

Remembered Reward Locations Restructure Entorhinal Spatial Maps

William N. Butler¹†, Kiah Hardcastle¹†, Lisa M. Giocomo¹*

¹Department of Neurobiology, Stanford University School of Medicine, Stanford, CA, USA.

*Corresponding author, e-mail: giocomo@stanford.edu.

5 †William N. Butler and Kiah Hardcastle contributed equally to this work.

Abstract: Ethologically relevant navigational strategies often incorporate remembered reward locations. While neurons in the medial entorhinal cortex provide a map-like representation of the external spatial world, it remains unknown if this map integrates information regarding learned reward locations. We compared entorhinal coding during a free foraging versus spatial memory task. Entorhinal spatial maps re-structured to incorporate a learned reward location, which in turn improved positional decoding near this location. This finding indicates that different navigational strategies drive the emergence of discrete entorhinal maps of space and points to a role for entorhinal codes in a diverse range of navigational behaviors.

One Sentence Summary: Training rats to perform a spatial memory task drives the emergence of a new entorhinal spatial map that incorporates a remembered reward location and improves positional decoding at that location.

The ability to recall and navigate to a remembered reward location is essential to survival. The hippocampus and medial entorhinal cortex (MEC) contain cells that provide representations of self-location and orientation within the local spatial environment (1-5). Initial experiments suggested a dissociation between representations in these regions: spatially-
5 modulated codes sensitive to contextual features in the hippocampus and context-independent codes for position, orientation and speed in MEC (2, 3, 5-10). In contrast, recent work has shown that MEC spatial codes are flexible and adaptive (6, 11-13). However, these MEC spatial coding features have primarily been observed during random foraging, whereas ethologically relevant strategies often employ more complex behaviors such as goal-directed navigation (14). While
10 MEC plays a critical role in navigation (15), the degree to which remembered reward locations influence MEC neural codes remains unknown.

We recorded neural activity in the MEC and surrounding cortical areas of seven rats as they explored two arenas (1.5 m x 1.5 m) (Fig. S1). In environment one (ENV1; black walls, lemon scent), rats foraged for randomly scattered crushed cereal (2-5, 12). In environment two
15 (ENV2; white walls, vanilla scent), rats navigated to a remembered, unmarked 20 cm x 20 cm zone in response to an auditory cue to receive a food reward (0.5-1 cereal units), and freely foraged for randomly scattered crushed cereal between trials (10) (Fig. 1A,B and S2). Reward trials (cue onset to reward zone entry) occurred ≥ 10 times per session (Fig. 1C). After training (mean # sessions to reach criterion = 15; range = 8-24), animals took rapid, direct paths to the
20 reward zone upon cue onset (Fig. 1D).

We considered the coding features of 778 cells recorded in both environments (Fig. S3). We identified cells as encoding position (P), head direction (H) or running speed (S), then further classified P-encoding cells as grid, border, or non-grid, non-border spatial cells (12). Between

environments, we observed equal proportions of grid and border cells, and cells encoding P, H, or S (Fig. S4A). Stability, information content, average and peak firing rates did not change between environments, apart from the firing rates of grid cells (Fig. S4C-E). Multiple features of local field potential theta oscillations (6 - 10 Hz) were also similar between environments (Fig. S5).

We next asked whether task-demands alter the structure of MEC firing patterns (6, 9). Grid cells' (n = 102) firing patterns re-organized between environments, despite their shared geometric shape and size (Fig. 1E, Table S1). First, the orientation of the grid pattern rotated (median absolute orientation change: 12.53° , $p=1.12 \times 10^{-12}$, Fig. 1F). These rotations varied across animals (mean rotation range: -27° to $+7^\circ$), and resulted in grid orientations that were less environmentally-aligned in ENV2 compared to ENV1 ($p=0.001$, Fig. 1I) (13). Second, there was a small decrease in grid spacing ($p=0.015$, Fig. 1G), but not in field size ($p=0.85$), in ENV2. Third, we observed less elliptical grid patterns in ENV2 ($p=0.006$, Fig. 1H). Finally, we observed a translation in the grid pattern (Fig. S6D-G) (16, 17). Co-recorded grid cells changed coherently and maintained their phase offsets (Fig. S6A). The observed grid orientation, scaling, and ellipticity changes also held for unpaired grid cell recordings clustered into modules (Fig. 1J,K) (18). Overall, 49/102 grid cells showed a statistically significant change on at least one measure (Fig. S6B), with changes largely conserved within animals (Fig. S6C-H). Critically, we observed grid pattern translation but not orientation, spacing or ellipticity changes when ENV1 and ENV2 had the same behavioral demand (random foraging, n = 3 rats), although there was no difference in the change in grid spacing between groups (Fig. 1F-H, right, Table S1) (11, 13, 16, 17).

Consistent with task demands re-structuring MEC representations, head direction, border,

and non-grid spatial cells reorganized between environments. Head direction (HD) cells coherently rotated their preferred direction within sessions and animals (both $p < 0.002$, Fig 2A, Fig. S7A-C), with 70/132 cells exhibiting significant changes in tuning. Rotations were consistent with the rotation in grid orientation (all HD-grid cell pairs: $r=0.45$, $p=0.02$; averaged within sessions: $r=0.70$, $p=0.02$; Fig. 2B-C). A majority (24/36) of border cells remapped between environments, primarily through rotations (Fig. 2D, E) (6). Lastly, 196/271 non-grid spatial cells significantly remapped between ENV1 and ENV2, with task-trained animals showing more remapping than free-foraging controls (task-trained mean correlation coefficient \pm SD: 0.32 ± 0.22 ; control: 0.41 ± 0.27 ; 49/100 control cells re-mapped, proportions test $p = 3 \times 10^{-5}$) (Fig. 2G, Fig. S7E). We observed no changes in speed cells (Fig. S7F,G).

We next examined whether spatial restructuring incorporated the remembered reward location. As running speed and spatial sampling differed between environments (Fig. S2), we first down-sampled the data to match in speed and position occupancy between environments (3, 5, 12). The relative activity of grid and non-grid spatial cells increased near the reward zone in ENV2 compared to ENV1 (signed-rank test, normalized activity vs distance slopes, grid: $p = 0.0025$; non-grid: $p = 5 \times 10^{-4}$) (Fig. 3A,B, Fig S8A-D). The robustness of this effect was reinforced by the observation of the same effect at the level of individual animals (Fig. S8A) and was not driven by increased occupancy near the reward zone (Fig S8E-H). Directional and non-directional grid cells showed comparable reward-related firing increases (Fig. S8I,J).

We next investigated how grid cells re-structure their firing toward the reward zone (Fig. S9A). Our observation of coordinated translations between simultaneously recorded grid cells (see Fig. S6A) eliminated the possibility that cells translate independently. Emergence of new grid fields, distortion of the grid pattern, and systematic re-shaping of grid fields were also

eliminated, as we did not observe changes in grid score (Fig. S4B), the number of fields, or the distance between the reward zone and closest field (Fig. S9B, Table S2). Moreover, we did not observe changes in field size or eccentricity as a function of fields' proximity to the reward zone (Fig. S9C). Finally, we examined whether grid field rate-remapping (19) shows reward specificity, such that fields near the reward zone exhibit higher firing rates. We did not observe significant changes in the overall field peak firing rates or coefficient of variation among field peak firing rates (Fig. S9D). However, the peak firing rate of grid fields closer to the reward zone was higher in ENV2 ($p=0.01$, Fig. 3D) and the distance from the reward zone to the grid field with the highest firing rate was smaller in ENV2 ($p=0.01$, Fig. 3E; see also Fig. S9E,F).

We then investigated how non-grid spatial cells ($n = 271$) remapped to support reward-localized changes in firing rates. Non-grid spatial cells did not extend their firing fields in a reward-specific manner as average field size, total field area, and number of fields did not change (Figure S10A). Instead, many cells ($n = 159$ cells) heterogeneously remapped to preferentially encode the reward location (Fig. 3F,G Table S3). First, some cells (Group I) exhibited coherent spatial tuning in both environments, with a firing field located closer to the reward zone in ENV2 ($p=2 \times 10^{-5}$). A second group of cells (Group II) exhibited coherent spatial tuning in ENV1, with the field farther from the reward than expected by chance ($p=0.02$). Third, a population of cells (Group III) had coherent spatial tuning only in ENV2, and this activity was closer to the reward zone than expected by chance ($p=0.002$). Finally, Group IV did not exhibit any coherent spatial fields but exhibited increased activity near the reward zone in ENV2. The proportion of cells exhibiting reward-preference did not depend on the group type (all $p>0.05$, Fig. 3G, bottom). Further, reward-preference and other coding features did not cluster (Figure S10B,C).

We next asked whether these changes reflected neural activity during the spatial task trials or were persistent throughout the ENV2 recordings. We analyzed two rate maps for each ENV2 session: one for task trajectories (tone onset to zone entry) and one for speed- and position-matched no-task trajectories (Fig. 4A,B). Grid cells' average firing rate did not differ between task and no-task, though non-grid spatial cells had higher firing rates during task times (Fig. S11). Critically, task and no-task maps both exhibited significant increases in normalized activity near the reward zone (Fig. 4C,D, Table S4), indicating that the reward influence was present throughout the session.

Finally, we asked how the task-associated changes in MEC representations could impact navigation. MEC representations can support vector navigation by providing unique combinations of spatial firing patterns, which downstream neurons may use to estimate the distance between an animal's position and a goal location (20). We estimated the animal's position using the activity from simultaneously recorded neurons in ENV1 and ENV2 (Fig. 4E,F). Using a Bayesian decoder, we observed that the decoding accuracy increased near the reward zone in ENV2 compared to ENV1 (ENV2 slope > ENV1 slope for 27/43 sessions, median slope difference = 1×10^{-3} , signed-rank $p = 0.042$) (Fig. 4G-I, S12). Moreover, the improved position decoding was highly localized to the reward zone, with a decrease in decoding error in ENV2 observed up to 30 cm from the reward zone center (Fig. 4J, Fig S12A-C). Reward-related decoding did not consistently co-vary with fluctuations in task performance (Fig. S12E,F).

Our understanding of how remembered reward locations mediate MEC navigational codes has lagged due to a lack of task diversity. Here, we report that the firing rate and spatial pattern of MEC representations restructure in response to changes in navigational strategy. This

restructuring did not reflect trajectory specific coding, as previously observed in MEC (21), suggesting task-relevant features of the two environments evoked separate long-term map representations (17). However, the precise parameters of MEC map restructuring may depend on experience and task familiarity, as recent work indicates (22). Combined, our data points to MEC as a region capable of dynamically altering its coding features to integrate relevant contextual features to support a range of navigational strategies.

References and Notes:

1. J. O'Keefe, J. Dostrovsky, The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res* **34**, 171-175 (1971).
2. T. Hafting, M. Fyhn, S. Molden, M. B. Moser, E. I. Moser, Microstructure of a spatial map in the entorhinal cortex. *Nature* **436**, 801-806 (2005).
3. E. Kropff, J. E. Carmichael, M. B. Moser, E. I. Moser, Speed cells in the medial entorhinal cortex. *Nature* **523**, 419-424 (2015).
4. T. Solstad, C. N. Boccara, E. Kropff, M. B. Moser, E. I. Moser, Representation of geometric borders in the entorhinal cortex. *Science* **322**, 1865-1868 (2008).
5. F. Sargolini *et al.*, Conjunctive representation of position, direction, and velocity in entorhinal cortex. *Science* **312**, 758-762 (2006).
6. G. W. Diehl, O. J. Hon, S. Leutgeb, J. K. Leutgeb, Grid and nongrid cells in medial entorhinal cortex represent spatial location and environmental features with complementary coding schemes. *Neuron* **94**, 83-92 (2017).
7. M. I. Anderson, K. J. Jeffery, Heterogeneous modulation of place cell firing by changes in context. *J Neurosci* **23**, 8827-8835 (2003).
8. R. U. Muller, J. L. Kubie, The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *J Neurosci* **7**, (1987).
9. S. Leutgeb *et al.*, Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science* **309**, 619-623 (2005).
10. C. G. Kentros, N. T. Agnihotri, S. Streater, R. D. Hawkins, E. R. Kandel, Increased attention to spatial context increases both place field stability and spatial memory. *Neuron* **42**, 283-295 (2004).
11. J. Krupic, M. Bauza, S. Burton, C. Barry, J. O'Keefe, Grid cell symmetry is shaped by environmental geometry. *Nature* **518**, 232-235 (2015).
12. K. Hardcastle, N. Maheswaranathan, S. Ganguli, L. M. Giocomo, A multiplexed, heterogeneous, and adaptive code for navigation in medial entorhinal cortex. *Neuron* **94**, 375-387 (2017).
13. T. Stensola, H. Stensola, M. B. Moser, E. I. Moser, Shearing-induced asymmetry in entorhinal grid cells. *Nature* **518**, 207-212 (2015).
14. B. a. K. J. Stephens, *Foraging Theory*. (Princeton University Press, Princeton, NJ, 1986).
15. M. Gil *et al.*, Impaired path integration in mice with disrupted grid cell firing. *Nature Neuroscience*, (In Press).

16. E. Marozzi, L. L. Ginzberg, A. Alenda, K. J. Jeffery, Purely translational realignment in grid cell firing patterns following nonmetric context change. *Cereb Cortex* **Jun 5**, Epub ahead of print (2015).
- 5 17. M. Fyhn, T. Hafting, A. Treves, M. B. Moser, E. I. Moser, Hippocampal remapping and grid realignment in entorhinal cortex. *Nature* **446**, 190-194 (2007).
18. H. Stensola *et al.*, The entorhinal map is discretized. *Nature* **492**, 72-78 (2012).
19. R. Ismakov, O. Barak, K. Jeffery, D. Derdikman, Grid cells encode local positional information. *Curr Biol* **27**, 2337-2343 (2017).
- 10 20. D. Bush, C. Barry, D. Manson, N. Burgess, Using grid cells for navigation. *Neuron* **87**, 507-520 (2015).
21. L. M. Frank, E. N. Brown, M. Wilson, Trajectory encoding in the hippocampus and entorhinal cortex. *Neuron* **27**, 169-178 (2000).
22. C. N. Boccara, M. Nardin, F. Stella, J. O'Neill, J. Csicsvari, The entorhinal cognitive map is attracted to goals. (Submitted).
- 15 23. A. Friedman, M. D. Keselman, L. G. Gibb, A. M. Graybiel, A multistage mathematical approach to automated clustering of high-dimensional noisy data. *Proc Natl Acad Sci U S A* **112**, 4477-4482 (2015).
24. N. Schmitzer-Tobert, J. Jackson, D. Henze, K. Harris, A. D. Redish, Quantitative measures of cluster quality for use in extracellular recordings. *Neuroscience* **131**, 1-11 (2005).
- 20 25. M. G. Campbell *et al.*, Principles governing the integration of landmark and self-motion cues in entorhinal cortical codes for navigation. *Nat Neurosci* **21**, 1096-1106 (2018).
26. K. Hardcastle, N. Maheswaranathan, S. Ganguli, L. M. Giocomo, A multiplexed, heterogeneous, and adaptive code for navigation in medial entorhinal cortex. *Neuron* **94**, 375-387 (2017).
- 25 27. L. Meng, M. A. Kramer, S. J. Middleton, M. A. Whittington, U. T. Eden, A unified approach to linking experimental, statistical and computational analysis of spike train data. *PLoS One* **9**, e85269 (2014).
28. T. Solstad, C. N. Boccara, E. Kropff, M. B. Moser, E. I. Moser, Representation of geometric borders in the entorhinal cortex. *Science* **322**, 1865-1868 (2008).
- 30 29. J. P. Bassett, M. L. Tullman, J. S. Taube, Lesions of the tegmentomammillary circuit in the head direction system disrupt the head direction signal in the anterior thalamus. *J Neurosci* **27**, 7564-7577 (2007).
30. E. Kropff, J. E. Carmichael, M. B. Moser, E. I. Moser, Speed cells in the medial entorhinal cortex. *Nature* **523**, 419-424 (2015).
- 35 31. R. F. Langston *et al.*, Development of the spatial representation system in the rat. *Science* **328**, 1576-1580 (2010).
32. W. E. Skaggs, B. L. McNaughton, M. A. Wilson, C. A. Barnes, Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus* **6**, 149-172 (1996).
- 40 33. M. M. Yartsev, M. P. Witter, N. Ulanovsky, Grid cells without theta oscillations in the entorhinal cortex of bats. *Nature* **479**, 103-107 (2011).
34. H. Stensola *et al.*, The entorhinal map is discretized. *Nature* **492**, 72-78 (2012).
35. T. Stensola, H. Stensola, M. B. Moser, E. I. Moser, Shearing-induced asymmetry in entorhinal grid cells. *Nature* **518**, 207-212 (2015).
- 45 36. C. S. Mallory, K. Hardcastle, J. S. Bant, L. M. Giocomo, Grid scale drives the scale and long-term stability of place maps. *Nature Neuroscience* **21**, 270-282 (2018).

Acknowledgments: We thank A Borryo and A Diaz for histology assistance. **Funding:** LMG is a New York Stem Cell Foundation–Robertson Investigator. This work was supported by funding from the New York Stem Cell Foundation, NIMH MH106475, NIDA DA042012, Office of Naval Research N000141812690, Simons Foundation 542987SPI and the James S
5 McDonnell Foundation awarded to LMG, and a Stanford Interdisciplinary Graduate Fellowship awarded to KH; **Author contributions:** LMG, WB and KH conceptualized experiments and analyses. WB and KH performed implantations, collected and analyzed data. All authors wrote the paper; **Competing interests:** Authors declare no competing interests; **Data and materials availability:** Data are available at <https://giocomolab.weebly.com/data.html>. Code is available
10 at <http://doi.org/10.5281/zenodo.2435310>

Supplementary Materials:

Materials and Methods

Figures S1-S12

15 Tables S1 to S4

References (23-36)

Fig. 1. Performance of a task induces grid rotation and rescaling. (A). Schematic of environments. (B) Trajectories (gray) from a paired session. Trial trajectories are highlighted. Mean trial circuitry and trial time noted below ENV2. Reward zone in red. (C) Histogram of inter-trial intervals. (D) Circuitry and trial time improved with training in individual animals (gray lines). Data aligned to each animal's first post-trained session (red line). (E) (Left) Grid cell rate maps in both environments; peak firing rate and grid score noted on top. (Middle) Corresponding autocorrelations; spacing and orientation noted on top. Red lines indicate grid axes; white text indicates ellipticity. (Right) Corresponding ENV1-ENV2 cross-correlations. Distance from the cross-correlation's center to the nearest peak noted on top. (F) (Left) Grid cell orientations, red lines indicate rotations equivalent to modulo 60. (Right) Histogram of grid orientation differences for experimental and control animals. (G) (Left) Grid cell spacing, red line indicates identical spacing. (Right) Histograms of grid spacing ratio. (H) (Left) Grid cell ellipticities, red line indicates identical ellipticity. (Right) Histograms of ellipticity ratio. (I) Scatter plots of the innermost six fields in each grid cell's autocorrelation. Orange lines represent north-south aligned axes; blue lines represent east-west aligned axes. (J) Unpaired grid cell recordings from four animals, clustered into modules according to spacing and orientation. (K) Mean orientations (Left) and spacings (Right) in each environment for each of the six modules in (J). Error bars indicate SEM.

Fig. 2. Performance of a task induces remapping in head direction, border and non-grid spatial cells. (A) Top row: Four co-recorded HD cells in each environment. Rightmost panel indicates each cell's rotation between environments. Bottom row: rotation angles observed

across sessions. Gray lines indicate boundaries between animals. **(B)** Co-recorded grid and HD cells. (Top) HD tuning curves. (Bottom) Grid cell autocorrelations with grid axes. Co-rotation of grid and HD signals shown by rotating the ENV1 grid axes by the rotation observed in co-recorded HD cell (blue dashed lines). **(C)** (Grey) HD cell orientation change (between environments) versus grid cell orientation change for all possible pairs of co-recorded HD and grid cells. (Blue) Same data, with all HD or grid cells recorded within the same session averaged together. **(D)** Border cell rate maps in ENV1 and ENV2. **(E)** Histograms of border cell rate map ENV1 versus ENV2 correlation coefficients (left) and rotation values (right). **(F)** Non-grid spatial cell rate maps in ENV1 and ENV2. **(G)** (Left) Histogram of non-grid spatial cell rate map ENV1 versus ENV2 correlation coefficients (black = cells with significant re-mapping, grey = non-significant re-mapping). (Right) Histogram of the difference in spatial stability between ENV1 and ENV2.

Fig. 3. Grid and non-grid spatial cells have localized firing rate changes near the reward.

(A) (Left) Mean normalized grid cell firing rate as a function of distance from the reward zone. Ribbon indicates SEM. (Right) Difference in grid cell firing rate (ENV2-ENV1). **(B)** (Left) Mean normalized non-grid spatial cell firing rate as a function of distance from the reward zone. (Right) Difference in firing rate (ENV2-ENV1). **(C)** (Left) Rate maps for three grid cells recorded in both environments. (Right) Corresponding field peak firing rates, plotted as a function of the field's distance from the reward zone. Best-fit lines shown; difference between the best-fit lines (slope ENV2 - slope ENV1) indicated in upper left. **(D)** (Top) Best fit slope values for each cell in ENV1 and ENV2. (Bottom) Histogram of slope differences for grid cells. **(E)** (Top) Distance from the reward zone to the highest FR field in each environment for each

cell. (Bottom) Histogram of distance differences. (F) Non-grid spatial cells that show reward preference in ENV2 correspond to four categories of remapping; two examples/group are shown. (G) (Top) Fraction of reward-preferring cells in each remapping category (of 159 total reward-preferring cells). (Bottom) Fraction of cells in each remapping category that show reward preference.

Fig. 4. Long-term changes in the spatial map support spatial decoding near the reward. (A)

Rate maps of the full ENV2 session (left), task (middle), and no-task trajectories (right) speed-matched for each position bin. (B) (Left) The cell in (A)'s average normalized firing rate as a function of distance from the reward zone for task (orange) and no-task trajectories (green). (Right) Mean running speed during task and no-task trajectories as a function of distance from the center of the reward zone, before and after speed-matching. (C-D) (Left panels) Average normalized firing rate for grid (C) and non-grid position (D) cells as a function of distance from the reward zone. (Right panels) The slopes of both task and no-task trajectories were significantly negatively distributed for grid (C) and non-grid position (D) cells. (E) Example decoding error maps for ENV1 (left), ENV2 (middle), and the normalized difference (ENV2-ENV1, right) from a single session (ENV1 n = 6 P-encoding cells, ENV2 n = 5 cells). (F) Normalized error (Left) and ENV2-ENV1 error difference (Right) as a function of distance from the reward zone for the example in (E). (G) Normalized error versus distance from reward zone for each environment, averaged over all decoding sessions (n = 43). (H) Average difference in error (ENV2-ENV1) for all sessions. (I) Distribution of slopes of ENV2-ENV1 tuning curves across sessions. (J) Across all sessions, decoding error within 30 cm of reward zone is lower in ENV2 than ENV1 (median difference in error = -4.3 cm, signed-rank p = 0.028).



5
Supplementary Materials for

Remembered reward locations restructure entorhinal spatial maps

William N. Butler, Kiah Hardcastle, Lisa M. Giocomo

10 Correspondence to: giocomo@stanford.edu

This PDF file includes:

15 Materials and Methods
 Figs. S1 to S12
 Tables S1 to S4
 References (23-36)

20

Materials and Methods

Subjects

All procedures were approved by the Institutional Animal Care and Use Committee at Stanford University School of Medicine. Each animal was handled daily for 2-3 weeks before experimental training or surgery. Implanted rats were singly housed and allowed free access to food and water for 1 week after surgery, after which they were mildly food deprived to ~90% of their starting body weight. All rats were housed on a 12-h light-dark cycle, and experiments were performed during the light phase. Due to the large size and weight of the implanted microdrives (~4 cm in diameter, ~10 g in weight), this study used male rats (rat n = 11, aged 4 - 8 months old at the start of experiments, weighing between 350 and 500 grams).

In vivo survival surgery

Surgical anesthesia was induced and maintained by inhalation of isoflurane gas (1-3%). Rats were unilaterally implanted with a 12- or 16-tetrode microdrive (Neuralynx, Inc.) containing 17- μm polyimide-coated platinum-iridium (90-10%) wires (plated to impedances of ~250 k Ω at 1 kHz). To target medial entorhinal cortex (MEC), tetrodes were implanted 0.5 mm anterior to the transverse sinus (~8.5 mm posterior to bregma), and 3.8-4.2 mm lateral from midline; in all animals, the electrode was implanted in the right hemisphere. In a subset of animals (n = 3), the microdrive was also angled so that the tetrodes pointed 5° laterally from vertical. The craniotomy was sealed with Kwik-Sil, and jeweler's screws, C&B-Metabond®, and dental cement were used to affix the microdrive to the skull. At the end of the surgery, after implantation, tetrodes were driven 2.5-3.5 mm ventrally into the brain.

In vivo single-unit data collection

Recorded signals were amplified and bandpass filtered between 0.6 and 6 kHz using a Digital Lynx SX data acquisition system (Neuralynx, Inc.). Signals that crossed a preset amplitude threshold (manually specified, 25-70 μV) were time-stamped and stored to disk at 32 KHz. LFP was also recorded from each electrode and amplified 3,000-10,000 times, bandpass filtered between 0.1 and 9,000 Hz, sampled at 1,524 Hz, and stored with the unit data. To track the rat's position, two light-emitting diodes attached to the head stage, one red and one green, were detected by an overhead camera (sampling rate 30 Hz). After each day of recording, individual tetrodes were each driven ventrally by 0-250 μm . To maximize the odds of recording the same cells across multiple environments, tetrodes were not moved between daily recording sessions (3-6 hours between recording sessions).

Histology and electrode localization

Upon completion of the final recording session, rats were anesthetized with 2-4% of isoflurane in O₂ and euthanized using an overdose of pentobarbital followed by transcardial perfusion with saline and 10% formalin. After the brain was stored in formalin for at least 24 hours, the hemisphere containing the implant was sliced into 25 μm sagittal sections using a cryostat microtome. Sections were mounted on slides and Nissl-stained with cresyl violet. Approximate recording sites were determined by calculating the distance corresponding to the amount of microdrive screw turns and following the tetrode track dorsally from the most ventral point (Fig. S1). Of the 778 cells recorded in both environments, 598 were recorded from MEC; the remaining 180 cells were recorded from surrounding cortical regions, including

parasubiculum and postsubiculum.

Location-based spatial task

Neural data was collected as rats freely foraged in a 150 cm x 150 cm box (ENV1). This box had black walls, a white cue card on one wall, and the floor was scented with an ethanol-diluted solution of lemon extract prior to each recording session. Animals only freely foraged for randomly scattered chocolate cereal (<0.02 g each) in ENV1. Each box was thoroughly cleaned with an alcohol solution after each recording session.

In addition, rats in the Experimental group (n = 8) were trained to perform a cued positional response task in a second, distinct 150 cm x 150 cm open field arena (ENV2). This box had white walls, a black cue card on one wall, and the floor was scented with an ethanol-diluted solution of vanilla extract prior to each recording session. All food rewards in this box were vanilla flavored frosted cereal. Both boxes were placed in the exact same physical location for recording, and thus the distal and extra-maze cues were identical between the two boxes. For each animal, a unique reward zone was defined as an unmarked 20 cm x 20 cm area whose borders were all at least 15 cm from the walls of the arena. The location of each reward zone was chosen pseudo-randomly to ensure reward zones were in dissimilar locations across animals. Each reward zone was saved to a configuration file and paired with a separate video tracking application (Zoned Video, Neuralynx, Inc.) to enable online detection of zone boundary crossings during recording.

Animals performed ~10-20 trials during each spatial task recording session, with each trial separated by ~1-3 minutes (Figure 1, Figure S2). During each trial, animals responded to a white noise tone by navigating to the unmarked reward zone. The tone was played continuously until the animal entered the reward zone. Successful completion of a trial was marked by silencing of the tone and delivery of a large food reward (0.5-1 pieces of cereal, ~0.06 - 0.12 g). Training occurred in two phases. In phase I (the initial training), the food reward was delivered in the reward zone at tone onset. Once animals began to anticipate reward delivery by navigating to the reward zone prior to the food being placed there, training progressed to phase II. In this phase, the food reward was only delivered after the animal entered the reward zone. All analyses presented here come from phase II recordings. During non-trial portions of the recording session, small pieces of cereal (<0.02 gram each) were scattered around the environment at random time intervals to promote spatial coverage (as in ENV1 recordings). On the same day neural data was collected in ENV1, each animal was also recorded performing the spatial task in ENV2. Animals were recorded once in each box per day, and the ordering of which box was recorded first was counterbalanced across days.

To quantify animals' performance in the spatial task, for each trial we calculated the amount of time it took for the animal to reach the reward zone after tone onset, and the circuitry of the path the animal took to reach the reward zone (shortest possible distance divided by actual distance traveled). For each session, we also computed a combined performance index, defined as the time to reach the reward zone divided by circuitry of path to reward. We considered animals trained once they were able to perform the phase II version of the task with a median performance index ≤ 60 for at least two consecutive sessions. One animal did not reach this performance criterion and therefore was excluded from all analyses. In comparing pre- and post-training performance (Fig. 1), we only compared phase II recordings before and after the training criteria was met. Animals performed the task significantly better after training (median trial time before training = 13.1 s, after training = 8.8 s, Wilcoxon rank sum $p = 6 \times 10^{-7}$; median circuitry before training = 0.18, after training = 0.33, Wilcoxon rank sum, $p = 1 \times 10^{-7}$).

We also recorded cells in a separate group of Control rats ($n = 3$). In this group, the two boxes were identical to the conditions described above. However, instead of spatial task training, the Control rats freely foraged for randomly scattered food rewards in ENV2 as well, identical to their behavior in ENV1. To further control for the non-behavioral aspects of the spatial task recording sessions, the white noise tone used as the trial cue in Experimental animals was played ten times/session for the control animals (10 s duration, ~ 2 min between each presentation). Importantly, the control tone onset/offset was never linked to the Control animals' behavior nor to the delivery of a large food reward.

Offline spike sorting

Spike sorting was conducted through a combination of automated clustering (SNAP sorter (23)) and graphical software (SpikeSort3D, Neuralynx Inc.). Automatically-defined clusters were inspected and further refined manually. Cluster quality was assessed by calculating Mahalanobis distances between each spike cluster and noise spikes (24). The cluster quality of cells recorded was comparable between Experimental and Control groups (Experimental: $n = 3,090$ unique cells, median Mahalanobis distance during free foraging = 11.48, during spatial task = 11.43, rank-sum test $p = 0.87$; Control: $n = 1,123$ cells, median Mahalanobis distance ENV1 = 11.53, ENV2 = 11.40, rank-sum test $p = 0.20$) (Fig. S3). Only cells with ≥ 100 spikes were used in further analyses. Position data was sorted into 5 cm x 5 cm bins (each 1.5 m box was 30 bins x 30 bins). Occupancy and spiking maps were smoothed separately using a quasi-Gaussian kernel. Firing rate maps were calculated from these smoothed maps by dividing spike number by the time spent in each bin. In all figures, rate maps are colored such that high firing rates are yellow and low firing rates are blue.

To identify single units recorded in different sessions, we calculated the Euclidean distances between the center-of-mass (COM) of each cell's waveform on the current tetrode and each cell's waveform on that same tetrode in the most recent session recorded in the other environment (Fig. S3) (25). Cells that were likely matches (at least 1 standard deviation below the mean COM distance between the current cluster and all clusters from the previous session) were then manually inspected across other features (waveform shape, interspike-interval histogram) in order to finalize whether or not to label the unit as a matching cell. If a cell was recorded more than once in the same environment, only the best recording in that environment was analyzed. The best recording in an environment was determined by two comparisons: first, a more stringent occupancy measure (≥ 0.5 s/bin) was quantified, and sessions with $< 75\%$ occupancy by this measure were eliminated. If multiple recordings still remained, the best recording was defined as the session with the highest isolation distance. For all analysis presented here, only cells in trained ENV2 sessions were considered, and they were matched with the best ENV1 recording.

Recording navigational information

Position and head direction samples that were not detected by the camera were marked as NaN values. Strings of less than five consecutive NaN values were filled in through linear interpolation. Speed for every time point was computed as an average of the forward and backward difference in position estimates divided by the time bin width (33 ms). Speeds over 200 cm/s were removed (set to NaN), and then were filled in through linear interpolation. Speed was then smoothed through convolution with a Gaussian filter (standard deviation = 2.5 bins, over a 5 bin window; all NaN values were filled in prior to filtering through linear interpolation,

and then replaced post-filtering). Speeds less than 2 cm/s were removed from analysis.

LNP model and cell classification

Cells were identified as encoding position (P), head direction (H), or speed (S), using a linear-nonlinear-Poisson (LNP) model following Hardcastle et al., 2017 (26). In this framework, the model estimates the spiking rate of each cell, \hat{r} , as a function of position, head direction, or running speed as:

$$\hat{r} = \frac{\exp(b + \sum_i X_i^T w_i)}{dt}$$

where r denotes a vector of firing rates over T time bins, b denotes the baseline input (a learned parameter), X_i^T denotes the input matrix for variable i ($i \in [P, H, S]$, X_i^T is a $T \times (\text{length of } w_i)$ matrix), w_i denotes the learned parameters for variable i , and dt is the time bin width (fixed to be 20 ms; all navigational variables are linearly interpolated to this sampling rate). Each row in the matrix X_i^T (referred to as $X_i^{T(t)}$) indicates the input related to variable i at time point t , and is computed using cardinal spline interpolation (meaning that the learned parameters w_i are the values for each control point) (27). Specifically, each variable (position, head direction, and speed; we will refer a variable in general as y_i), is first binned into n_i bins ($n_P = 12 \times 12$ bins, $n_H = 10$ bins, $n_S = 8$ bins). These bins create a vector of control point positions x_i , which correspond to w_i . For each time point t , and value of each variable y_i at time point t (denoted $y_i^{(t)}$), we then find the control point positions that lie on either side of $y_i^{(t)}$, $x_{i,j}$

and $x_{i,j+1}$. From this, we compute a value $\alpha^{(t)} = \frac{y_i^{(t)} - x_{i,j}}{x_{i,j+1} - x_{i,j}}$, which is the fractional distance between control points. We then compute a vector $p^{(t)T} = [\alpha^{(t)3} \quad \alpha^{(t)2} \quad \alpha^{(t)1} \quad 1]$

$\begin{bmatrix} -s & 2-s & s-2 & s \\ 2s & s-3 & 3-2s & -s \\ -s & 0 & s & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix}$. For analysis in this paper, we chose $s = 0.5$. This vector $p^{(t)}$

corresponds to the input vector related to the control points $w_{i,j-1}, w_{i,j}, w_{i,j+1}, w_{i,j+2}$. Specifically, if a model contains only the baseline parameter and a single variable, then $\hat{r}^{(t)} * dt = \exp\left(b + \begin{bmatrix} -s & 2-s & s-2 & s \\ 2s & s-3 & 3-2s & -s \\ -s & 0 & s & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} w_{j-1} \\ w_j \\ w_{j+1} \\ w_{j+2} \end{bmatrix}\right) = \exp\left(b + p^{(t)T} \begin{bmatrix} w_{j-1} \\ w_j \\ w_{j+1} \\ w_{j+2} \end{bmatrix}\right) = \exp\left(b + [0 \dots p^{(t)T} \dots 0] \begin{bmatrix} w_1 \\ \vdots \\ w_j \end{bmatrix}\right)$, where J is the total number of parameters.

Thus, $X_i^{T(t)}$ simply a zero-padded $p^{(t)T}$, such that the j^{th} index of $X_i^{T(t)}$ corresponds to the j^{th} index of w_i .

The model was fit assuming that spikes follow a Poisson noise distribution. Learned

parameters $\hat{w} = \begin{bmatrix} b \\ w_1 \\ \vdots \\ w_j \end{bmatrix}$, where M is the total number of external variables in a model, were

computed by maximizing the log-likelihood of observing the data given the model:

$$\hat{w} = \operatorname{argmax}_{w_i, b} \sum_t \log P(n_t | \exp\left(b + \sum_i X_i^T w_i\right) - \frac{1}{2} \beta_i \sum_i \sum_j w_{i,j}^2)$$

where n_t is the number of spikes for in time bin t , and β_i is the weight of the L2 penalty. This was set to 1 for all variables. This maximization problem was carried out using `fminunc` in MATLAB. Models were fit and performance was assessed using 10-fold cross-validation (following (26)). For each fold, models were fit on 90% of the data set (10% were held out, selected as random 20-second long chunks), and then model performance was computed on the corresponding testing data set. Each testing data set was non-overlapping with the other testing data sets. Model performance was computed as the log-likelihood increase from a mean-firing rate model, divided by the number of spikes observed in that testing data fold.

To determine which variables each cell encodes, we used the forward search method described in Hardcastle et al., 2017 (26). In brief, we first fit single-variable models (e.g. a model with just position information), and then determined which model had the highest model performance. We then fit all possible two-variable models that include the original highest-performing single variable, and determine whether adding additional information significantly improved model performance (using a signed-rank test across the 10 folds, and a $p = 0.05$ significance criteria). This process was repeated until there was no significant improvement in model performance. At the end of this process, we determined whether the highest-performing model (on held-out data) was significantly better than a mean firing rate model. If not, the cell was recorded as not encoding a navigational variable. Mixed selective cells were defined as cells that encoded two or more navigational variables.

Tuning curves and classification of score-defined cell types

We calculated each cell's grid score by taking a circular sample of the spatial rate map autocorrelation centered on the central peak and compared it to rotated versions of the same circular sample (60° and 120° versus 30° , 90° and 150°). The grid score was defined as the mean correlation at 60° and 120° minus the mean correlation at 30° , 90° and 150° . A border score was computed following previous work by calculating $\frac{CM-DM}{CM+DM}$, where CM is the proportion of high firing rate bins along one wall, and DM is the normalized mean product of the firing rate of each bin and its distance to the nearest wall (28). To calculate head direction tuning, we plotted the firing rate of a neuron as a function of the rat's directional heading, divided into 10° bins and smoothed with a 5-bin Gaussian filter. The head direction score was defined as the length of the mean Rayleigh vector for the circular distribution of firing rate x head direction (29). Speed scores were computed as the Pearson correlation coefficient between instantaneous running speed and firing rate, the latter of which was smoothed according to (30).

To identify grid cells, border cells, HD cells, and speed cells, we shuffled the spike trains with respect to the behavioral data, re-computed the grid and border scores, generated a null distribution of each score value by pooling the shuffled score distributions for all cells, and then computed each score threshold as the 95th percentile of the shuffled distribution (31). In the shuffling procedure, each cell's spike train was circularly time shifted by a random interval between 20 s and T-20s, where T is the length of the recording. Each spike train was shuffled 50 times. This procedure was carried out separately for the recordings in the two environments (ENV1 and ENV2). Cells were assigned to a specific class (e.g., grid cell) if their score exceeded the shuffled population score threshold in either environment, and if the cell also encoded the relevant variable (e.g. position). Stronger criteria were used for head direction and border cells

when quantifying re-mapping: only head direction cells with a score > 0.4 (above 95th percentile) were considered, and only border cells with a border score of 0.55 and a spatial coherence measure above 2.25 in either environment were considered.

5

Tuning curve stability, information content, spatial coherence

Stability for a given tuning curve was computed as the Pearson correlation coefficient between the tuning curves generated from the first half of the session and the second half of the session. Position tuning curves were vectorized prior to computing the correlation coefficient.

10

Information content was computed as:

$$I = \sum_i \frac{1}{N} \frac{\lambda_i}{\lambda} \log_2 \frac{\lambda_i}{\lambda}$$

15

where i is the bin number, λ_i is the mean firing rate in the i th bin, λ is the overall mean firing rate, and N is the number of bins (assuming equal probability in each position bin) (following (32)). Spatial coherence was computed as the Fisher z-transformation (the arctangent) of the correlation between the firing rate of each spatial bin and the average firing rate of the neighboring 3-8 spatial bins.

Theta measures

20

Theta power and theta frequency of the local field potential (LFP) were computed as a function of running speed for each session. All measures were computed using the tetrode with the highest theta power (theta power = power in 6-10 Hz range / power in 1-50 Hz range) for a given session. The power spectrum was computed using the pwelch function in MATLAB. We used a window size of 10 times the length of the LFP time bins (LFP was sampled at 1524 Hz), an overlap of half the window size, and $2^{(\text{window size})}$ as the number of NFFT points. To compute theta power and frequency across running speeds, we first binned the session into 2-second segments. For each segment, we computed the theta power, the average running speed, and the frequency with maximum power in the theta-band (defined here 6-10 Hz). We then computed the Pearson correlation coefficient between theta power and running speed, as well as theta frequency and running speed. In addition, we found the slope and intercept of the line that best fit the running speed and theta power (or theta frequency) data (Figure S5).

25

30

35

To compute the influence of theta on single-cell activity, we computed theta locking, theta modulation (through computing a theta index), and the degree of theta skipping for each cell. To compute theta locking, we first band-pass filtered the LFP for 6-10 Hz (using the butter and filtfilt MATLAB functions), computed the phase of each time point through the Hilbert transform, and then recorded the theta phase of each spike. The theta locking score was computed as the mean vector length of the distribution of spike phases (similar to the head direction score). Theta modulation was quantified through the theta index score following (33), which was calculated by first computing the power spectrum of the spike-time autocorrelation (the normalized histogram of all inter-spike-intervals), finding the frequency f_{max} of maximum power within the theta band (6-10 Hz), and then computing the ratio between the average power within 1 Hz of f_{max} and the average power between 1 and 50 Hz. To quantify the degree of theta skipping observed in a spike train, we first smoothed the spike time autocorrelation with a Gaussian window (2.5 bin standard deviation with 5 bin width). We then quantified the first theta peak as the maximum in the autocorrelation between 100 and 160 ms and the second theta peak as the maximum in the autocorrelation between 200 and 320 ms. The theta skipping score was

40

45

then calculated as the difference between the second and first theta peaks, divided by the maximum of the two peaks.

Speed- and position-matching for comparison of cell activity between session types

As differences in running speed have been shown to influence coding properties in entorhinal neurons (26), we implemented a speed- and position-matching down-sampling protocol in order to compare changes in activity near the reward zone between ENV1 and ENV2. In this protocol, we first binned position into a 10 x 10 grid (15 cm bin width). Within each position bin, we then binned speed into 10 bins (15 cm/s bin width), and computed the number of samples observed in each position-by-speed bin. The number of samples was matched between the two sessions by randomly sampling from the session with greater occupancy for that bin. Due to the stochastic nature of this process, this procedure was repeated 50 times for each cell, and the down-sampled rate maps were then averaged to obtain a single map for each session.

Grid cell analyses

Grid cell field spacing, grid orientation, ellipticity, and grid module analyses were all performed according to methods described in detail by (34). Briefly, field spacing and grid orientation were both determined by applying a threshold to each cell's autocorrelation to identify fields using an extended-maxima transform (scalar = 0.1). Grid spacing was defined as the mean of the distances of the six closest fields from the center peak in the autocorrelation. Grid orientation was determined by fitting three axes through each of the three closest pairs of fields, then taking the circular mean of the angular orientation of these three axes. Grid ellipticity was found by fitting an ellipse to the six closest fields and then computed as the length of the ellipse's major axis divided by the length of the ellipse's minor axis. Translation of the grid pattern was computed by first creating a cross-correlation of the rate maps in both environments. Because grid orientation rotated between the environments, the rate map for ENV2 was rotated to align grid orientations prior to making the cross-correlation. The distance to the peak nearest to the center of the cross-correlation was used as the amount of grid translation. Absolute angle to the nearest wall (A_{min}) was calculated as in (35) using the equation:

$$A = 45 - \text{abs}([X \text{ mod } 90] - 45)$$

where X is orientation of a single grid axis. A_{min} was then computed as the smallest angle value among the axes of each cell. Grid axes across all cells (Fig. S6C) were estimated on an animal-by-animal basis by creating histograms of the grid axes values in each environment and finding the three largest peaks, with the additional requirement that the angle between peaks must be at least 30 deg and could be no more than 90 deg.

Previous work has shown that grid cells form discrete clusters or modules based on orientation and field spacing (34). For this analysis, we removed the requirement that cells had to be recorded in both environments, and instead considered all unique units with a grid score >0.5 ($n = 193$ cells in ENV1, $n = 216$ cells in ENV2). Grid cells were clustered by a k-means algorithm according to their spacing and orientation (in modulo 60) within each animal, with the number of clusters estimated by the number of peaks in a kernel smoothed density estimate of grid spacing. The k-means clustering process was repeated 300 times for each animal and the iteration with the greatest mean silhouette coefficient was selected as the optimal clustering solution. The silhouette coefficient was computed for each point i as:

$$s(i) = \frac{b(i) - a(i)}{\max\{a(i), b(i)\}}$$

where $a(i)$ is the average distance between i and each point within its cluster, and $b(i)$ is the average distance between i and each point within the nearest neighboring cluster. Clusters with at least four cells were examined further, and mean orientation and mean spacing were calculated for each cluster. Clusters were determined to be corresponding modules between the two environments if they had mean spacings that were within 15% of one another (for example, if a cluster in ENV1 had a mean spacing of 100 cm, a corresponding module in ENV2 would have to have a mean spacing between 85 cm and 115 cm).

To determine how many cells exhibit significant changes in orientation, spacing, and ellipticity, we compared each difference metric (“across-session values”) to a null distribution of “within session values”. The null distribution was generated by randomly splitting each ENV1 session into two equal halves, re-computing the rate-maps for each partition, and then computing the difference in the measure of interest between rate map patterns. The random split was done by randomly choosing a breakpoint in the session; times after the breakpoint (wrapping around to the beginning the session if necessary) were placed into one half, while all times before the breakpoint (wrapping around to the end the session if necessary) were assigned to the other half. This procedure was repeated 100 times. If the across-session values were greater than the 95th percentile, the cell was considered to have significantly changed along that metric.

Quantification of grid cell translation

Translation was only quantified on cells with high grid scores in both environments (grid score > 0.35 in both environments, $n = 50$). First, the ENV1 grid cell pattern was rotated so that it matched in orientation to the ENV2 pattern. The cross-correlation was then computed by dragging the ENV1 pattern over the ENV2 pattern, and computing the correlation between the overlapping regions of pixels. The translation was quantified by the nearest local-maximum to the origin (local maxima had to be at least 75% of the maximum value of cross-correlation values observed within $0.8d$ of the center of the cross-correlation, where d = the minimum grid spacing between ENV1 and ENV2). To determine how many cells exhibit significant translations, the translation observed in each cell was compared to a null distribution of translations, which was generated estimating the degree of translations observed within each session. Specifically, each ENV1 and ENV2 session was split randomly into two (splits were done in 1-minute long chunks), rate-maps computed for each partition, and then the translation between rate map patterns was computed and recorded. This procedure was repeated 100 times, and the actual translation was compared to the combined null distributions of ENV1 and ENV2. If the actual translation was greater than the 95th percentile, the cell was considered to have significantly translated.

Quantification of head direction cell re-mapping and rotation

Remapping was quantified by computing the correlation coefficient of the head direction tuning curves derived from the recording in ENV1 compared to that in ENV2. To compute the fraction of HD cells that remap, we compared the across-session correlation coefficients to a distribution of within-session correlation coefficients. Within-session correlation coefficients were computed by first randomly dividing each ENV1 session into two halves and then

computing the correlation coefficient between tuning curves derived from each half. The random split was done by randomly choosing a breakpoint in the session; time after the breakpoint (wrapping around to the beginning the session if necessary) were placed into one half, while all times before the breakpoint (wrapping around to the end the session if necessary) were assigned to the other half. We did this 100 times for each session type (ENV1 and ENV2), and then combined the resulting null distributions. A cell was then considered to have significantly remapped if its correlation was lower than the lowest value in the null distribution.

Rotations were quantified through first identifying all task sessions in which at least two head direction cells were simultaneously recorded ($n = 30$ sessions), and then identifying the corresponding ENV1 head direction tuning curve for each cell. We computed the rotation between HD cells recorded in ENV1 and ENV2 by computing the change in angle in the ENV2 tuning curve that, when rotated by that amount, would produce that maximum correlation between the ENV1 and ENV2 tuning curves. Coherency of within-session rotations were quantified by computing the range of rotation angles observed within simultaneously recorded cells. We determined whether rotation angles were more coherent within sets of simultaneously recorded cells than chance by comparing the median rotation angle observed (median across 30 sessions) to a null distribution of median rotation angles. This distribution was generated by shuffling cells across sessions, and re-computing the median angle range across sessions, 1000 times (Figure S7B). Similar analyses were carried out to determine the significance of rotation coherence within animals (Figure S7B).

Identification of significant re-mapping in border and non-grid spatial cells

To determine how many cells exhibit significant changes in tuning, we compared the correlation coefficient between tuning curves generated from ENV1 and ENV2 (“across-session” coefficients) to a null distribution of “within session” correlation coefficients. The null distribution was generated by randomly splitting each ENV1 session into two equal halves, re-computing the rate-maps for each partition, and then computing the correlation coefficients between rate maps. The random split was done by randomly choosing a breakpoint in the session; times after the breakpoint (wrapping around to the beginning the session if necessary) were placed into one half, while all times before the breakpoint (wrapping around to the end the session if necessary) were assigned to the other half. This procedure was repeated 100 times. If the across-session correlation coefficients were less than 5th percentile of the distribution, the cell was considered to have significantly re-mapped.

Identification of fields in grid and non-grid position cells

Fields were identified in each speed- and position-matched rate map in order to compute field properties as a function of distance from the reward zone. First, we converted the rate map to a log-normalized scale. This log-normalized rate map was then thresholded at a value of 0.8 (0.85 for non-grid, non-border position cells) to identify fields. Firing fields were defined as regions of at least six contiguous spatial bins, and field size was computed as the total area of the spatial bins comprising that field. Field centers were defined as the center of mass of the firing field. Field eccentricity was calculated as the ratio of the major and minor axes of the best-fit ellipse to the field perimeter. For grid cells, fields were considered to be relatively complete (as opposed to falling partially outside the environment) if their area was at least 75% of the area of the central peak in the cell’s autocorrelation. Only grid cells with at least two complete fields in

each environment were analyzed to examine how changes in grid field structure varied (Fig. 3, Fig. S9).

5 Generating a null distribution of nearest field distances to the reward zone center

The distance of any field to the reward zone was computed by taking the Euclidean distance from the center of mass of the field to the center of the reward zone. To compare whether the distance of the nearest field to the zone center was significantly different from chance for non-grid spatial cells, we pooled all ENV1 non-grid spatial maps with a spatial coherence > 2.25, and generated a pool of reward zone centers based on the statistics (e.g. distance from boundaries) of the actual reward zones. We then randomly sampled from the pool of spatial maps and reward zone centers and computed the distance of the field closest to the reward zone. We repeated this process N times (with N = number of cells categorized into each remapping group, i.e. Group I, II, III or IV), took the mean, and then repeated the entire procedure 1000 times to generate a distribution of means with which to compare the mean of the data.

15 Comparison of task and no-task related trajectories

To compare activity near the reward zone during the task with no-task-related times, we created two rate maps for every ENV2 recording session: one consisting of only of task-related data (e.g. the concatenated sections of time beginning of tone onset to zone entry), and another consisting of only no-task-related times. As the number of data points and the position occupancy differed significantly between the two maps, we used the position-by-speed down-sampling procedure described above to match the data. To compute differences in activity as a function of the distance from the reward zone, we computed the tuning curve for the average activity vs distance from reward zone center for each map, range-normalized each tuning curve separately (each curve has min = 0, max = 1), and then computed the slope for each pair of tuning curves.

In addition, we computed the center of mass (COM) of the normalized activity, with respect to the distance from the reward zone, in each rate map. Smaller COM values correspond to more activity near the reward zone. To compute this measure, we used the normalized tuning curves (e.g. that in Figure 4B). For every cell, we computed:

$$COM = \frac{1}{M} \sum_{i=1}^{N=10} FR_i * dist_i$$

where $N = 10$ is the number of binned points in the tuning curve, FR_i is the normalized firing rate of bin i in the tuning curve, $dist_i$ is the distance of bin i from the reward zone center, and $M = \sum_i FR_i$. To compare this to what one would expect if there were no influence of the reward zone on cell activity, we computed a null distribution of COM values for each condition (grid-task data, grid-no task data, etc.). To compute each null distribution, we first generated a simulated dataset by randomly sampling a rate map, randomly sampling a zone location (> 30 cm from the zone location originally associated with that rate map), computing the normalized tuning curve of firing rate activity versus distance from the reward zone, and then computing the COM from the tuning curve. By doing this $n = 102$ times for grid cells, or $n = 271$ times for non-grid spatial cells, we can compute an estimate of the median COM values for that dataset. By repeating this process 500 times, we can build a null distribution of median COM values for each condition (Fig S11E).

45 To determine the degree of translation observed between grid cell ratemaps during task and

no-task times, an identical procedure to that described in “Quantification of grid cell translation” (although maps were not rotated). In addition, only cells with a grid score > 0.35 in ENV2, and with 50% coverage post-down-sampling were considered ($n = 35$).

5

Position decoding

To decode the animal’s position, we first identified pairs of sessions (1 pair = 1 ENV1 session and 1 ENV2 session) with at least 5 cells that were recorded simultaneously in both sessions. For each session in a session pair, we then determined the variables encoded by cells recorded during that session by fitting a series of LNP models. We only considered position and spike history as input variables, and we used the forward search procedure (as previously described) to determine the least complex model that could best describe the data. Similar to the previous fitting of LNP models, variables were modeled using cardinal spline functions. The position spline is identical to that described in the LNP model section, while the spline for the spiking history followed Meng et al., 2014. Specifically, a model containing only spike history input, along with a baseline parameter, is given by:

$$\hat{r}^{(t)} = \frac{\exp\left(b + \sum_{\tau=1}^M \Delta N_{t-\tau} X_{sh}^{T(t)} w_{sh}\right)}{dt}$$

where $\hat{r}^{(t)}$ is the estimated firing rate for time point t , dt is the time bin width (20 ms), b is the baseline input, τ denotes time bins in the past, $\Delta N_{t-\tau}$ is the number of spikes in the time bin given by $t - \tau$, $X_{sh}^{T(t)}$ denotes the input (computing using the splines), and w_{sh} are the learned parameters (values of the control point locations). $X_{sh}^{T(t)}$ is computed analogously to the procedure described in the LNP model section. Control points are placed at 20, 40, 60, 80, 100, 140, and 200 ms in the past. Models were fit in the same manner as described in the LNP modeling section (using 10-fold cross-validation). Prior to model fitting, 5 minutes worth of data was randomly chosen and excluded to be used later to test decoding accuracy.

Once the variables encoded, and associated parameters, were learned for each cell in a session, we then used all cells that significantly encoded position to decode the animal’s position. Knowledge of the past spiking history was used when decoding position. Position was binned into a 50 x 50 grid (of width 3 cm), and we estimated which position bin x the animal occupied for all decoding points t . To do this, we employed the principle of maximum-likelihood, and found the binned position that maximized the product of the likelihood (of the position given the observed spiking, under the LNP modeling framework) and prior on position occupancy (a quantity proportional to the posterior):

$$\hat{X}^{(t)} = \operatorname{argmax}_x \prod_{i=1}^N P\left(\sum_{\tau=0}^L \Delta N_i^{(t-\tau)} \mid \sum_{\tau=0}^L \hat{r}_i^{(t-\tau)}(x)\right) * P(x)$$

where $\hat{X}^{(t)}$ is the estimated position bin for time bin t , $\Delta N_i^{(t-\tau)}$ is the number of spikes that occurred for neuron i during time bin $t - \tau$ (so $\sum_{\tau=1}^L \Delta N_i^{(t-\tau)}$ is the total number of spikes emitted from neuron i from $t - L$ to $t - 1$), $\hat{r}_i^{(t-\tau)}(x)$ is the estimated rate for neuron i during time bin $t - \tau$ under the assumption that the animal is occupying position bin x , and $P(x)$ is the probability that the animal is in position bin x . We computed $P(x)$ empirically, and we took $L = 25$ bins (500 ms). The value of $\hat{r}_i^{(t-\tau)}(x)$ was computed for each time bin, position, and neuron using the selected model and the associated learned model parameters for each cell. For example,

if a cell encoded both previous spiking history and position, then the estimated number of spikes observed during the decoding time bin for each position would be:


$$\hat{r}_i(x) = \exp(b_i + X_p^T(x)w_p) \left(\sum_{\tau=1}^L \exp(\Delta N_i^{(t-\tau)} X_{sh}^{T(t-\tau)} w_{sh}) \right)$$

Decoding error was quantified by computing the Euclidean distance between the actual position and estimated position.

Down-sampling spikes prior to position decoding

To investigate the degree to which firing rate changes influence the reduced decoding error near the reward location, we used a downsampling procedure that matched the total firing rate (summed across neurons) across position for each ENV1 and ENV2 session pair. Specifically, for each ENV1 and ENV2 session pair, we binned position into a 30 x 30 grid and matched the average firing rate observed across all spatially-modulated cells in each position bin. Spatially-modulated cells were determined to be the cells that significantly encoded position according to the model-selection step of the previously-described decoding analysis. To match the firing rates, we randomly rejected spikes from cells recorded in the session containing more spikes. The number of spikes rejected varied from 0-70%. Once the spike-rejection procedure terminated, we then proceeded with the same down-sampling procedure described above, with the exception that it was assumed that all cells encoded position.

Fig. S1

**Fig. S1.**

Histology for all rats. (A) Histological examples for each of the seven experimental rats. Images show sagittal brain slices stained for cresyl violet. Approximate locations of tetrode tips

5

are indicated by red circles. **(B)** Histological examples from three of the experimental animals verifying that no cells were recorded from lateral entorhinal cortex (LEC). Across all animals, no tetrode tracks were observed further lateral than the medial-lateral locations shown.

Fig. S2

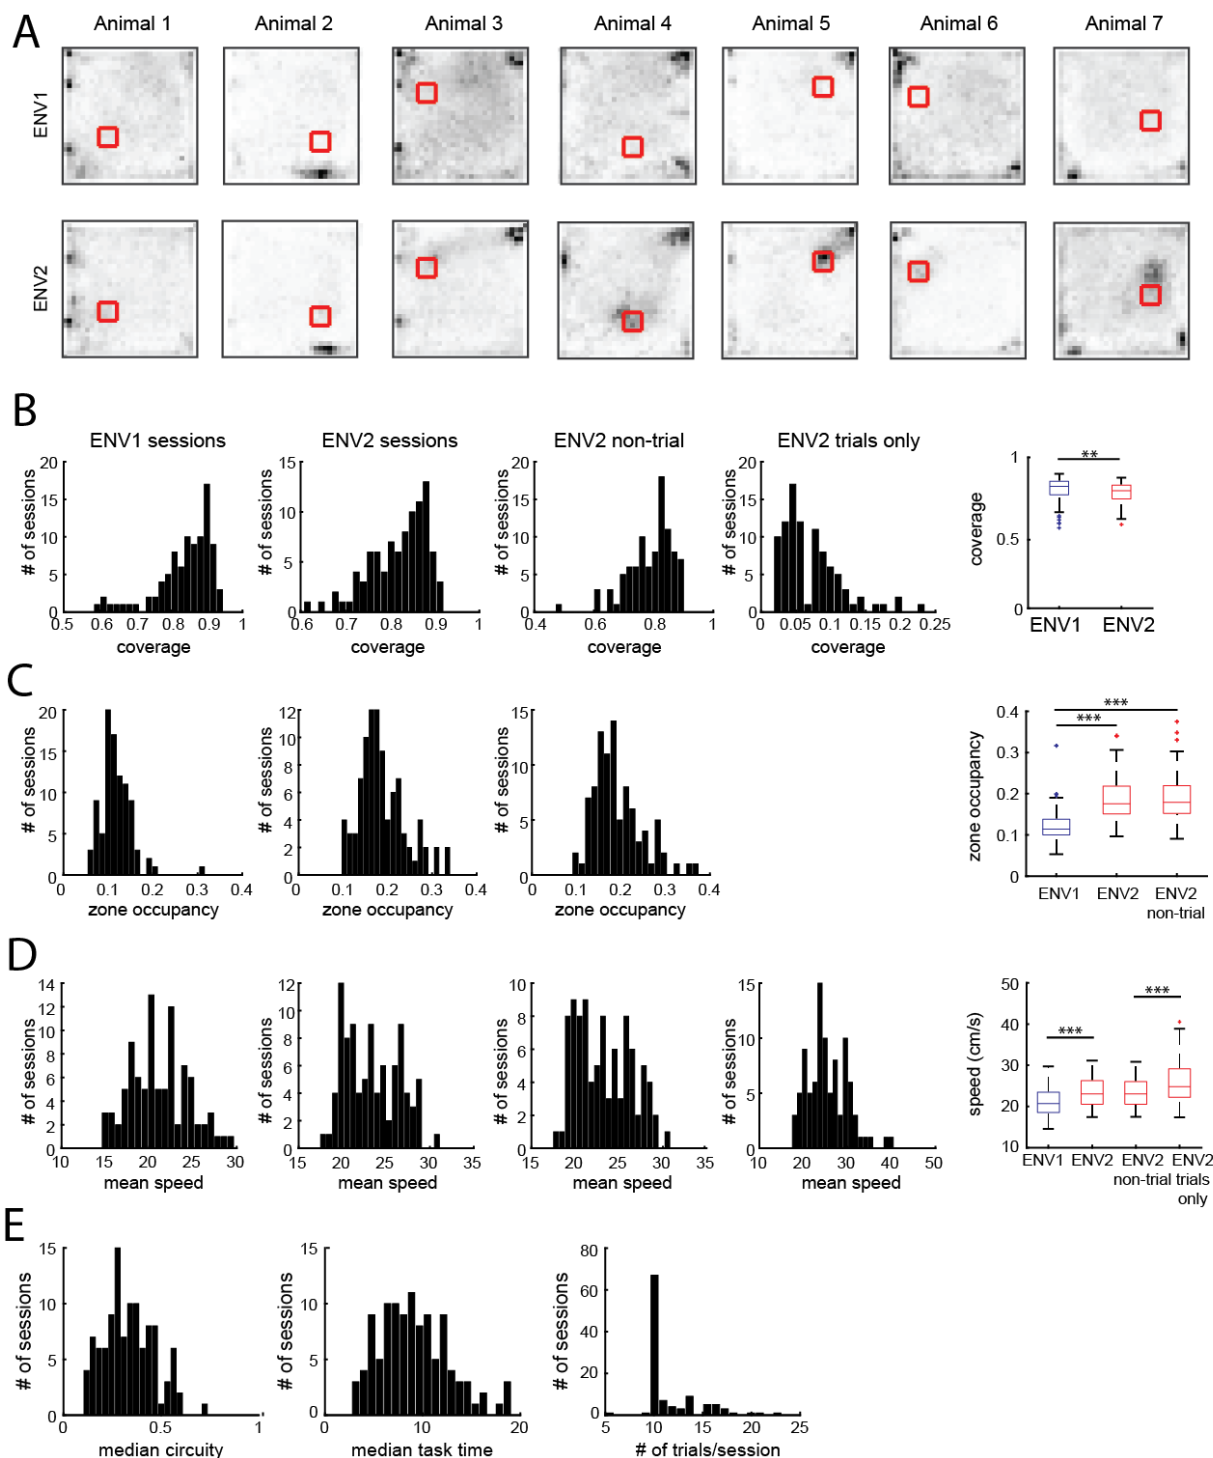


Fig. S2.

Behavior during the spatial task. (A) Average occupancy across pairs of ENV1 sessions (top row) and ENV2 sessions (bottom row) for the seven experimental animals. Only pairs of sessions recorded on the same day were considered (n = 93 sessions total, n sessions per animal left to right: n = 13, n = 12, n = 20, n = 10, n = 7, n = 16, n = 15). Red square indicates reward zone location. White = no occupancy, black = high occupancy. While animals tended to spend

5

more time near the boundaries in ENV1, animals also spent a significant portion of time near the reward zone in ENV2 (see panel C), indicating that the presence of a remembered reward location resulted in context-dependent differences in navigational behavior. **(B)** Left panels: Histograms of the coverage (computed as the fraction of 900 position bins with > 500 ms of data) for each pair of recording sessions (after training) across all rats. “ENV2 non-trial” denotes all times during ENV2 sessions after removing the trial times, while “ENV2 trials only” indicates all times during which the animal was performing the task in ENV2. Far right panel: box plot comparing the coverage for ENV1 and ENV2 sessions. Coverage for ENV1 sessions was significantly higher than ENV2 sessions (signed-rank $p = 0.003$), but ENV1 sessions were significantly longer than ENV2 sessions by an average of 3.5 minutes (median (2nd quartile, 3rd quartile) coverage for ENV1 = 33.7 (29.4, 36.3) minutes, ENV2 = 30.2 (26, 30.5) minutes, signed-rank test $p = 4e-9$). **(C)** Left panels: Histograms of the zone occupancy (computed as the fraction of the total session time spent within 30 cm of the reward zone center) for each pair of recording sessions. ENV2 trials-only is not shown because trials ended when the animal entered the reward zone (resulting in nearly zero zone occupancy). Far right panel: box plot comparing the zone occupancy for ENV1 and ENV2 sessions, as well as the non-trial component of the ENV2 sessions. Zone occupancy was significantly higher in ENV2 than ENV1 (signed-rank $p = 1e-14$), and significantly higher in ENV2 non-trials than ENV1 (signed-rank $p = 1e-14$), indicating that the animals’ behavior changes in ENV2 in a way that reflects the saliency of the reward zone. **(D)** Left panels: Histograms of the average running speed for each type of recording session. Far right panel: plot comparing the average running speeds for ENV1, ENV2, ENV2 minus task trajectory times, and task trajectory times only. Animals ran faster overall in ENV2 than ENV1 (signed-rank $p = 6e-11$), and faster during task trials than no-task trial times in ENV2 (signed-rank $p = 8e-7$). Given these differences in running speeds, we matched running speeds (and occupancy) when comparing reward-based tuning changes between ENV1 and ENV2. **(E)** Median circuitry/session (left), median task trial time/session (middle), and the number of trials per session (right) across all ENV2 sessions ($n = 109$).

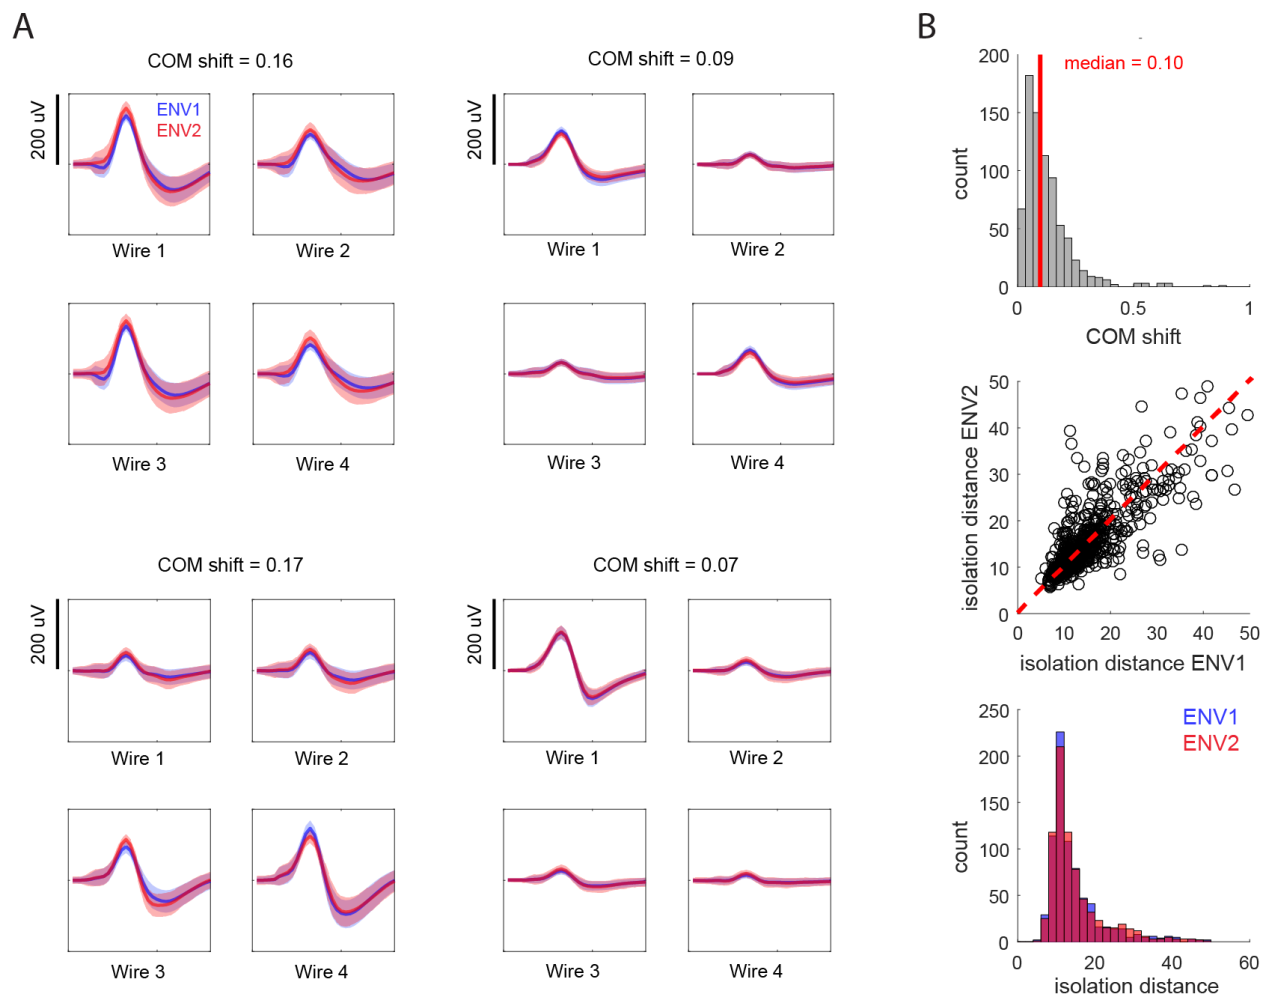


Fig. S3.

Matching recordings between ENV1 and ENV2. (A) Tetrode waveforms for four examples of cells that were recorded in both environments. Blue: mean waveform ± standard deviation (SD) in ENV1; red: mean waveform ± SD in ENV2. The calculated center-of-mass (COM) shift between the two environments is indicated above each set of waveforms. (B) Top: Histogram for the COM shifts of all 778 neurons recorded in both environments (ENV1: n = 110 sessions; ENV2: n = 108 sessions) in the experimental group. Red line indicates the median COM shift (0.10). Middle: Each cell's isolation distance in ENV2 plotted against its isolation distance in ENV1. High correlation indicates that there were no overall changes in isolation between environments (correlation coefficient $r = 0.85$). Bottom: Histogram of all isolation distance values observed in each environment. Blue: ENV1; red, ENV2. There was no difference between individual cells' isolation distances in ENV1 and ENV2 (signed-rank test, $p = 0.55$).

Fig. S4

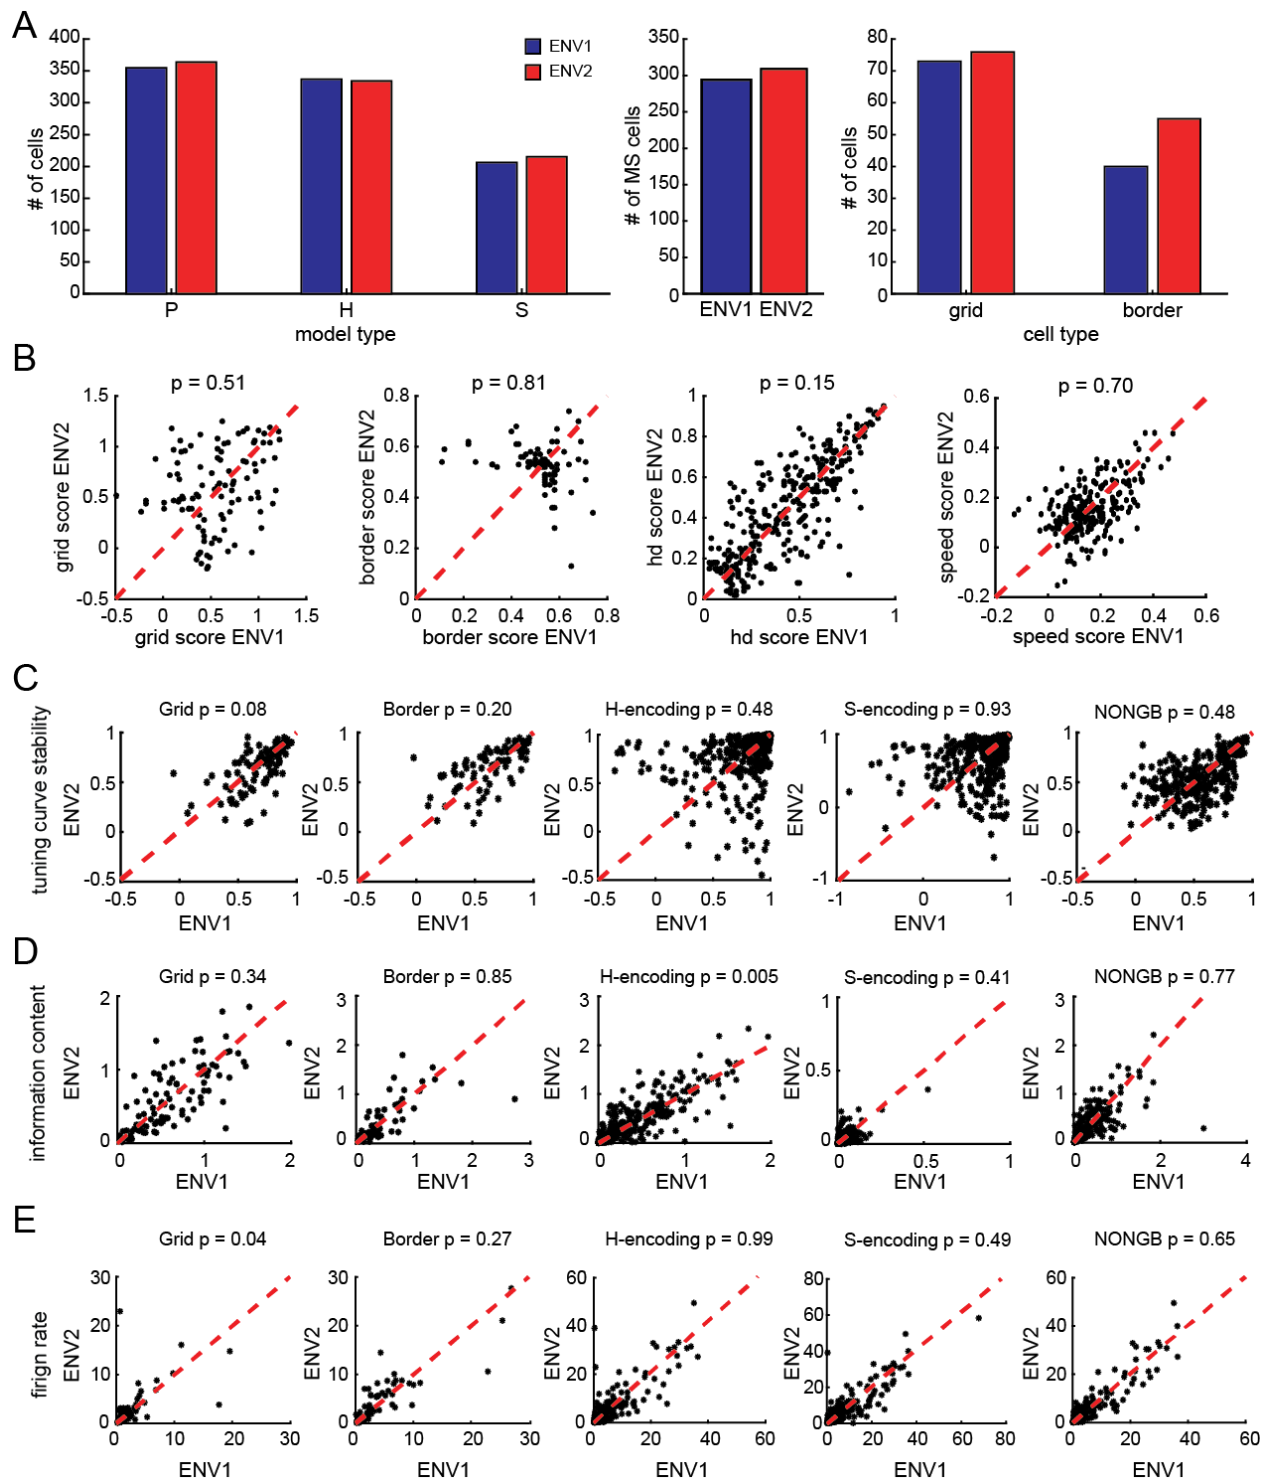


Fig. S4.

Changes in coding content, tuning scores, stability, information, and firing rate of cells between ENV1 and ENV2. (A) Left: The number of cells classified as position (P), head direction (H), or speed (S) encoding did not vary between the environments (proportion tests: P: $z = 0.46$, $p = 0.65$; H: $z = 0.15$, $p = 0.88$; S: $z = 0.51$, $p = 0.61$). Middle: The number of mixed-

5

selective cells (cells that encoded two or more navigational variables) did not significantly change between environments (proportion test, $z = 0.58$, $p = 0.56$). Right: The number of cells classified as grid or border in each environment did not change (proportion test, grid: $z = 0.26$, $p = 0.80$; border: $z = 1.59$, $p = 0.11$). **(B)** There was no significant change in the grid score of grid cells, border score of border cells, head direction score (mean vector length) of head direction cells, or speed score of speed cells (signed-rank tests, grid cell $n = 102$, $p = 0.51$; border cell $n = 79$, $p = 0.81$; HD cell $n = 304$, $p = 0.15$; speed cell $n = 234$, $p = 0.7$). For this analysis, head direction cells were defined as cells that exceeded the 95th percentile head direction score threshold (p_{95} in ENV1 = 0.115, p_{95} in ENV2 = 0.120), and encoded H (according to the LNP model) in at least one environment. Similarly, speed cells were defined as cells that exceeded the 95th percentile speed score threshold (p_{95} in ENV1 = 0.047, p_{95} in ENV2 = 0.048) and encoded S (according to the LNP model) in at least one environment. Because animals ran faster in ENV2 than ENV1 (Fig. S2), speed scores were re-computed on speed-matched data (i.e. sessions were down-sampled to match coverage across speed bins between ENV1/ENV2). **(C)** Tuning curve stability was not significantly different between grid cells, border cells, H-encoding cells, S-encoding cells, or non-grid spatial (NONGB) cells recorded in ENV1 and ENV2 (all signed-rank tests $p > 0.05$, H-encoding cells $n = 408$, S-encoding cells $n = 303$, non-grid spatial cells $n = 271$). In this analysis, H-encoding and S-encoding cells were all cells that significantly encoded H or S in either environment according to the LNP model. **(D)** Information content (measured in bits/spike) was not significantly different between grid cells, border cells, speed cells, or non-grid spatial cells recorded in ENV1 and ENV2 (all signrank $p > 0.3$). We observed a slight increase in the information content of H-encoding cells in ENV2 compared to ENV1 (median increase in score = 0.004, signrank $p = 0.005$). **(E)** Firing rate (spikes / length of recording in seconds) was not significantly different between ENV1 and ENV2 for border cells, H-encoding cells, speed cells, or non-grid spatial cells (all signed-rank tests $p > 0.2$). Grid cells fired at slightly higher rates in ENV2 than ENV1 (median firing rate, (2nd quartile, 3rd quartile): ENV1 = 1.18, (0.69, 1.9), ENV2 = 1.4, (0.78, 2.50), signed-rank test $p = 0.036$).

Fig. S5

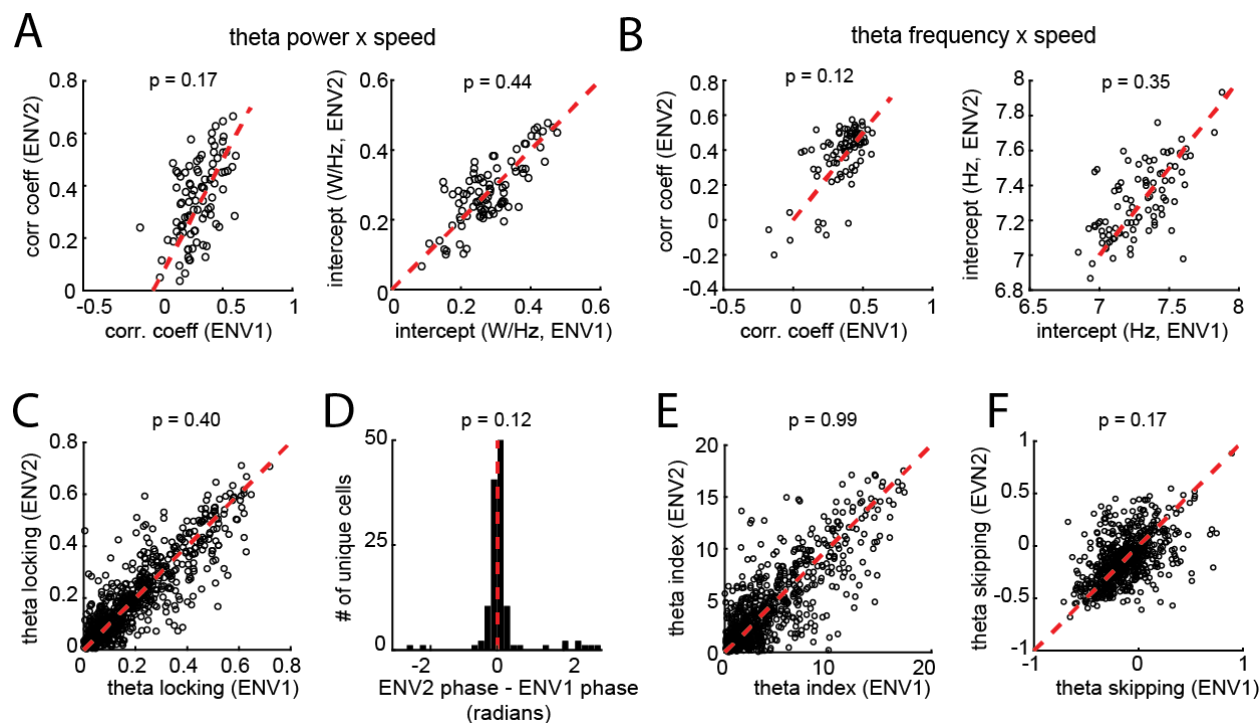


Fig. S5.

Measures of theta are unchanged between environments. (A) Between ENV1 and ENV2, there was no significant difference in the Pearson correlation between theta power and running speed (left), or the intercept of the theta power-running speed best-fit line for each session (right) ($n = 93$ sessions, median difference in correlation = 0.029 (ENV2-ENV1), signed-rank test $p = 0.17$; median difference in intercept = -0.002 W/Hz (ENV2-ENV1), signed-rank test $p = 0.44$). (B) Between ENV1 and ENV2, there was no significant difference in the Pearson correlation between theta frequency and running speed (left), or the intercept of the theta frequency-running speed best-fit line for each session (right) ($n = 93$ sessions, median difference in correlation = 0.024 (ENV2-ENV1), signed-rank test $p = 0.12$, median difference in intercept = 0.002 Hz (ENV2-ENV1), signed-rank test $p = 0.35$). (C) There was no significant difference in theta-locking score between ENV1 and ENV2 ($n = 778$ cells, median difference in theta-locking score = 0.002 (ENV2-ENV1), signed-rank test $p = 0.4$). This still held true when considering cells with a theta-locking score above 0.3 in either environment ($n = 196$ cells, median difference in theta-locking score = -0.005, signed-rank test $p = 0.3$). (D) There was no significant difference between environments in the preferred phase of the theta frequency for theta-locked cells (theta-locking score > 0.3 in both environments, $n = 122$ cells, median phase angle difference = 0.03 radians (ENV2-ENV1), signed-rank $p = 0.13$) (Materials and Methods). (E) There was no significant difference in the theta index between ENV1 and ENV2 (median difference in theta index = -0.04 (ENV2-ENV1), signed-rank test $p = 0.99$) (Materials and Methods). This lack of difference held true even when considering cells with a theta index above 5 ($n = 298$, median difference in theta index = 0.43 (ENV2-ENV1), signed-rank test $p = 0.43$). (F) There was no significant difference in theta skipping between ENV1 and ENV2 (median difference in theta skipping = -0.02 (ENV2-ENV1), signed-rank test $p = 0.17$) (Materials and Methods). This lack

of difference also held true when only considering cells with a theta index score above 5 (n = 298, median theta skipping difference = -0.01 (ENV2-ENV1), signed-rank test p = 0.53).

Fig. S6

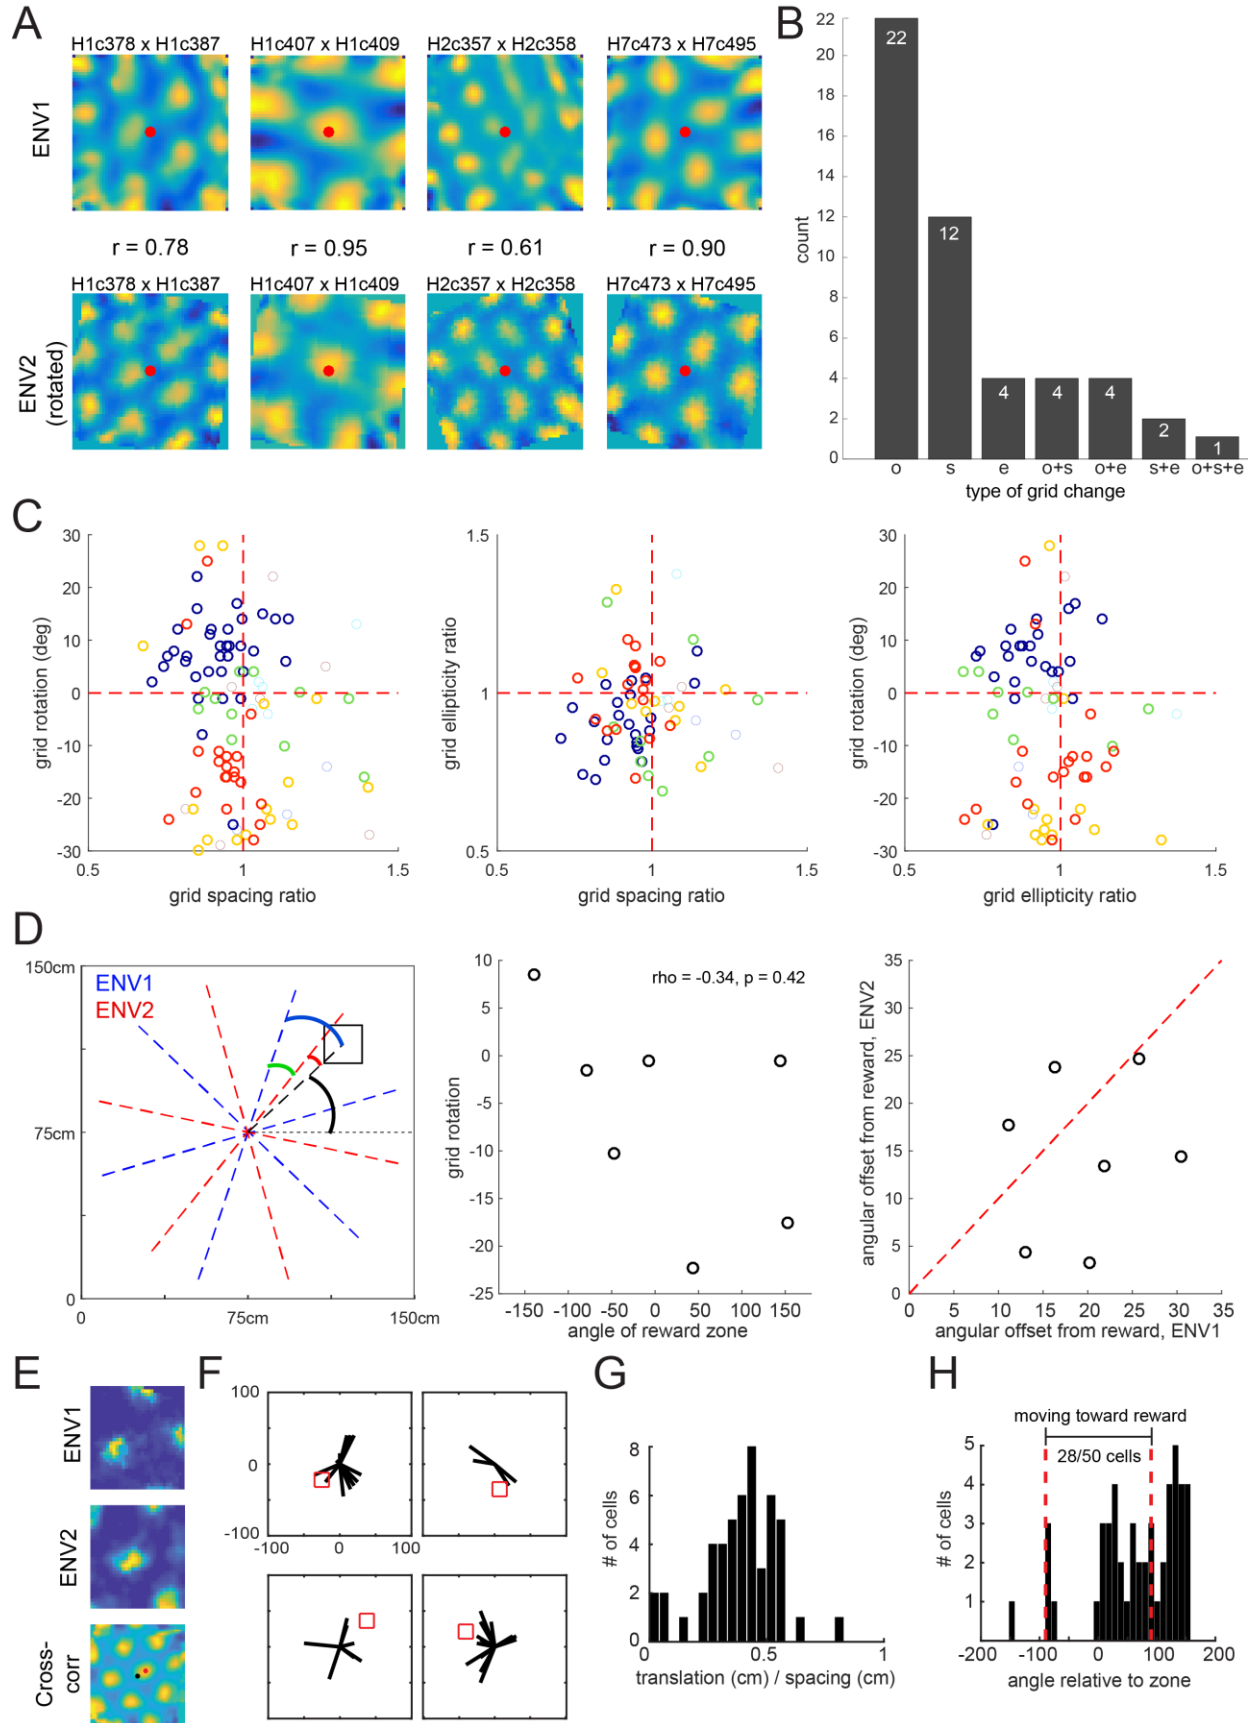


Fig. S6.

Changes in grid structure are coherent across grid cells and are unrelated to the reward zone. (A) Four cross-correlations of pairs of grid cells that were co-recorded in both environments. (Top) Cross-correlation for a pair of grid cells co-recorded in ENV1. (Bottom) Cross-correlation for the same pair of cells in ENV2. Animal and cell identities are indicated above each cross-correlation. Because grid cells' orientations changed between the two environments, the ENV2 cross-correlation was rotated to produce the maximal correlation with ENV1; the value of this correlation (r) is indicated between each pair of cross-correlations. The high degree of similarity between the ENV1 and ENV2 cross-correlations indicate that grid cells maintain their phase offset and translate in concert with one another between the two environments, even when the orientation of the grid pattern changes. Red points indicate the center of the cross-correlations. (B) Number of grid cells that showed a significant change on each grid measure, or on a combination of grid measures. o: orientation; s: spacing; e: ellipticity. (C) All 102 experimental grid cells, plotted according to their conjunctive changes in orientation, spacing, or ellipticity. Each cell is colored according to animal identity. Cells recorded from animals with at least 10 grid cells are shown with heavier lines. Grid cell changes cluster within animals. The observed changes in grid spacing may be due to attentional or arousal-mediated modulation of ion channels (36). (D) (Left) Example grid axes from a single animal, showing grid cells' overall orientation in ENV1 and ENV2 for that animal. The four measures derived from the grid axes and plotted in the middle and right panels are shown with bolded lines. The black square represents the location of the reward zone in a top-down view of the environment. (Middle) The observed changes in grid orientation between environments (green arc in left panel) were not related to the angle of the reward zone in the environment (black arc in left panel; circular-circular correlation, $\rho = -0.34$, $p = 0.42$). (Right) The angle between the reward zone angle and the closest grid axis in ENV1 (blue arc, left panel) and the closest grid axis to the reward zone in ENV2 (red arc, left panel) did not significantly differ (signed-rank test, $p = 0.16$). (E) To examine grid cell translation as a function of the reward zone location, grid patterns were cross-correlated after rotating the ENV1 rate map to match in orientation. Example rotated firing pattern in ENV1 (top), example firing pattern in ENV2 for the same cell (middle), and the cross-correlation (bottom). The black star marks the center, while the red star marks the observed translation. (F) Grid translations observed for 4 animals with > 3 grid cells recorded. The axes in the top left hold for all plots. The red box denotes the reward zone for that animal. The angle and direction of each line denotes the translation for a single grid cell recorded for that animal. (G) Distribution of the normalized translation (translation in centimeters / grid scale) observed for each grid cell. In total we observed that 45/50 grid cells exhibited significant translations (see Materials and Methods). (H) Distribution of the angle of translation relative to the reward zone. We observed that 28/50 cells translated towards the reward zone in ENV2.

Fig. S7

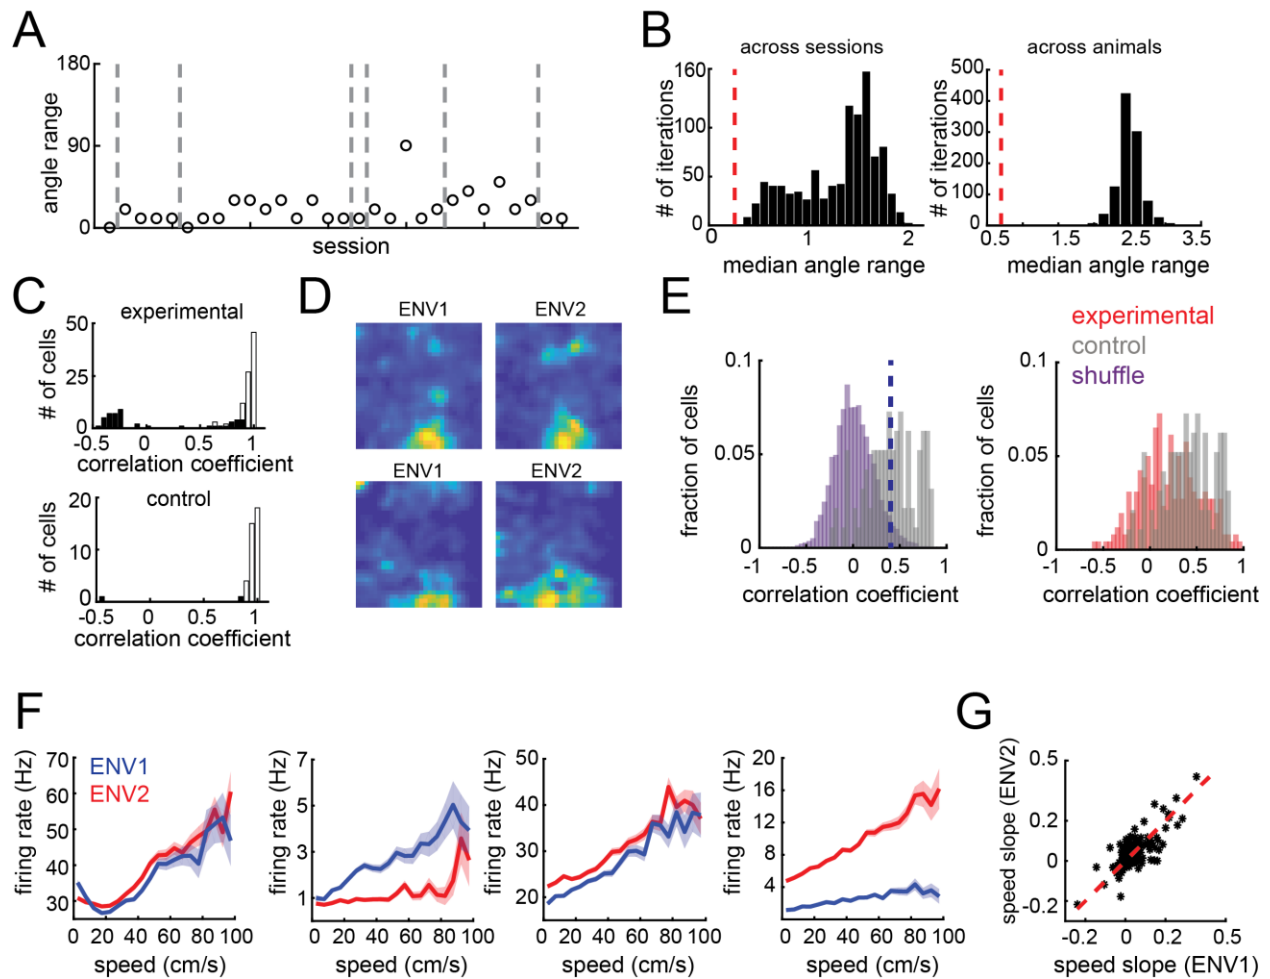


Fig. S7.

HD, non-grid spatial cells, and speed cells in ENV1 and ENV2. (A) Range of rotation angle (computed as shown in Fig. 2) observed across HD cells for each session (n = 30). Gray lines denote boundaries between animals (similar to Fig. 2A). (B) Null distribution of median angle ranges when cells are shuffled across sessions (left) or across animals (right). The experimentally-measured median angle range across sessions (left) or across animals (right) is shown as a red dashed line. In both cases, the actual median lay to the left of all points in the null distributions, suggesting that HD cells rotate coherently with one another. (C) Histogram of correlation coefficients between HD tuning curves recorded from ENV1 and ENV2 in the experimental dataset (top) or the control dataset (bottom). For both histograms, black bars correspond to cells that significantly re-mapped, while white bars indicate cells that did not significantly re-map. For the experimental dataset, 48/132 cells significantly re-mapped, while 2/39 cells remapped in the control dataset. Further, the control dataset had significantly lower correlation coefficients overall compared to the experimental dataset (ranksum $p = 0.0018$). (D) Example non-grid spatial cells recorded from control animals. Note that cells largely retain their tuning across environments. (E) Left: In control animals, a large fraction (54/100) of non-grid spatial cells significantly retain their tuning between ENV1 and ENV2 (gray). Cells were considered to significantly retain their tuning if they showed a correlation above the 95th percentile (denoted by the dotted black line, $P95 = 0.38$) of the shuffled distribution (purple),

which was created by computing the correlations between all cells recorded in each environment. Right: non-grid spatial cells recorded in the control animals (gray) showed less remapping between ENV1 and ENV2 than those recorded in the experimental animals (red) (control cell $n = 100$, experimental cell $n = 271$, rank-sum test $p = 2e-7$). **(F)** Firing rate (FR) x speed tuning curves for four S-encoding cells (that encoded S in at least one environment) recorded in both ENV1 and ENV2. Speed tuning curve shape is largely unchanged between the environments. **(G)** There was no significant difference in S-encoding cells' speed slope (computed between running speed and firing rate) between ENV1 and ENV2 ($n = 303$, median (2nd quartile, 3rd quartile) slope ENV1 = 0.013 (-3e-4, 0.03), ENV2 = 0.01 (0.002, 0.03), signed-rank test $p = 0.67$). There was also no significant difference when considering only S-encoding cells that surpassed the speed score threshold ($n = 234$, median (2nd quartile, 3rd quartile) slope ENV1 = 0.019 (0.004, 0.04), ENV2 = 0.02 (0.007, 0.05), signed-rank test $p = 0.3$).

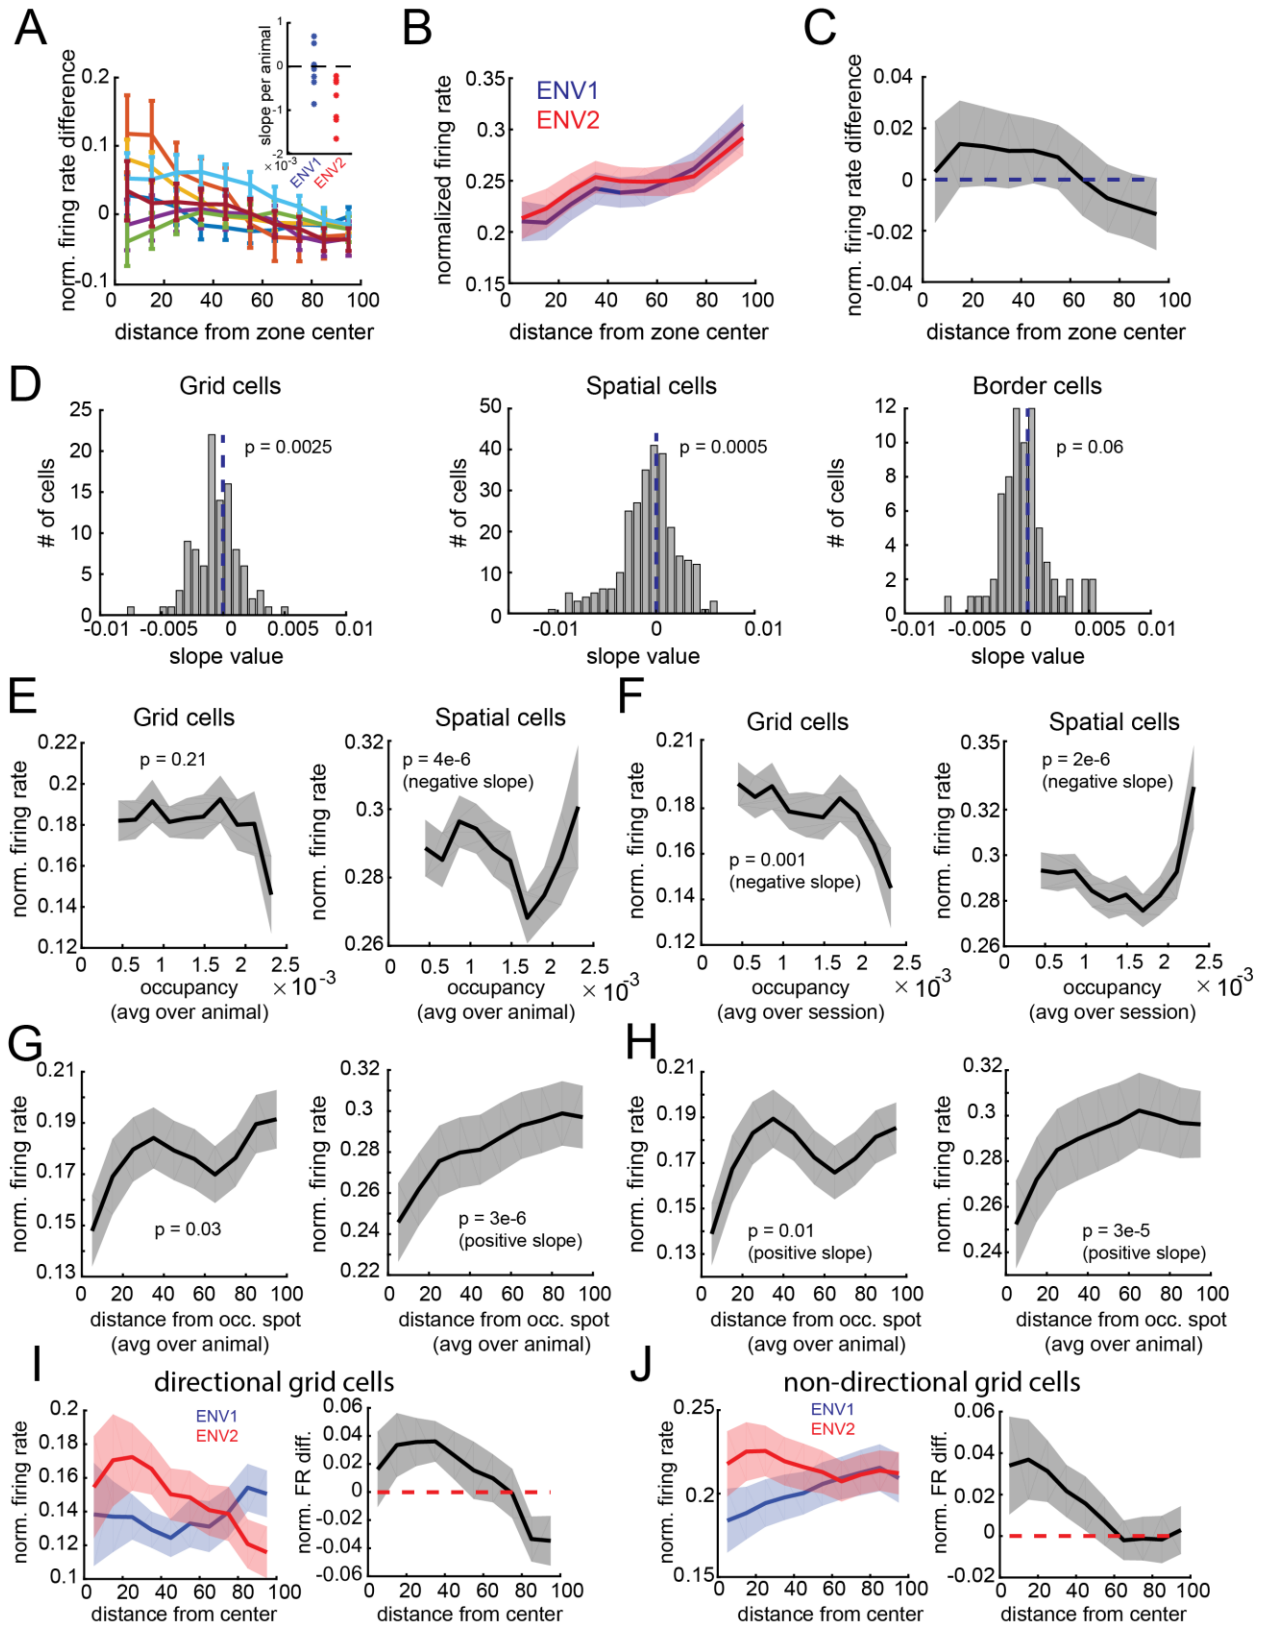


Fig. S8.

Cells remap to preferentially encode reward zone location in ENV2. (A) The normalized firing rate difference vs distance from reward zone center for non-grid spatial cells in the seven experimental animals. Six out of the seven animals showed a negative slope (increased firing as distance to reward decreased). Individual animals shown in different colors. (Inset) Slopes for ENV1 and ENV2 for each animal. All ENV2 slopes were negative (signed-rank $p = 0.01$), while ENV1 slopes across animals were distributed around 0 (signed-rank $p = 0.8$; slopes between environments were different with signed-rank $p = 0.03$). (B) The average normalized firing rate vs distance from reward zone center in ENV1 (blue) and ENV2 (red) for border cells recorded in both environments. Shading indicates the standard error of the mean across cells. The increase in firing rate at greater distances from the reward zone can be attributed to the zone not being located near the borders of the box. (C) The average difference between the normalized firing rate vs distance from reward zone center tuning curves (i.e. the difference between each cell's ENV1 and ENV2 tuning curves), across all border cells. (D) Distribution of the difference in activity vs distance slope values (ENV2 slope – ENV1 slope). Each value in the histogram is computed from the slopes of the normalized firing rate vs distance from reward zone center tuning curves in ENV1 and ENV2. The slopes in ENV2 are significantly more negative than the slopes in ENV1 for grid cells (ENV1 median slope = $5e-4$, ENV2 median slope = $8e-5$, median slope difference = $-6e-4$, signed-rank test $p = 0.0025$), and non-grid spatial cells (ENV1 median slope = $3e-4$, ENV2 median slope = $-4e-4$, median slope difference = $-5e-4$, signed-rank test $p = 5e-4$), and a non-significant effect for border cells (ENV1 median slope = $6e-4$, ENV2 median slope = $7e-4$, median slope difference = $-5e-4$, signed-rank test $p = 0.06$). (E) The average normalized firing rate versus average occupancy tuning curves computed from data collected in ENV1 for grid cells (left) and non-grid spatial cells (right). We chose to focus only on ENV1 as occupancy was biased toward the reward location in ENV2. Occupancy was computed by taking the average occupancy across all sessions for that animal (as in Figure S2, position binned into a 30×30 grid). Tuning curves for each cell were created by first plotting the normalized firing rate of each spatial bin versus the average occupancy for that bin, and then taking the average normalized rate along binned occupancy values ($n = 10$ bins, distributed uniformly across the observed occupancy values). Grid cells did not show any change in normalized activity with occupancy (signed-rank test $p = 0.21$), while non-grid spatial cells showed a slight decrease in activity with increasing occupancy (signed-rank test $p = 4e-6$). (F) The average normalized firing rate vs occupancy tuning curves computed for grid cells (left) and non-grid spatial cells (right). In these panels, the occupancy was computed for each ENV1 session individually. Thus, tuning curves were created by plotting the normalized firing rate of each position bin versus the occupancy during that session, and taking the average normalized firing rate along binned occupancy values (binned as in panel E). Both grid and non-grid spatial cells showed a slight decrease in activity with increasing occupancy (grid: signed-rank test $p = 0.001$, spatial: $p = 2e-6$). In both cases, this indicates that activity decreases with occupancy. (G) The average normalized firing rate vs the distance from the most occupied location in the ENV1 environment for grid cells (left) and non-grid spatial cells (right). The highest-occupied spot in ENV1 was computed by finding the location with the highest average occupancy across all sessions for that animal. Tuning curves were created in a manner identical to that in (B), with the exception that the recordings were from ENV1 only, and the reference location was the highest-occupied location instead of the reward zone center. Both grid and non-grid spatial cells showed an increased in activity with increasing distance from the reward zone center, which is the opposite

effect observed when considering the reward zone center as the target location (grid: signed-rank test $p = 0.03$, spatial: $p = 3e-6$). **(H)** Same as **(G)**, but the most occupied spot was computed for each session individually, instead of averaged within an animal. Again, grid and non-grid spatial cells showed an increased in activity with increasing distance from the reward zone center, which is the opposite effect observed when considering the reward zone center as the target location. **(I)** (Left) Average normalized activity for ENV1 (blue) and ENV2 (red) as a function of reward zone distance for directional grid cells. (Right) Difference in firing (ENV2-ENV1). Directional grid cells showed a significant increase in activity near the reward zone in ENV2 (median slope < 0 with signed-rank $p=0.047$). Directional grid cells were defined as grid cells with an average MVL (averaged across ENV1 and ENV2) of greater than 0.2. **(J)** (Left) Average normalized activity for ENV1 (blue) and ENV2 (red) as a function of reward zone distance for non-directional grid cells. Non-directional grid cells were grid cells with an average MVL (averaged across ENV1 and ENV2) less than 0.2. (Right) Difference in firing (ENV2-ENV1). Similar to directional grid cells, non-directional grid cells showed a significant increase in activity near the reward zone in ENV2 (median slope < 0 with signed-rank $p=0.03$).

Fig. S9

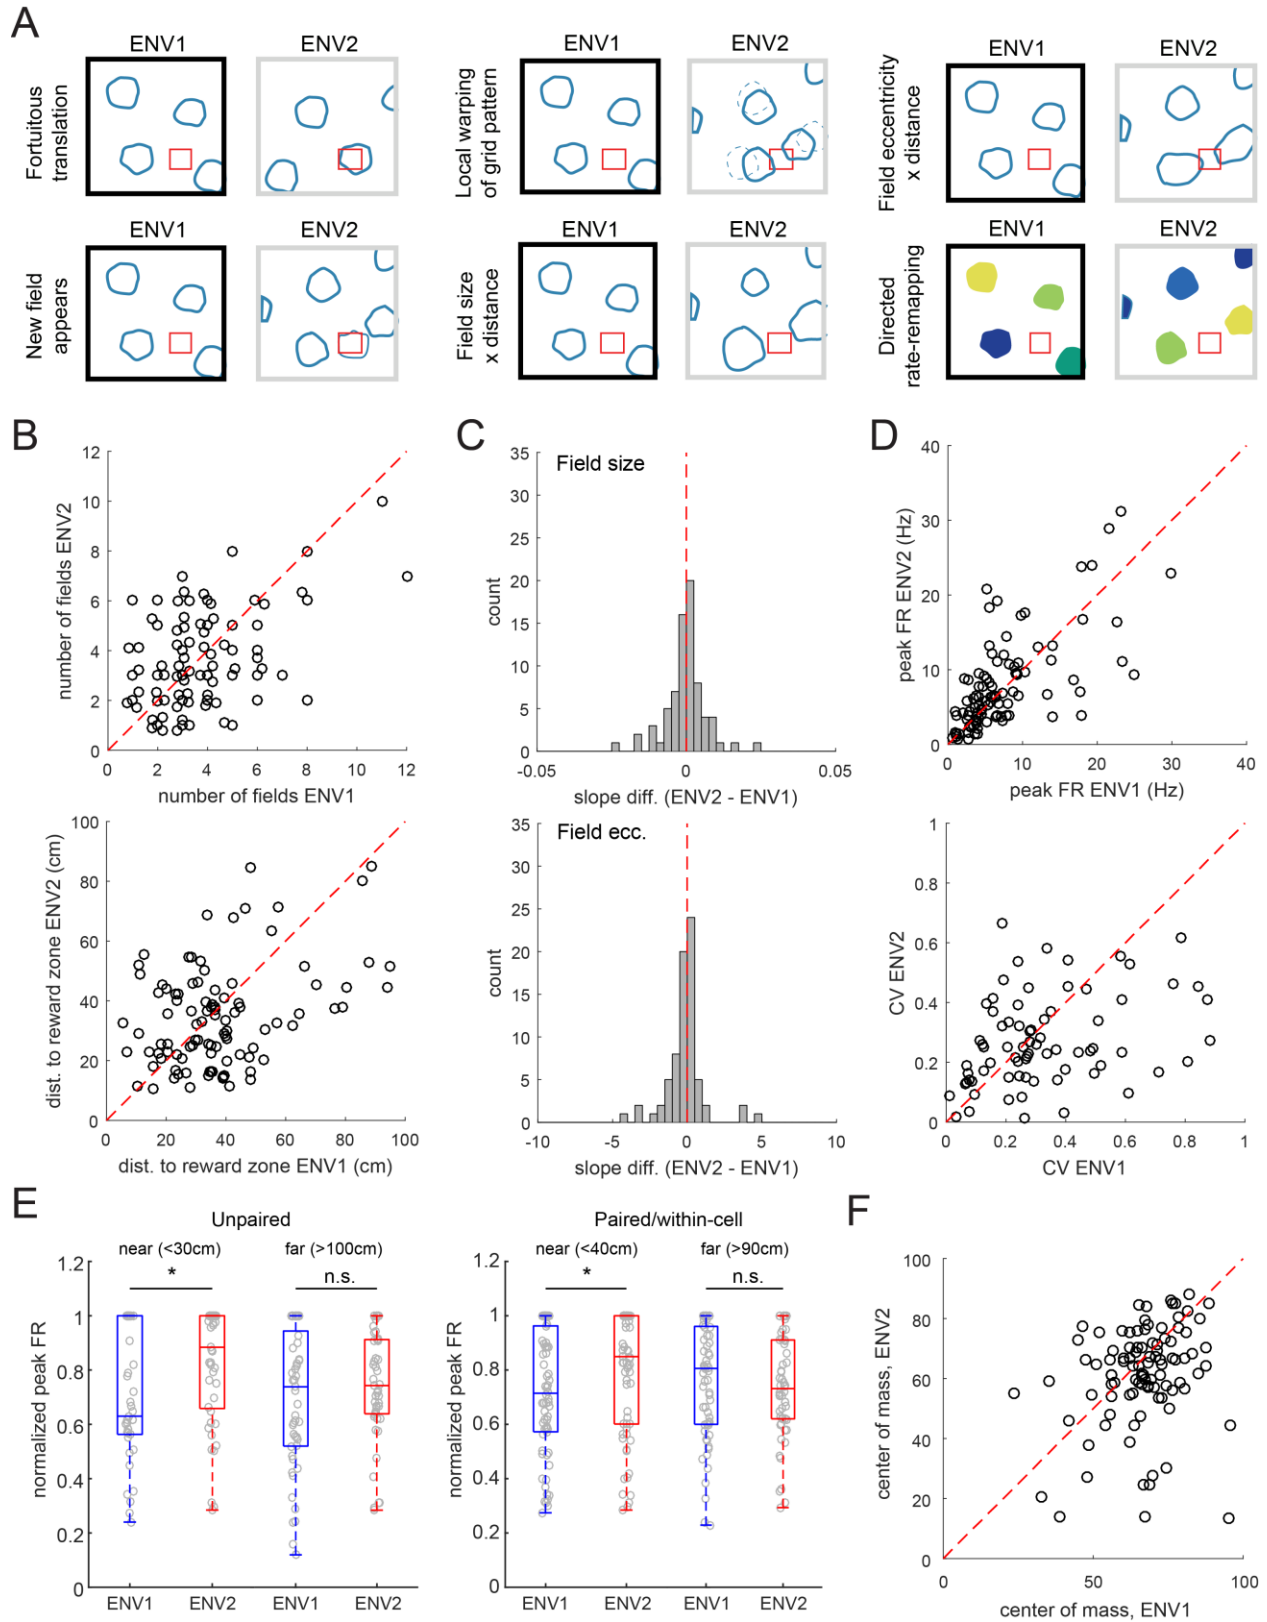


Fig. S9.

Other measures of grid fields are unaltered. (A) Schematic representation of possible grid cell structural changes that would result in an increased firing rate near the reward zone in ENV2. *Fortuitous translation*: individual grid cells translate independently to align a firing field near the reward zone. *New field appears*: a new grid field is formed at the reward zone. *Local warping*: grid fields are ‘drawn’ toward the reward zone. *Field size x distance*: fields change size as a function of their distance from the reward zone. *Field eccentricity x distance*: fields change shape as a function of their distance from the reward zone. *Directed rate-remapping*: field locations remap between environments without regard for the reward zone, but the firing rates within each field remap non-randomly such that fields nearer to the reward zone have higher firing rates. (B) (Top) The number of grid fields observed in each grid cell did not differ between ENV1 and ENV2 (signed-rank, $p = 0.59$). Red line indicates unity; points are jittered to minimize overlap since field count is always a whole number. (Bottom) The distance to the center of the closest grid field from the center of the reward zone did not differ between ENV1 and ENV2 (signed-rank, $p = 0.14$). (C) (Top) Histogram of the within-cell differences in slope of two best-fit lines: grid field size vs. distance from the reward zone in ENV2, and grid field size vs. distance from the reward zone in ENV1. (Bottom) Histogram of the within-cell differences in slope of two best-fit lines: grid field eccentricity vs. distance from the reward zone (ENV2), and grid field eccentricity vs. distance from the reward zone (ENV1). Neither measure showed a significant differential effect of distance between the two environments (signed-rank tests, field size and field eccentricity slopes: $p = 0.31$ and $p = 0.97$, respectively). (D) (Top) The peak field FR observed across all grid fields did not differ between ENV1 and ENV2 (signed-rank test, $p = 0.83$). (Bottom) The coefficient of variation (CV; standard deviation of field peak FRs/mean of field peak FRs) for each grid cell did not differ between ENV1 and ENV2 (signed-rank test, $p = 0.07$). (E) (Left) Across cells, fields close (<30cm) to the reward zone have higher firing rates in ENV2 than ENV1 (rank-sum test, $p=0.028$), while fields far (>100cm) from the reward zone show no difference in firing rates between the two environments (rank-sum test, $p=0.43$). (Right) Within grid cells with close or far fields in both environments, fields close (<40cm) to the reward zone have a higher firing rate in ENV2 (signed-rank test, $p=0.02$), while fields far (>90cm) from the reward zone have similar firing rates between the two environments (signed-rank test, $p = 0.81$). (F) For each cell, the firing field center of mass was computed as

$$COM = \frac{1}{M} \sum_{i=1}^N FR_i * dist_i$$

where N is the number of firing fields, FR_i is the firing rate of firing field i , and $dist_i$ is the distance of firing field i from the reward zone location and M is the sum of firing field rates. The center of mass (COM) for all firing fields was significantly closer to the reward zone in ENV2 than in ENV1 (signed-rank test, $p=0.009$).

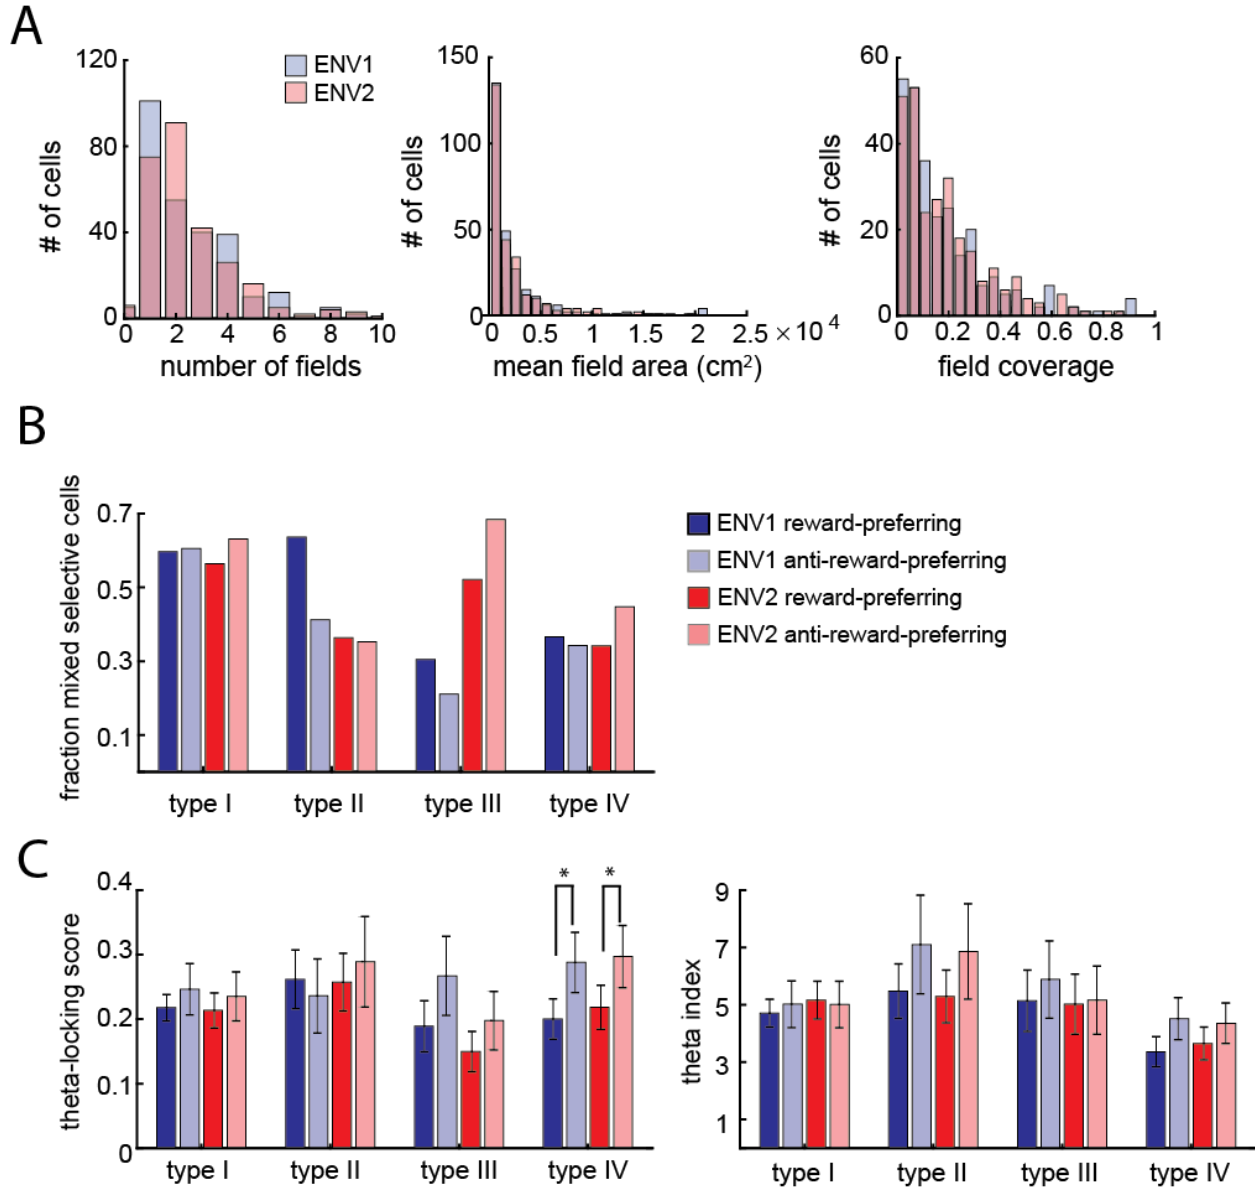


Fig. S10.

Non-grid spatial cells don't change field metrics and exhibit heterogeneous remapping. (A)

In non-grid spatial cells, the number of fields (left), average field size (middle), and the total field coverage (right; fraction of environment covered by fields) did not significantly change between ENV1 and ENV2 (field number: median field number for ENV1 = 2, ENV2 = 2, median difference in fields = 0, signed-rank test $p = 0.81$; field size: median field size for ENV1 = 1225 cm², ENV2 = 1106 cm², median difference in size = 5 cm², signed-rank test $p = 0.35$, total field coverage: median field coverage for ENV1 = 0.13, ENV2 = 0.14, median difference in coverage = 0.001, signed-rank test $p = 0.84$). Further, we did not observe any difference in field properties when considering only cells with high spatial coherence in both environments (spatial coherence > 2.25, $n = 100$, field number: median field number for ENV1 = 1, ENV2 = 2, median difference in fields = 0, signed-rank test $p = 0.37$, field size: median field size for ENV1 = 1958 cm², ENV2 = 1969 cm², median difference in size = 383 cm², signed-rank test $p = 0.13$; total

field coverage: median field coverage for ENV1 = 0.15, ENV2 = 0.17, median difference in coverage = 0.008, signed-rank $p = 0.70$). **(B)** For each non-grid spatial cell re-mapping type (I-IV), we computed the fraction of mixed selectivity for cells that exhibited reward-preferring remapping (activity increased more near reward zone in ENV2 than ENV1, dark red and dark blue) and anti-reward-preferring re-mapping (activity decreased more near reward zone in ENV2 than ENV1, light red and light blue). Mixed selectivity was computed based on the encoded variables observed for each cell when it was recorded in ENV1 (blue colors) or ENV2 (red colors). Across all types of re-mapping, and for each way of computing mixed selectivity, we did not see a significant difference in the degree of mixed selectivity between cells that did or did not show reward-preferring remapping (all proportions significance test; type I/ENV1: $p = 0.93$, type I/ENV2 = 0.51; type II/ENV1: $p = 0.13$, type II/ENV2 = 0.94; type III/ENV1: $p = 0.49$, type III/ENV2 = 0.29; type IV/ENV1: $p = 0.83$, type IV/ENV2 = 0.34). **(C)** Left: for each remapping type and environment, we computed the theta-locking score for cells that exhibited reward-preferring remapping and anti-reward-preferring remapping. For types I-III, we saw no difference in the theta locking score between cells that did or did not show reward-preferring remapping (all rank-sum test; type I/ENV1: $p = 0.35$, type I/ENV2 = 0.28; type II/ENV1: $p = 0.71$, type II/ENV2 = 0.49; type III/ENV1: $p = 0.20$, type III/ENV2 = 0.28). However, for type IV remapping cells, we observed a lower theta-locking score for cells that showed a reward-preferring remapping (type IV/ENV1: $p = 0.03$, type IV/ENV2 = 0.04). Right, same as the left, but computed for theta index. For all types of re-mapping, and for the theta index calculation in both environments, we did not see a significant difference in the magnitude of the modulation between cells that did or did not show reward-preferring remapping (all rank-sum test; type I/ENV1: $p = 0.97$, type I/ENV2 = 0.69; type II/ENV1: $p = 0.26$, type II/ENV2 = 0.20; type III/ENV1: $p = 0.36$, type III/ENV2 = 0.60; type IV/ENV1: $p = 0.16$, type IV/ENV2 = 0.40).

25

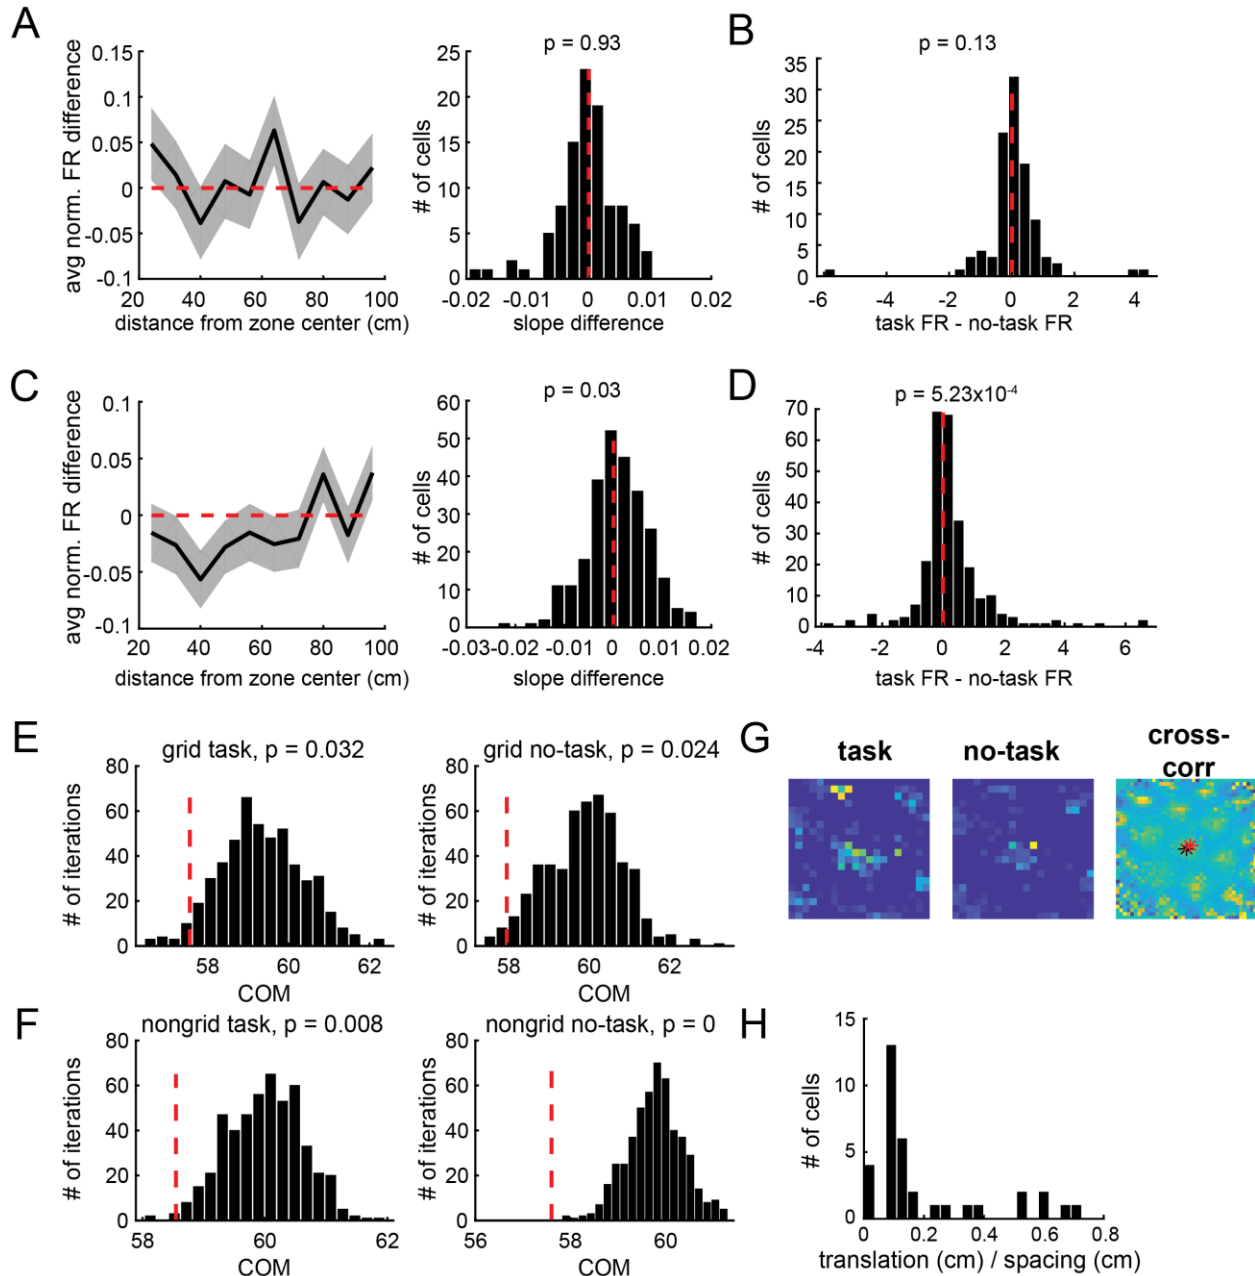


Fig. S11.

Firing rate during task-related and no-task trajectories. (A) Left: the average difference in the normalized firing rate by distance tuning curves (task-related trajectories - no-task-related trajectories) across all grid cells (analogous to Figure 3A,B; $n = 102$ grid cells). Right: Histogram of the distribution of difference in slopes computed from the task-related and no-task-related activity vs distance tuning curves (in Figure 4). There was no significant difference from zero (median task slope = $-1e-3$, median no-task slope = $-2e-3$, median difference in slopes = $-3e-5$, signed-rank test $p = 0.93$). (B) Distribution of the difference in mean firing rate, computed across down-sampled spatial bins, for grid cells. There was no significant difference in spatially-averaged firing rate (median task FR = 1.7 Hz, median no-task FR = 1.6 Hz, median difference in FR = 0.06 Hz, signed-rank test $p = 0.13$). In addition, we did not observe a difference when

mean firing rate was computed over time (average spikes/recording length in seconds, averaged across down-sampled iterations) (median task FR = 1.5 Hz, median no-task FR = 1.4 Hz, median difference = 0.06 Hz, signed-rank test $p = 0.06$). **(C)** Left: the average difference in the normalized firing rate by distance tuning curves (task-related trajectories - no-task-related trajectories) across all non-grid spatial cells. Right: Histogram of the distribution of slopes, computed from data in the left panel. There was a significant difference from zero (median task slope = -0.001, median no-task slope = -0.003, median difference in slopes = $7e-3$, signed-rank test $p = 0.03$), where no-task-related trajectories showed a steeper slope than task-related trajectories. **(D)** Distribution of the difference in mean firing rate, computed across down-sampled spatial bins, for non-grid spatial cells. Cells tended to fire more while the animal performed the task (median task FR = 1.83 Hz, median no-task FR = 1.77 Hz, median difference in FR = 0.07 Hz, signed-rank test $p = 5e-4$). In addition, we also observed a difference when mean firing rate was computed over time (average spikes/recording length in seconds, averaged across down-sampled iterations) (median task FR = 1.8 Hz, median no-task FR = 1.7 Hz, median difference = 0.06 Hz, signed-rank test $p = 0.002$). **(E)** The center of mass (COM, computed from all spatial bin firing rates) of grid cell activity is closer to the reward zone than would be expected if there was no relationship between activity and the reward zone center for both task times (left) and no-task times (right). The black bar correspond to the null distribution of median COM values across 500 null datasets, while the red line is the median COM values for grid cells during task (right) or no-task times (left). **(F)** Same as (E), but for non-grid spatial cells. **(G)** Example of rate-maps for a single cell, computed during task (left) and no-task times (middle), and the cross-correlation between them (right). The black star indicates the center of the cross correlation, while the red star indicates the translation. **(H)** A histogram of the normalized translations (translation / grid spacing) observed across 35 grid cells. There was no evidence of translation of the grid pattern between task- and no-task trajectories.

Fig. S12

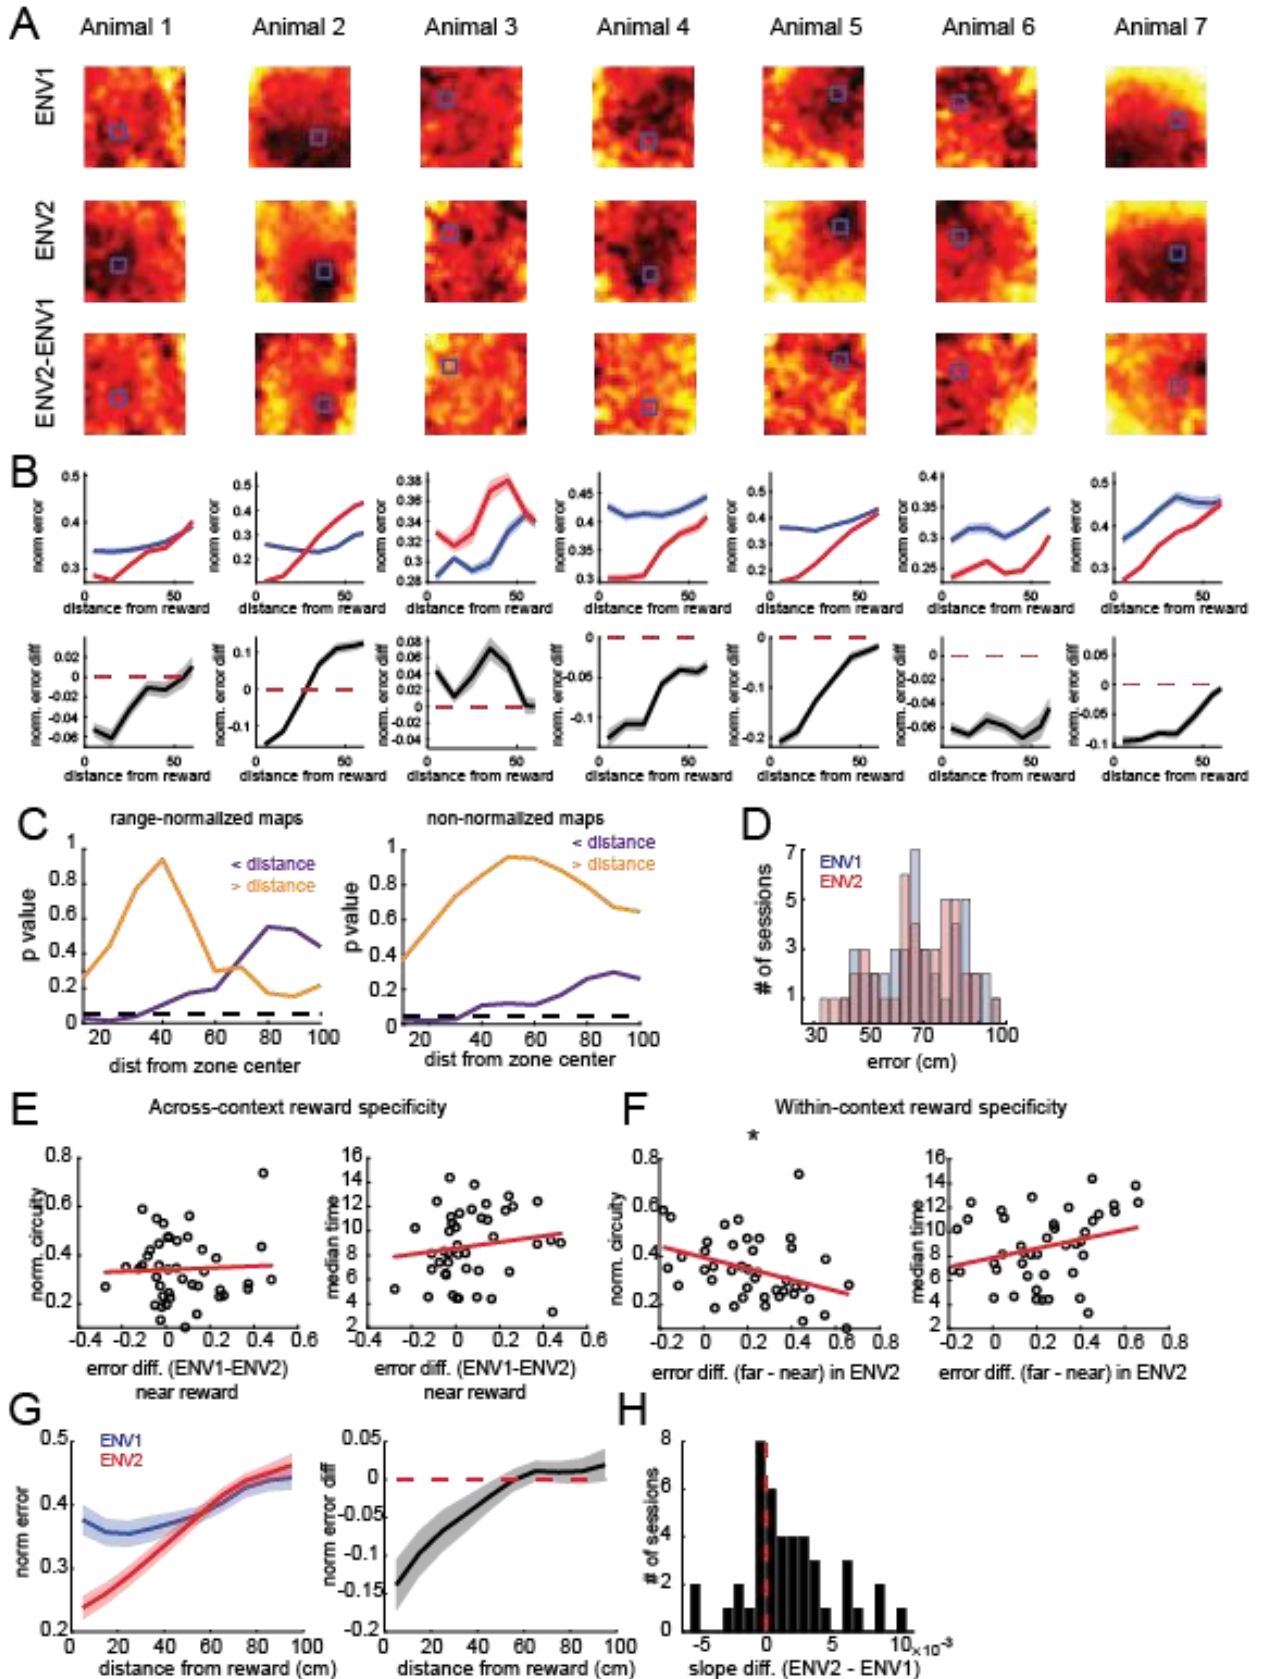


Fig. S12.

Decoding accuracy increases with proximity to reward zone. (A) Top row of images: average decoding error across all ENV1 sessions, by animal. Yellow colors indicate high error, while dark colors indicate low error. Reward zone is denoted in blue. Middle row of images: average decoding error across all ENV 2 sessions, by animal. Bottom row of images: Difference in decoding error (ENV2 - ENV1). Yellow colors indicate regions in which ENV1 accuracy is greater than ENV2 accuracy, while darker colors indicate regions in which ENV2 accuracy is greater than ENV1 accuracy. (B) Top row of plots: the average normalized decoding error as a function of distance from the reward zone center for each environment. Blue is ENV1, while red is ENV2. Bottom row of plots: the average difference between the average normalized decoding error vs distance tuning curves. Negative numbers indicate that ENV2 is more accurate than ENV1. Note that in the majority of animals, ENV2 becomes more accurate than ENV1 as distance to the reward zone center decreases. (C) To investigate the range over which ENV2 decoding is more accurate than ENV1 decoding, we computed signed-rank tests on the average decoding error between sessions. We considered the average decoding error within a certain distance from the reward zone center (orange), and also average decoding error a certain distance away from the reward zone center (purple). We also considered the differences computed from the range-normalized error maps (left), and non-normalized error maps (right). Each plot shows the p-value for the signed-rank test computed at each distance. Note that in both plots only decoding error within 30 cm of the reward zone is significantly lower in ENV2 than ENV1. (D) There was no significant difference in the overall decoding error between environments averaged across all spatial bins ($n = 43$ sessions, median decoding error in ENV1 = 68.1 cm, median decoding error in ENV2 = 68.5 cm, median difference in decoding error = 0.7 cm, signed-rank test $p = 0.33$). (E, F) These plots investigate whether there is a relationship between relative decoding accuracy near the reward zone and the animal's performance during that task session. We took each of the 43 sessions in which we decoded position in ENV1 and ENV2, computed error maps" (as in Figure 4) and range-normalized each error map to go between 0 and 1. We then computed the average error near the reward zone in ENV2 (<30 cm, shown in the paper to be the region in which there is significantly higher decoding accuracy in ENV2), as well as the equivalent zone in ENV1, and average normalized decoding error far (>100 cm) from the reward zone in ENV2. We then computed two measures: across-context reward specificity (decoding error near reward in ENV1 – decoding error near reward in ENV2, panel E), and within-session reward specificity (decoding error far from reward in ENV2 – decoding error near reward in ENV2, panel F). For across-context specificity, we did not see a significant correlation between decoding error and task performance, measured by either circuitry or task time (E left: $\rho = 0.05$, $p = 0.8$; E right: $\rho = 0.15$, $p = 0.34$). Further, for the within-context specificity, we did not observe a significant correlation between decoding error difference and median task time (F right: $\rho = 0.28$, $p = 0.07$). However, we did observe a significant correlation between decoding error difference and the median normalized circuitry, although error was higher near reward, compared to far from it, for low circuitry (low task performance) (F right: $\rho = -0.37$, $p = 0.01$). (G) Difference in decoding error once spikes were downsampled to match average firing rate across position for each environment (see Methods). Left: Normalized error versus distance from reward zone for each environment, averaged over all decoding sessions ($n = 43$). Right: Average difference in error (ENV2-ENV1) for all sessions. (H) Distribution of slopes of ENV2-ENV1 tuning curves across sessions. Histogram was right shifted with signed-rank $p = 0.002$.

Table S1

Changes in grid cell structure (Fig. 1)

All values are shown as median (25th, 75th quantiles)

Paired grid cells (n=102)

5	<i>measure</i>	<i>ENV1</i>	<i>ENV2</i>	<i>test</i>	<i>p-value</i>
	abs. orientation change (ENV2 - ENV1)		12.53 (7.10, 24.82)	signed-rank	1.12x10 ⁻¹²
	spacing	99.14 (85.85,107.37)	93.07 (82.22,104.77)	signed-rank	0.015
	ellipticity	1.24 (1.18, 1.44)	1.17 (1.08, 1.25)	signed-rank	0.006
	mean field size (cm ²)	1420 (925, 2013)	1325 (1027, 1863)	signed-rank	0.85
10	autocorr. peak size (cm ²)	675 (525, 975)	625 (525, 875)	signed-rank	p=0.16

Unpaired grid cell recordings (ENV1 n = 339, ENV2 n = 362)

15	<i>measure</i>	<i>ENV1</i>	<i>ENV2</i>	<i>test</i>	<i>p-value</i>
	wall offset (A _{min})	7.66 (3.65, 12.44)	10.35 (5.49, 14.02)	rank-sum	0.001

Experimental grid cells (n=102) vs Control grid cells (n=27)

20	<i>measure</i>	<i>experimental</i>	<i>control</i>	<i>test</i>	<i>p-value</i>
	abs(orient_change)	12.53 (7.10, 24.82)	3.82 (1.62, 7.35)	rank-sum	5.88x10 ⁻⁵
	spacing ratio	0.95 (0.88, 1.05)	1.01 (0.96, 1.96)	rank-sum	0.19
	ellipticity ratio	0.94 (0.84, 1.02)	1.07 (0.96, 1.16)	rank-sum	0.003

25

Table S2
Measures of grid cell field restructuring (Figs. 3, S9)

All measures are comparing ENV1 v ENV2

<i>measure</i>	<i>test</i>	<i>p-value</i>
5 # of fields	signed-rank	p=0.59
distance between reward and closest field	signed-rank	p=0.14
slope, field size x distance	signed-rank	p=0.31
slope, eccentricity x distance	signed-rank	p=0.97
10 slope, field peak FR x distance	signed-rank	p=0.01
field peak FR	signed-rank	p=0.83
coefficient of variation, field peak FR	signed-rank	p=0.07

15

Table S3
Measures of non-grid spatial cell remapping (Figs. 3, S10)

5	<i>Group #</i>	<i>n cells</i>	<i>proportion of reward-preferring cell population</i>	
	I	62	39%	
	II	33	20%	
	III	23	14%	
	IV	41	26%	
10	<i>measure</i>	<i>median value</i>	<i>test</i>	<i>p-value</i>
	distance to closest field, Group I (ENV1, ENV2)	59, 36	signed-rank	2×10^{-5}
	distance to closest field, Group II, ENV1	60+-26cm	shuffling test	0.02
	distance to closest field, Group III, ENV2	35+-24	shuffling test	0.002
15	proportion effect, cells in each category (Fig. 3G) (baseline = 0.58)			
	<i>Group #</i>		<i>test</i>	<i>p-value</i>
20	I		binomial test	0.28
	II		binomial test	0.18
	III		binomial test	0.75
	IV		binomial test	0.91

Table S4
Measures of task/no-task FRxdistance slopes (Fig. 4)

	<i>Group</i>	<i>time</i>	<i>median slope (25th, 75th quantiles)</i>	<i>test</i>	<i>p-value</i>
5	grid	task	-0.0095 (-0.0059, 0.0029)	signed-rank	0.04
	grid	no-task	-0.0015 (-0.0052, 0.0027)	signed-rank	0.03
	non-grid task		-0.001 (-0.006, 0.0032)	signed-rank	0.001
	non-grid no-task		-0.0027 (-0.007, 0.0025)	signed-rank	2x10 ⁻⁶

10