were continuously sampled at 1 kHz. The rat's position was tracked with an overhead camera recording a circular array of light emitting diodes (red and blue) positioned over the head of the rat and a 13 cm extension behind the head with additional diodes (green) at 30 Hz.

Tetrodes were independently advanced by small increments every day for approximately 3 weeks. After entering the CA1 layer, tetrodes were advanced at ~40–148 μm (the larger movements occurred after leaving CA1) each day for an additional 300 μm. For tetrodes targeting DG, advancement was significantly reduced to 10–20 μm per day once gamma activity and dentate spikes in the LFP were detected (Bragin et al., 1995a, 1995b). A tetrode was no longer advanced after it detected well-isolated units that fired during behavior. Tetrodes that did not show active cells during behavior (even though they may have had active cells during the quiet BL sessions) were advanced by 10 μm. This procedure continued until at least five putative DG cells that fired during behavior were present on any combination of tetrodes. For tetrodes targeting CA3, tetrodes were advanced daily by ~50 μm in an attempt to enter the CA3 layer at the same time that DG units were detecting cells. No attempt was made to track cells through the experiment; therefore, some of the same units may have been recorded over multiple days. However, our primary analyses of individual ensembles do not depend on the number of cells; moreover, because the data are analyzed separately for different mismatch amounts, any unit is counted at most twice.

Unit Isolation

Multiple waveform characteristics (i.e., spike amplitude peak, area under the waveform, and valley depth) recorded simultaneously on the four wires of a tetrode were used to isolate single units offline with a custom, interactive software program. A cell's isolation quality was rated on a subjective scale from 1 (very good) to 5 (poor), depending on the distance each cluster was separated from other clusters and from background noise. Cluster isolation was judged prior to examining any of the behavioral firing correlates of the cells. All cells rated as fair or better (categories 1, 2, and 3) were potentially included in all analyses.

Data Analysis

Analyses were performed on data restricted to times when the animal's head was within the boundaries of the track and traveling with a velocity greater than 1 cm/s. The circular, 2D data for each cell were transformed into a 1D linear representation by converting the rat's Cartesian position into units of degrees on the track, and the mean firing rate for every degree of the track was calculated. A Gaussian smoothing algorithm ($\sigma = 5.34°$) was applied to the linearized firing rate maps.

Population correlation matrices were created by forming normalized firing rate vectors for the sample of cells at each 1° bin of the track and correlating these vectors with the vectors for every location in a comparison session (Figure S4). A band of high correlation along one of the diagonals of this matrix indicates a high degree of coherence of the representation between the two sessions. The correlation matrix was reduced to 1D polar plots by averaging the correlation values along each of the 360 diagonals of the matrix. Bootstrap procedures were used to estimate the 95% confidence intervals of the mean vectors of the polar correlation plots (Efron and Tibshirani, 1991). For each region and mismatch angle, the sample of cells was randomly resampled with replacement to generate a new sample of the same number of data points as the original. These data points were used to calculate 2D correlation matrices and polar plots, and the mean vector was calculated. The procedure was repeated 1,000 times, and the 2.5th and 97.5th percentile values were taken as the limits of the 95% confidence interval. To test for statistical differences between CA3 and DG, we first looked at overall differences collapsed across mismatch angles based on the mean vectors from the bootstrapped samples as follows:

$$MV_{CA3\text{-}DG} = \sum_{i} \left(MV(i)_{CA3} - MV(i)_{DG}\right),$$

where MV is the mean vector, and i is a member of the set [45°,90°, 135°, and 180°]. The $MV_{CA3\text{-}DG}$ was calculated each time for 1,000 random resamplings with replacement of the original sample, and a p value was assigned as the number of times that the $\Delta MV_{CA3\text{-}DG}$ was <0. For example, if $MV_{CA3\text{-}DG} > 0$ in all 1,000 bootstraps, the p value associated with this was $p \leq 0.001$. To look at differences between brain regions for individual mismatch angles, we created bootstrapped distributions based on the difference in the MV for each particular mismatch angle and pair of brain regions.

The rotation angle and direction that each cell's rate map rotated between consecutive Std and Mis sessions was determined for every cell that met the inclusion criteria in both sessions. The linearized rate map in the Std session was correlated with the linearized rate map for the Mis session. The Mis session rate map was then shifted in increments of 5° and correlated with the Std session rate map at each increment. A cell's rotation angle was assigned to the shift producing the maximum correlation.

Statistical test were calculated in Excel (Microsoft Corp.), Matlab (MathWorks), or Statistica (StatSoft, Inc.). Functions from the Matlab circular statistics toolbox were used to determine circular statistics. Every statistical analysis was two-tailed and considered significant at $p < 0.05$, unless a different significance level was indicated.

Histological Procedures

For a subset of tetrodes, 10 μA of positive current was passed for 10 s to generate lesions used to aid in identifying the location of the tetrodes during histological reconstruction. Lesions were made 0–10 days after finishing the double rotation experiments. The day after lesioning, rats were euthanized with formalin perfused through the heart. Two rats (227 and 232) were euthanized immediately after making the lesions. Brains were sliced (40 μm) in the coronal plane with a freezing microtome, mounted on microscope slides, and stained with cresyl violet. A Moticam 2000 camera (Motic Instruments Inc., Richmond) or IC Capture DFK 41BU02 camera (The Imaging Source, Charlotte) attached to a Motic SMZ-168 stereoscope was used to image the brain slices. All tetrode tracks were identified, and the lowest point of the track was used to determine the recording location. Recording tips that were located in or near the CA3a and CA3b pyramidal layers were assigned to CA3. Consistent with published examples from other groups (Jung and McNaughton, 1993; Leutgeb et al., 2007), tetrode tracks that terminated in the granule cells layer or at the interface between the granule cell layer and the polymorphic layer were considered (for descriptive purposes) to be located in the granule cell layer. Tetrode tracks that terminated below the granule layer but above the CA3 pyramidal layer were considered to be located in the hilus. As described in the Results, the final histological location of the recording tips is not necessarily a valid indication of the identity of the recorded cells. However, all recordings assigned to the DG came from tetrodes that showed no indication of having encroached upon the CA3 pyramidal layer.

SUPPLEMENTAL INFORMATION

Supplemental Information includes six figures and can be found with this article online at http://dx.doi.org/10.1016/j.neuron.2013.11.017.

ACKNOWLEDGMENTS

This work was supported by National Institutes of Health Grants R01 NS39456 and T32 NS07467. We thank M.R. Neunuebel, M. Hussain Shuler, and M. Yassa for comments on an earlier version of this manuscript and F. Savelli, S.S. Deshmukh, S. Mihalas, and C.E. Connor for helpful discussions.

Accepted: November 5, 2013
Published: January 22, 2014

REFERENCES

Alme, C.B., Buzzetti, R.A., Marrone, D.F., Leutgeb, J.K., Chawla, M.K., Schaner, M.J., Bohanick, J.D., Khoboko, T., Leutgeb, S., Moser, E.I., et al. (2010). Hippocampal granule cells opt for early retirement. Hippocampus *20*, 1109–1123.

Bakker, A., Kirwan, C.B., Miller, M., and Stark, C.E. (2008). Pattern separation in the human hippocampal CA3 and dentate gyrus. Science *319*, 1640–1642.

Bonnevie, T., Dunn, B., Fyhn, M., Hafting, T., Derdikman, D., Kubie, J.L., Roudi, Y., Moser, E.I., and Moser, M.B. (2013). Grid cells require excitatory drive from the hippocampus. Nat. Neurosci. *16*, 309–317.

Bragin, A., Jandó, G., Nádasdy, Z., Hetke, J., Wise, K., and Buzsáki, G. (1995a). Gamma (40-100 Hz) oscillation in the hippocampus of the behaving rat. J. Neurosci. *15*, 47–60.

Bragin, A., Jandó, G., Nádasdy, Z., van Landeghem, M., and Buzsáki, G. (1995b). Dentate EEG spikes and associated interneuronal population bursts in the hippocampal hilar region of the rat. J. Neurophysiol. *73*, 1691–1705.

Brandon, M.P., Bogaard, A.R., Libby, C.P., Connerney, M.A., Gupta, K., and Hasselmo, M.E. (2011). Reduction of theta rhythm dissociates grid cell spatial periodicity from directional tuning. Science *332*, 595–599.

Burgess, N., Barry, C., and O'Keefe, J. (2007). An oscillatory interference model of grid cell firing. Hippocampus *17*, 801–812.

Efron, B., and Tibshirani, R. (1991). Statistical data analysis in the computer age. Science *253*, 390–395.

Eichenbaum, H. (2004). Hippocampus: cognitive processes and neural representations that underlie declarative memory. Neuron *44*, 109–120.

Ge, S., Yang, C.H., Hsu, K.S., Ming, G.L., and Song, H. (2007). A critical period for enhanced synaptic plasticity in newly generated neurons of the adult brain. Neuron *54*, 559–566.

Gilbert, P.E., Kesner, R.P., and Lee, I. (2001). Dissociating hippocampal subregions: double dissociation between dentate gyrus and CA1. Hippocampus *11*, 626–636.

Gold, A.E., and Kesner, R.P. (2005). The role of the CA3 subregion of the dorsal hippocampus in spatial pattern completion in the rat. Hippocampus *15*, 808–814.

Gothard, K.M., Hoffman, K.L., Battaglia, F.P., and McNaughton, B.L. (2001). Dentate gyrus and ca1 ensemble activity during spatial reference frame shifts in the presence and absence of visual input. J. Neurosci. *21*, 7284–7292.

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.

Hargreaves, E.L., Rao, G., Lee, I., and Knierim, J.J. (2005). Major dissociation between medial and lateral entorhinal input to dorsal hippocampus. Science *308*, 1792–1794.

Hargreaves, E.L., Yoganarasimha, D., and Knierim, J.J. (2007). Cohesiveness of spatial and directional representations recorded from neural ensembles in the anterior thalamus, parasubiculum, medial entorhinal cortex, and hippocampus. Hippocampus *17*, 826–841.

Hasselmo, M.E., and Wyble, B.P. (1997). Free recall and recognition in a network model of the hippocampus: simulating effects of scopolamine on human memory function. Behav. Brain Res. *89*, 1–34.

Jung, M.W., and McNaughton, B.L. (1993). Spatial selectivity of unit activity in the hippocampal granular layer. Hippocampus *3*, 165–182.

Kesner, R.P., Gilbert, P.E., and Wallenstein, G.V. (2000). Testing neural network models of memory with behavioral experiments. Curr. Opin. Neurobiol. *10*, 260–265.

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., Lee, I., and Hargreaves, E.L. (2006). Hippocampal place cells: parallel input streams, subregional processing, and implications for episodic memory. Hippocampus *16*, 755–764.

Kropff, E., and Treves, A. (2008). The emergence of grid cells: intelligent design or just adaptation? Hippocampus *18*, 1256–1269.

Langston, R.F., Ainge, J.A., Couey, J.J., Canto, C.B., Bjerknes, T.L., Witter, M.P., Moser, E.I., and Moser, M.B. (2010). Development of the spatial representation system in the rat. Science *328*, 1576–1580.

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.

Leutgeb, J.K., Leutgeb, S., Moser, M.B., and Moser, E.I. (2007). Pattern separation in the dentate gyrus and CA3 of the hippocampus. Science *315*, 961–966.

Manns, J.R., and Eichenbaum, H. (2006). Evolution of declarative memory. Hippocampus *16*, 795–808.

Marr, D. (1969). A theory of cerebellar cortex. J. Physiol. *202*, 437–470.

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

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

McHugh, T.J., Jones, M.W., Quinn, J.J., Balthasar, N., Coppari, R., Elmquist, J.K., Lowell, B.B., Fanselow, M.S., Wilson, M.A., and Tonegawa, S. (2007). Dentate gyrus NMDA receptors mediate rapid pattern separation in the hippocampal network. Science *317*, 94–99.

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. (1990). 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: Erlbaum), pp. 1–63.

McNaughton, B.L., Battaglia, F.P., Jensen, O., Moser, E.I., and Moser, M.B. (2006). Path integration and the neural basis of the 'cognitive map'. Nat. Rev. Neurosci. *7*, 663–678.

Monaco, J.D., and Abbott, L.F. (2011). Modular realignment of entorhinal grid cell activity as a basis for hippocampal remapping. J. Neurosci. *31*, 9414–9425.

Nakazawa, K., Quirk, M.C., Chitwood, R.A., Watanabe, M., Yeckel, M.F., Sun, L.D., Kato, A., Carr, C.A., Johnston, D., Wilson, M.A., and Tonegawa, S. (2002). Requirement for hippocampal CA3 NMDA receptors in associative memory recall. Science *297*, 211–218.

Neunuebel, J.P., and Knierim, J.J. (2012). Spatial firing correlates of physiologically distinct cell types of the rat dentate gyrus. J. Neurosci. *32*, 3848–3858.

Neunuebel, J.P., Yoganarasimha, D., Rao, G., and Knierim, J.J. (2013). Conflicts between local and global spatial frameworks dissociate neural representations of the lateral and medial entorhinal cortex. J. Neurosci. *33*, 9246–9258.

O'Keefe, J., and Burgess, N. (2005). Dual phase and rate coding in hippocampal place cells: theoretical significance and relationship to entorhinal grid cells. Hippocampus *15*, 853–866.

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

O'Reilly, R.C., and McClelland, J.L. (1994). Hippocampal conjunctive encoding, storage, and recall: avoiding a trade-off. Hippocampus *4*, 661–682.

Rolls, E.T., and Kesner, R.P. (2006). A computational theory of hippocampal function, and empirical tests of the theory. Prog. Neurobiol. *79*, 1–48.

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

Santoro, A. (2013). Reassessing pattern separation in the dentate gyrus. Front. Behav. Neurosci. *7*, 96.

Sargolini, F., Fyhn, M., Hafting, T., McNaughton, B.L., Witter, M.P., Moser, M.B., and Moser, E.I. (2006). Conjunctive representation of position, direction, and velocity in entorhinal cortex. Science *312*, 758–762.

Savelli, F., and Knierim, J.J. (2010). Hebbian analysis of the transformation of medial entorhinal grid-cell inputs to hippocampal place fields. J. Neurophysiol. *103*, 3167–3183.

Skaggs, W.E., McNaughton, B.L., Wilson, M.A., and Barnes, C.A. (1996). Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. Hippocampus *6*, 149–172.

Solstad, T., Moser, E.I., and Einevoll, G.T. (2006). From grid cells to place cells: a mathematical model. Hippocampus *16*, 1026–1031.

Squire, L.R., Stark, C.E., and Clark, R.E. (2004). The medial temporal lobe. Annu. Rev. Neurosci. *27*, 279–306.

Suzuki, W.A., Miller, E.K., and Desimone, R. (1997). Object and place memory in the macaque entorhinal cortex. J. Neurophysiol. *78*, 1062–1081.

Vivar, C., Potter, M.C., Choi, J., Lee, J.Y., Stringer, T.P., Callaway, E.M., Gage, F.H., Suh, H., and van Praag, H. (2012). Monosynaptic inputs to new neurons in the dentate gyrus. Nat. Commun. *3*, 1107.

Wills, T.J., Cacucci, F., Burgess, N., and O'Keefe, J. (2010). Development of the hippocampal cognitive map in preweanling rats. Science *328*, 1573–1576.

Yassa, M.A., and Stark, C.E. (2011). Pattern separation in the hippocampus. Trends Neurosci. *34*, 515–525.

Yoganarasimha, D., Yu, X., and Knierim, J.J. (2006). Head direction cell representations maintain internal coherence during conflicting proximal and distal cue rotations: comparison with hippocampal place cells. J. Neurosci. *26*, 622–631.

Zhang, K. (1996). Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: a theory. J. Neurosci. *16*, 2112–2126.

Zugaro, M.B., Berthoz, A., and Wiener, S.I. (2001). Background, but not foreground, spatial cues are taken as references for head direction responses by rat anterodorsal thalamus neurons. J. Neurosci. *21*, 1–5.

Neuron, Volume *81*

Supplemental Information

CA3 Retrieves Coherent Representations from Degraded Input: Direct Evidence for CA3 Pattern Completion and Dentate Gyrus Pattern Separation

Joshua P. Neunuebel and James J. Knierim

Supplemental Figures

Figure S1

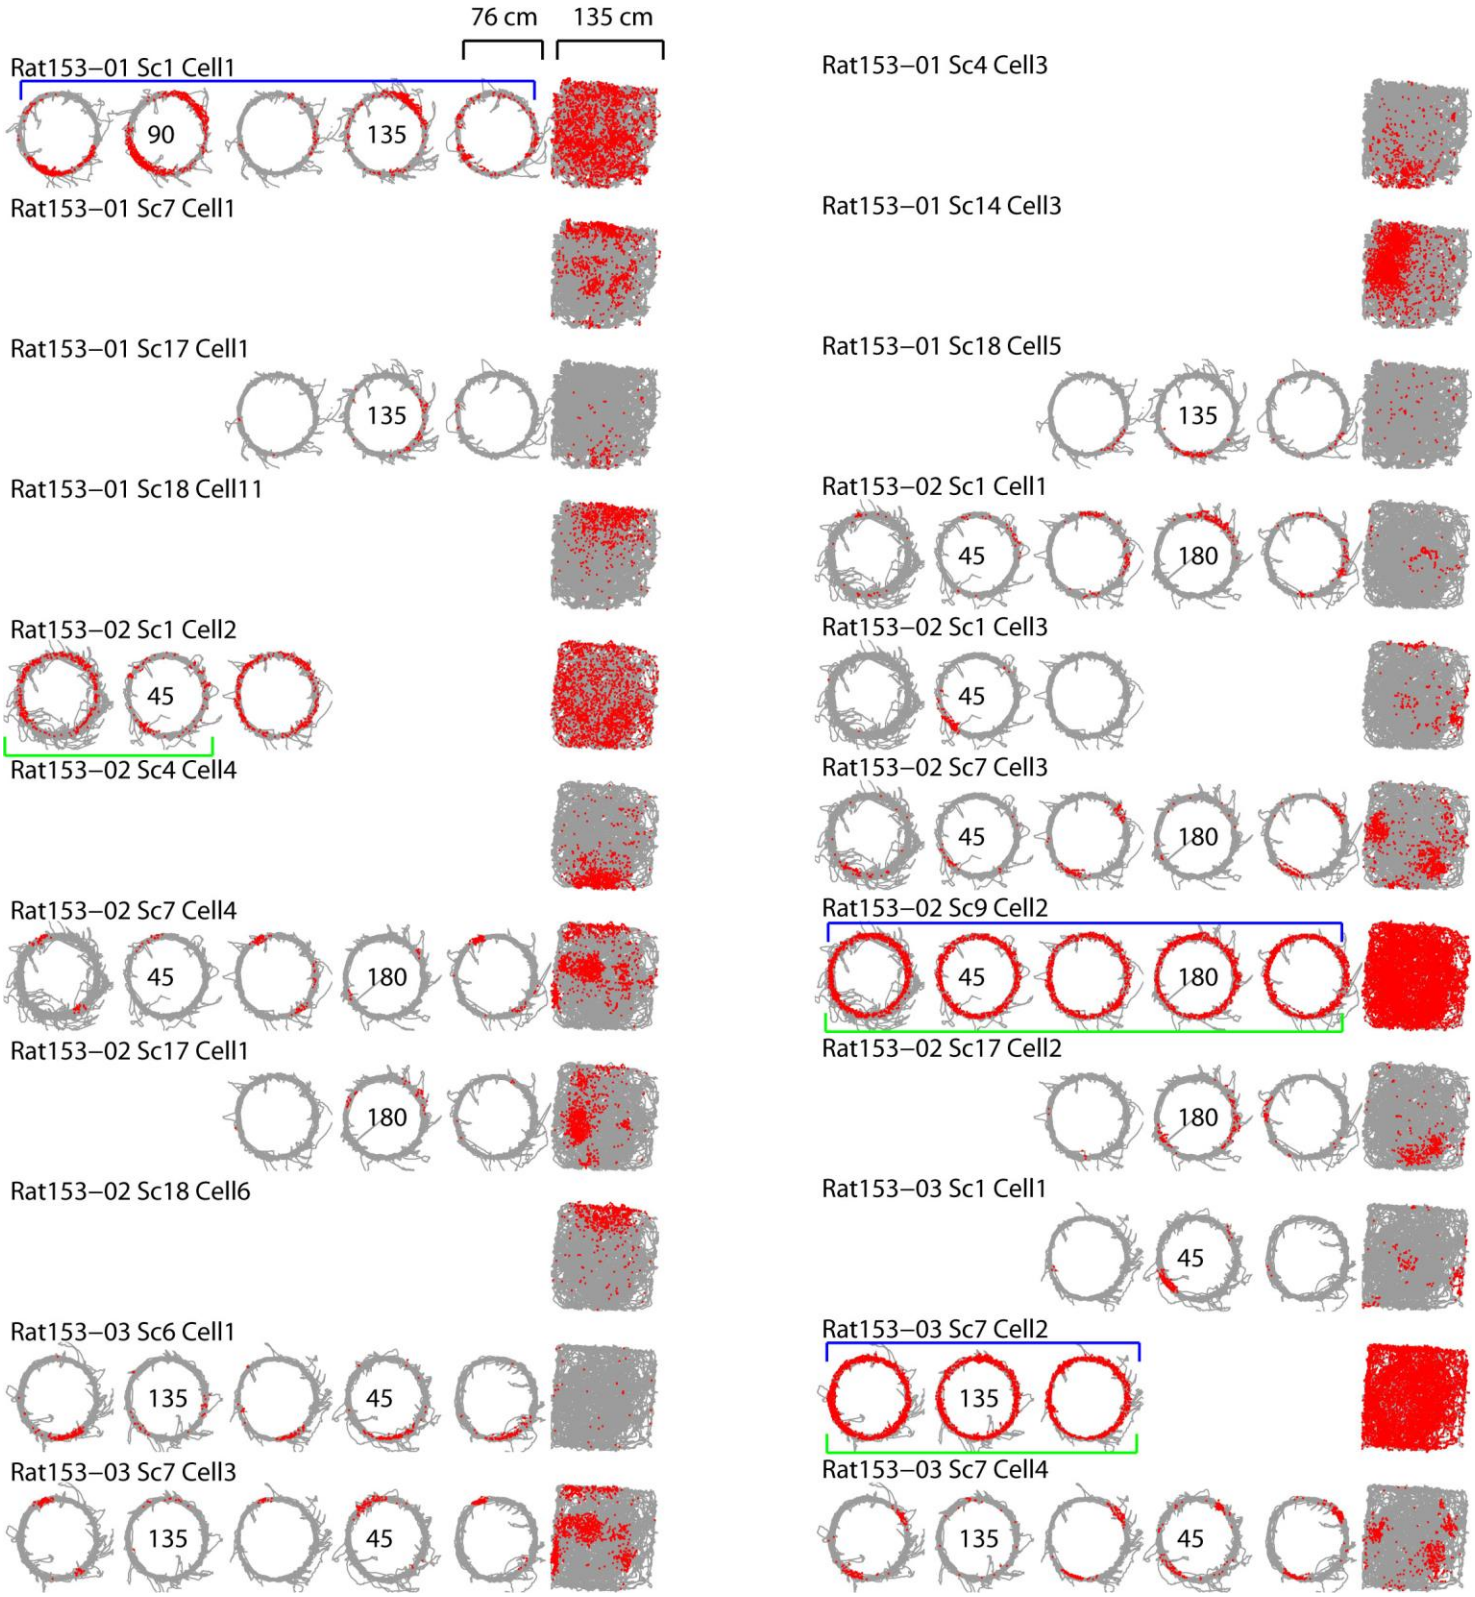

Figure S1 (cont.)

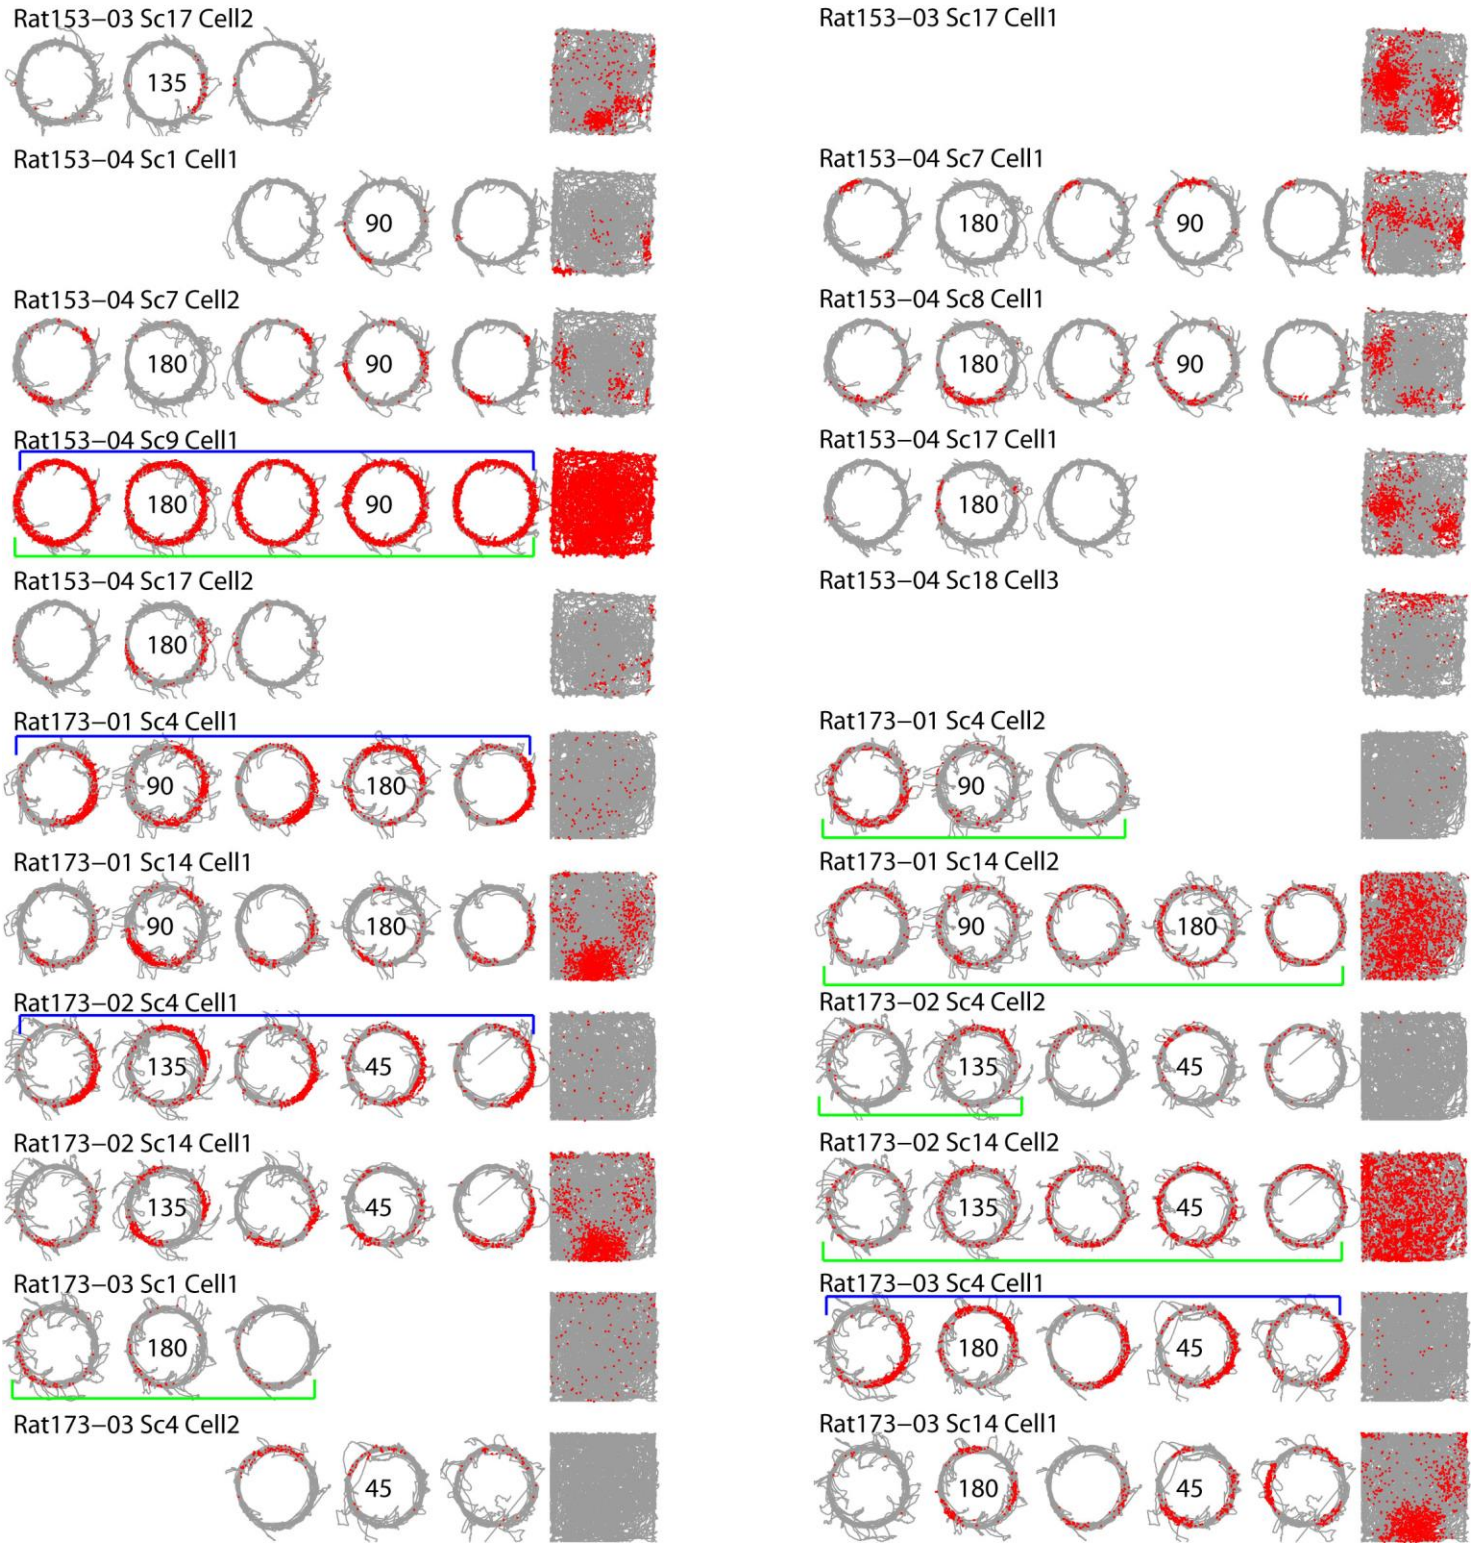

Figure S1 (cont.)

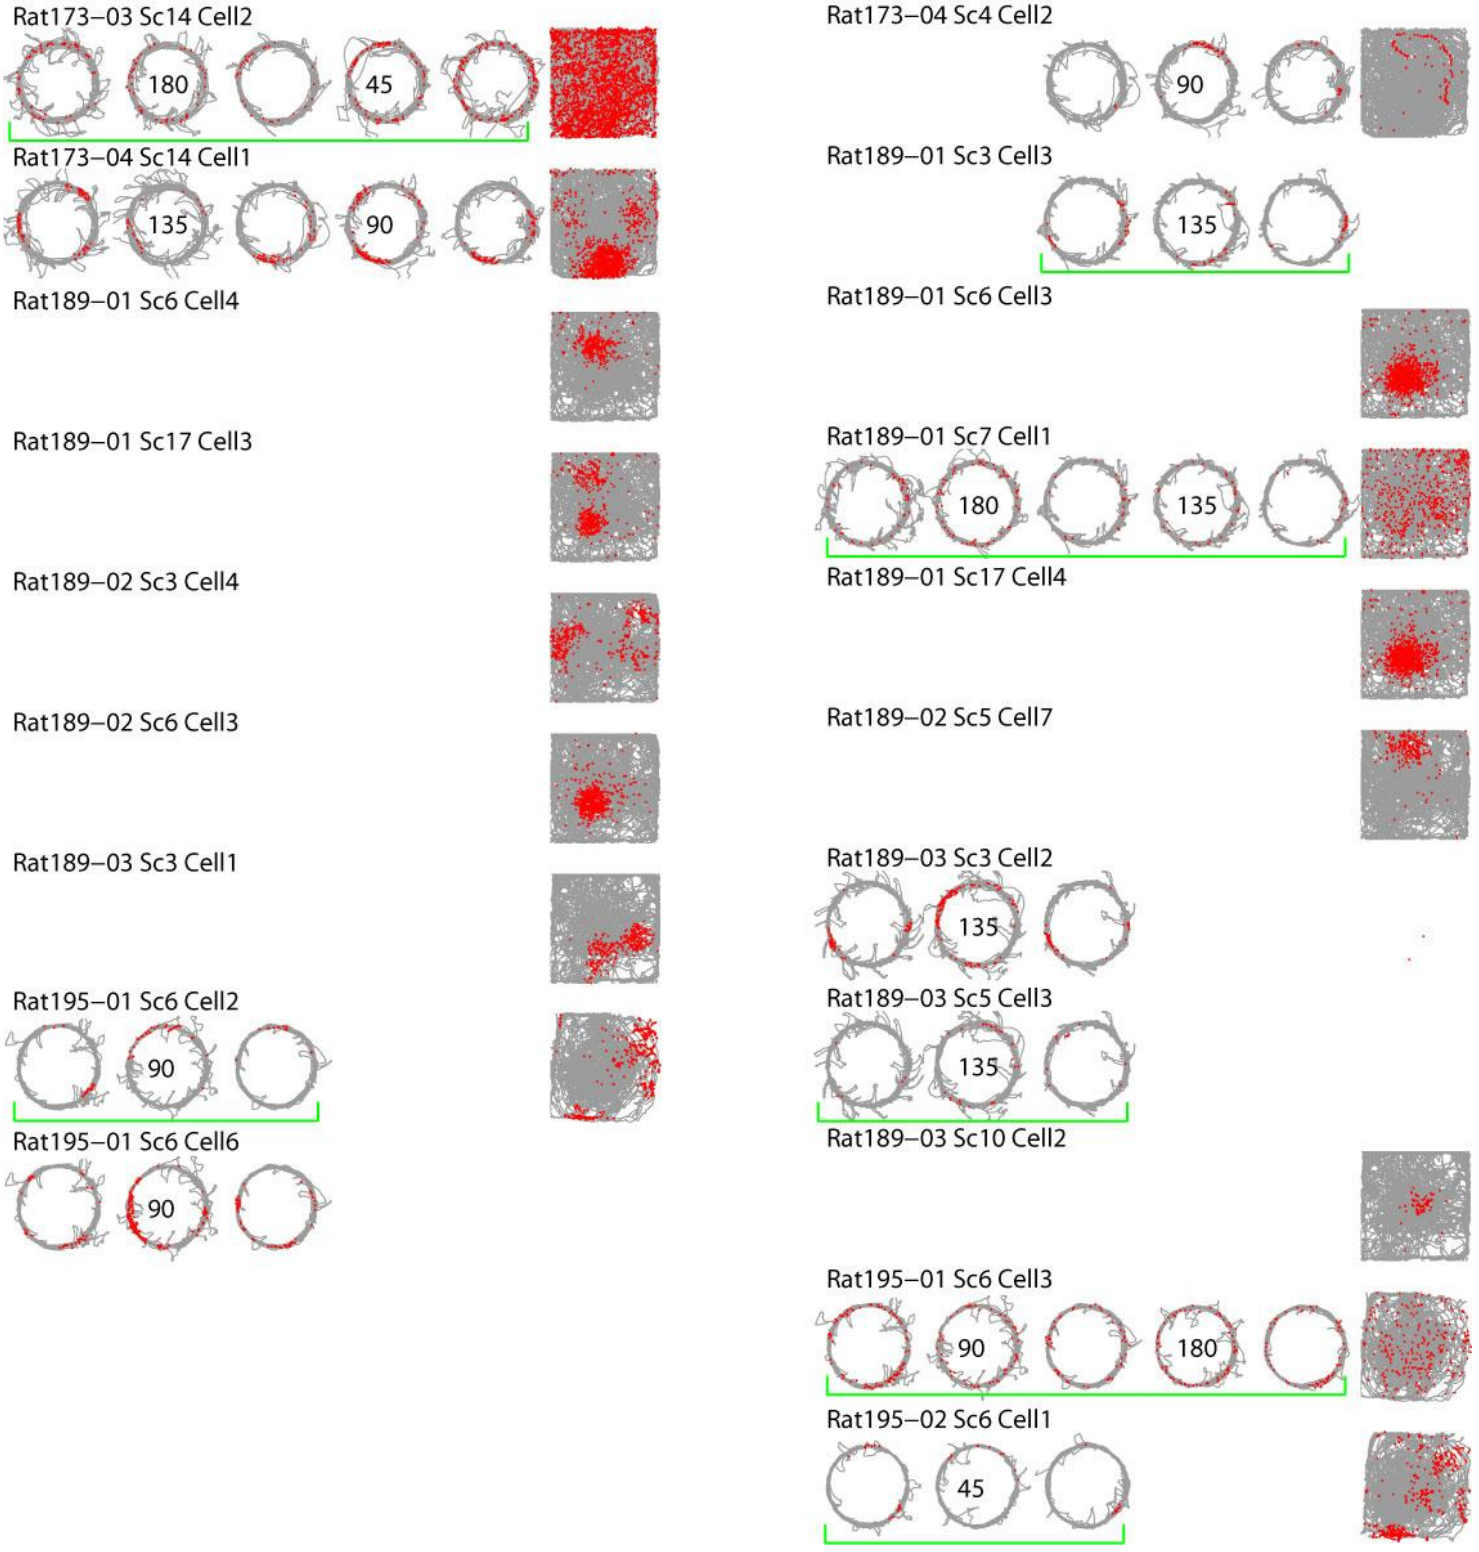

Figure S1 (cont.)

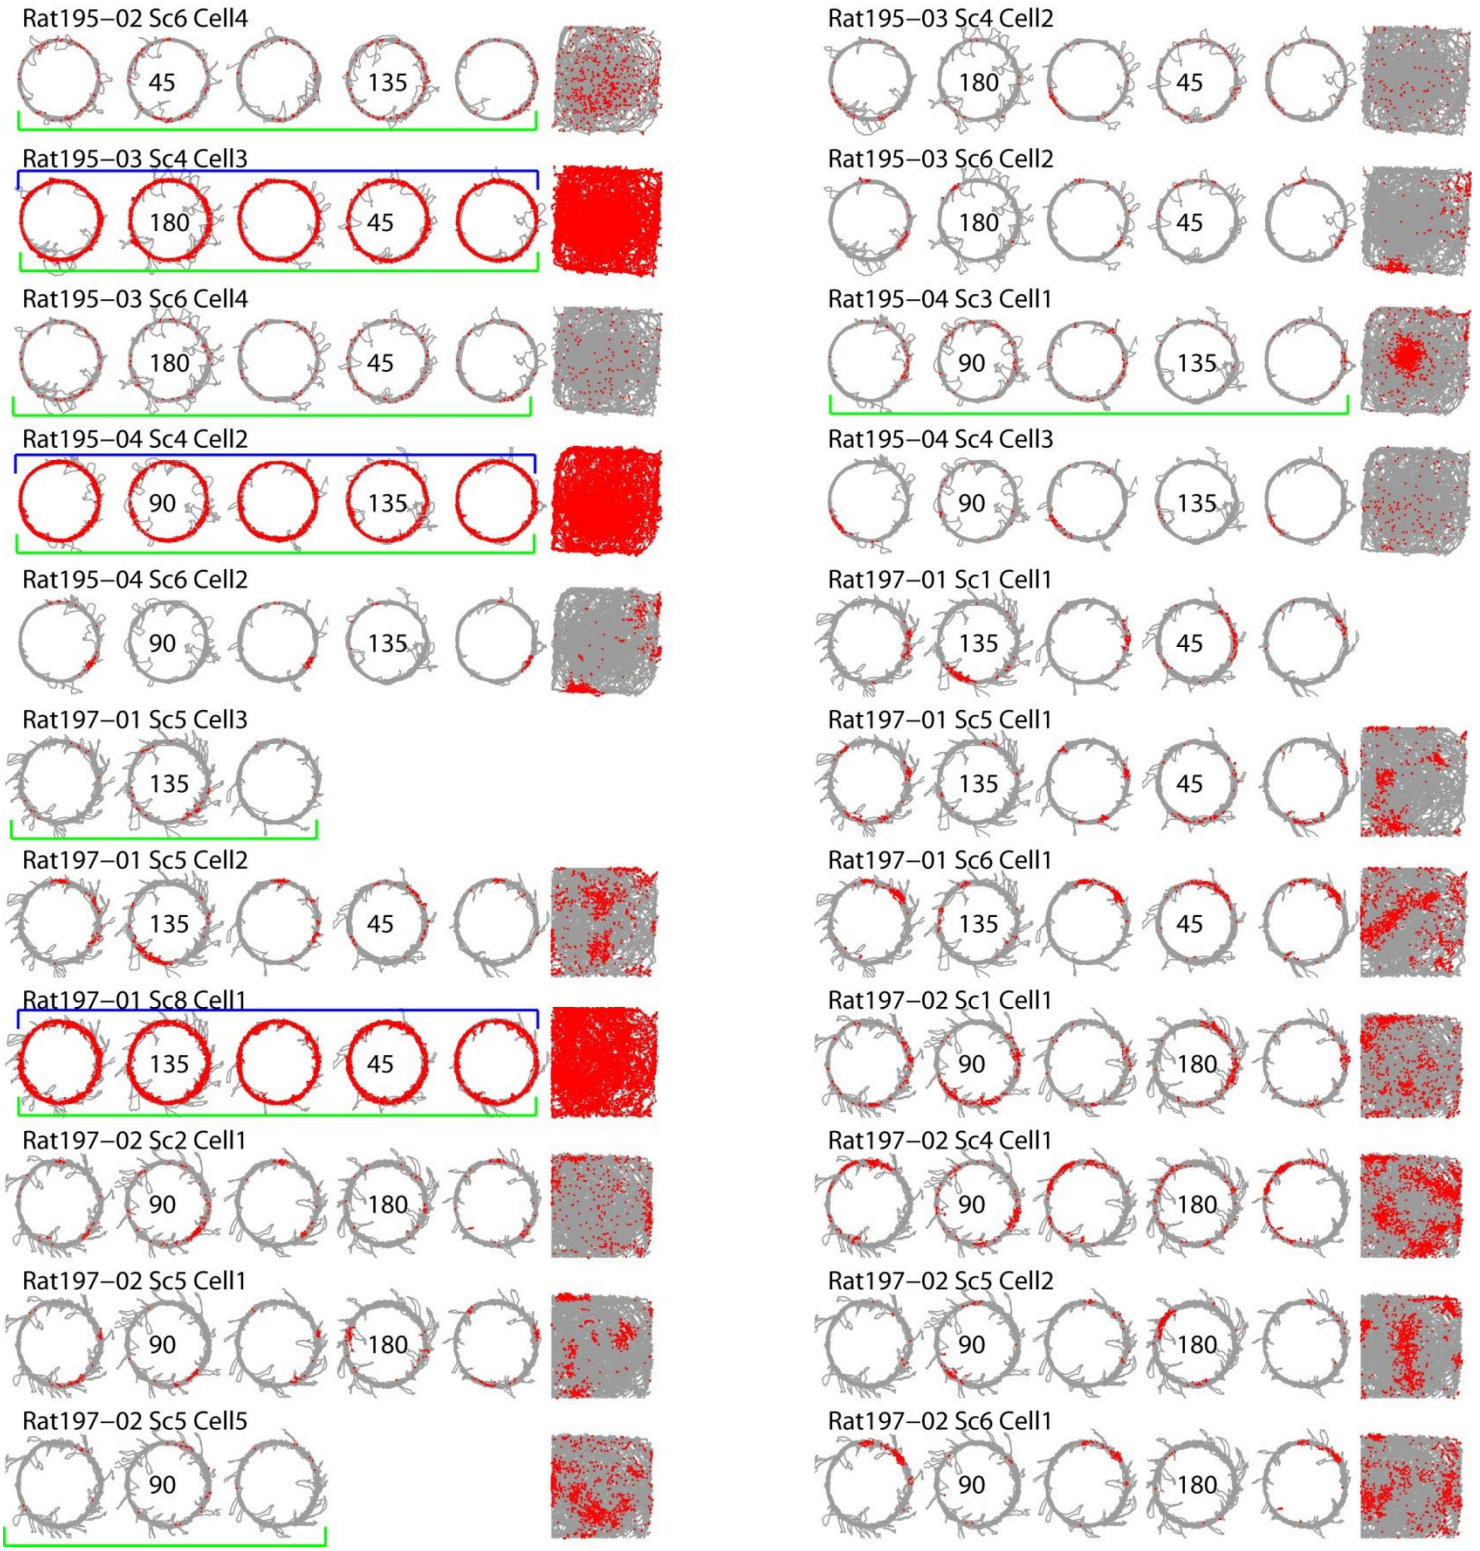

Figure S1 (cont.)

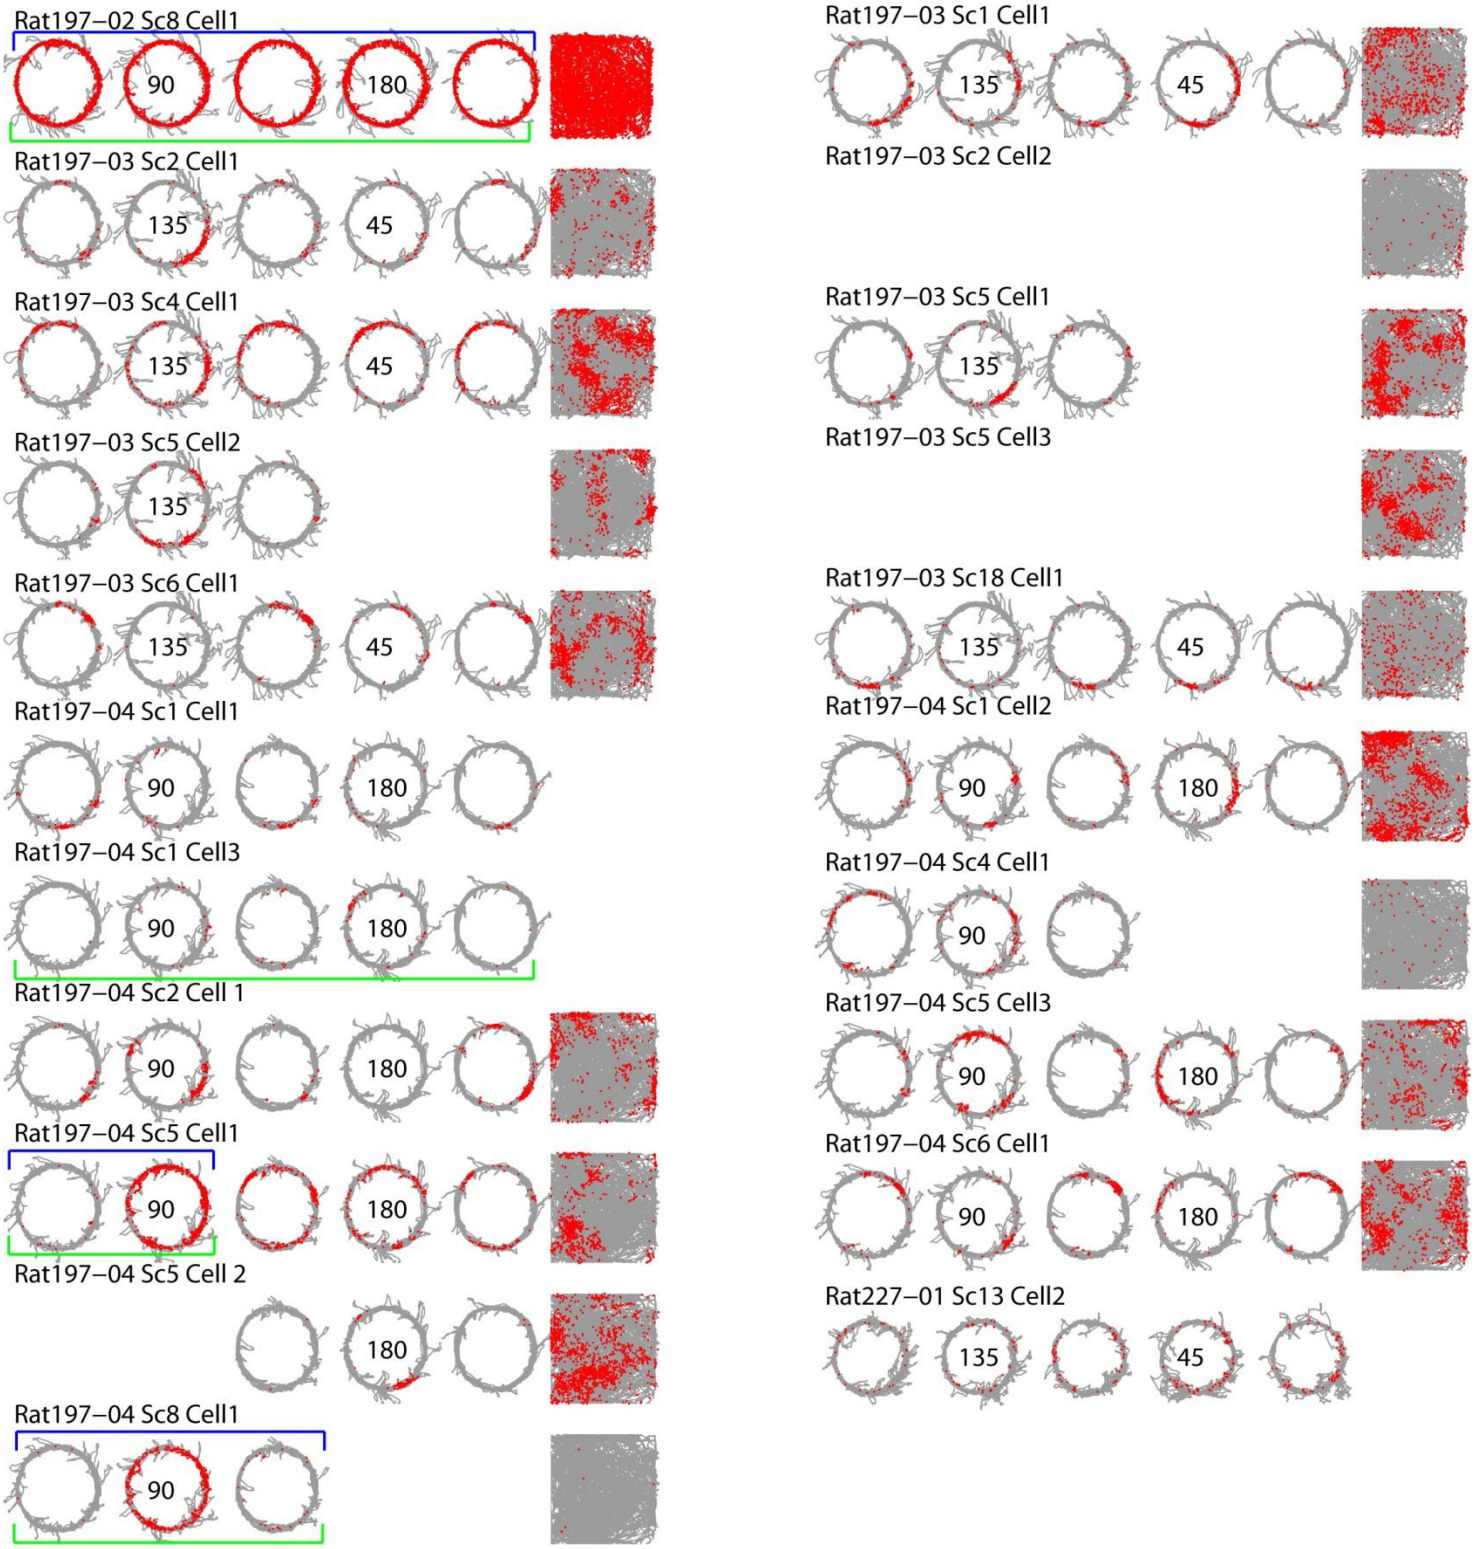

Figure S1 (cont.)

Figure S1. Spatial firing plots for the DG cells that met activity criteria, related to Figure 3. This figure shows the rat's trajectory (gray lines) and locations where the spikes fired (red dots) for the 5 circular track sessions and the open-field foraging session for cells that met minimum activity criteria (≥ 20 spikes, mean firing rate < 10 Hz) for the main analyses of the double rotation manipulation (Figure 4, columns 3-4). All of the circular track data that went into the primary analyses are shown in this figure; the spikes that occurred when the rat's head was off the track were removed. The first, third, and fifth columns are standard sessions; the second and fourth columns are mismatch sessions (the mismatch angle is shown in the center of the plot). For many cells, only 3 of the five rate maps are shown. In most of these cases, the cells did not meet the activity criteria in the missing sessions, so they are not plotted here. The lack of strong firing in these sessions was not due to recording instability, as the cells usually fired a few spikes during those sessions; in most cases in which the cell fired no spikes in the missing sessions, the recording stability was verified by identifying the cell in either baseline sleep or behavioral sessions before or after the missing sessions. For the purposes of the main analyses, the recording stability is shown clearly by the Std vs. Std correlation matrices and polar plots (Figure 4; Figures S5 and S6), which show strong spatial correlations in the two standard sessions that bracket the mismatch sessions. For other cells, only the foraging sessions are shown; these are cells that did not meet activity criteria in any of the circular track sessions but met the following criteria in the foraging session: > 75 spikes, info > 0.5 bits/spike, and p value ≤ 0.01. Not shown are any of the cells that did not meet activity criteria in any of the 6 sessions of the day, but were active during the baseline sleep sessions (Neunuebel and Knierim, 2012). Blue lines on top of the circular track maps indicate the moderate-rate cells that fired between 2 and 10 Hz; these cells were not included in the analyses of Figure 4 (columns 5-6) and Figure S6. Green lines at the bottom of the circular track data indicate the cells that did not meet the spatial criteria (information score > 0.5 bits/spike, $p < 0.01$); these cells were not included in the analyses of Figure S5. None of the open-field foraging sessions are shown for Rats 227 and 232, because the behavior of both rats in these sessions was poor on all days of recording and the box was not adequately sampled. In a few other cases the foraging session is not shown because technical issues (e.g., recording noise) made the cells poorly isolated in those sessions. Note that many of the cells fired in multiple locations in the open-field sessions, as reported by Leutgeb and colleagues (2007) (also Jung and McNaughton, 1993). Many other cells fired in single locations, like classic place cells, or were virtually silent, as reported by Neunuebel and Knierim (2012). It is not known whether these cells correspond to different DG cell types (e.g., mature, developmentally born granule cells; mature, adult-born granule cells; immature, adult-born granule cells; or cells of the hilus that may be recorded from tetrodes potentially in the granule cell layer). There were no conclusive differences in the double rotation analyses between the cells that fired in multiple subfields vs. cells that fired in single fields or were silent (not shown), although the small numbers of cells in each group when the data are divided in this way limit the power of this negative result. Note that the circular track and the open field are not drawn to scale. Cell identification key: RatX-Y ScM CellN, where X = rat number, Y = recording day, M = tetrode number, and N = cell number.

Figure S2

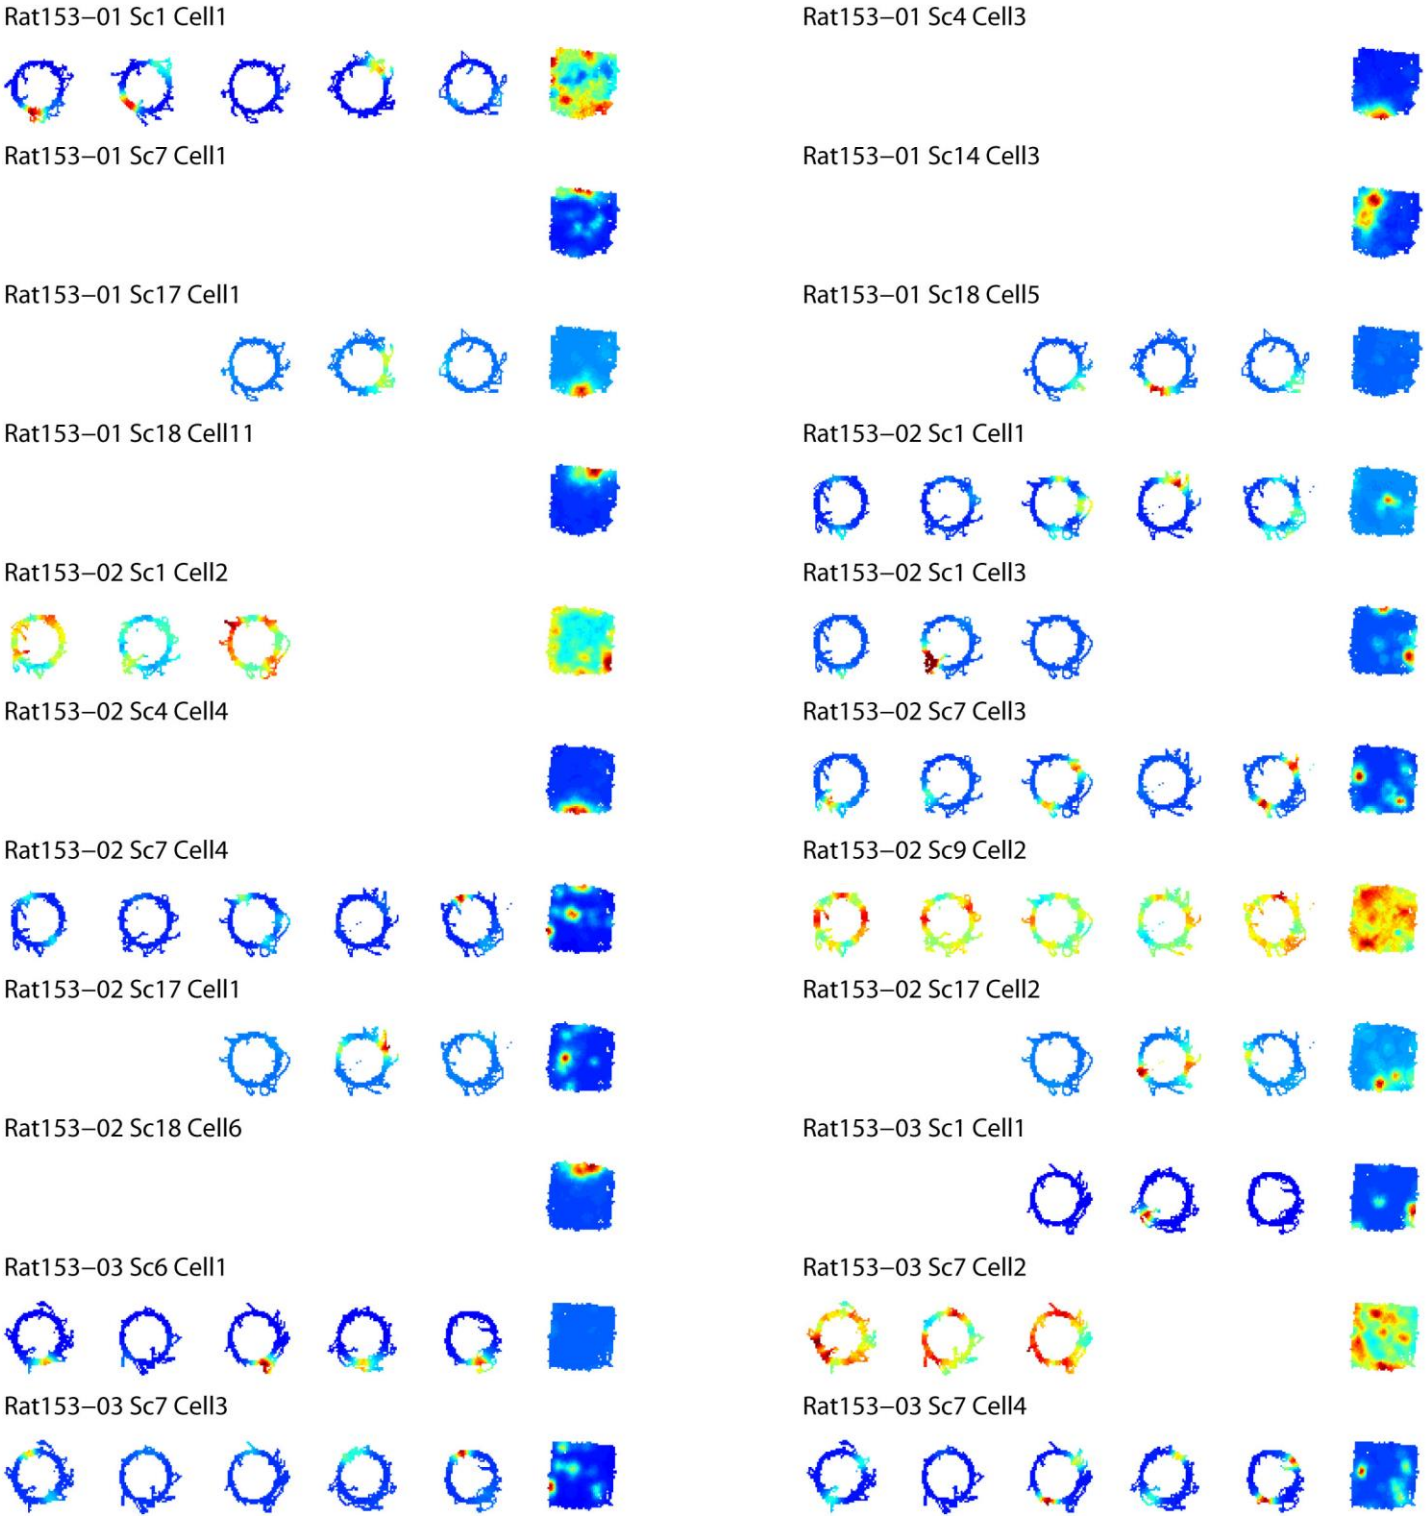
Rat153–01 Sc1 Cell1
Rat153–01 Sc4 Cell3
Rat153–01 Sc7 Cell1
Rat153–01 Sc14 Cell3
Rat153–01 Sc17 Cell1
Rat153–01 Sc18 Cell5
Rat153–01 Sc18 Cell11
Rat153–02 Sc1 Cell1
Rat153–02 Sc1 Cell2
Rat153–02 Sc1 Cell3
Rat153–02 Sc4 Cell4
Rat153–02 Sc7 Cell3
Rat153–02 Sc7 Cell4
Rat153–02 Sc9 Cell2
Rat153–02 Sc17 Cell1
Rat153–02 Sc17 Cell2
Rat153–02 Sc18 Cell6
Rat153–03 Sc1 Cell1
Rat153–03 Sc6 Cell1
Rat153–03 Sc7 Cell2
Rat153–03 Sc7 Cell3
Rat153–03 Sc7 Cell4

Figure S2 (cont.)

Rat153–03 Sc17 Cell2

Rat153–03 Sc17 Cell1

Rat153–04 Sc1 Cell1

Rat153–04 Sc7 Cell1

Rat153–04 Sc7 Cell2

Rat153–04 Sc8 Cell1

Rat153–04 Sc9 Cell1

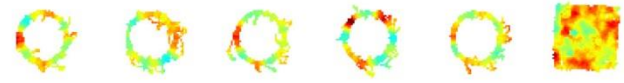

Rat153–04 Sc17 Cell1

Rat153–04 Sc17 Cell2

Rat153–04 Sc18 Cell3

Rat173–01 Sc4 Cell1

Rat173–01 Sc4 Cell2

Rat173–01 Sc14 Cell1

Rat173–01 Sc14 Cell2

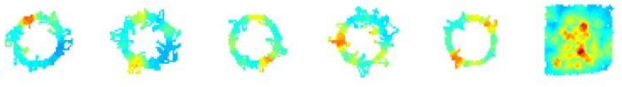

Rat173–02 Sc4 Cell1

Rat173–02 Sc4 Cell2

Rat173–02 Sc14 Cell1

Rat173–02 Sc14 Cell2

Rat173–03 Sc1 Cell1

Rat173–03 Sc4 Cell1

Rat173–03 Sc4 Cell2

Rat173–03 Sc14 Cell1

Figure S2 (cont.)

Rat173–03 Sc14 Cell2

Rat173–04 Sc14 Cell1

Rat189–01 Sc6 Cell4

Rat189–01 Sc17 Cell3

Rat189–02 Sc3 Cell4

Rat189–02 Sc6 Cell3

Rat189–03 Sc3 Cell1

Rat195–01 Sc6 Cell2

Rat195–01 Sc6 Cell6

Rat173–04 Sc4 Cell2

Rat189–01 Sc3 Cell3

Rat189–01 Sc6 Cell3

Rat189–01 Sc7 Cell1

Rat189–01 Sc17 Cell4

Rat189–02 Sc5 Cell7

Rat189–03 Sc3 Cell2

Rat189–03 Sc5 Cell3

Rat189–03 Sc10 Cell2

Rat195–01 Sc6 Cell3

Rat195–02 Sc6 Cell1

Figure S2 (cont.)

Rat195–02 Sc6 Cell4

Rat195–03 Sc4 Cell3

Rat195–03 Sc6 Cell4

Rat195–04 Sc4 Cell2

Rat195–04 Sc6 Cell2

Rat197–01 Sc5 Cell3

Rat197–01 Sc5 Cell2

Rat197–01 Sc8 Cell1

Rat197–02 Sc2 Cell1

Rat197–02 Sc5 Cell1

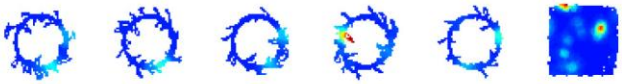

Rat197–02 Sc5 Cell5

Rat195–03 Sc4 Cell2

Rat195–03 Sc6 Cell2

Rat195–04 Sc3 Cell1

Rat195–04 Sc4 Cell3

Rat197–01 Sc1 Cell1

Rat197–01 Sc5 Cell1

Rat197–01 Sc6 Cell1

Rat197–02 Sc1 Cell1

Rat197–02 Sc4 Cell1

Rat197–02 Sc5 Cell2

Rat197–02 Sc6 Cell1

Figure S2 (cont.)

Rat197–02 Sc8 Cell1

Rat197–03 Sc2 Cell1

Rat197–03 Sc4 Cell1

Rat197–03 Sc5 Cell2

Rat197–03 Sc6 Cell1

Rat197–04 Sc1 Cell1

Rat197–04 Sc1 Cell3

Rat197–04 Sc2 Cell 1

Rat197–04 Sc5 Cell1

Rat197–04 Sc5 Cell 2

Rat197–04 Sc8 Cell1

Rat197–03 Sc1 Cell1

Rat197–03 Sc2 Cell2

Rat197–03 Sc5 Cell1

Rat197–03 Sc5 Cell3

Rat197–03 Sc18 Cell1

Rat197–04 Sc1 Cell2

Rat197–04 Sc4 Cell1

Rat197–04 Sc5 Cell3

Rat197–04 Sc6 Cell1

Rat227–01 Sc13 Cell2

Figure S2 (cont.)

Rat227–01 Sc15 Cell3

Rat227–02 Sc4 Cell1

Rat227–02 Sc13 Cell1

Rat227–02 Sc13 Cell2

Rat227–02 Sc15 Cell1

Rat227–02 Sc16 Cell1

Rat227–03 Sc4 Cell1

Rat227–03 Sc4 Cell3

Rat227–03 Sc13 Cell1

Rat227–03 Sc13 Cell2

Rat227–03 Sc15 Cell1

Rat227–03 Sc16 Cell1

Rat227–04 Sc13 Cell1

Rat227–04 Sc15 Cell2

Rat232–01 Sc12 Cell1

Rat232–03 Sc12 Cell1

Rat232–04 Sc1 Cell1

Rat232–04 Sc12 Cell1

Rat232–05 Sc12 Cell1

Rat232–05 Sc17 Cell1

Rat232–05 Sc17 Cell3

Rat232–05 Sc18 Cell3

Figure S2. Firing rate maps for the DG cells that met activity criteria, related to Figure 3. The cells are arranged in the same order as in Figure S1. These rate maps provide a standard representation of the data normalized by occupancy. They are provided here as a complement for the spike-trajectory plots of Figure S1. These maps are especially useful to demonstrate that some of the moderate-rate cells identified with the blue bars in Figure S1 show a degree of spatial selectivity that is consistent across sessions. For each cell, the firing rate scale (red = maximal firing rate; blue = no firing) is the same for each of the 6 sessions.

Figure S3

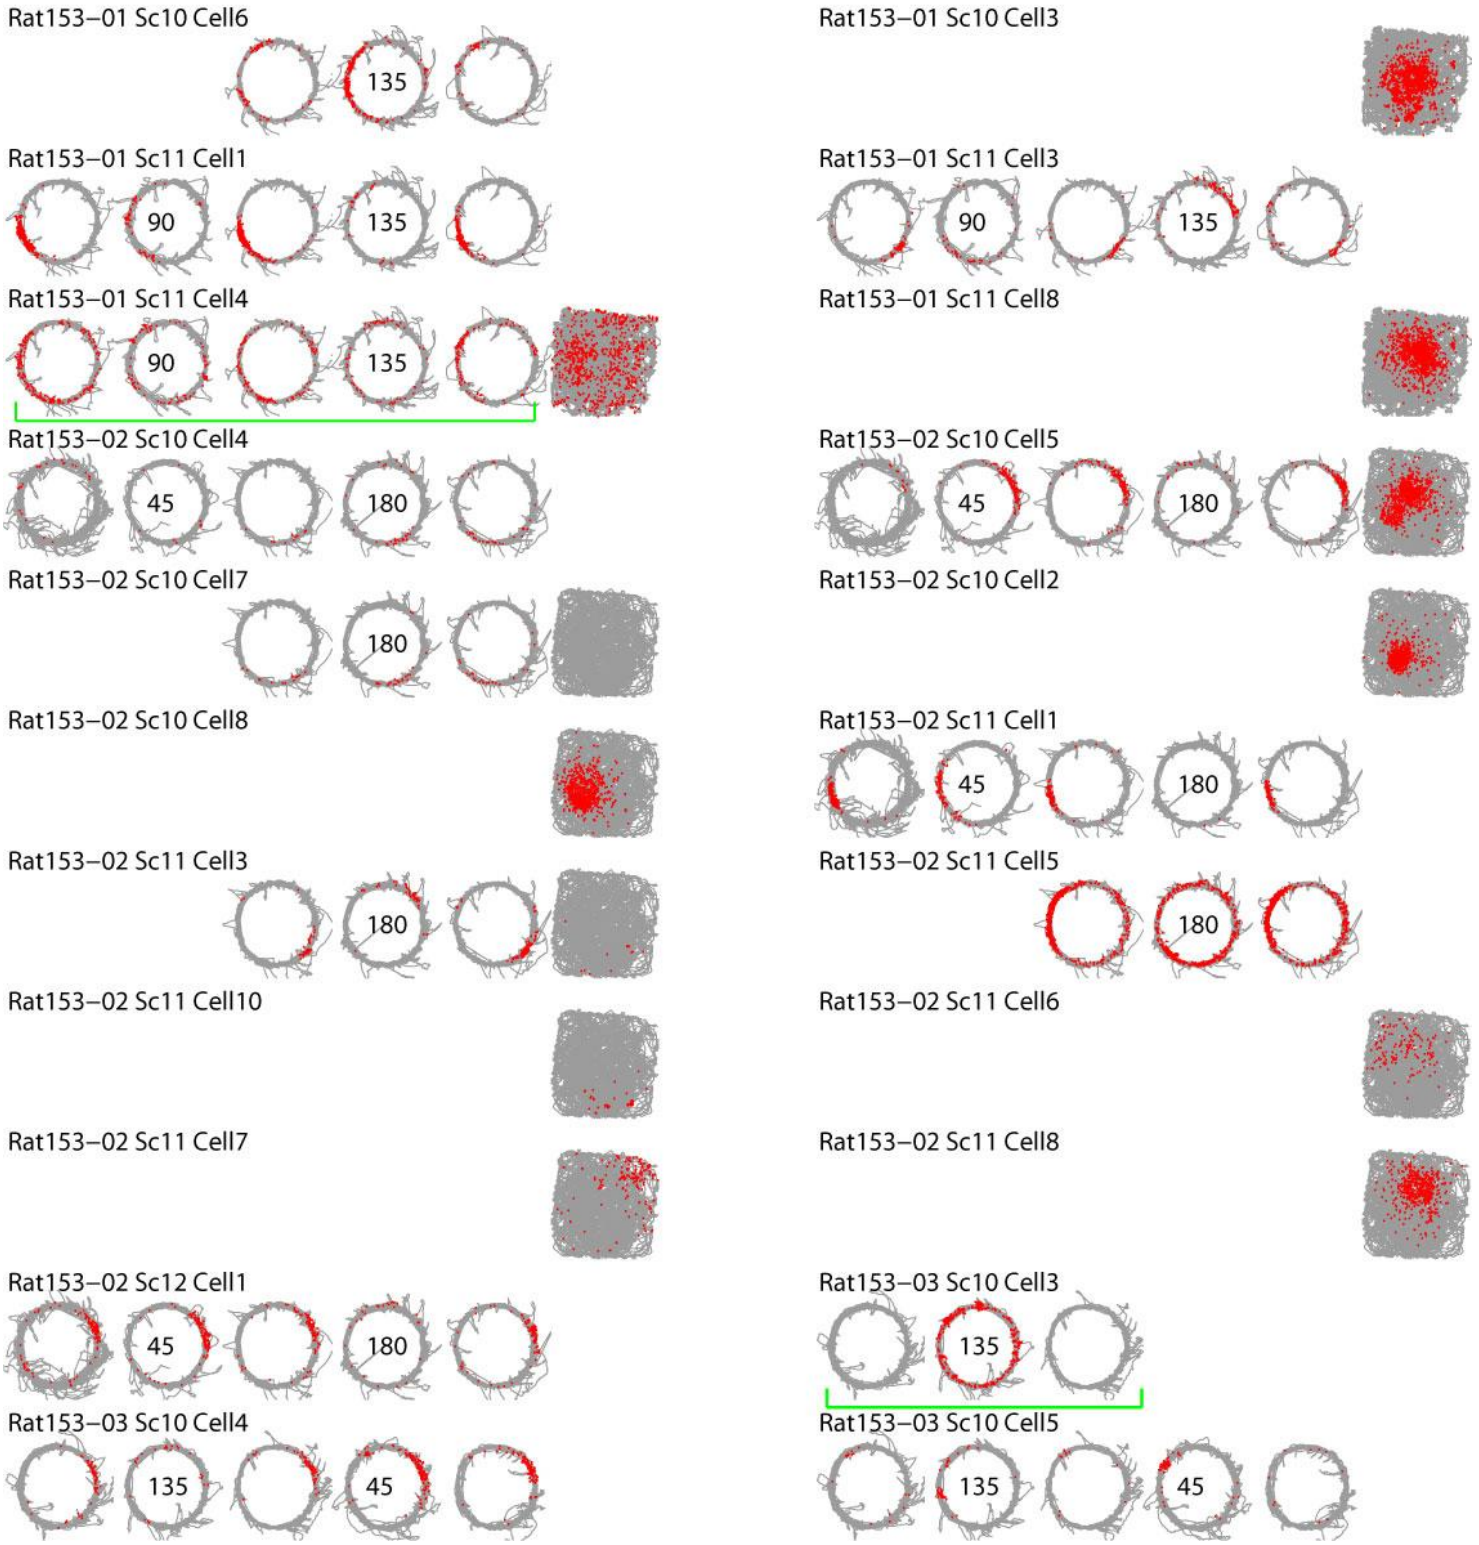
Rat153–01 Sc10 Cell6
135
Rat153–01 Sc10 Cell3
Rat153–01 Sc11 Cell1
90
135
Rat153–01 Sc11 Cell3
90
135
Rat153–01 Sc11 Cell4
90
135
Rat153–01 Sc11 Cell8
Rat153–02 Sc10 Cell4
45
180
Rat153–02 Sc10 Cell5
45
180
Rat153–02 Sc10 Cell7
180
Rat153–02 Sc10 Cell2
Rat153–02 Sc10 Cell8
Rat153–02 Sc11 Cell1
45
180
Rat153–02 Sc11 Cell3
180
Rat153–02 Sc11 Cell5
180
Rat153–02 Sc11 Cell10
Rat153–02 Sc11 Cell6
Rat153–02 Sc11 Cell7
Rat153–02 Sc11 Cell8
Rat153–02 Sc12 Cell1
45
180
Rat153–03 Sc10 Cell3
135
Rat153–03 Sc10 Cell4
135
45
Rat153–03 Sc10 Cell5
135
45

Figure S3 (cont.)

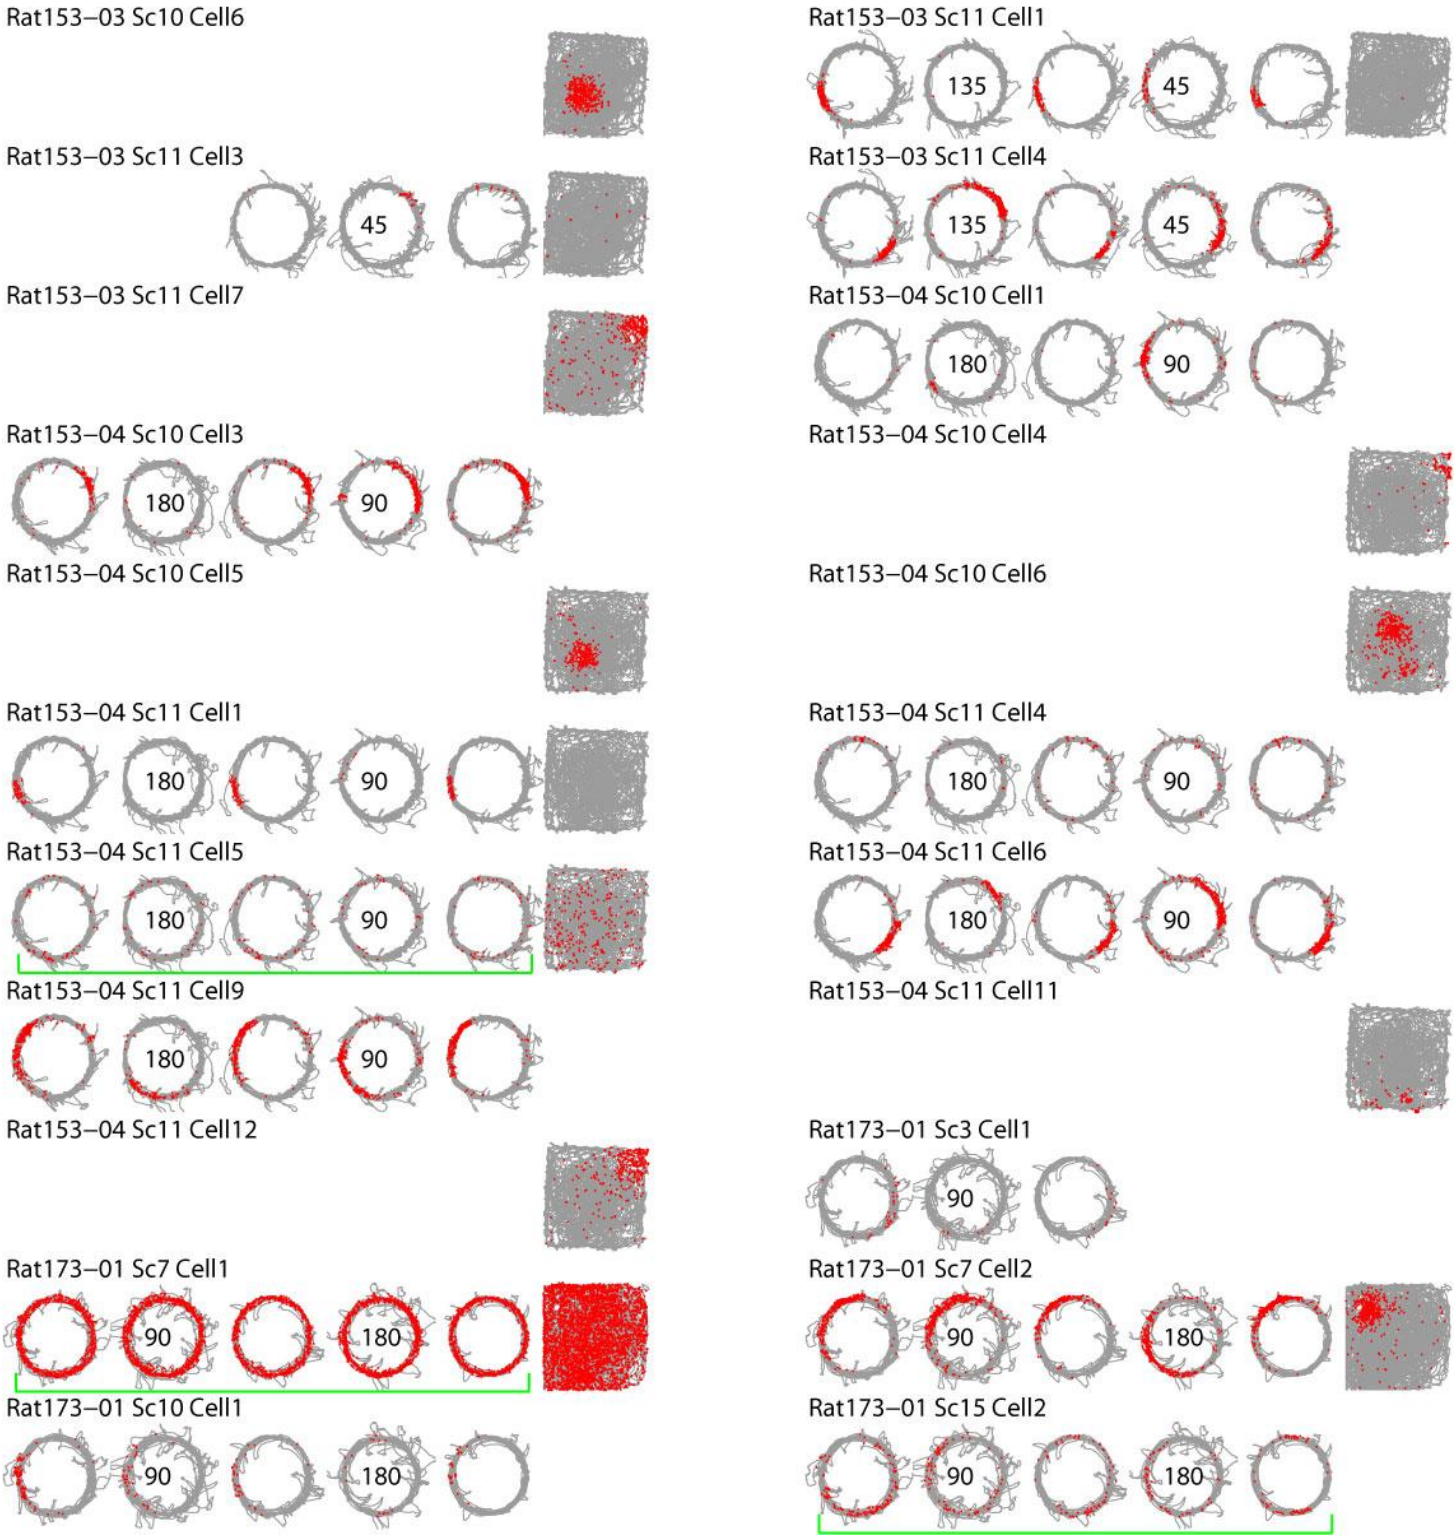

Figure S3 (cont.)

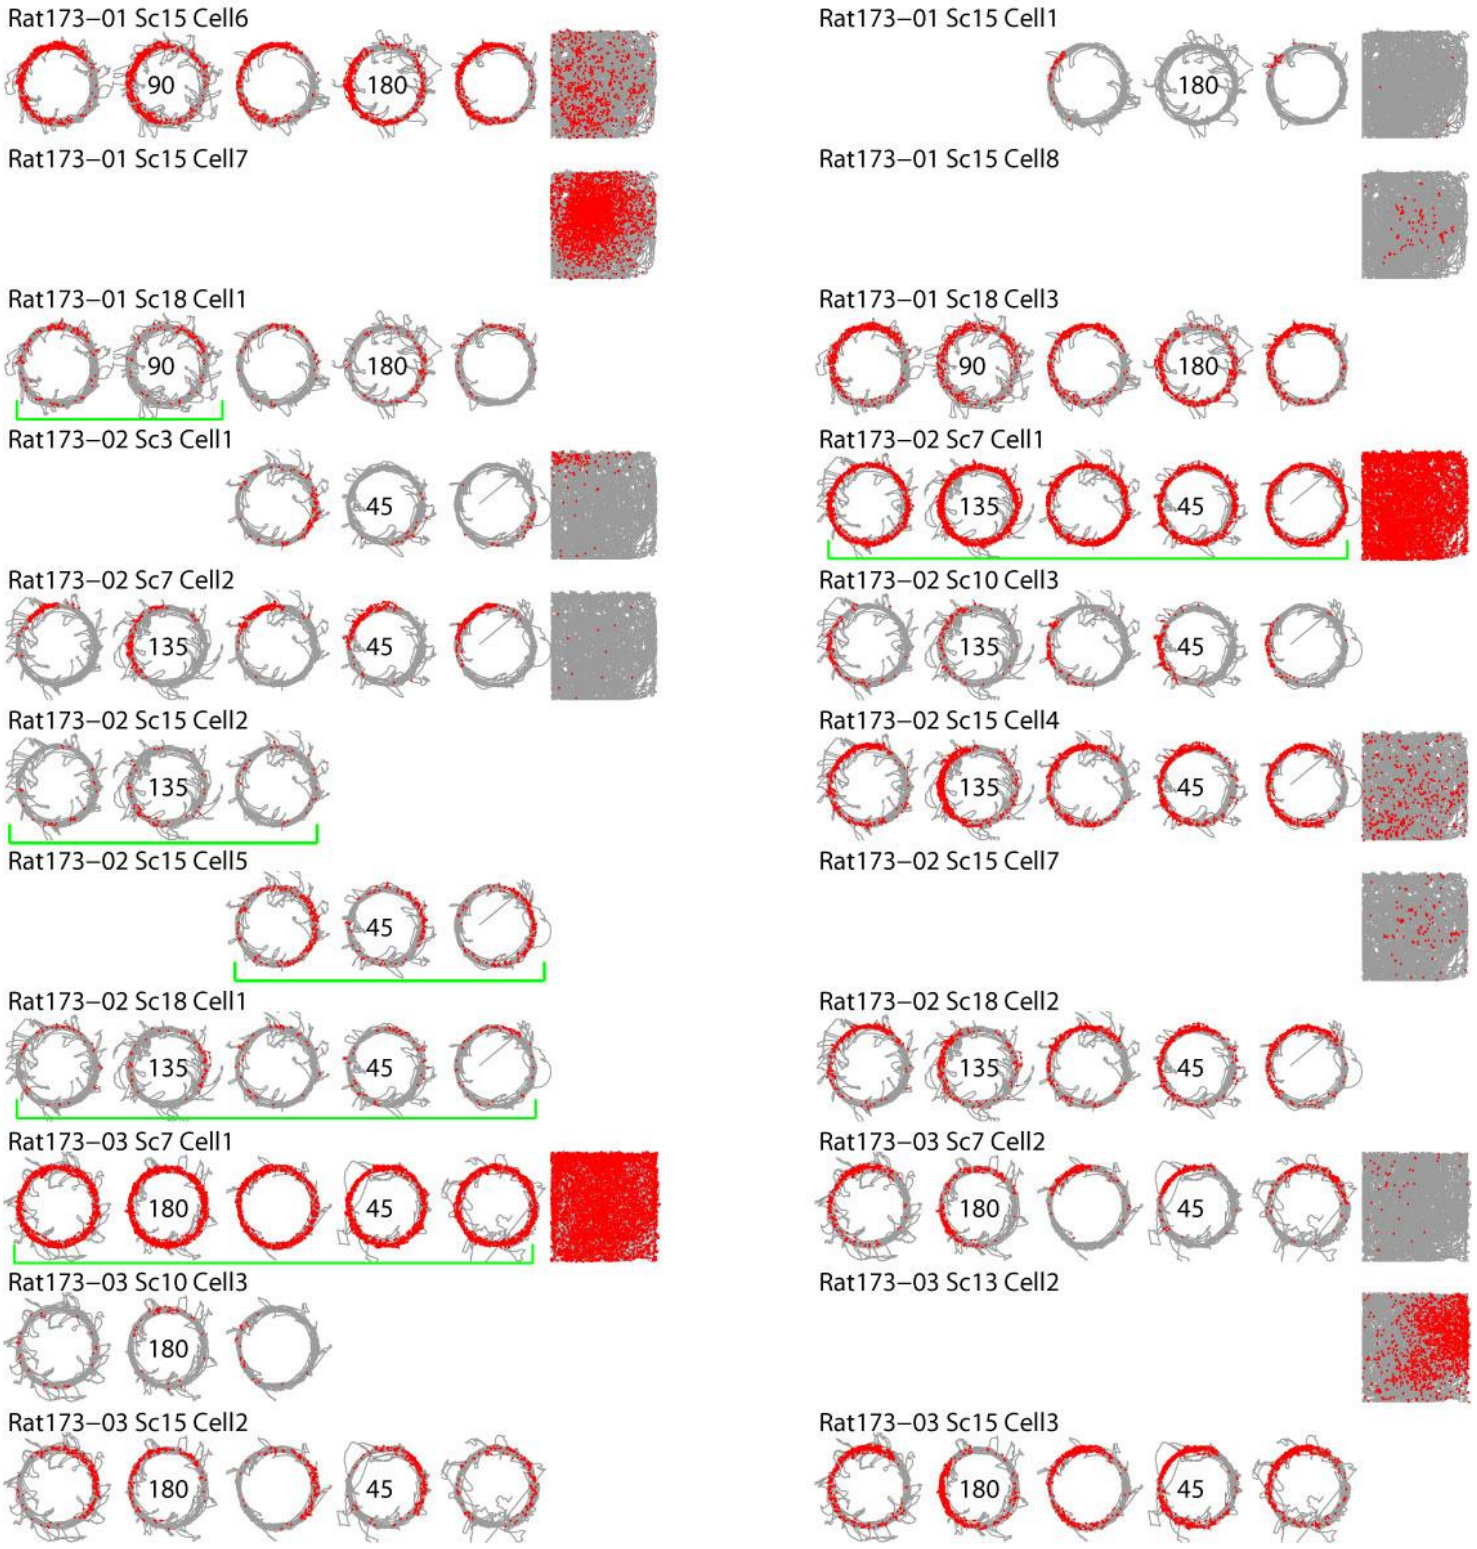

Figure S3 (cont.)

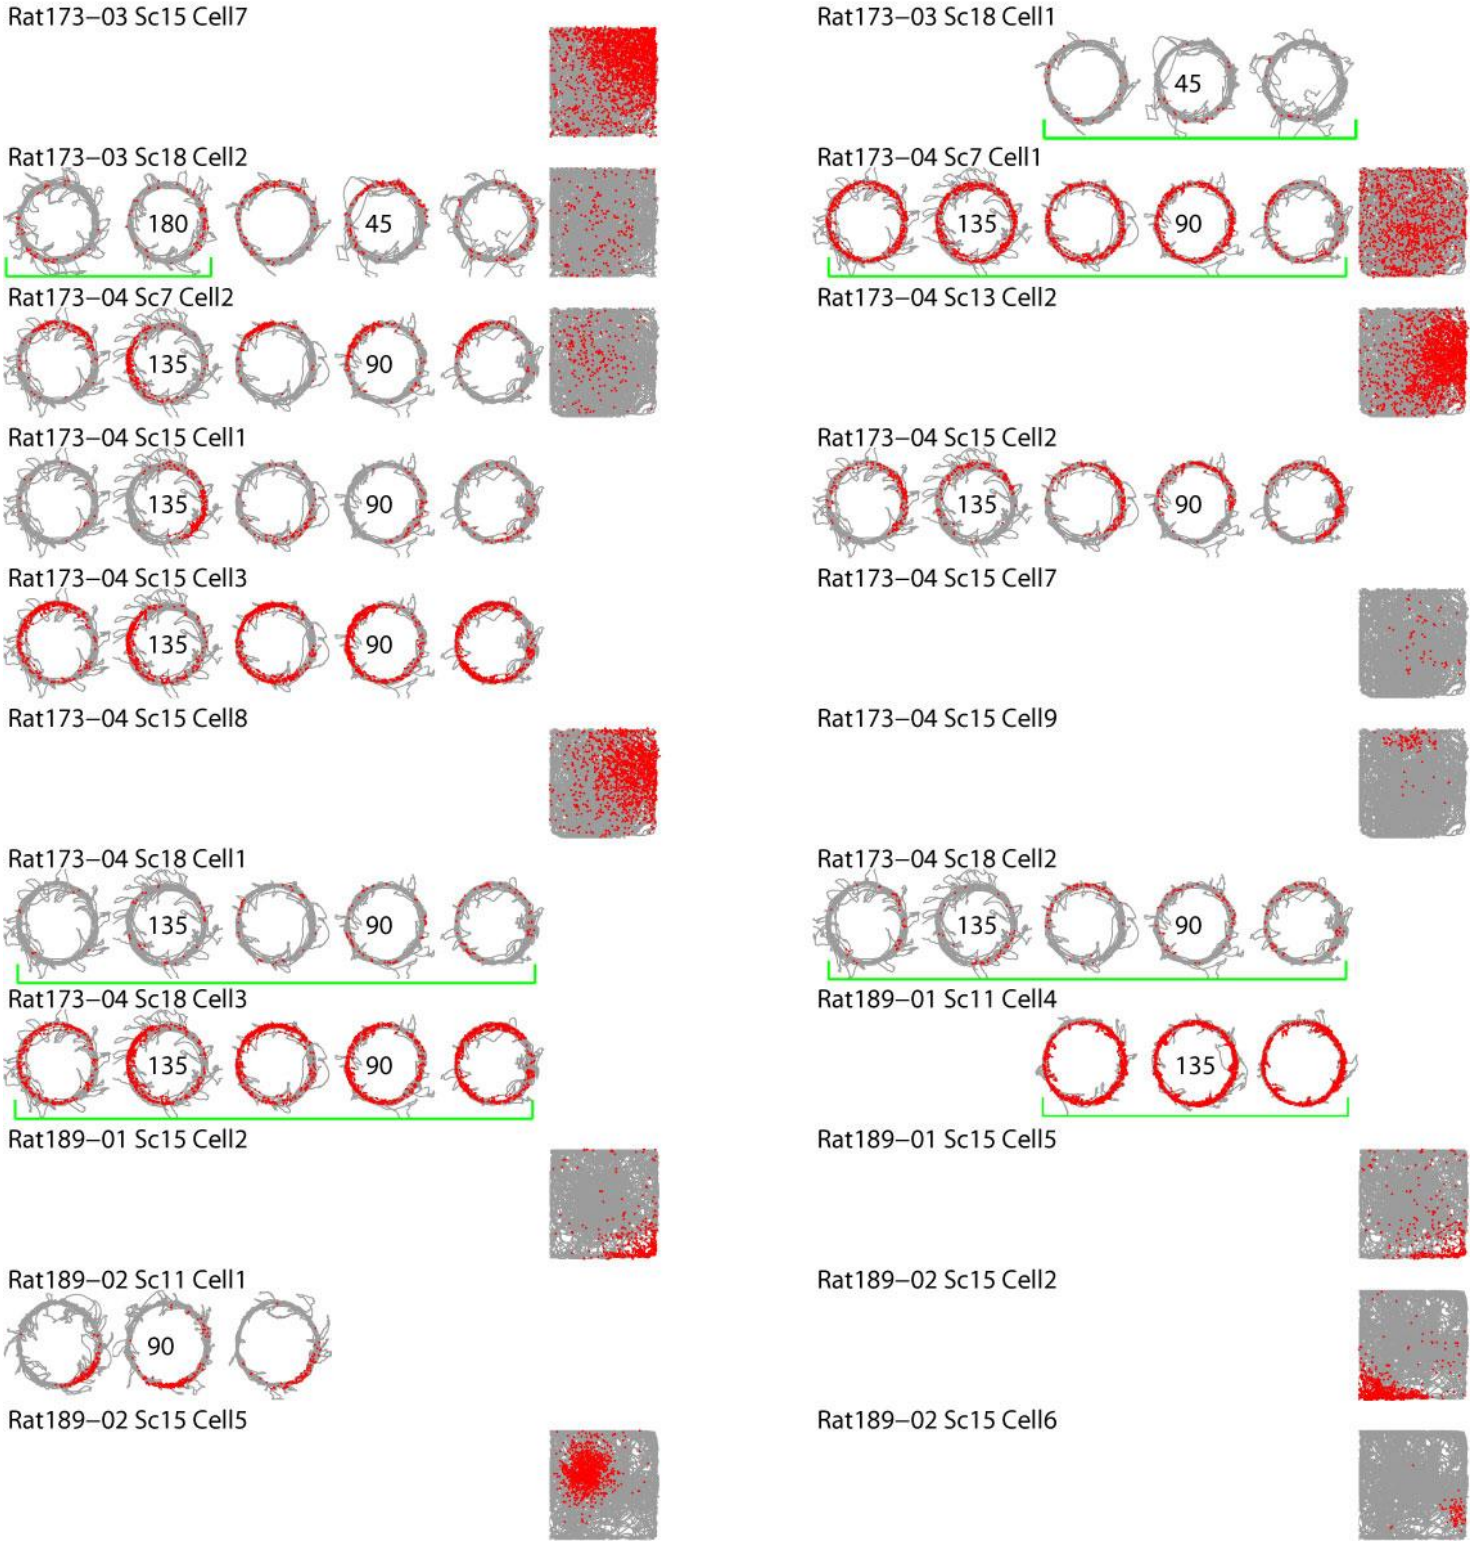

Figure S3 (cont.)

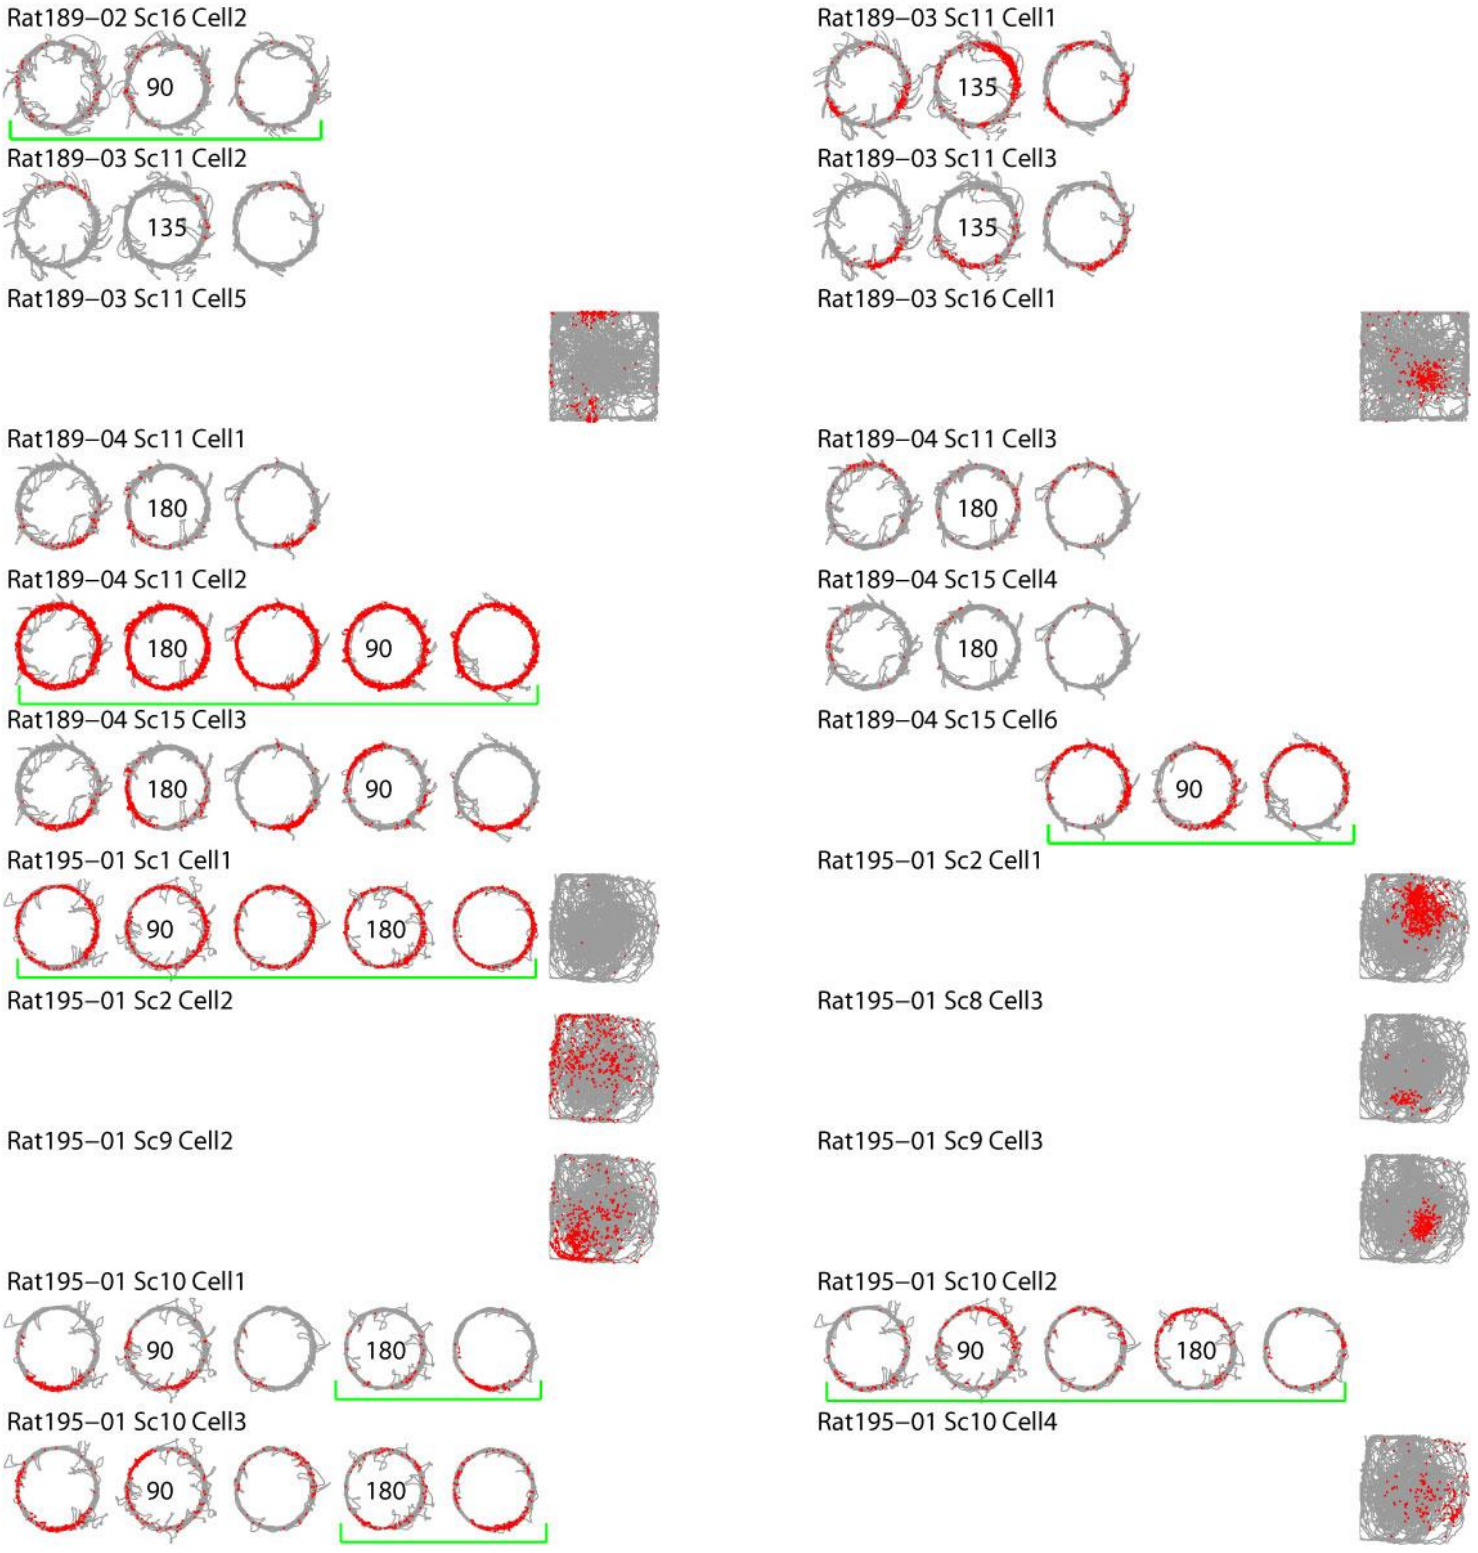

Figure S3 (cont.)

Figure S3 (cont.)

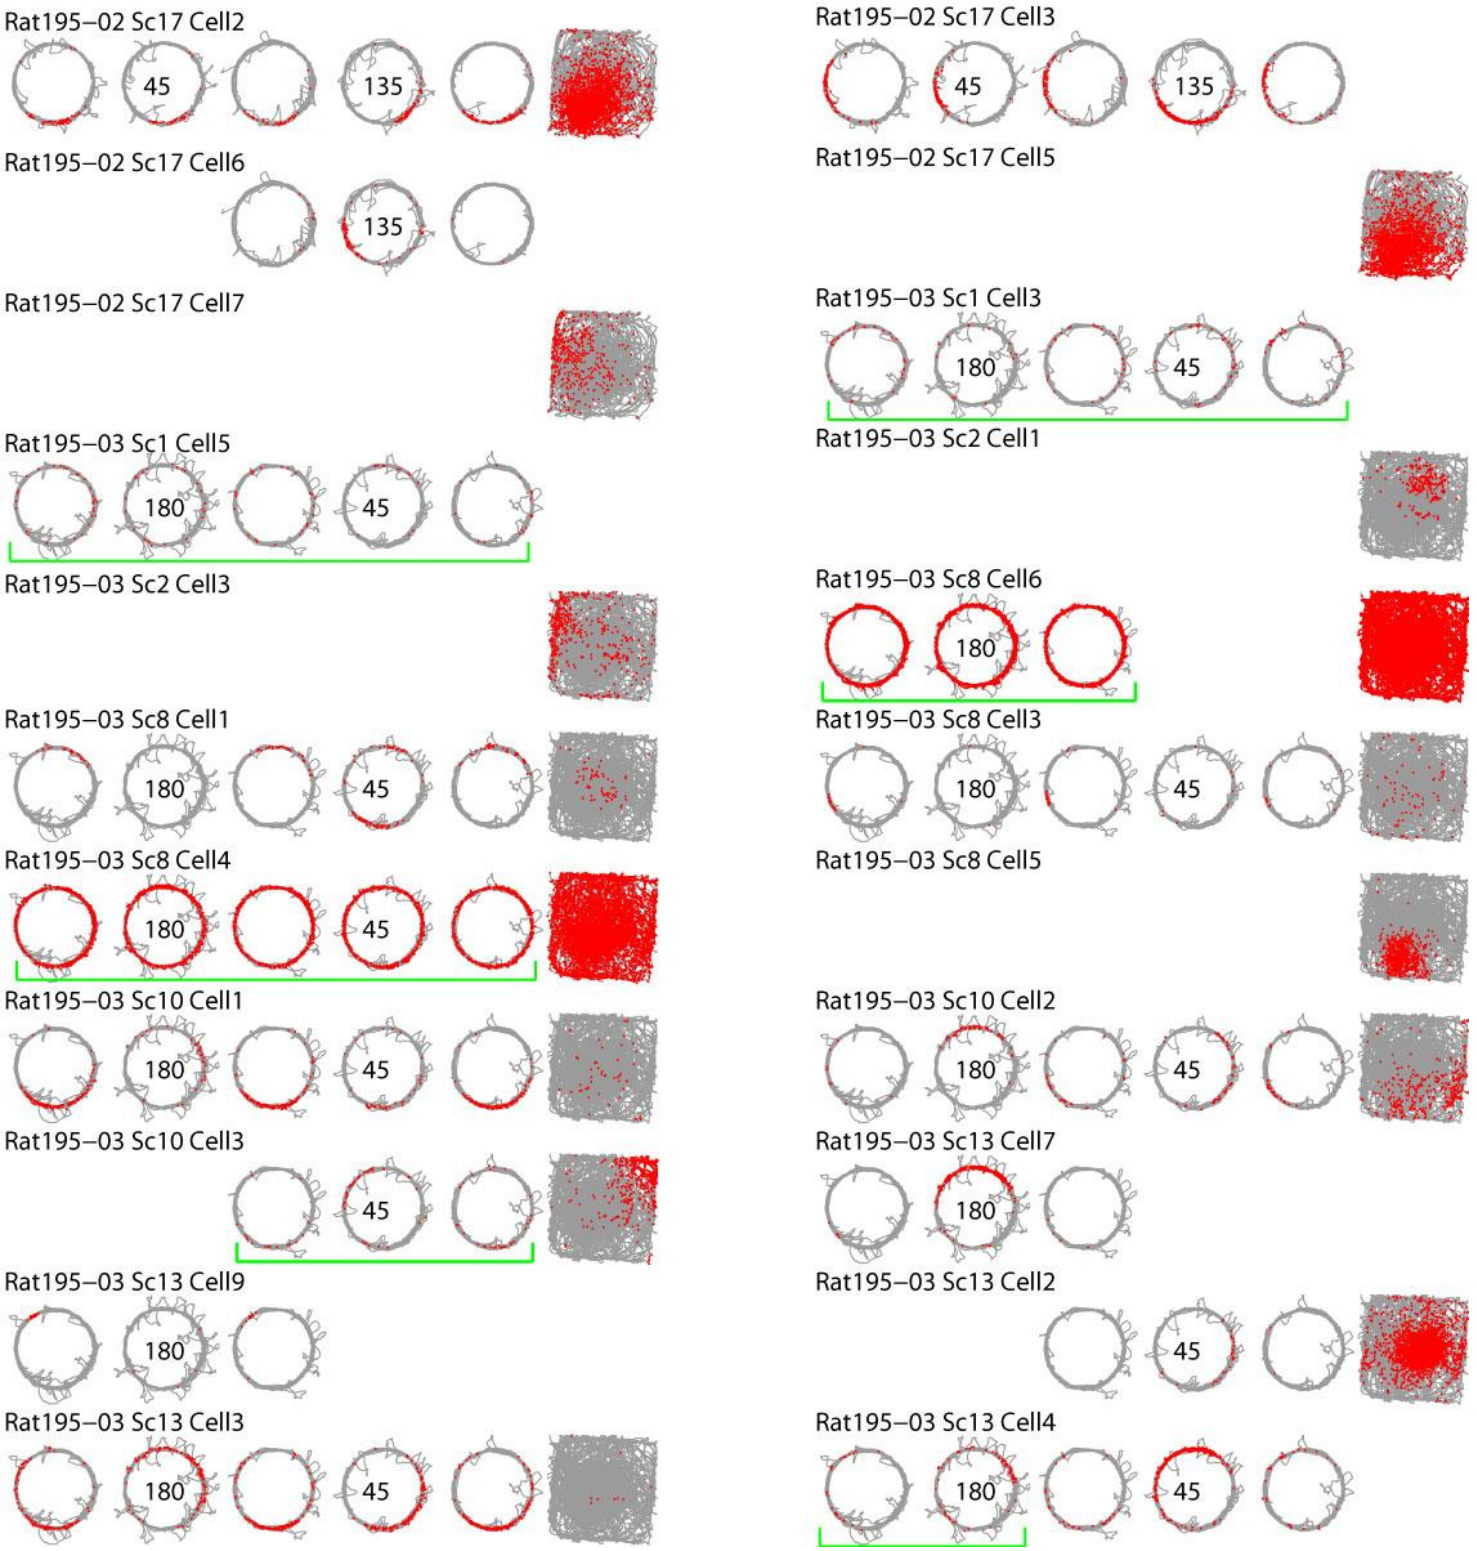

Figure S3 (cont.)

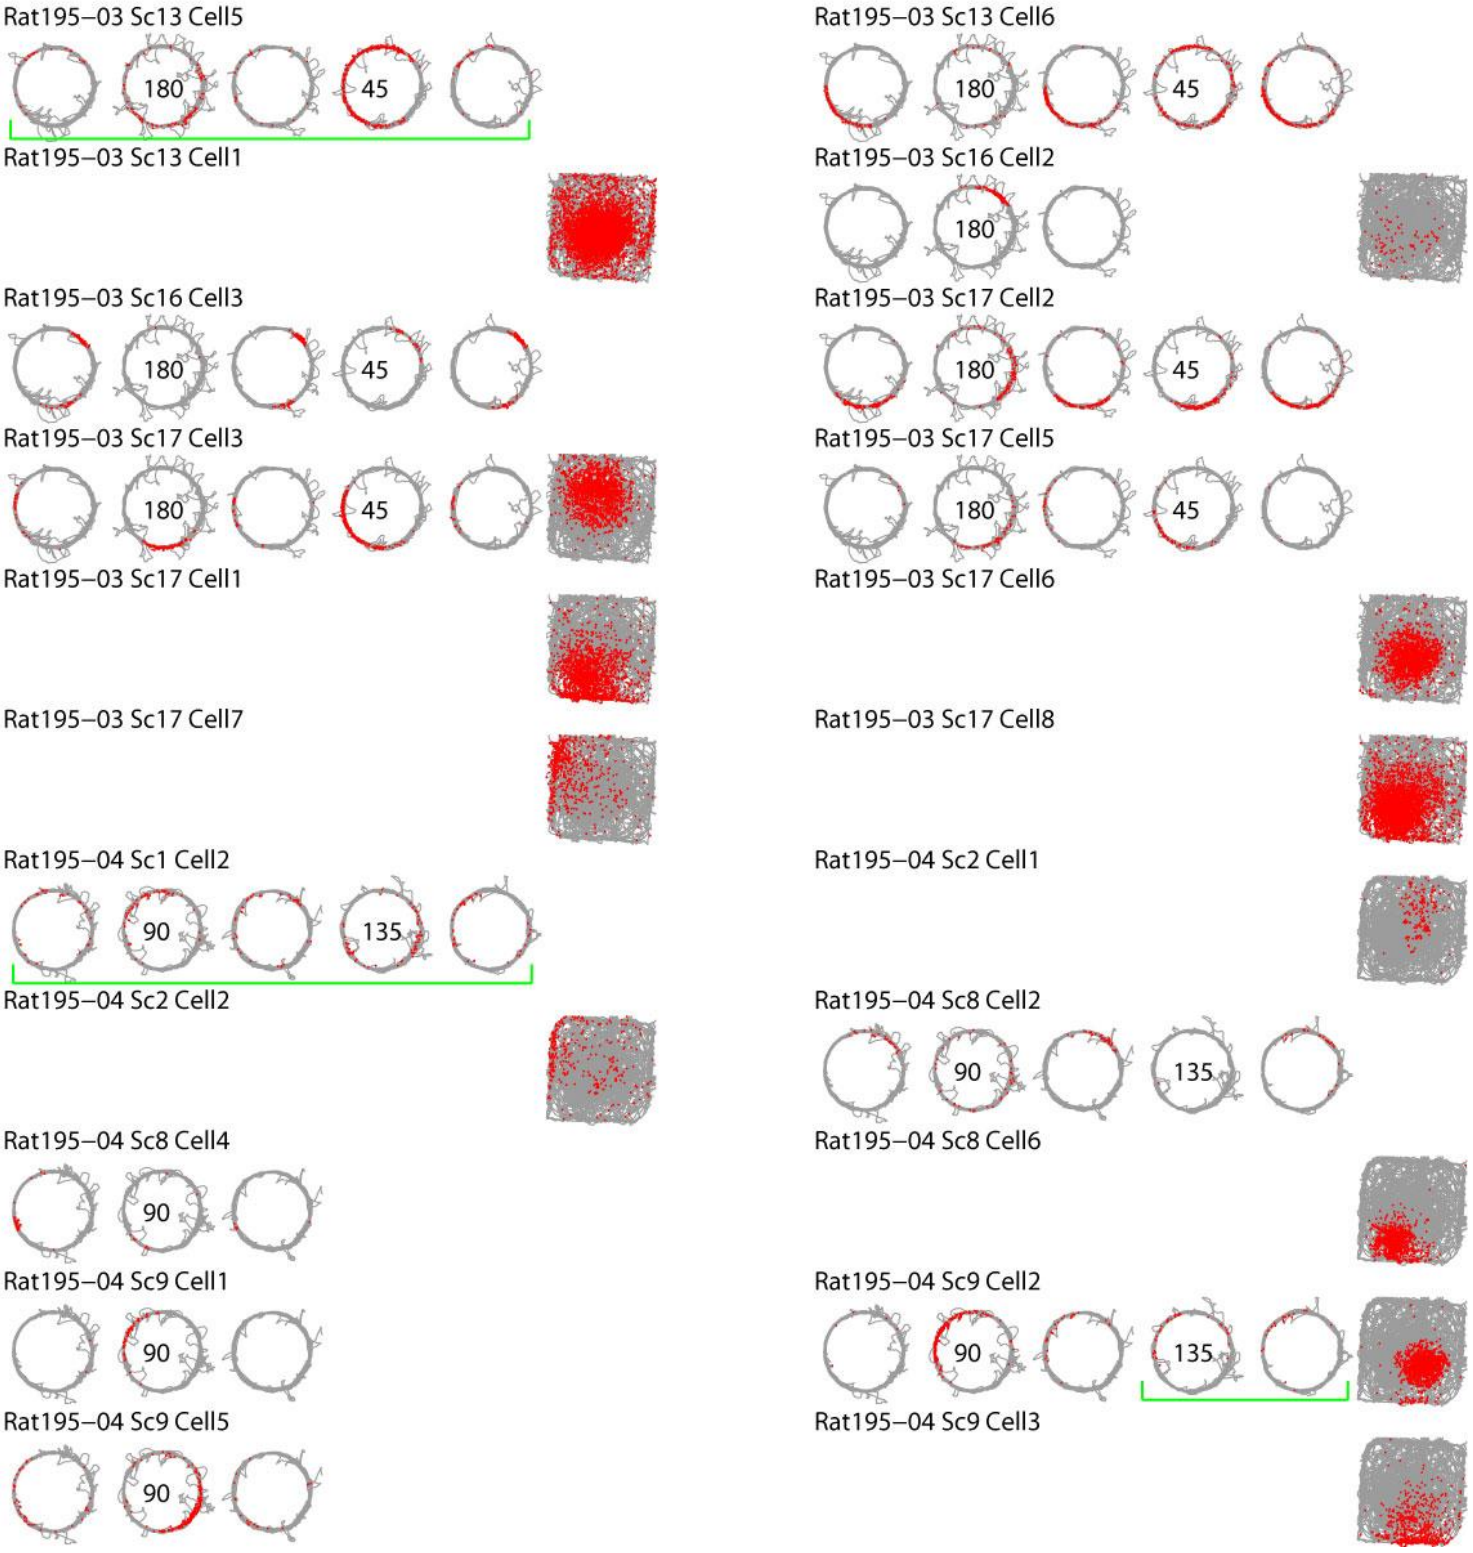

Figure S3 (cont.)

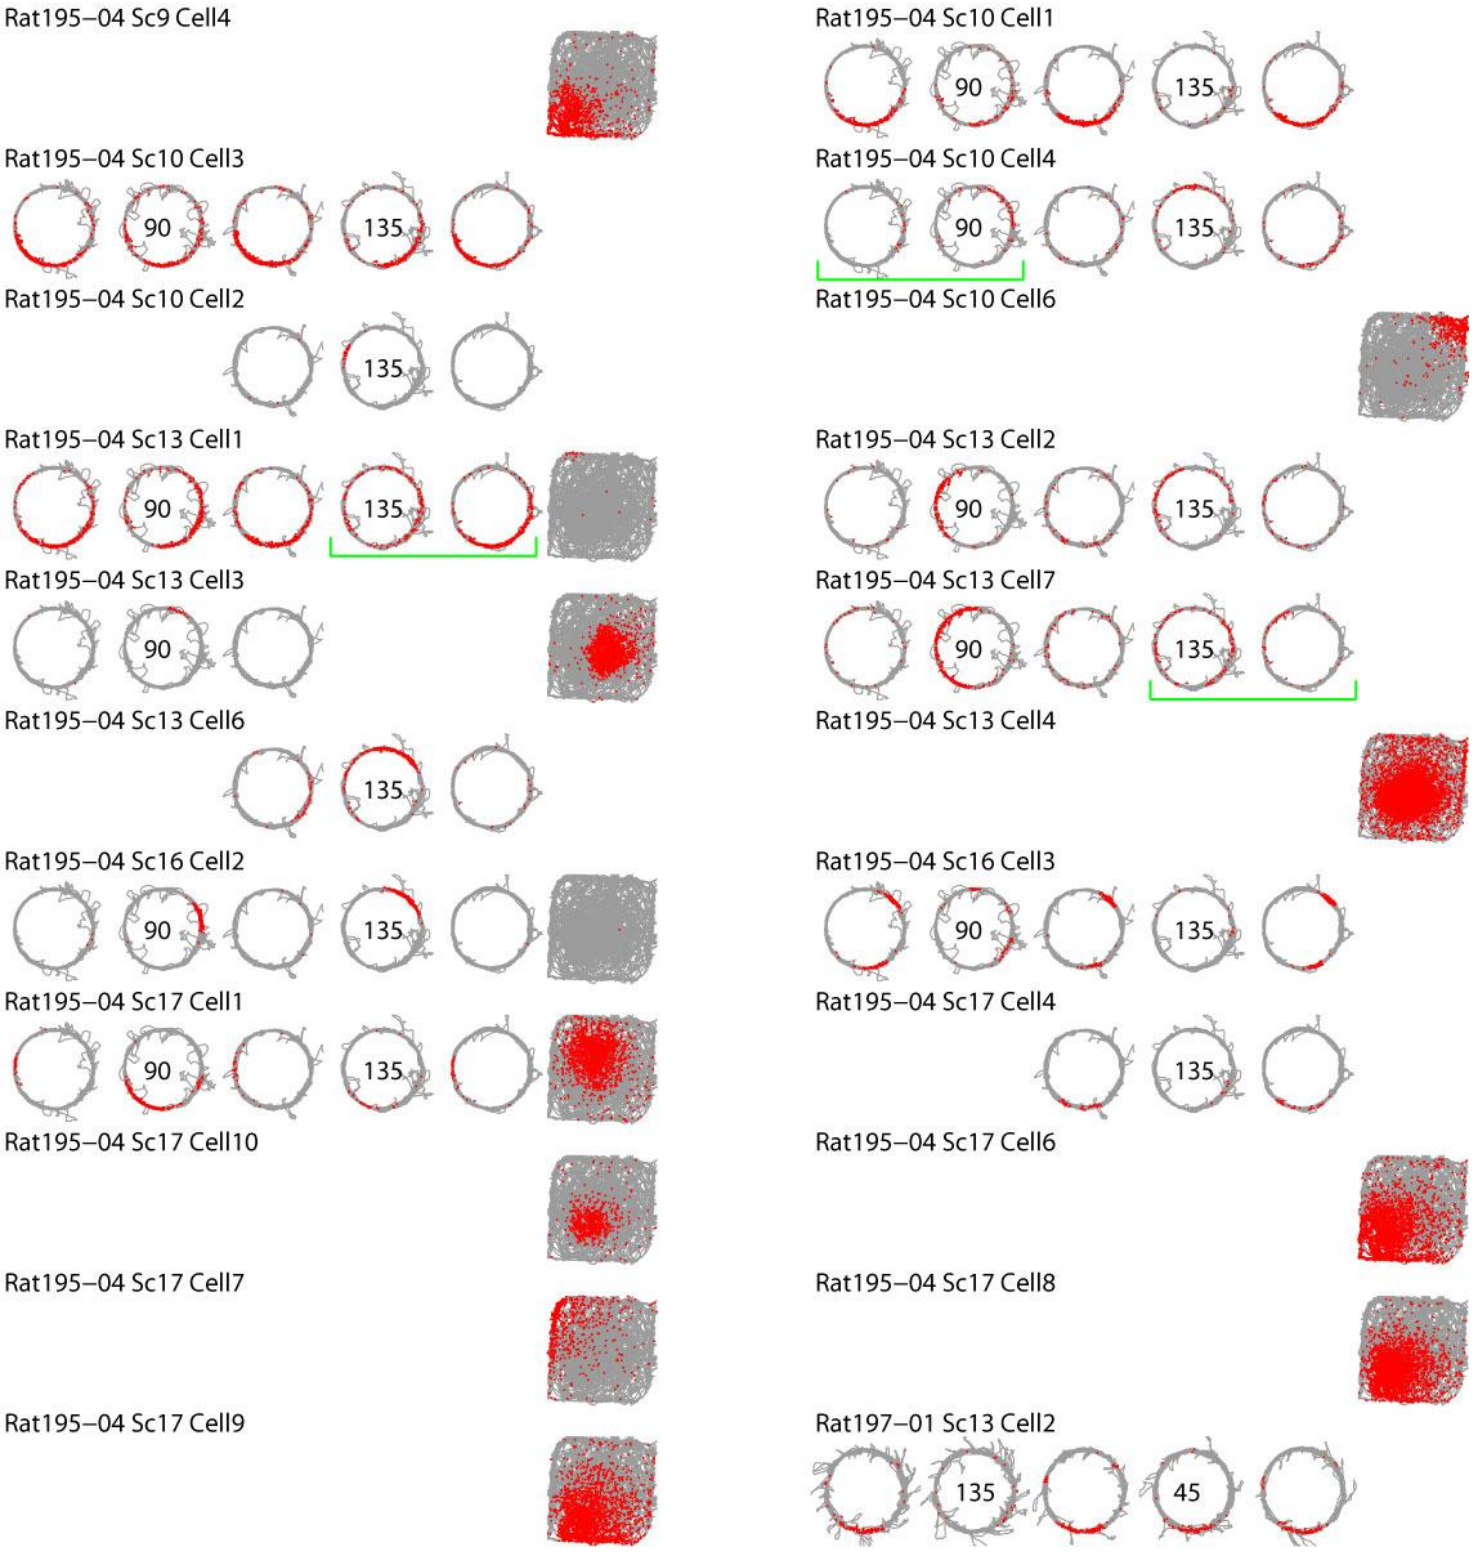

Figure S3 (cont.)

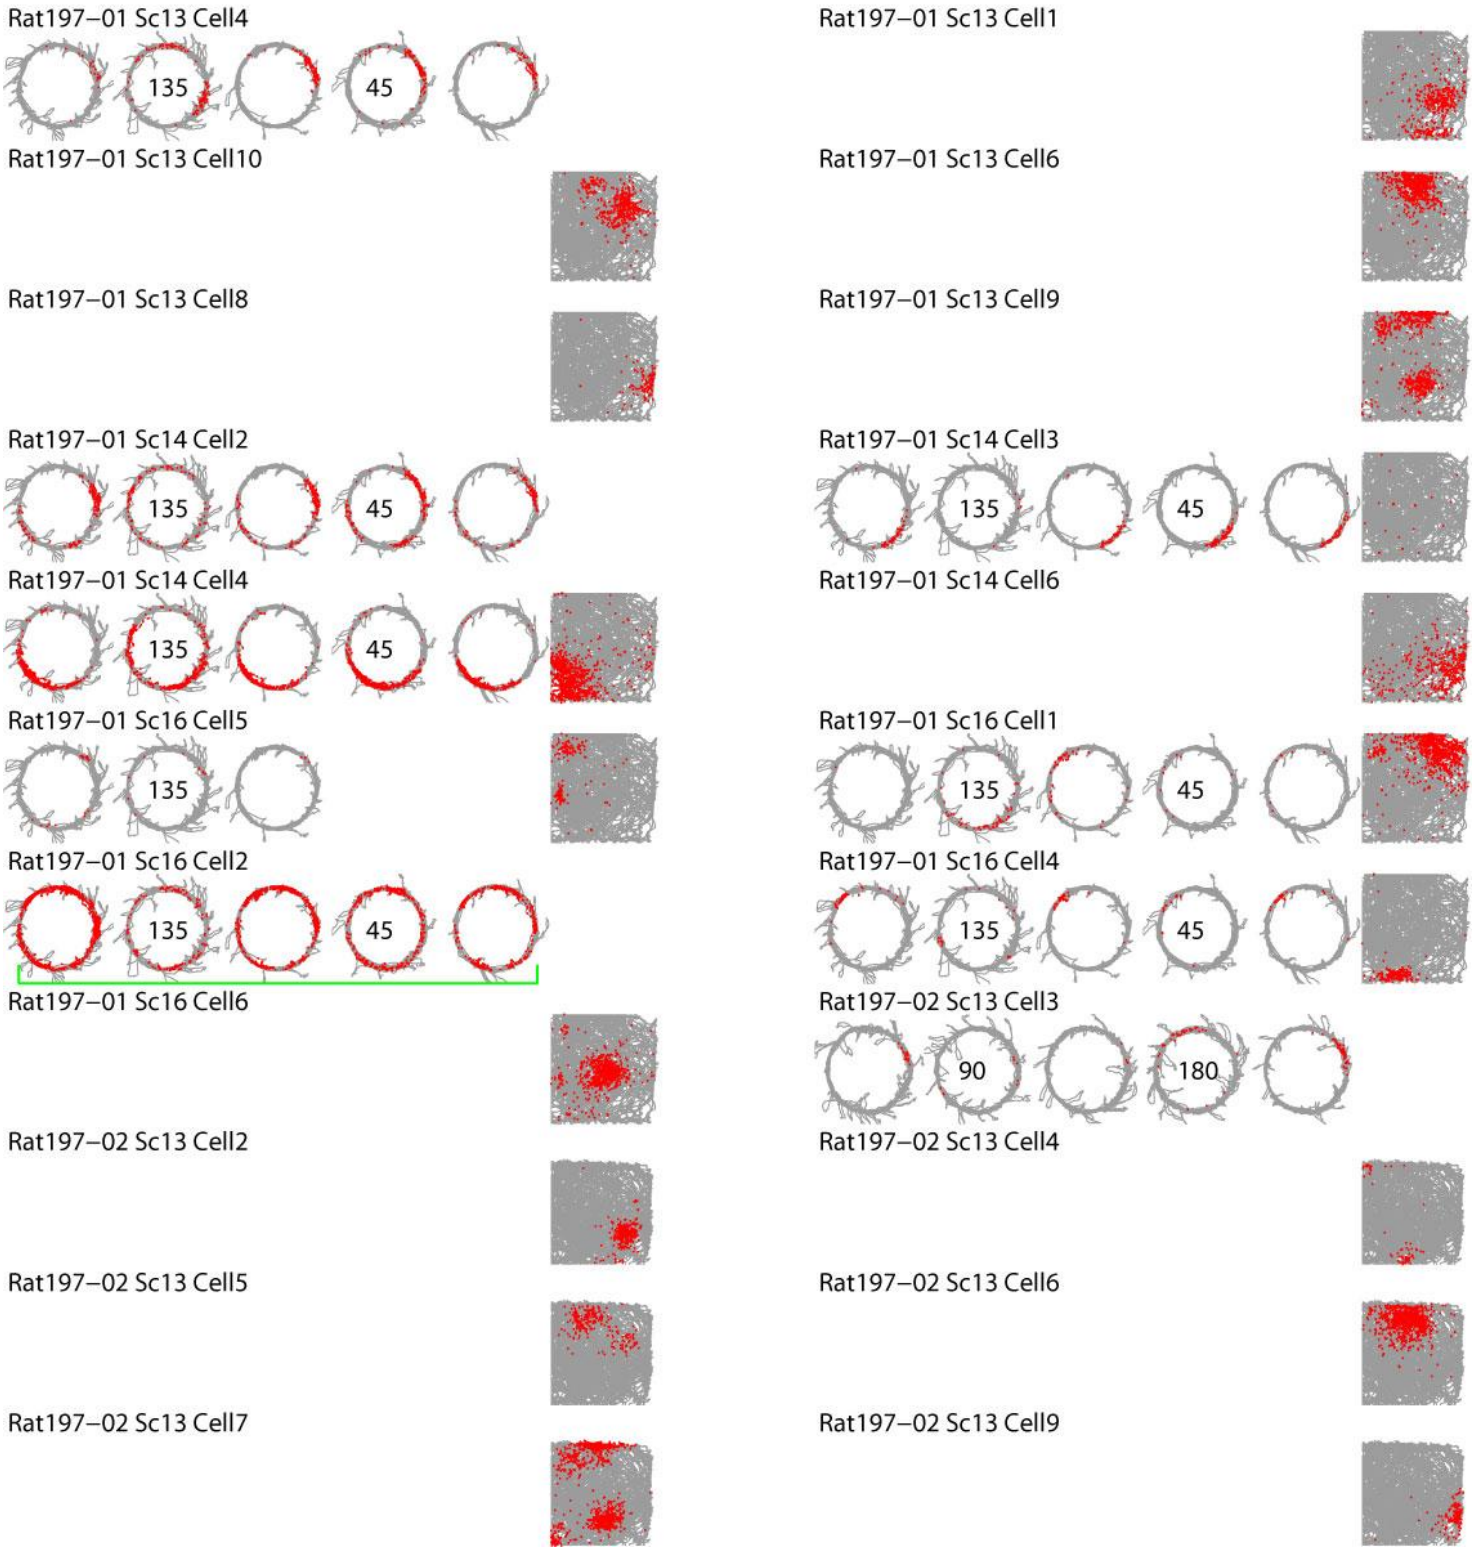

Figure S3 (cont.)

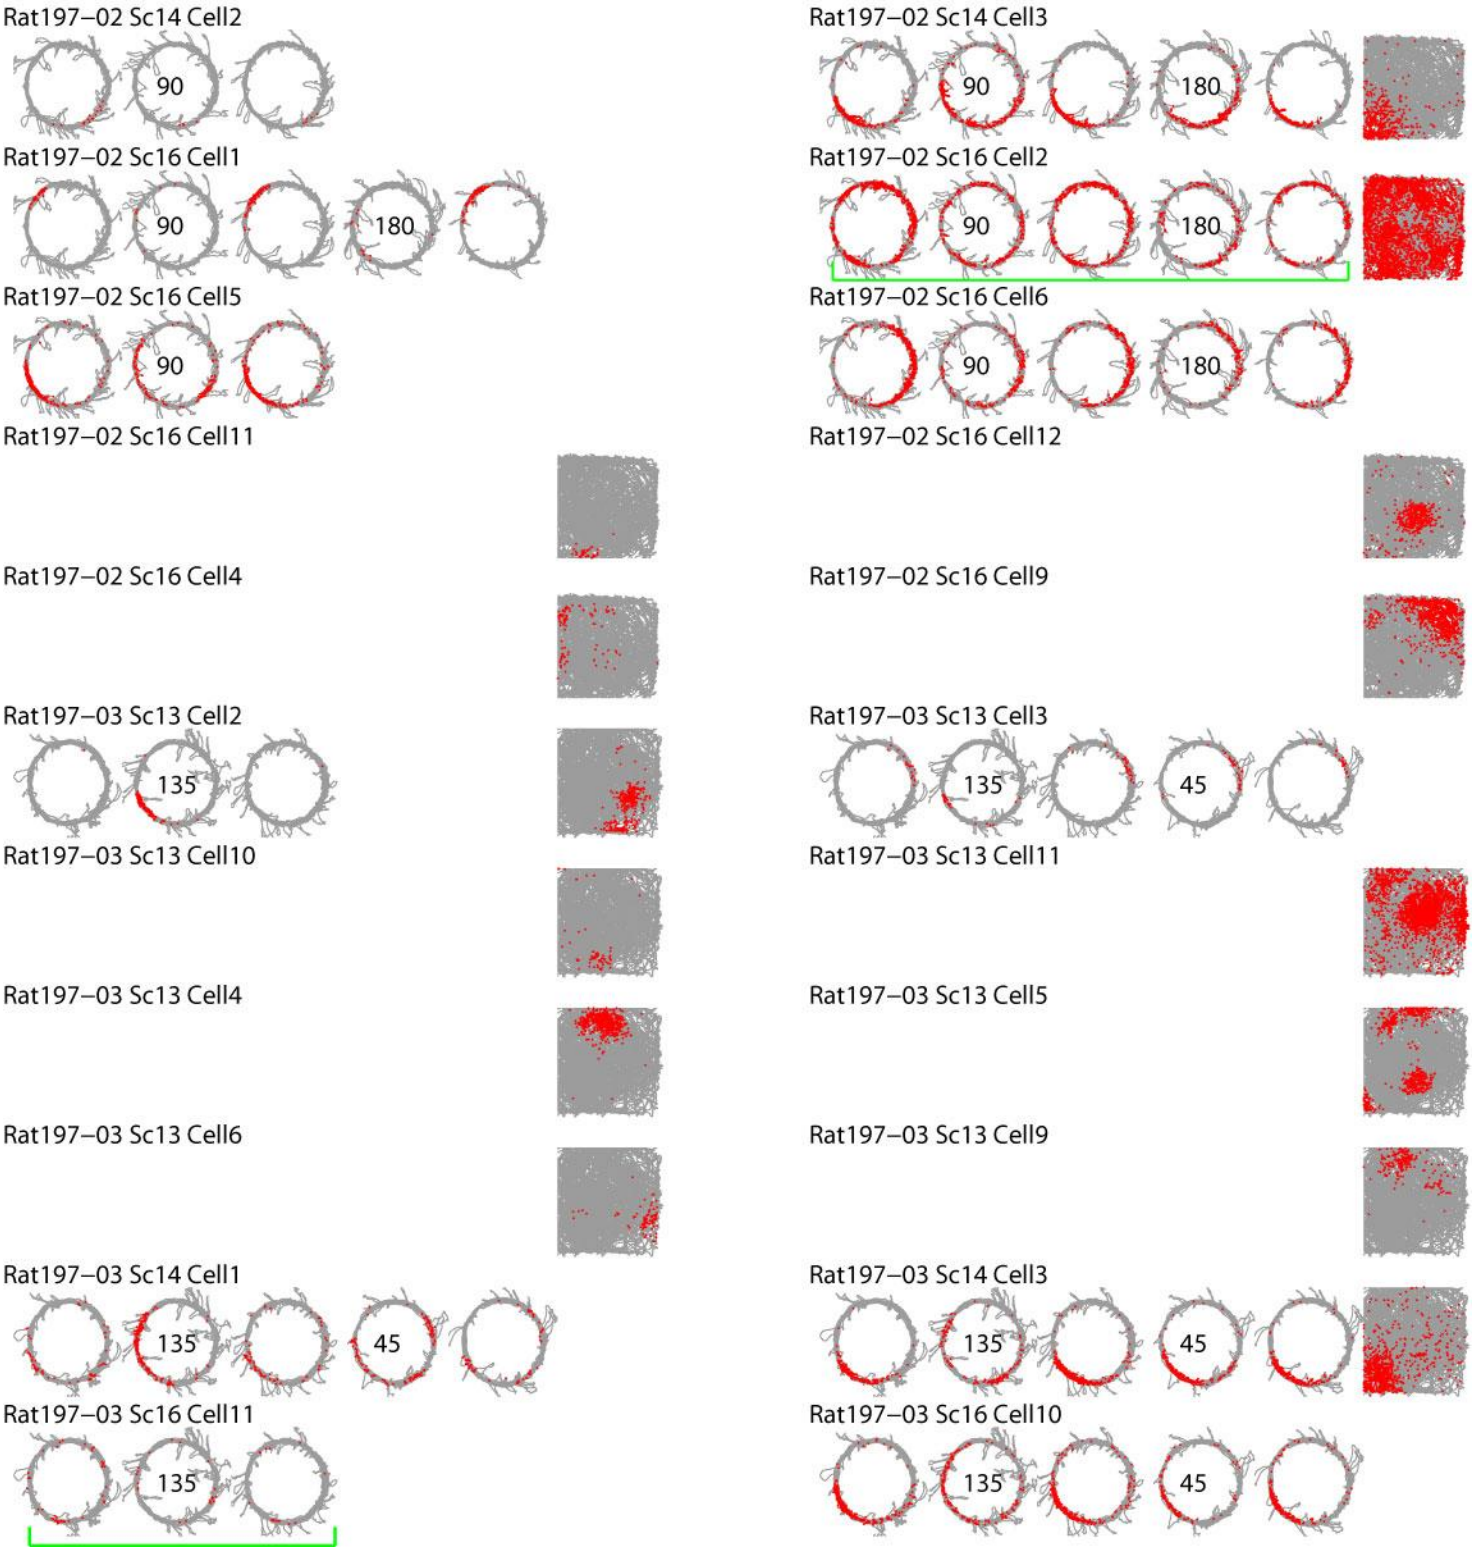

Figure S3 (cont.)

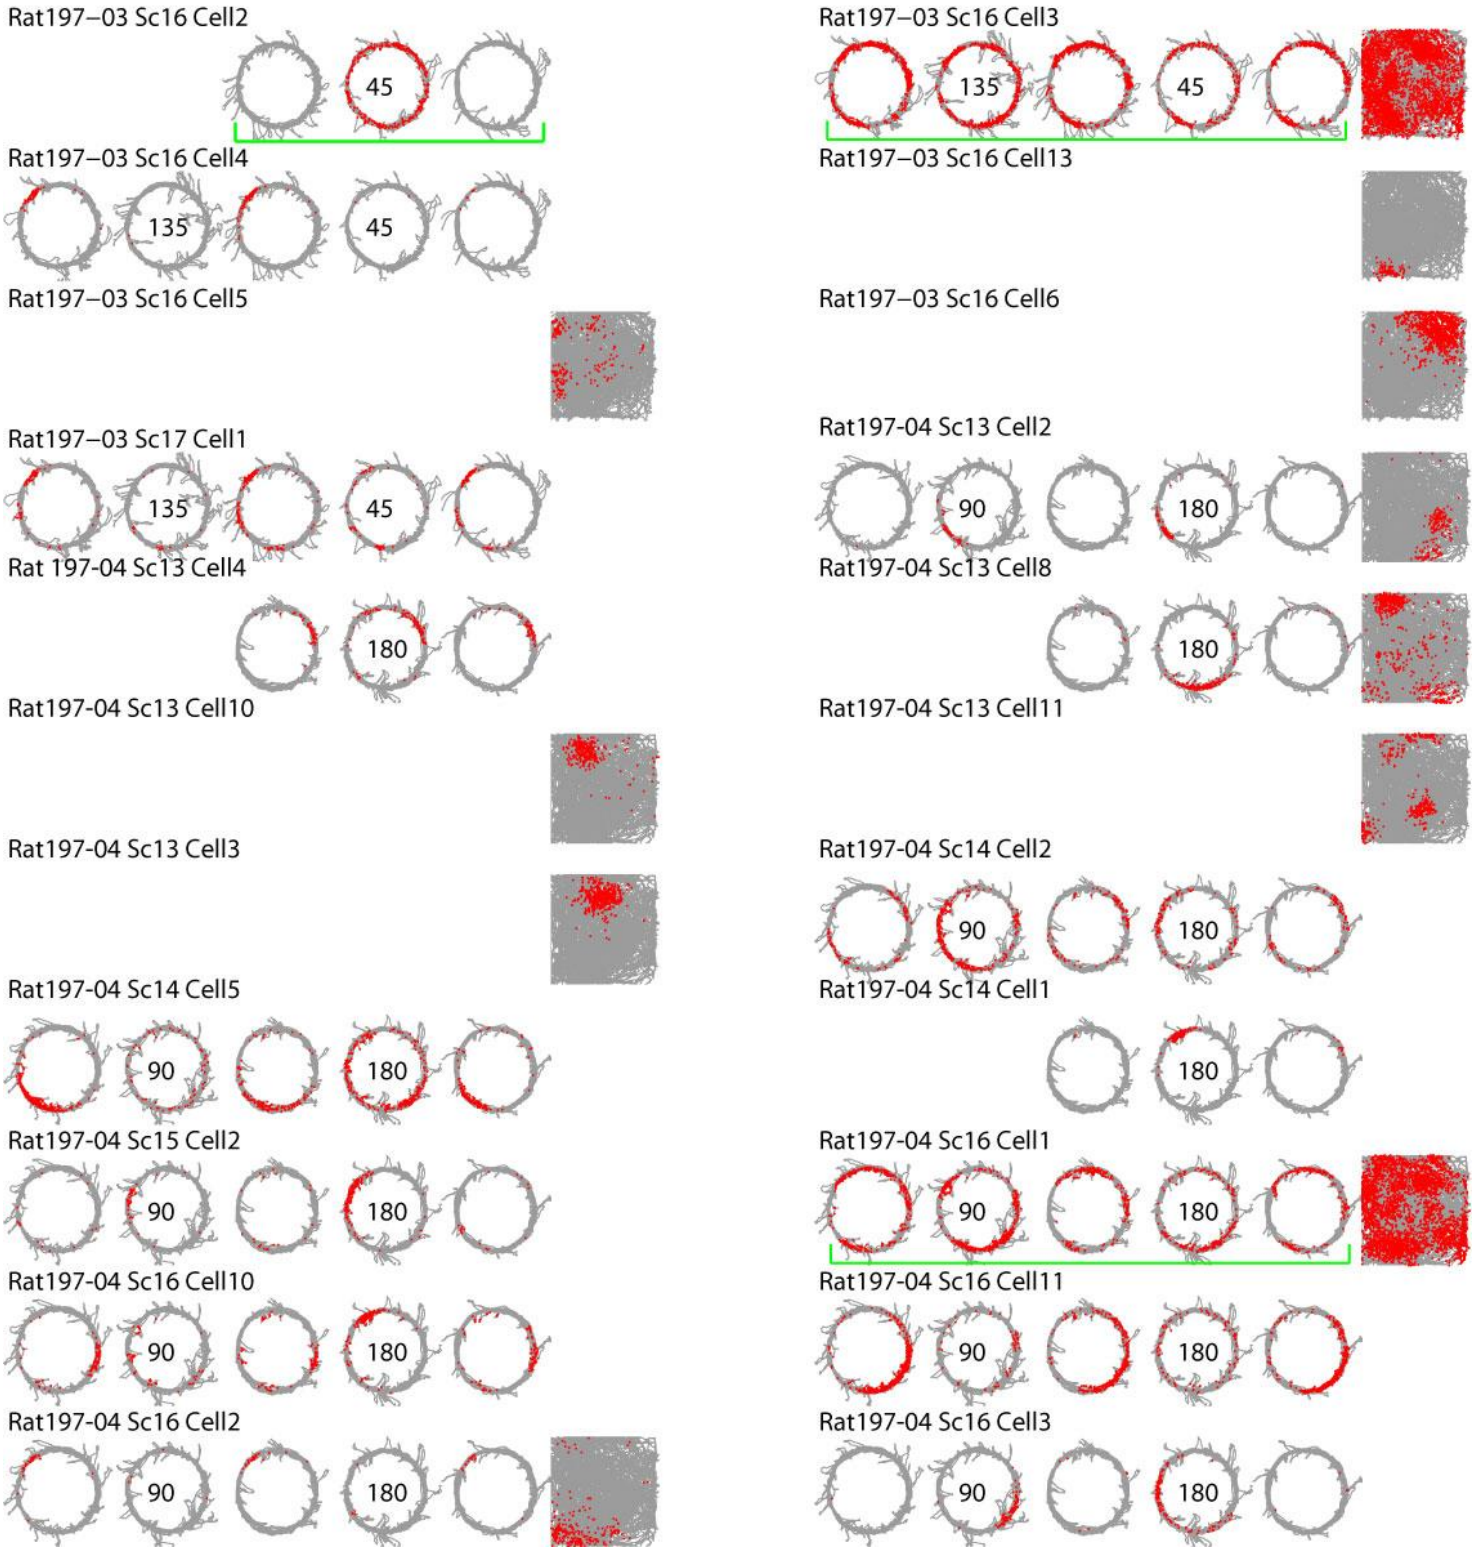

Figure S3 (cont.)

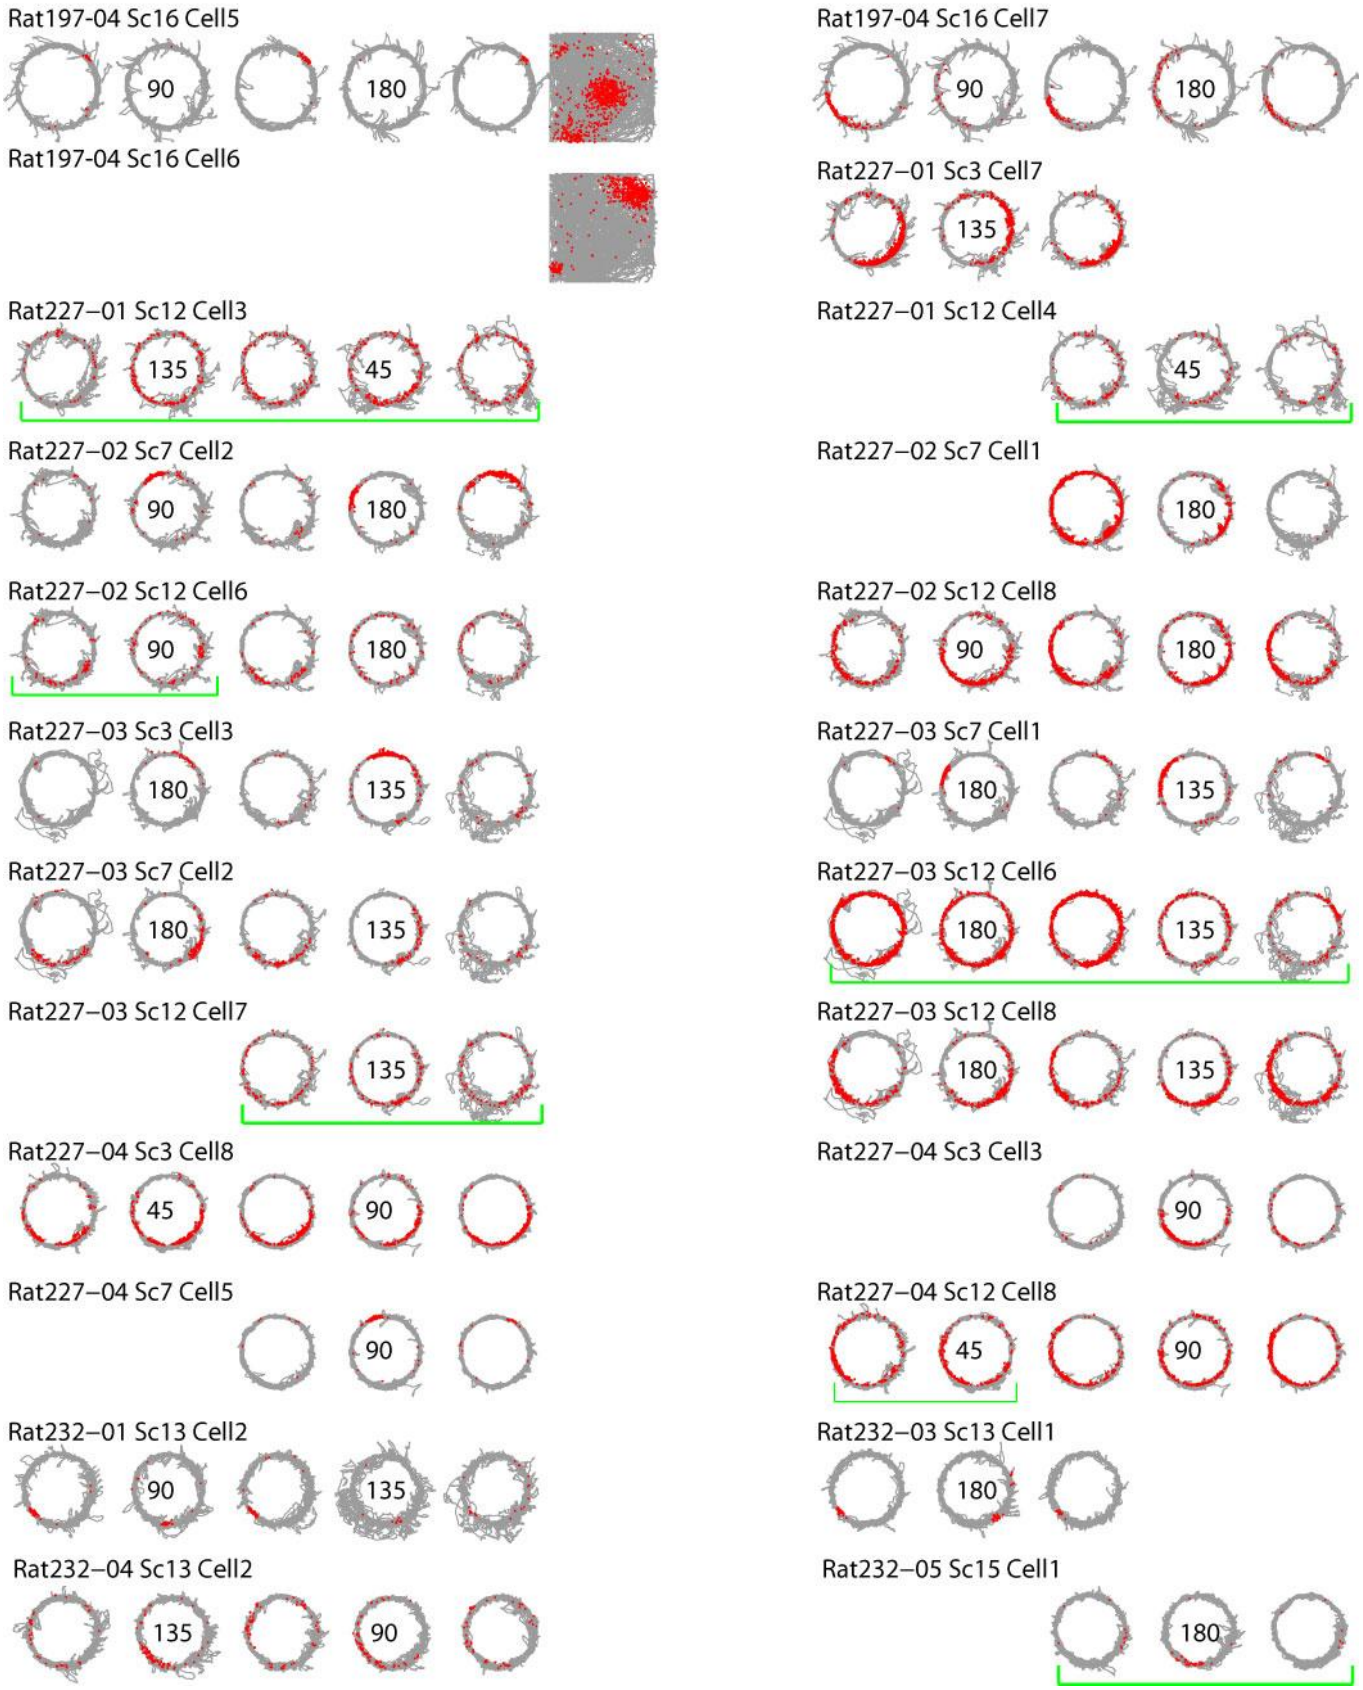

Figure S3. Spatial firing plots for the CA3 cells that met activity criteria, related to Figure 2. This figure is organized identically to Figure S1, which showed similar data for the sample of DG cells. The figure shows the rat's trajectory (gray lines) and locations where the spikes fired (red dots) for the 5 circular track sessions and the open-field foraging session for cells that met minimum activity criteria (≥ 20 spikes, mean firing rate < 10 Hz) for the main analyses of the double rotation manipulation (Figure 4, columns 1-2). For other cells, only the foraging sessions are shown; these are cells that did not meet activity criteria in any of the circular track sessions but met the following criteria in the foraging session: > 75 spikes, info > 0.5 bits/spike, and p value ≤ 0.01. Not shown are any of the cells that did not meet activity criteria in any of the 6 sessions of the day, but were active during the baseline sleep sessions. (See caption for Figure S1 for further details.) Green lines at the bottom of the circular track data indicate the cells that did not meet the spatial criteria for the circular track sessions (information score > 0.5 bits/spike, $p < 0.01$); these cells were not included in the analyses of Figure S5. To reduce the size of the supplemental material, we did not plot the firing rate maps as we did in Figure S2 for the DG cells. Note that the circular track and the open field are not drawn to scale (see Fig. S1). Cell identification key: RatX-Y ScM CellN, where X = rat number, Y = recording day, M = tetrode number, and N = cell number.

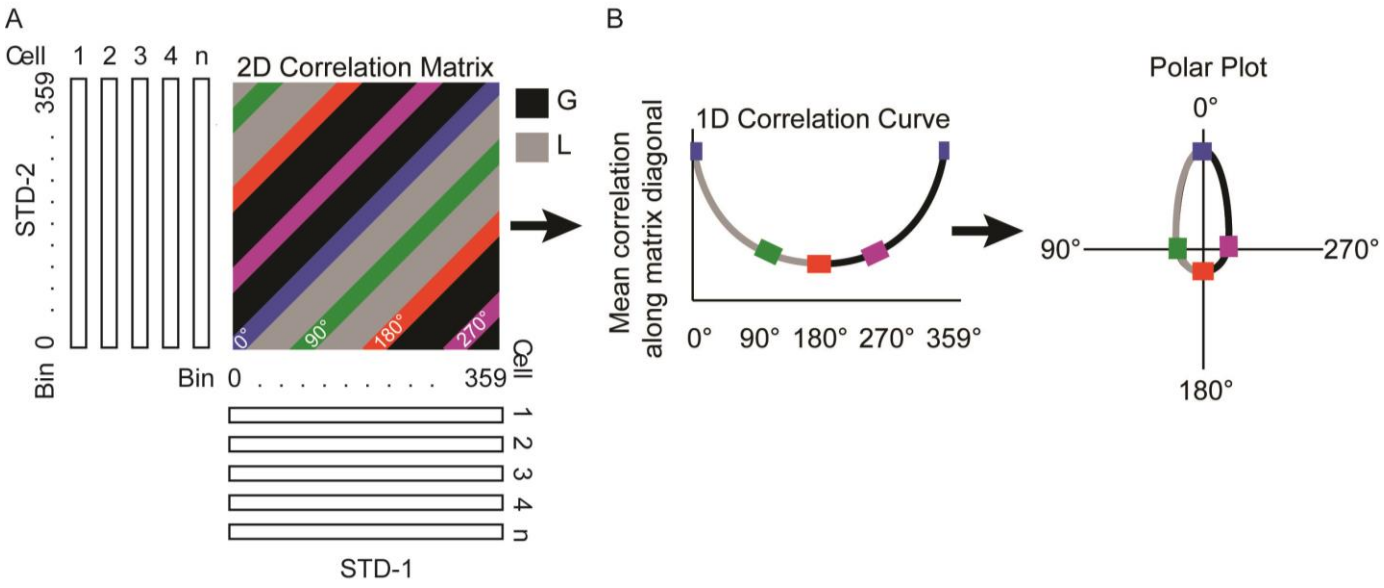

Figure S4. Creation of population correlation matrices, related to Figure 4. (A) Normalized population firing rate vectors are created for the sample of cells at each 1° bin of the track. The firing rate of each cell in each bin is normalized to that cell's peak firing rate. The correlation matrix contains the Pearson product moment correlation values for each of the 360 x 360 firing rate vectors. A band of high correlation along the 0° main diagonal (blue) indicates that the representation maintained coherence and did not rotate in the world-centered (i.e., video camera) frame of reference. A band of high correlation along the 90° diagonal (green) indicates a coherent representation that rotated 90° CCW, whereas a band of high correlation along the 270° diagonal (magenta; note that the band wraps around the matrix) indicates that the representation rotated 90° CW. Because the local cues always were rotated CCW and the global cues CW, bands that fall in the gray or black areas of the matrix indicate that the representations were controlled more strongly by the local cues or global cues, respectively. (B) The 2D correlation matrices were transformed to 1D linear plots by averaging the correlations along each diagonal of the matrix. The linear plots were then represented in polar coordinates in order to visualize the rotation angles. All figures are reproduced from Neunuebel et al. (2013).

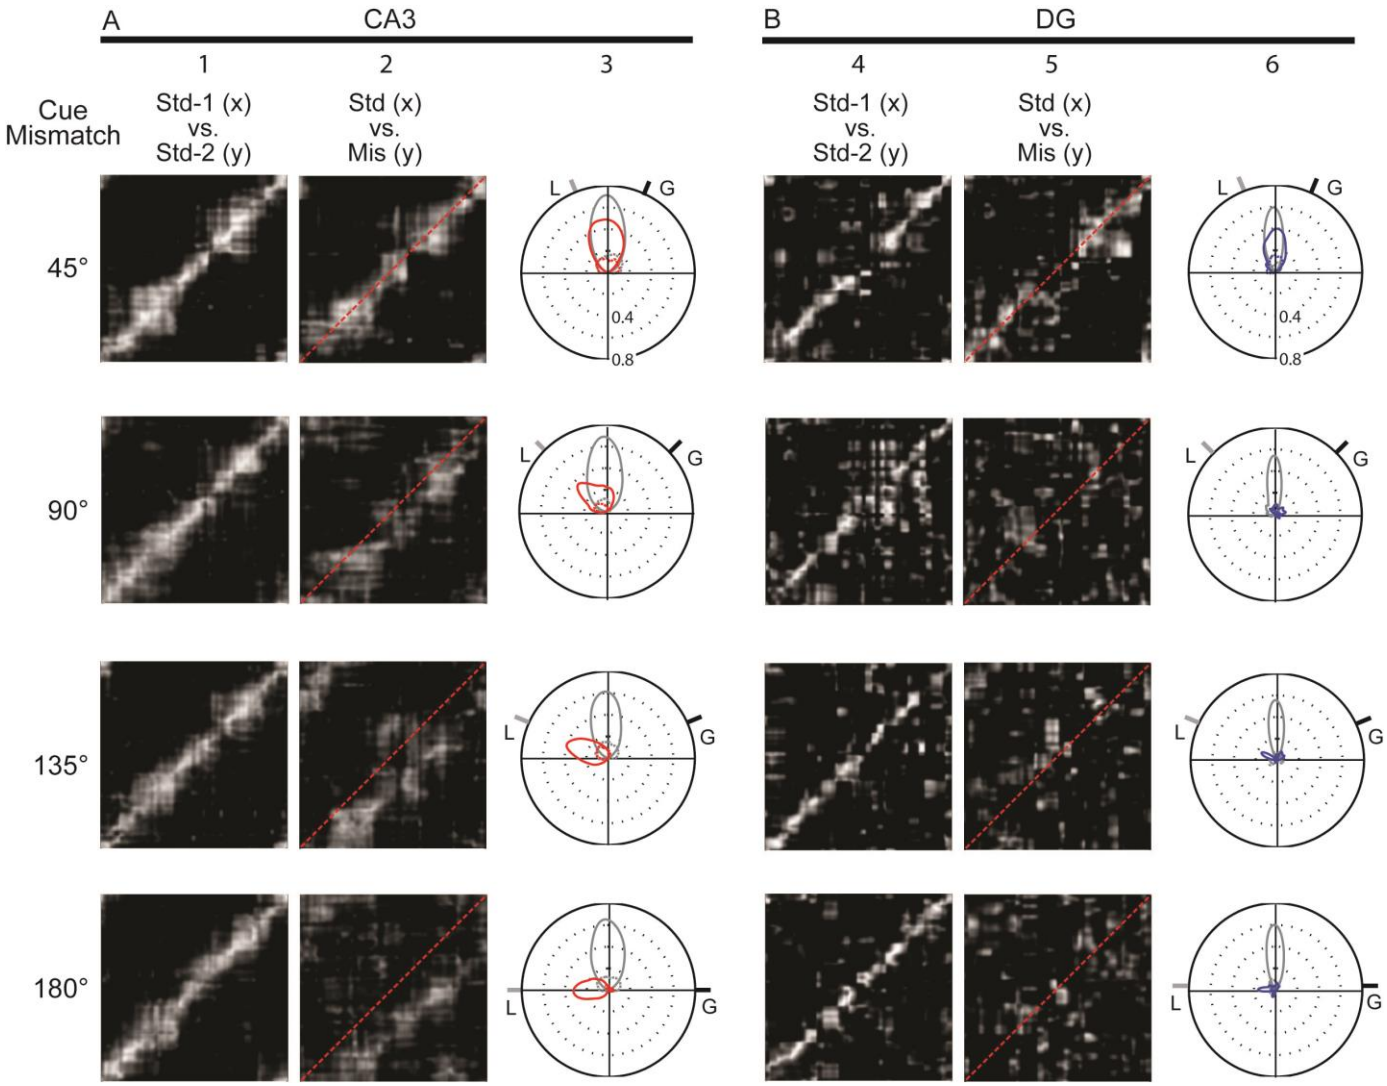

Figure S5. Population responses of CA3 and DG cells that met spatial information selection criteria, related to Figure 4. This figure presents the population correlation analyses for cells that met the inclusion criteria for the main analyses plus two additional criteria to ensure that only spatially selective cells were analyzed. The full criteria were that, in at least one of the sessions being compared: (a) the cell fired more than 20 spikes; (b) the mean firing rate was < 10 Hz; (c) the spatial information score was > 0.5 bits/spike; and (d) the information score was statistically significant ($p < .01$). (A) CA3 correlations. Column 1 shows the correlation between the population vectors for the Standard 1 vs. Standard 2 conditions. See Figure S4 for description of how these plots are generated. Column 2 shows the correlation between the population vectors for the Standard vs. Mismatch conditions. The correlation plots look very similar to those in the main text (Figure 4), except that the results are cleaner due to the exclusion of cells that do not meet the spatial firing criteria. Column 3 shows the 1-D polar plot representations of the 2-D correlation matrices. The gray curves represent the Std-1 vs. Std-2 conditions and the red curves represent the Std vs. Mis conditions. Note that in

these correlation plots, a number of the diagonals showed negative correlations, which are indicated by the dashed lines. These plots reinforce strongly that the CA3 representation maintained a large degree of coherence in the mismatch session, and that the representation was primarily controlled by the local cues (at least for the mismatch sessions > 45°). (B) DG correlations. The format is the same as Part A, except that the polar plots for the Std vs. Mis conditions are colored blue. These results confirm the conclusions of the main analysis, with cleaner correlation plots due to the exclusion of the spatially nonselective neurons. Note that the Std-1 vs. Std-2 correlations are much cleaner than the plots shown in Figure 4 of the main text (evidenced clearly by the narrow tuning curve on the polar plots). For the Std vs Mis 45° plot, the correlations are also high and centered near 0°, although the tuning curve is broader than the Std-1 vs Std-2 curve. For the mismatch angles > 45°, the DG correlation breaks down almost completely. These data provide strong evidence that the DG representation is severely disrupted by the mismatch manipulation, whereas the CA3 representation is able to reproduce the representation of the standard environment with a great deal of accuracy even with the severely corrupted DG input.

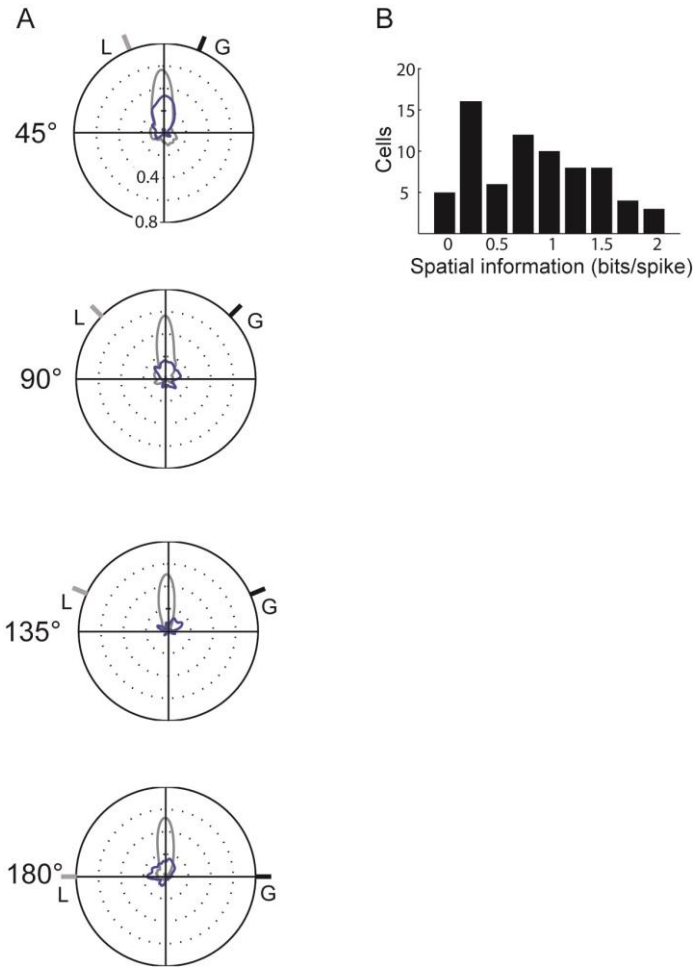

Figure S6. Population responses of low rate cells from the DG to cue-mismatch manipulations, related to Figure 4. (A) The analyses in the main text used all cells that fired < 10 Hz, for consistency with the criteria for CA3 and the criteria in our prior publication of the MEC and LEC (Neunuebel et al., 2013). However, the DG showed 3 classes of cells based on firing rate, a low-rate group (< 2 Hz), a medium-rate group (2-10 Hz), and a high-rate group (> 10 Hz) (Fig. 1). The medium-rate group showed spatial firing rate maps that were more active and on average less spatially specific than the low-rate cells (although some of them showed strong spatial tuning; Figure S2)**.** Only 15 medium-rate cells met inclusion criteria for the analyses, but it was important to determine if excluding these cells changed the pattern of results. The correlation matrices were shown in Figure 4 of the main text. Here we show the polar plots that were created from the spatial correlation matrices. The population activity between standard 1 vs. standard 2 (grey lines) show that the low-rate cells had a very tight distribution of high correlations centered on the 0° rotation angle. This result indicates high stability of the spatial firing between the standard sessions. The correlations between the standard and mismatch sessions (blue lines) were much smaller than the Std-Std correlations, especially for the mismatch angles > 45°. These correlation plots demonstrate the dramatic decorrelation of the DG representation between the standard

and mismatch sessions. The lines at the edges of the plots show the amount of local (L; grey) and global (G; black) cue rotations. (B) Histogram of information scores for the DG cells that fired < 2 Hz. The median spatial information of this sample is not statistically different from the sample of CA3 cells shown in Figure 1F (Mann-Whitney U-test, $Z = -0.7$, $p = 0.488$).