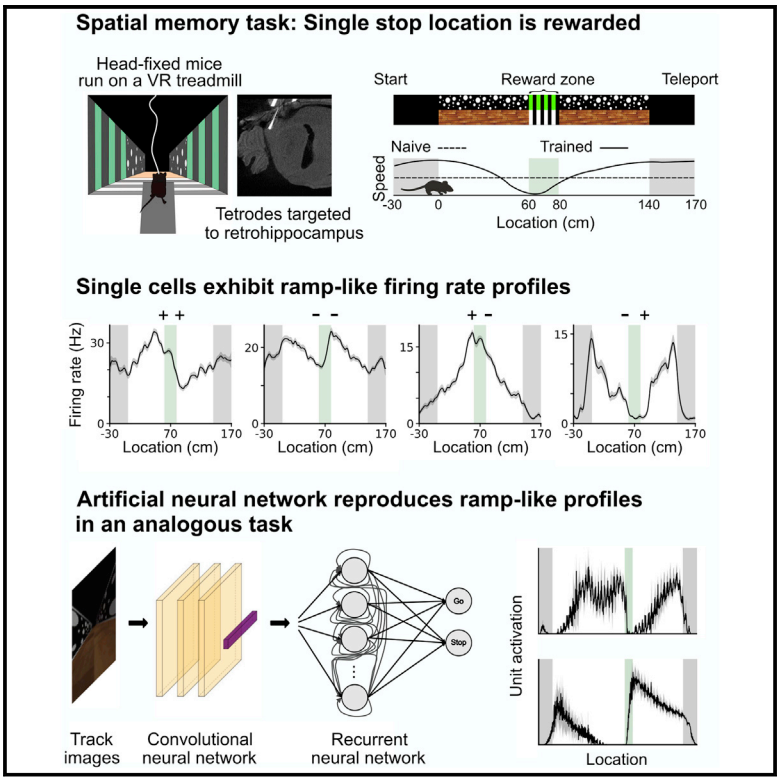
Article

Current Biology

Spatial representation by ramping activity of neurons in the retrohippocampal cortex

Graphical abstract



Highlights

- Neurons in the retrohippocampal cortex use ramping activity to encode location
- Ramping activity is consistent with recall of learned models of the environment
- Neurons with ramping activity are observed more frequently than grid or border cells
- Ramping emerges in artificial recurrent neural networks that learn a similar task

Authors

Sarah A. Tennant, Harry Clark, Ian Hawes, ..., Klara Z. Gerlei, Emma R. Wood, Matthew F. Nolan

Correspondence

mattnolan@ed.ac.uk

In brief

The retrohippocampal cortices are crucial for spatial memory. Tennant et al. show that neurons in the retrohippocampal cortices represent learned environments through ramping activity profiles. These profiles are distinct from well-studied discrete spatial firing fields and may be well suited to readout learned models of space.

Tennant et al., 2022, Current Biology *32*, 4451–4464
October 24, 2022 © 2022 The Author(s). Published by Elsevier Inc.
https://doi.org/10.1016/j.cub.2022.08.050

CellPress

Current Biology

CellPress
OPEN ACCESS

Article

Spatial representation by ramping activity of neurons in the retrohippocampal cortex

Sarah A. Tennant,[1,5] Harry Clark,[1,5] Ian Hawes,[1,5] Wing Kin Tam,[1,2] Junji Hua,[1] Wannan Yang,[1,4] Klara Z. Gerlei,[1] Emma R. Wood,[1,2] and Matthew F. Nolan[1,2,3,6,*]

[1]Centre for Discovery Brain Sciences, University of Edinburgh, Edinburgh, UK
[2]Simons Initiative for the Developing Brain, University of Edinburgh, Edinburgh, UK
[3]Centre for Statistics, University of Edinburgh, Edinburgh, UK
[4]Present address: Center for Neural Science, New York University, New York, NY, USA
[5]These authors contributed equally
[6]Lead contact
*Correspondence: mattnolan@ed.ac.uk
https://doi.org/10.1016/j.cub.2022.08.050

SUMMARY

Neurons in the retrohippocampal cortices play crucial roles in spatial memory. Many retrohippocampal neurons have firing fields that are selectively active at specific locations, with memory for rewarded locations associated with reorganization of these firing fields. Whether this is the sole strategy for representing spatial memories is unclear. Here, we demonstrate that during a spatial memory task retrohippocampal neurons encode location through ramping activity that extends across segments of a linear track approaching and following a reward, with the rewarded location represented by offsets or switches in the slope of the ramping activity. Ramping representations could be maintained independently of trial outcome and cues marking the reward location, indicating that they result from recall of the track structure. When recorded in an open arena, neurons that generated ramping activity during the spatial memory task were more numerous than grid or border cells, with a majority showing spatial firing that did not meet criteria for classification as grid or border representations. Encoding of rewarded locations through offsets and switches in the slope of ramping activity also emerged in recurrent neural network models trained to solve a similar spatial memory task. Impaired performance of model networks following disruption of outputs from ramping neurons is consistent with this coding strategy supporting navigation to recalled locations of behavioral significance. Our results suggest that encoding of learned spaces by retrohippocampal networks employs both discrete firing fields and continuous ramping representations. We hypothesize that retrohippocampal ramping activity mediates readout of learned models for goal-directed navigation.

INTRODUCTION

Spatial memory relies on complex brain network dynamics, with neurons in the hippocampus and associated retrohippocampal cortices (medial entorhinal cortex, pre- and para-subiculum) playing essential roles.[1] Many neurons in these brain areas have receptive fields selective for location (Figure 1A). For example, place cells in the hippocampus and grid cells in retrohippocampal cortices have been widely studied[1–3] and are central to theories of navigation and memory.[4–9] However, theoretical studies suggest that continuous coding schemes in which neuronal firing rate is directly proportional to location (Figure 1B) may also be suited to control of navigational behaviors.[10–17] While in principle such codes could be generated downstream of place or grid cells,[10–17] it is unclear whether neurons in spatial memory circuits represent information using continuous codes or if coding schemes of this kind can be engaged by recall of memories.

Representations during learned spatial behaviors serve the purposes of tracking location along routes to a goal and of indicating whether the expected goal has been reached.[18] Goal locations are associated with an increase in the density of hippocampal place fields[19–23] and with reorganization of the firing fields of grid cells in the entorhinal cortex.[24,25] Under these coding schemes, locations along the track and at goals are typically represented by different neurons, with the current location reflected in the identity of the currently active neurons. In contrast, under continuous coding schemes in which position is represented by firing rate, neural firing rates could ramp upward or downward during movement through an environment. If the initiation and termination of the ramp is then anchored to salient landmarks, which could include goal locations, then positions during movement between the anchoring landmarks may be represented by continuous changes in the firing rates of the ramping neurons. In this case, reaching a landmark could be indicated by cessation or reconfiguration of the ramping activity (Figure 1B). Coding schemes of this kind appear to be used by neurons in the striatum, which generate spatial ramping activity in anticipation of reward locations.[26–32] Whether similar strategies are employed by neurons in the hippocampal formation is unclear.

This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

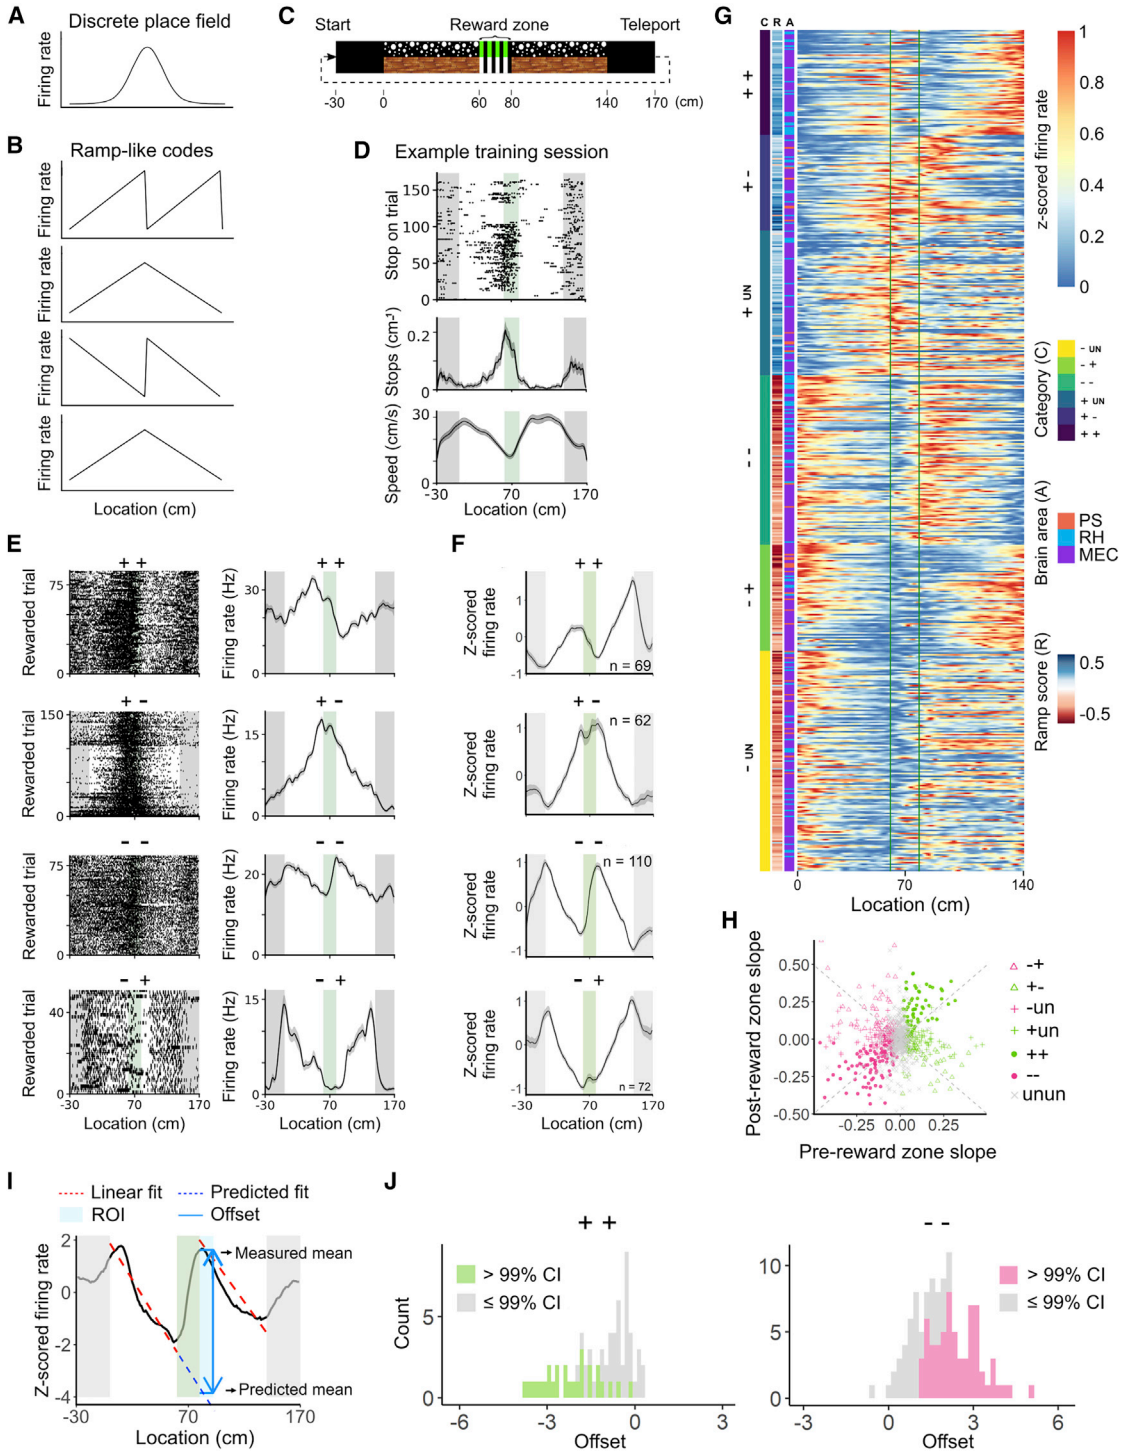


Figure 1. Ramping neural activity approaching and following a learned reward location

(A and B) Conceptualization of discrete (A) and continuous (B) firing fields.

(C) Schematic of the virtual track in the beaconed trial configuration. Numbers below the track indicate distance (cm) relative to the start of the track.

(D) Example behavior of a trained mouse on beaconed trials, showing a raster plot of stopping locations (upper), mean stops (per cm) (middle), and average running speed (lower) each as a function of position on the track (solid lines indicate mean values with standard error indicated in gray). The vertical green bar indicates the location of the reward zone.

(E) Spike rasters (left) and the corresponding mean firing rate (right) as a function of track position, for example, neurons (solid lines indicate mean values with standard error indicated in gray). The first and third neurons are recorded from the training session shown in (D). Symbols above the plots indicate classification of the slope of the neuron's ramping activity before and after the reward zone, e.g., ++ indicates neurons with positive slope ramps on each side of the reward zone.

(legend continued on next page)

To investigate coding schemes used by retrohippocampal neurons during recall of locations in learned environments, we used a memory task in which external sensory cues can be controlled on a trial by trial basis, enabling separation of cue-driven and memory-dependent spatial representations.[33] The task requires recall of the location of a reward zone on a virtual linear track (Figure 1C).[33] On cued trials the reward zone is marked by a prominent visual cue but otherwise the track lacks unique spatial cues. With training, mice efficiently obtain rewards by running directly to the reward zone (Figures 1D and S1).[33] This spatial behavior is maintained when the cues marking the reward zone are removed, indicating that mice can solve the task by comparing online estimates of their distance from the start of the track, obtained using a path integration strategy, with estimates stored in memory.[33] The path integration process may be driven either by sensory input or internal command signals. However, path integration alone is insufficient for accurate stopping unless the mouse also remembers where to stop. Thus, a memory for location on the track is critical to efficiently obtain rewards.[33] For tasks with this structure, previously hypothesized continuous coding schemes[10–17] lead to predictions for neural representations that are distinct from and complementary to coding using classical discrete firing fields. First, under continuous coding schemes, location may be represented by ramping neural activity. Second, the initiation and termination of ramping activity should be anchored to salient locations including the reward zone as well as the start and end of the track. Third, if ramping activity is driven by a neural path integrator, then it should be maintained in the absence of location-specific cues. Fourth, if the locations that anchor the initiation or termination of ramping activity are recalled from memory, then transitions in firing at these locations should be maintained in the absence of cues that indicate the transition point.

We show here that in the retrohippocampal cortices of mice that have learned the location memory task, neurons encode location through ramping activity that either resets or switches polarity at the reward zone. Resetting and switches in ramp polarity were maintained after removal of reward zone cues, indicating that they can be driven by recall of the reward location. Neurons that generated ramping activity were observed more frequently than classical grid or border cells, and in open arenas a majority of these neurons had spatial firing fields that did not meet the criteria for classification as grid or border representations. Finally, we show that artificial recurrent neural networks can be trained to solve a similar location memory task and in doing so generate similar ramping representations that are interrupted at reward locations. Because successful activity patterns emerge in these artificial networks through reinforcement rather than by design, and as successful training did not require networks to be constrained by the specific connectivity of the retrohippocampal cortex, this suggests that ramping activity profiles that reset or change direction at salient landmarks may provide a general solution to the problem of representing remembered locations.

RESULTS

We used tetrodes targeted to retrohippocampal cortices, including the medial entorhinal cortex, presubiculum, and parasubiculum (Data S1), to record neural activity from mice carrying out the location memory task. Following spike sorting we obtained 1,395 isolated single units (126.8 units/mouse, range 20–314; 11 mice) recorded from trained mice during sessions that contained a minimum of 30 correct trials (196.5 ± 7.7 trials/session; 187 sessions, 17 sessions/mouse, range 6–25). The tetrodes were usually lowered by ~50–100 μm following each session.

Encoding of location by ramping activity

We focus our initial analyses on trials in which the animal stops in the reward zone, which we refer to as successful trials, and in which the reward zone was visible, which we refer to as beaconed trials. On these trials, a majority of recorded retrohippocampal neurons generated ramping activity within discrete regions bounded by the start of the track and the reward zone, and by the reward zone and the end of the track (n = 836/1,395 neurons) (Figures 1E–1G). For many neurons movement through the reward zone was associated with either a change in the slope of the ramp or resetting of the spike rate (Figures 1E–1G). Changes in ramp slope were observed from positive before the reward to negative afterward and vice versa (indicated as +− and −+ in Figures 1E–1G). Resetting of the spike rate occurred as a step like reduction in spike rate of neurons with a positive ramp slope before and after the reward zone, and a step like increase for neurons with negative slopes before and after the reward zone (indicated as ++ and −− in Figures 1E–1G).

To quantitatively evaluate ramp firing profiles, we classified neurons on the basis of their firing rate slopes in the track

(F) *Z* scored firing rate as a function of position averaged across neurons classified by their ramp slope before and after the reward zone (symbols above each panel indicate classification).

(G) Normalized firing rates as a function of location along the track (0–140 cm) for all neurons that have ramp-like firing in the region of the track before the reward zone. The start and end of the reward zone region is marked with vertical green lines. Neurons are organized along the y axis by the direction of their slope before and after the reward zone (e.g., ++, −−, etc.) and then ordered within each category according to their ramp slope before the reward zone. Columns to the left indicate the classification of activity based on fit of the linear model (C), the ramp score (R), and the location within the retrohippocampal region (L).

(H) Slope of the fit to the firing rate on the track segment before the reward zone as a function of slope on the track segment after the reward zone. Dashed lines indicate points for which the slope would be identical on both track segments. Symbols indicate classification of individual neurons.

(I) The ramp slope was estimated by fitting spiking activity as a function of location separately for the track segments before the reward zone (0–60 cm) and after the reward zone (80–140 cm). Offsets in firing rate were calculated by subtracting the actual firing rate in the first 5 cm after the reward zone with the firing rate predicted by extrapolating the fit from the track segment before the reward zone.

(J) Distribution of offsets as calculated in (I) for neurons classified on beaconed trials as having ++ (left) or −− (right) activity patterns (++, $t = -9.2$, $df = 67$, $p = 1.9 \times 10^{-13}$; −−, $t = 19.6$, $df = 109$, $p < 2.2 \times 10^{-16}$, t test versus mean = zero). Colored bars indicate individual neurons for which the offset was outside the 99% confidence of intervals of the neuron's predicted value.

For related analyses, see also Figures S1–S5. For documentation of tetrode locations, see Data S1.

segments before and after the reward zone (Figures 1E–1G). While there was a slight bias toward neurons with negative ramp slopes on the track segment before the reward zone, a continua of ramp slopes were observed and ramp slope on the track segment before the reward zone did not predict the ramp slope after the reward zone (Figure 1H). When neurons maintained a positive or negative ramp slope on both sides of the reward zone, their firing rate usually reset as the animal passed through the reward zone (Figures 1I and 1J). Classifications were similar when ramping activity was identified by fitting a linear model ($p < 0.01$ after correction for multiple comparisons; STAR Methods) or using a "ramp score" based on correlations between position and firing rate (Figure 1G and S2A–S2C). Ramp-like activity was present in similar proportions of neurons from all animals (mean 60.7% ± 4.99%, range 45.8–77.4), had firing rate peaks for ramps with positive slopes (or minima for ramps with negative slopes) that were strongly biased toward locations adjacent to the reward zone (Figures S2D and S2E), was independent of theta modulation of spiking activity (Figure S3), was similar in medial entorhinal cortex and pre/para-subiculum (Figures S2F and S2G), and was present in only 2.02% of shuffled datasets (n = 28,235/1,395,000 shuffles) (Figure S4). For ramping neurons, ramp slopes and track wide firing profiles were correlated between the first and second half of each recording session, indicating that ramp representations are stable at this timescale (ramp slope before the reward zone, $r^2 = 0.88$; ramp slope after the reward zone, $r^2 = 0.77$; track wide correlation, $r^2 = 0.47 \pm 0.28$). The corresponding correlation scores ($r^2 = 0.25$, $r^2 = 0.21$, and $r^2 = 0.23 \pm 0.26$) were substantially lower for non-ramping neurons (respectively $p < 2 \times 10^{-16}$, $F = 164.4$, $df = 1$, ANOVA; $p = 5 \times 10^{-15}$, $F = 63.0$, $df = 1$, ANOVA; $t = 14.6$, $df = 1{,}103$, $p < 2 \times 10^{-16}$, t test).

To test whether ramping activity reflects position or other kinematic variables, we first fit firing rates in the track segment before the reward zone with generalized linear mixed effect models that included position, speed, and acceleration as fixed effects (Figure 2A). For most neurons with ramping activity before the reward zone, their activity could best be accounted for by encoding of position, either alone or conjunctively with position as the dominant coefficient alongside speed and/or acceleration (n = 340/546 neurons with ramping activity before the reward zone) (Figures 2B–2D). Ramping activity was also better accounted for by encoding of position rather than time (Figure S5). As a further test, we evaluated neural firing rate profiles on trials in which mice failed to stop in the reward zone (Figures 3 and S6). We subdivided these miss trials into try trials, in which mice slowed down in the reward zone but did not stop, and run trials, in which mice maintained a high running speed within the reward zone (Figures 3A and 3B). Firing rate profiles of positionally modulated ramping neurons were similar for each trial outcome (Figures 3C, 3D, and S6A), with firing rate slopes independent of trial outcome (Figures 3E and 3F) and offset behavior of ++ and −− neurons maintained although somewhat reduced compared with hit trials (Figure 3G), indicating that for most neurons ramping activity, along with switches in slope and resetting of the firing rate in the region of the reward zone, is independent from changes in running speed associated with different trial outcomes.

Together, these data indicate that firing rate profiles in which ramping activity encodes location within discrete track segments and reconfigures at a reward location are a common feature of neural activity in retrohippocampal cortices. These activity patterns contrast with discrete firing fields that are typical of place and grid cell fields recorded previously on linear virtual tracks.[34–36]

Ramp-like firing rate profiles in the absence of reward zone cues

To test whether changes in the slope or resetting of the rate of ramp firing at the reward zone results from recall of a memory of the track structure, we investigated activity on probe trials in which rewards and the visible cues indicating the reward zone were omitted (1/10 trials were probe trials, interleaved with 4/5 trials that had the standard cue configuration and 1/10 trials that are rewarded but lack the visible cue) (Figures 4A and 4B). We focus on neurons with firing rate profiles that depend on track position, either alone or conjunctively with speed or acceleration, and for which there were sufficient numbers of probe trials (STAR Methods) (n = 97/240 neurons with positional ramping before the reward zone, 38 sessions, 9 mice).

The effects of cue removal differed between positional ramping neurons. For many neurons, encoding of the reward zone, by changes in the slope of ramping activity or by resetting of the spike rate, was maintained on probe trials (Figures 4C and S6B). This stability was manifest in similarity between average firing rate profiles obtained on beaconed and probe trials from populations of positional neurons grouped according to their slopes before and after the reward zone (e.g., ++, +−, −+, and −−) (Figure 4D). To assess the effects of cue removal across the neuronal population we determined the effects on ramp slopes and offsets. Ramp slopes showed variable responses to removal of the cues, with some neurons maintaining a similar slope, which is consistent with their activity being driven by a path integration mechanism, and others showing a reduced slope, which is consistent with the reward zone cues driving their activity (Figures 4E and 4G). The latter group of neurons drove a reduction in the mean ramp slope on probe compared with beaconed trials, although the population slopes remained substantially above chance (Figures 4E and 4G), and slopes on probe trials were correlated with slopes on beaconed trials (pre-reward zone, adjusted $R^2 = 0.50$, $p = 3.9 \times 10^{-14}$; post-reward zone, adjusted $R^2 = 0.45$, $p = 2.9 \times 10^{-12}$) (Figures 4F and 4H), which is consistent with the maintained ramping profiles in the average activity. The effects of cue removal on the offset in ramping activity was also consistent with variation between neurons in the influence of path integration and cue-driven mechanisms. Thus, offsets in ramping activity of ++ and −− neurons were maintained on probe trials, although were reduced compared with standard trials (Figure 4I) but again were strongly correlated between the two trial types (adjusted $R^2 = 0.39$, $p = 2.2 \times 10^{-10}$) (Figure 4J).

The maintenance of ramping profiles in the absence of reward zone cues indicates that recall of the track structure is sufficient for ramping activity and its interruption at the reward zone. Reduced ramping and offset scores of some neurons suggests that cues associated with the reward zone also play a role in establishing firing rate profiles. These observations are consistent with activity of many neurons in the medial entorhinal cortex

CellPress
OPEN ACCESS

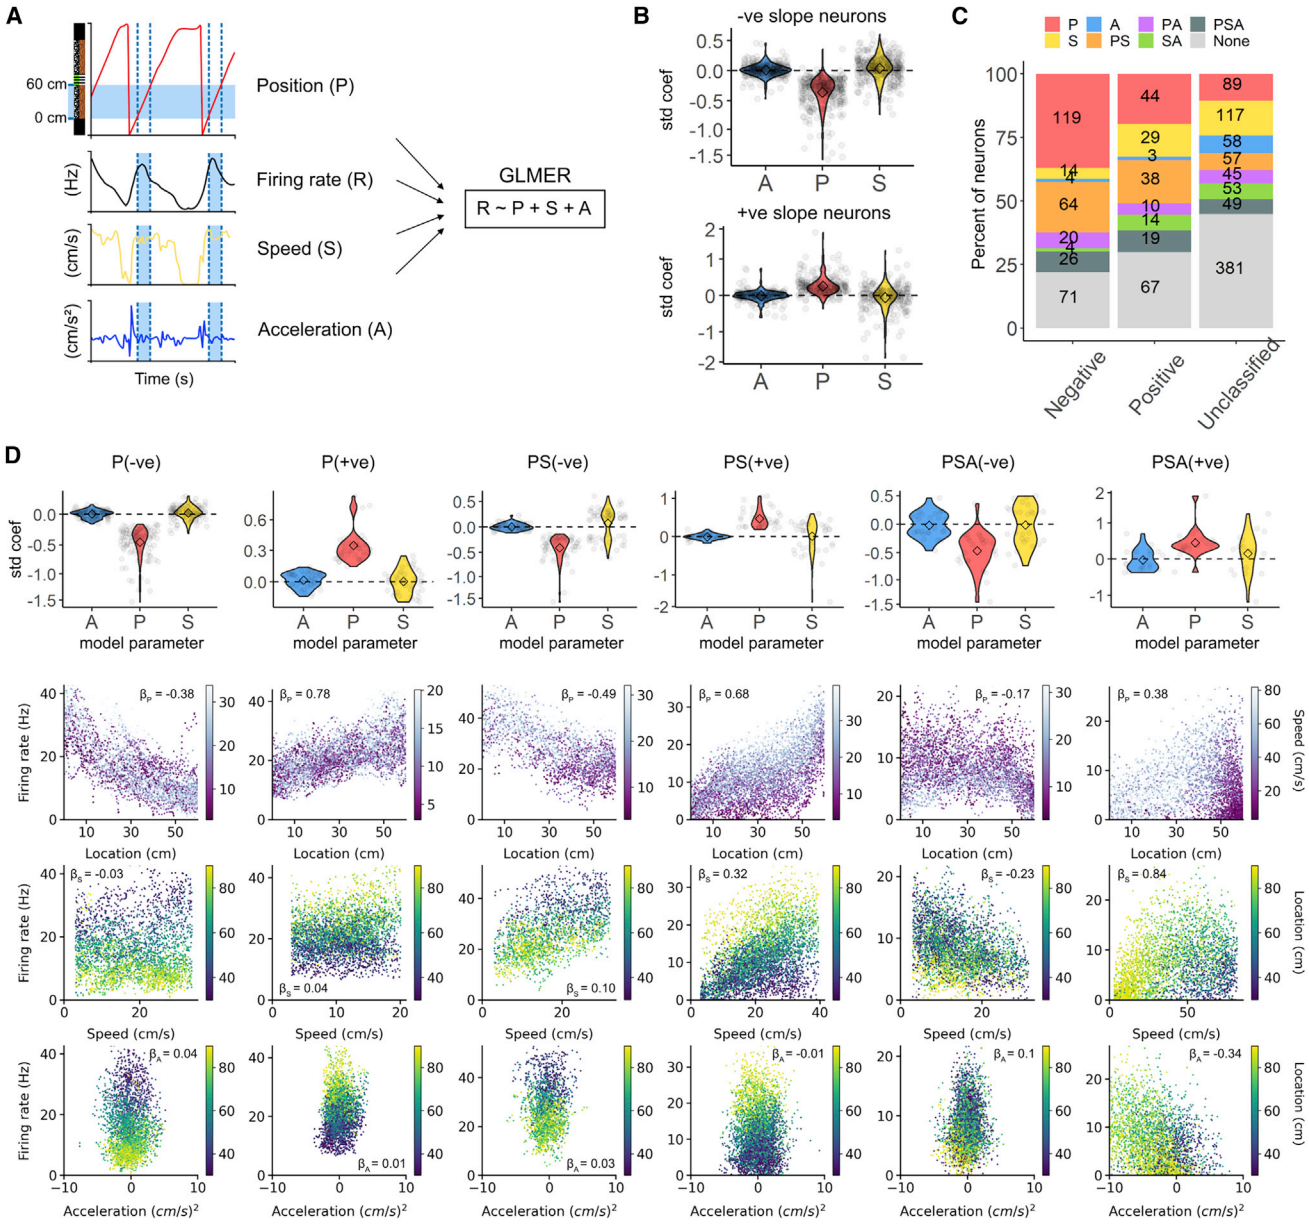


Figure 2. Differential influence of position, speed, and acceleration on ramp firing

(A) Generalized linear mixed effect models that include speed (S), acceleration (A), and position (P) as fixed effects were fit to the firing rate in the track region approaching the reward zone (0–60 cm; see Figure 1C). The left panels show firing rate, position, speed, and acceleration as a function of time for two consecutive trials from an exemplar experiment. The track region over which the model was fitted is indicated by the blue shaded bar. Speed and acceleration were calculated as the first and second derivatives with respect to time of the track position. Trial number was included as a random effect.

(B) Standardized coefficients, which index the relative strength of each fixed effect variable, obtained from the model fits for neurons with negative (upper) and positive (lower) ramp slopes on the track segment before the reward zone. ANOVA indicates that the relative amplitude of coefficients for speed, acceleration, and position differ (+ve slope, $p < 2 \times 10^{-16}$, df = 2, 446, f = 93.18; −ve slope, $p < 2 \times 10^{-16}$, df = 2, 642, f = 421.3; repeated measures ANOVA). For neurons in both groups the position coefficients were larger than those for acceleration ($p < 2 \times 10^{-16}$, Bonferroni-corrected paired t test) and speed ($p < 2 \times 10^{-16}$, Bonferroni-corrected paired t test).

(C) Proportions of neurons for which each combination of P, A, and S coefficients were significant at a threshold of $p < 0.01$ (Wald's chi-squared test) for neurons classified as having a positive slope, negative slope, or unclassified.

(D) Distribution of standardized coefficients (upper row) and example data (lower rows) for neurons with positive (+ve) or negative (−ve) ramp slopes and classified as modulated solely by position (P), by position conjunctively with speed (PS), and position conjunctively with speed and acceleration (PSA).

Example data show firing rate as a function of position (second row), running speed (third row), and acceleration (lower row). Firing rates as a function of position are color coded by speed (cm/s) while firing rates as a function of speed or acceleration are color coded by position along the track (0–60 cm).

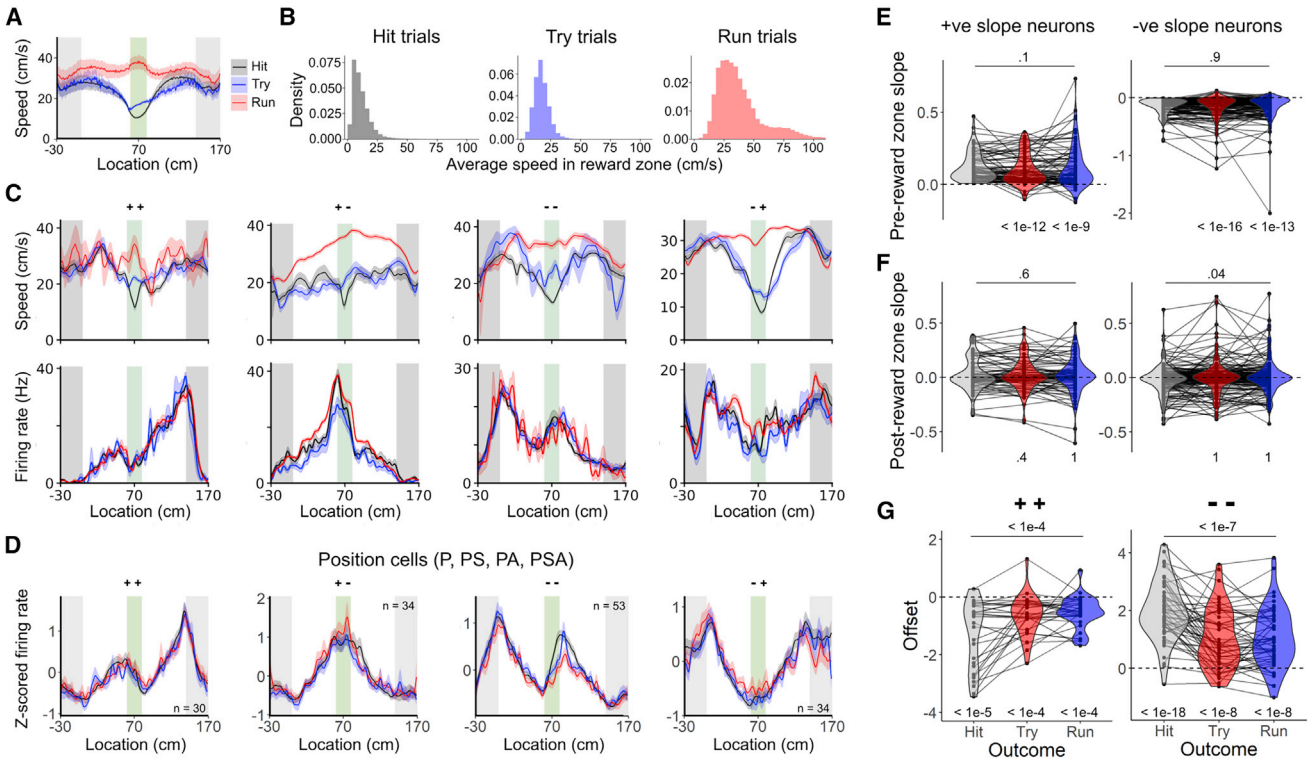


Figure 3. Relationship between task performance and ramping activity

(A) Population average of mean running speed as a function of position for each trial outcome across all animals (n = 11). Trials were classified as "hit trials" when mice stopped in the reward zone region, as "try trials" when mice slowed down to <95th percentile of reward zone speeds in hit trials but did not stop, and as "run trials" when mice maintained a speed >95th percentile of reward zone speeds in hit trials. Shaded regions correspond to the standard error of the mean across animals. Note that given the reversal of the speed profile between hit and try groups (reduced in the reward zone) compared with the run group (higher in the reward zone), neurons with activity determined exclusively by running speed should show a similar reversal.

(B) Trial level distribution of mean running speeds within the reward zone for hit, try, and run trial outcomes. The histograms are calculated from all sessions (n = 187) and all trials with a given outcome (hit, n = 12,130; try, n = 1,534; run, n = 1,111).

(C) Examples of running speed (upper) and firing rate (lower) as a function of track position. Note that for some animals the average running speed in the hit group can be greater than zero because the stopping location varies between trials. Shaded regions correspond to the standard error of the mean across trials.

(D) Population averaged firing rates as a function of position on hit, try, and run through trials for positionally modulated neurons (P, PS, PA, and PSA groups). The neurons were classified on hit trials as having +−, −+, ++, and −− firing rate profiles. Shaded regions correspond to the standard error of the mean across cells.

(E) Slope of positive and negative ramping activity on the track segment before the reward zone as a function of trial outcome. There was no significant effect of trial outcome (+ve slopes, df = 2, 164, F = 2.1, p = 0.12; −ve slopes, df = 2, 308, F = 0.069, p = 0.93; repeated measures ANOVA) and for all outcomes the slope values differed significantly from zero ($p < 2 \times 10^{-10}$, Bonferroni-corrected t test).

(F) Slope of positive and negative ramping activity on the track segment after the reward zone as a function of trial outcome. There was no significant effect of trial outcome (+ve slopes, df = 2, 308, F = 3.2, p = 0.04; post-reward zone −ve slopes, df = 2, 164, F = 0.427, p = 0.65; repeated measures ANOVA).

(G) Offset between the predicted and actual firing rate following the reward zone as a function of trial outcome for neurons with positive slopes before and after the reward zone (++) and for neurons with negative slopes before and after the reward zone (−−). The extent of the offsets varied according to trial outcome (++, df = 2, 56, F = 12.43, $p = 3.4 \times 10^{-5}$; −−, df = 2, 102, F = 21.91, 1.2×10^{-8}), but for all outcomes the offsets differed from zero ($p < 7 \times 10^{-5}$, Bonferroni-corrected t test). Differences between offsets may reflect increased trial to trial variation in the location at which resetting occurs in some neurons (see also Figure S6A).

reflecting readout of an internally generated learned model of the track, with other neurons being driven directly by visual cues indicating the reward zone.

Ramping neurons have spatial firing fields in open environments

Retrohippocampal cortices are widely associated with grid, border, and head direction representations.[3,37–40] However, because ramp firing patterns that we report here are recorded during a task in which recall of the track locations is necessary for reward, whereas firing properties of retrohippocampal neurons are usually evaluated in the absence of location-specific reward contingencies, it is unclear which of the previously described functional cell types ramping neurons correspond to. To address this, in a subset of behavioral sessions after recording from neurons in the location memory task we recorded from the same neurons in an open arena (Figure 5A). We used a hierarchical scheme to classify cells,[37] with rank grid > border > other spatial > head direction > non-spatial. In this scheme, cells were assigned to the highest ranked category for which their identity score was greater than the 99th percentile of their corresponding shuffled data (STAR Methods). For most neurons with positional ramping activity in the location memory task, their firing in the open arena was classified in the other spatial category

CellPress
OPEN ACCESS

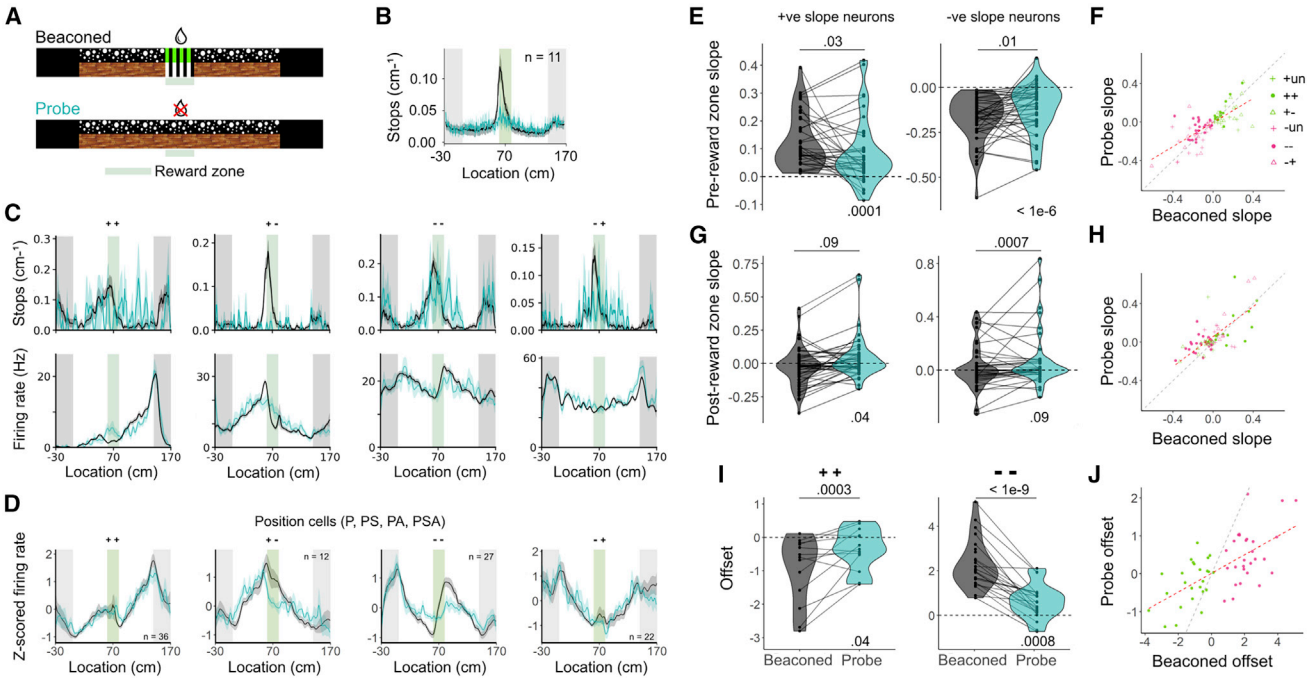


Figure 4. Maintenance of ramping activity profiles following removal of reward zone cues

(A) Schematic of track configuration.

(B) Population mean of behavior from trained mice on beaconed and probe trials, showing the mean stops (per cm) as a function of position on the track. Shaded regions correspond to the standard error of the mean across animals.

(C) Examples from beaconed trials (black) and probe trials (blue) of stop histogram (upper row) and mean firing rate (lower row, shaded areas indicate the SEM) as a function of track position. The neurons were classified (++, +−, −−, −+) according to their firing rate profiles on beaconed trials. See also Figure S6B.

(D) Mean firing rates as a function of track position on beaconed trials (black) and probe trials (blue) for the populations of neurons classified with ++, +−, −−, −+ firing rate profiles.

(E and G) Slopes of ramping activity on the track segment before the reward zone (E) and after the reward zone (G) differed between beaconed and probe trials (pre-reward zone +ve slopes, $t = 2.3$, $df = 36$, $p = 0.028$; pre-reward zone −ve slopes, $t = -2.6$, $df = 45$, $p = 0.012$; post-reward zone +ve slopes, $t = -1.7$, $df = 36$, $p = 0.093$; post-reward zone −ve slopes, $t = -3.7$, $df = 45$, $p = 0.00067$; paired Student's t tests). Slopes before the reward zone on probe trials nevertheless remained above/below zero on for +ve/−ve slope neurons, respectively (+ve, $t = 4.3$, $df = 36$, $p = 0.00011$; −ve, $t = 2.1$, $df = 36$, $p = 0.038$; one-sample t test versus mean of zero). Lines connect points from the same neuron.

(F and H) On track segments before (F) and after (H) the reward zone, ramp slopes on probe trials varied as a function of the beaconed trial slope (pre-reward zone, adjusted $R^2 = 0.50$, $F = 83.9$ on 1 and 81 DF, $p = 3.9 \times 10^{-14}$; post-reward zone, adjusted $R^2 = 0.45$, $F = 67.3$ on 1 and 81 DF, $p = 2.9 \times 10^{-12}$).

(I) Firing rate offsets for ++ and −− neurons, measured as in Figure 1H, were reduced on probe compared with beaconed trials (++, $t = -4.3$, $df = 23$, $p = 0.00027$; −−, $t = 9.1$, $df = 30$, $p = 4.0 \times 10^{-10}$; paired t test), but remained greater than expected from linear extrapolation of ramping activity before the reward zone (++, $t = -2.19$, $df = 23$, $p = 0.039$; −−, $t = 3.7$, $df = 30$, $p = 0.00078$; one-sample t test).

(J) Firing rate offsets on probe trials were correlated with offsets on beaconed trials (adjusted $R^2 = 0.3938$, $F = 52.97$ on 1 and 79 DF, $p = 2.2 \times 10^{-10}$).

(n = 145/201) (Figures 5B–5D). Positional ramping neurons that had border, grid or pure head direction firing patterns in the open field were found relatively infrequently (n = 10, 7, and 18/201, respectively). The proportion of border, grid, and pure head direction cells was similar among ramping and non-ramping neurons (Figures 5C, S7A, and S7B), and the proportion of ramping and non-ramping neurons was similar among each open arena cell type (Figure 5D). The distribution of speed modulated activity in the open arena was also similar among ramping and non-ramping neurons (Figure S7B).

We also evaluated the correspondence between the activity of ramping neurons and features used to distinguish principal cells from interneurons. The majority of ramping neurons had firing properties characteristic of principal cells, while a small proportion had firing properties that are typical of interneurons (Figure S7C). Ramp firing of putative interneurons usually had a negative slope on the track segment before the reward zone, with few putative interneurons having positive ramp slopes (Figure S7D).

Together, these data indicate that the population of neurons with ramping activity is largely, but not exclusively, non-overlapping with grid, head direction, and border cell populations, while many neurons that in an open arena lack well-defined spatial firing fields may participate in representation of learned information by generating ramp-like representations that interrupt around recalled reward locations.

Recurrent neural network models learn to estimate location using ramp-like representations

If initiation and termination of ramping representations at behaviorally important locations is an effective representation for guiding learned behaviors, then similar activity profiles may emerge in artificial neural networks trained to solve location memory tasks. To investigate this, we trained a network model that

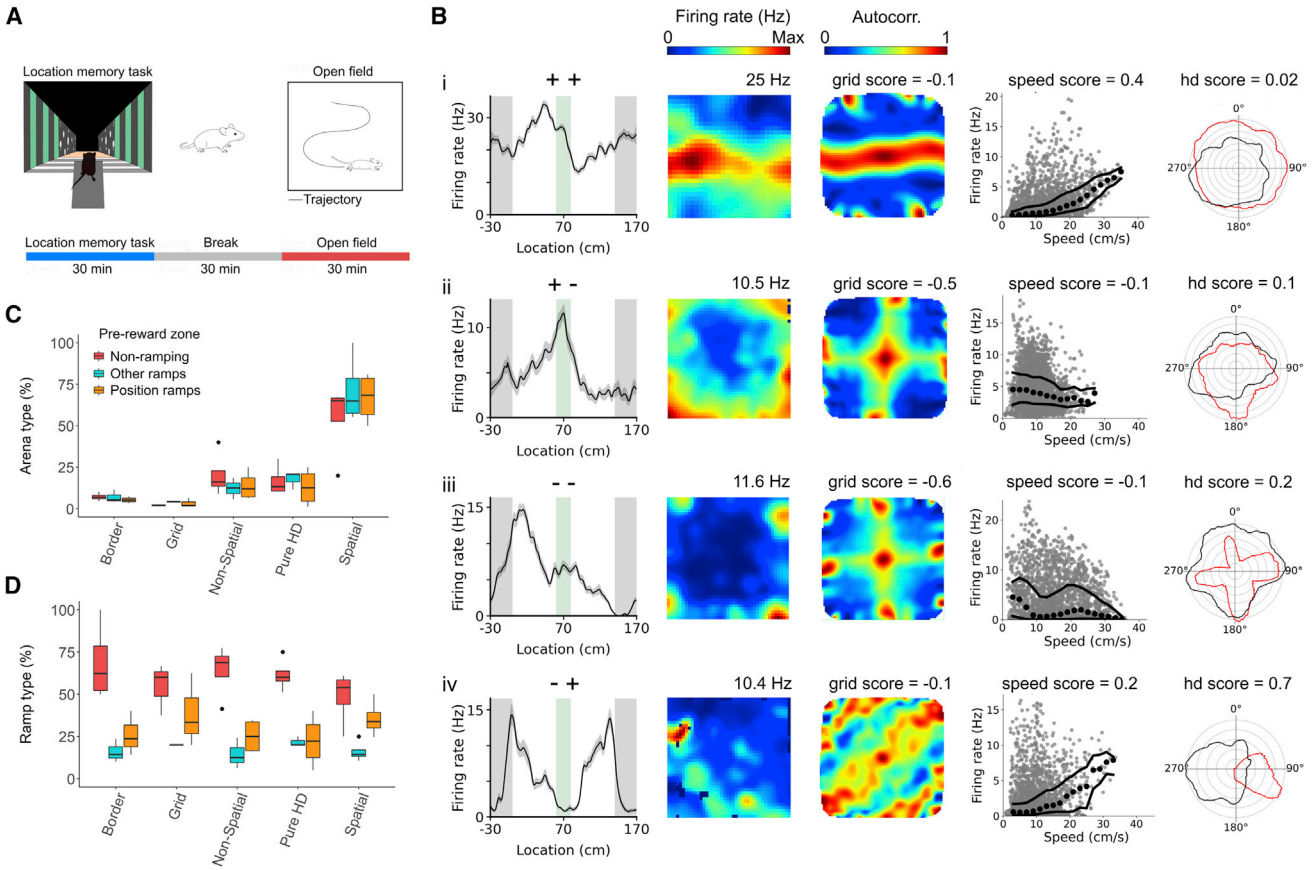


Figure 5. Ramping neurons have diverse spatial firing patterns in an open arena

(A) Animals performed the location estimation task for ~30 min, were placed in a holding arena for ~30 min, and then explored the open arena for ~30 min.

(B) Example activity in the open arena of cells with ramping activity on the linear track characterized by +ve slopes approaching and following the reward zone (Bi), +ve slope approaching and −ve slope following the reward zone (Bii), −ve slopes approaching and following the reward zone (Biii), and −ve slope approaching and +ve slope following the reward zone (Biv). Activity in the open arena was classified as either spatial (Bi and Biv) or border (Bii and Biii). Panels show (from left to right) mean firing rate on the virtual track as a function of position; spatial heat maps of firing rates in the open arena; spatial autocorrelograms of the open arena firing rate; firing rate in the open arena as a function of speed; and polar plots of average firing rate in the open arena as a function of head direction (red) and movement (black). See also Figure S7.

(C) Proportion of each open arena cell type (border, grid, head direction, non-spatial, other spatial) among neurons with activity on the linear virtual track classified as position-dependent ramping, other ramping, and non-ramping.

(D) Proportion of neurons classified on the linear virtual track as position-dependent ramping, other ramping, and non-ramping among neurons classified as border, grid, head direction, non-spatial, other spatial.

In (C) and (D), boxes show the median, 25th and 75th percentiles, and whiskers extend to the largest value 1.5 times the interquartile range from the box. Points outside this range are plotted individually.

received visual input corresponding to a virtual track with layout similar to the experimental track. To promote cue-independent strategies, we trained the network solely with tracks in which the visual reward zone cues were absent. The model network generated stop or go signals to control movement along the track. The model was organized such that visual signals were processed by a three-layer convolutional neural network (CNN) with the output fed into a recurrent neural network that generated movement control signals (Figure 6A). We chose this network architecture because recurrent connections are a feature in many handcrafted models of the entorhinal cortex[41] and we focused on visual inputs as visual cues delineate the segments of the experimental track and visual signals are encoded by neurons within the entorhinal cortex.[42] Networks trained with a reinforcement learning algorithm[43] successfully learned the task (Figure 6B). Before training, the agent controlled by the network stopped frequently on the track with no particular spatial stopping pattern (Figure 6C). After training, the agent stopped at the reward zone on non-beaconed (Figure 6D) and probe trials (Figure 6E), indicating that the network solves the task using a distance-based strategy. The agent's stopping patterns were remarkably similar to those of mice before and after learning the task.[33]

As in experimental recordings from retrohippocampal neurons, units in the trained recurrent network showed ramping activity that was interrupted at the reward zone (Figures 6F and 6J). The ramping activity of units in the recurrent network was clearly distinguishable from activity patterns generated by shuffled datasets (Figures 6G and 6H), and had either a positive or a negative slope (Figures 6F, 6I, and 6J), with the slope on the region of

CellPress
OPEN ACCESS

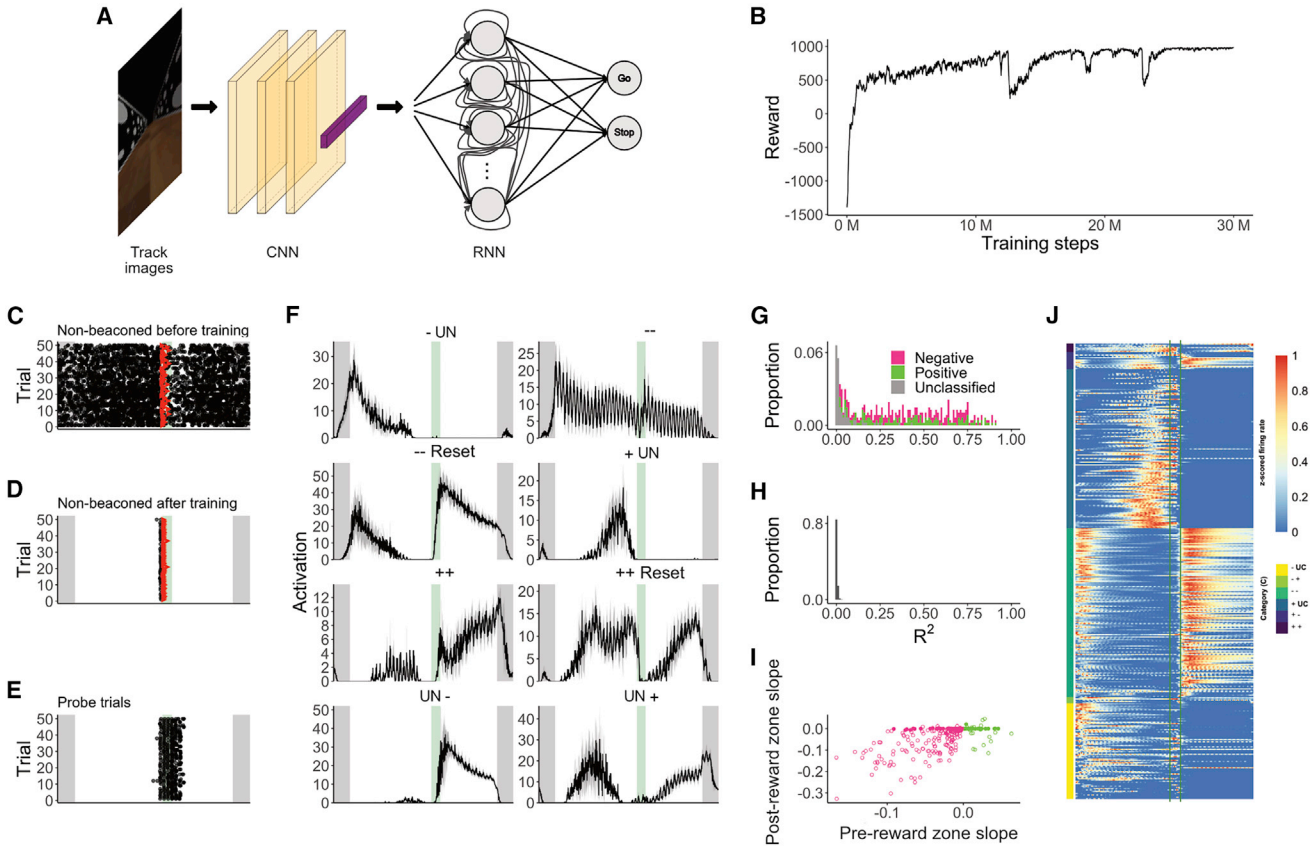


Figure 6. Recurrent network models trained to navigate to a rewarded location generate ramp-like representations

(A) Schematic of the network. Images of the track were played into a convolutional neural network (CNN) that sends outputs to a recurrently connected network (RNN) that sets the activity of movement control neurons. On each simulation step the network can either select "Go," which advances the agent, or "Stop," which causes the agent to remain at the same location. During training, a reinforcement learning algorithm modifies the network weights according to whether the agent stops at the reward zone.

(B) Reward obtained as a function of training step. If the agent stops in a non-goal location, it receives a reward of −1, whereas moving forward is a neutral action with a reward of 0 and stopping in the goal locations yields a reward of 1,000. Dips in the reward occur when the agent tests new (unsuccessful) navigation strategies. After developing a successful strategy during training, the agent typically takes ~250 simulation steps to reach the end of the track.

(C and D) Rasters of stopping locations on non-beaconed trials before (C) and after (D) the agent has learned the task.

(E) Rasters of stopping locations on probe trials after the agent has learned the task.

(F) Examples of RNN unit activity as a function of track location after the agent has learned the task.

(G) R^2 values for fits of activity of RNN units as a function of position in the track segment before the reward zone. Units are classified as in Figure 1.

(H) As for (G) but for shuffled data.

(I) Slopes of the firing rate of RNN units after the reward zone as a function of slopes before the reward zone. On the track segment before the reward zone 43.4% of neurons were classified as having a negative ramp slope and 29.3% a positive ramp slope. On the track segment after the reward zone 33.0% were classified as having a negative slope and 4.7% a positive slope.

(J) Normalized firing rate as a function of position on the track for all units in the RNN that demonstrate ramping activity on the track segment before the reward zone.

the track approaching the reward zone not predicting the slope after the reward zone (Figure 6I). Also similar to the experimental data, ramps generated by the trained recurrent network more frequently had a negative rather than a positive slope, although in the recurrent networks negative slopes were typically steeper than positive slopes (Figure 6I).

The presence in artificial recurrent neural networks and in the retrohippocampal cortices of activity profiles characterized by ramping activity and its interruption at reward locations is consistent with these activity patterns providing an effective readout of location. To test whether ramping activity contributes to successful completion of the task, we explored the consequences of perturbing neurons in the recurrent network model targeted according to their activity pattern (Figure 7A). We first tested blocking of the output from either ramping or non-ramping neurons to other neurons in the recurrent network. In contrast to the trained network, which stops at the reward location on all trials (Figure 7B), networks in which recurrent outputs from either ramping neurons (Figure 7C) or non-ramping neurons (Figure 7D) were inactivated did not show a spatial stopping strategy, indicating that activity of both neuron types is critical for the network to function. Because inactivation might be expected to cause substantial behavioral deficits simply through disrupting recurrent network activity, we evaluated two further manipulations.

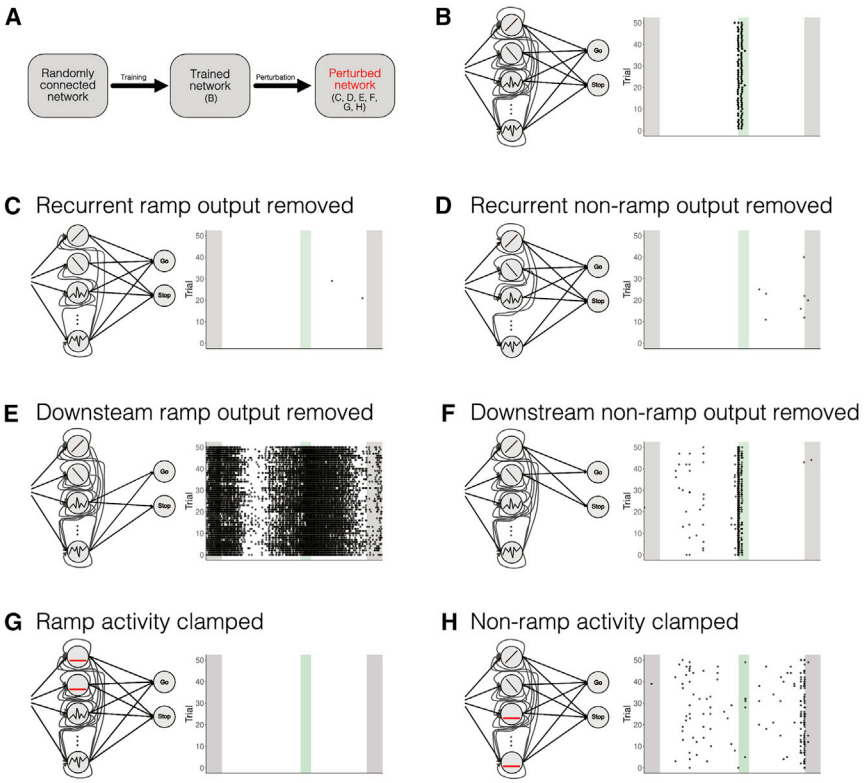


Figure 7. Perturbation of ramping activity disrupts control of spatial stopping by trained recurrent network models

(A) Perturbations were introduced after training randomly connected recurrent networks in the linear location estimation task.

(B–H) Raster plots show stopping locations of the trained network in control conditions (B), after blocking recurrent outputs from ramping neurons (C) and non-ramping neurons (D), after blocking the outputs to movement control neurons from ramping neurons (E) and non-ramping neurons (F), and after clamping the activity of ramping neurons (G) and non-ramping neurons (H) to its mean value. Criteria for selection of ramping and non-ramping neurons are in the methods.

Selective block of outputs from ramping neurons to downstream movement control neurons reduced the spatial selectivity of the stopping behavior (Figure 7E), whereas the spatial selectivity was largely maintained when the corresponding outputs from non-ramping neurons were blocked (Figure 7F). Clamping the activity of either ramping or non-ramping neurons to their mean activity level across a trial also abolished spatial stopping behavior, resulting in either abolished stopping when the activity of ramping neurons was clamped (Figure 7G), or a greater frequency and modified distribution of stops when activity of non-ramping neurons was clamped (Figure 7H).

Together, these analyses show that network configurations that solve the location estimation task are discoverable by reinforcement rather than being designed a priori, and that ramping activity is critical to the success of these network configurations. This is consistent with ramping activity patterns being able to play causal roles in recall of location memories.

DISCUSSION

Our results show that retrohippocampal neurons employ ramping neural activity to generate structured representations of learned environments, with reward location associated with switches in the ramp slope and resetting of the firing rate of ramping neurons. These representations could be maintained in the absence of location-specific cues, indicating that they result from recall of a spatial memory. As the direction of ramping activity could differ between segments of the behavioral track, representations of position may be specific to individual segments, potentially enabling ramp codes to combinatorially represent location and local context, for example by using ascending and descending ramp slopes to encode locations before and after the reward zone. These representations differ from the global metric provided by grid cells[3,44] but are reminiscent of robotic control strategies that segment the environment.[45] The ramp-like firing within a track segment is also reminiscent of spatial firing of neurons in the ventral striatum in anticipation of rewarded locations,[26–32] and of asymmetric firing of place cells in familiar track environments.[46] However, retrohippocampal activity differs in the diverse combinations of firing rate slopes approaching and following the reward location (Figure 1), and in their insensitivity to whether reward is delivered (Figures 3 and 4). These distinct features of ramp-like firing of retrohippocampal neurons may be suited to stable representation of the structure of the learned environments and location within them relative to salient landmarks, which in the task here are the reward zone and the start and end of the track.

Our results extend the representational strategies through which neurons in the hippocampal formation might encode goals. In the hippocampus, learning about rewarded goal locations has been associated with an increase in the density of place fields,[19–22] with the emergence of distinct discrete goal related firing fields,[20,23] or when goal and reward locations are distinct with additional discrete firing fields at the goal location.[47–49] In bats, discrete firing fields may also encode direction to goals and the distance from goals.[50] In the medial entorhinal cortex, learning about rewarded goal location causes a shift in the location of nearby grid firing fields.[24,25] In contrast, ramping activity that we describe here provides information about the distance to a goal, while resetting of ramps and changes in their polarity are associated with the goal location. More generally, this could be considered as encoding event boundaries. Similar to discrete directional and distance codes in the bat hippocampus,[50] ramping activity profiles did not require cues associated with the goal to be present. Thus, memory recall appears to be able to generate both types of representation.

What signals initiate, maintain, and terminate ramping activity? Initiation of ramping activity was associated with the start of the virtual track suggesting that it must be determined by visual signals. The progression of ramping activity was then best explained by the animal's track position alone or in conjunction with its speed and/or acceleration (Figure 2). Several observations indicate that for most neurons modulation of their firing rate by speed or acceleration is unlikely to be the primary drive of ramping activity. First, position rather than speed or acceleration was usually the dominant coefficient explaining their firing (Figure 2B). Second, pure speed and acceleration coding neurons were identified more frequently among non-ramping then ramping neurons (Figure 2C), and speed and acceleration profiles during the task differed from ramping activity profiles (Figures 2A and 3A). Third, evaluation of activity profiles on trials in which mice failed to stop for rewards enables further dissociation of the contributions of position, speed, and acceleration (Figure 3). Despite the speed profile being reversed on these trials ramping profiles were in general maintained. Similar considerations argue that changes in speed or acceleration are unlikely to account for resetting of ramping activity or switches in ramp slope in the regions of the reward zone. An alternative possibility is that reward expectation signals rather than position per se drive ramping activity (cf. Hok et al.,[47] Duvelle et al.,[48] and Takamiya et al.[49]); in this scenario, activity profiles on probe trials must still reflect recall of the structure of the track and task. While ramping activity on the track segment after the reward zone is inconsistent with reward expectation, dissociating reward expectation from location coding will require tasks in which the stopping location and reward delivery location differ (cf. Hok et al.[47]). It is important to note here that our results suggest considerable heterogeneity between ramping neurons, with subpopulations appearing to depend on kinematic variables rather than position, and with some positional neurons primarily driven by the visual cue rather than a path integration-dependent mechanism. Furthermore, our analyses focus on representations in trained animals, and it is possible that during learning kinematic or visual cues will play greater roles.

According to current theories for spatial computation, the location task we use here could be solved using path integration by grid cells to drive sequential activation of place cells along the track. In this scenario visual cues anchor the grid representations to the start of the track and to the reward zone. However, it is challenging for neural circuits to read out metric representations of distance to the stopping position from discrete spatial firing fields without using graded intermediate representations.[10,11,17] Our results are consistent with proposed solutions to this problem in which the structure of the environment is stored in the weights of synaptic inputs from place or grid cells onto ramping neurons.[10,11,17] In this case, learning modifies synaptic weights so that the discrete spatial inputs from upstream neurons drive ramping activity profiles across the length of the track. Ramping activity within a track segment could also be generated without place or grid cell input through integration of velocity signals by recurrent connectivity (Figures 6 and 7) or intrinsic dynamics of entorhinal neurons,[51–54] or could be driven by inputs from postrhinal neurons.[55] The intrinsic dynamics of recurrent networks could also mediate interruption of ramping activity at reward locations (Figures 6 and 7). In this case, a key difference from models in which grid or place inputs drive ramp firing is that the recurrent network performs a path integration-like computation, and the reward zone location is encoded by plasticity in the synaptic weights within the recurrent network.

Our analysis of artificial neural networks that solved the location memory task is consistent with the idea that ramping activity contributes to control of learned behaviors (Figures 6 and 7). The ramping activity emerged without a priori design to make the artificial networks resemble the retrohippocampal cortices, other than through the selection of an initially unstructured recurrent architecture. Thus, ramping codes may be a general solution for recall of location in track-like environments in the sense that their emergence is not tied specifically to the initial details of circuit organization. Further investigation will be required to establish whether this applies across other learning algorithms and to determine the extent to which the learned synaptic weights resemble connectivity in the retrohippocampal cortices. Because different spatial coding schemes emerge in artificial neural networks trained with different tasks (e.g., Banino et al.[56] and Sorscher et al.[57]), it will also be of interest to explore how task contingencies interact with learning rules to shape emergent representations. Impaired spatial localization after manipulation of the outputs from ramping neurons in the artificial networks suggests that ramping activity can play a necessary role in control of spatial behaviors. Similar tests of causality in animal models will require strategies for selective manipulation of ramping neurons during behaviors.

Given that deep layers of entorhinal cortex have extensive projections throughout the telencephalon,[58–60] ramping activity may be used by motor control neurons receiving input from the retrohippocampal cortices, or as a template for ramping activity of neurons in striatal circuits that receive direct inputs from retrohippocampal neurons. Thus, our results suggest a model for memory recall in which a neural path integrator drives ramping neural activity that reconfigures at salient locations and that may be suitable for control of behaviors by downstream brain structures.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- RESOURCE AVAILABILITY
  - Lead contact
  - Materials availability
  - Data and code availability
- EXPERIMENTAL MODEL AND SUBJECT DETAILS
- METHOD DETAILS
  - Microdrive fabrication
  - Surgical procedures
  - Handling and habituation
  - Food restriction
  - Experimental design overview
  - Location memory task
  - Open arena exploration
  - Post recording anatomical assessment of tetrode location

- QUANTIFICATION AND STATISTICAL ANALYSIS
  - Spike sorting
  - Classification of ramp-like activity
  - Ramp score calculation
  - Generalized linear mixed effect models for position, speed and acceleration
  - Classification of trial outcomes
  - Calculation of average speed
  - Calculation of spiking theta index
  - Power spectral analyses of local field potentials
  - Spike half width calculation
  - Classification of putative inhibitory and excitatory cells
  - Classification of cell types in the open arena
  - Neural network model

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at https://doi.org/10.1016/j.cub.2022.08.050.

ACKNOWLEDGMENTS

We thank Richard Morris for feedback on an earlier version of the manuscript, Robert Wallace for assistance with micro-CT imaging, Holly Stevens for technical assistance, and Gülşen Sürmeli, Ian Duguid, and members of the Sürmeli, Duguid, and Nolan labs for helpful discussions. This work was supported by grants to M.F.N. from the Wellcome Trust (200855/Z/16/Z), the BBSRC (BB/L010496/1), and the Simons Initiative for the Developing Brain; by a College of Medicine and Veterinary Medicine PhD Studentship, funded by the Thomas Work Fellowship, to K.Z.G.; by the Wellcome Trust (108890/Z/15/Z) Translational Neuroscience PhD programme to I.H.; and by the Medical Research Council Precision Medicine PhD programme to H.C. This work made use of resources provided by the Edinburgh Compute and Data Facility. For the purpose of open access, the author has applied a CC BY public copyright license to any Author Accepted Manuscript version arising from this submission.

AUTHOR CONTRIBUTIONS

Conceptualization, S.A.T., K.Z.G., E.R.W., and M.F.N.; formal analysis, visualization, and data curation, S.A.T., H.C., I.H., W.K.T., and M.F.N.; investigation, S.A.T., H.C., I.H., W.K.T., J.H., and W.Y.; resources, software, and validation, K.Z.G., S.A.T., H.C., I.H., and W.K.T.; writing – original draft, M.F.N. and S.A.T.; writing – review & editing, all authors; funding acquisition, supervision, and administration, M.F.N.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: January 14, 2022
Revised: July 5, 2022
Accepted: August 17, 2022
Published: September 12, 2022

REFERENCES

1. Moser, E.I., Moser, M.-B., and McNaughton, B.L. (2017). Spatial representation in the hippocampal formation: a history. Nat. Neurosci. *20*, 1448–1464.
2. O'Keefe, J., and Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. Brain Res *34*, 171–175.
3. Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., and Moser, E.I. (2005). Microstructure of a spatial map in the entorhinal cortex. Nature *436*, 801–806.
4. O'Reilly, R.C., and Rudy, J.W. (2001). Conjunctive representations in learning and memory: principles of cortical and hippocampal function. Psychol. Rev. *108*, 311–345.
5. McClelland, J.L., McNaughton, B.L., and O'Reilly, R.C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. Psychol. Rev. *102*, 419–457.
6. Whittington, J.C.R., Muller, T.H., Mark, S., Barry, C., Burgess, N., and Behrens, T.E.J. (2019). The Tolman-Eichenbaum Machine: unifying space and relational memory through generalisation in the hippocampal formation. Preprint at bioRxiv. https://doi.org/10.1016/j.cell.2020.10.024.
7. Eichenbaum, H. (2000). A cortical-hippocampal system for declarative memory. Nat. Rev. Neurosci. *1*, 41–50.
8. Eichenbaum, H. (2017). The role of the hippocampus in navigation is memory. J. Neurophysiol. *117*, 1785–1796.
9. O'Keefe, J., and Nadel, L. (1978). The Hippocampus as a Cognitive Map (Clarendon Press).
10. Stemmler, M., Mathis, A., and Herz, A.V.M. (2015). Connecting multiple spatial scales to decode the population activity of grid cells. Sci. Adv. *1*, e1500816.
11. Bush, D., Barry, C., Manson, D., and Burgess, N. (2015). Using grid cells for navigation. Neuron *87*, 507–520.
12. Stachenfeld, K.L., Botvinick, M.M., and Gershman, S.J. (2017). The hippocampus as a predictive map. Nat. Neurosci. *20*, 1643–1653.
13. Foster, D.J., Morris, R.G., and Dayan, P. (2000). A model of hippocampally dependent navigation, using the temporal difference learning rule. Hippocampus *10*, 1–16.
14. Gershman, S.J. (2014). Dopamine ramps are a consequence of reward prediction errors. Neural Comput *26*, 467–471.
15. Gustafson, N.J., and Daw, N.D. (2011). Grid cells, place cells, and geodesic generalization for spatial reinforcement learning. PLoS Comput. Biol. *7*, e1002235.
16. Burgess, N., and O'Keefe, J. (1996). Neuronal computations underlying the firing of place cells and their role in navigation. Hippocampus *6*, 749–762.
17. Fiete, I.R., Burak, Y., and Brookings, T. (2008). What grid cells convey about rat location. J. Neurosci. *28*, 6858–6871.
18. Nyberg, N., Duvelle, É., Barry, C., and Spiers, H.J. (2022). Spatial goal coding in the hippocampal formation. Neuron *110*, 394–422.
19. Dupret, D., O'Neill, J., Pleydell-Bouverie, B., and Csicsvari, J. (2010). The reorganization and reactivation of hippocampal maps predict spatial memory performance. Nat. Neurosci. *13*, 995–1002.
20. Jarzebowski, P., Hay, Y.A., Grewe, B.F., and Paulsen, O. (2022). Different encoding of reward location in dorsal and intermediate hippocampus. Curr. Biol. *32*, 834–841. e5.
21. Lee, J.S., Briguglio, J.J., Cohen, J.D., Romani, S., and Lee, A.K. (2020). The statistical structure of the hippocampal code for space as a function of time, context, and value. Cell *183*, 620–635. e22.
22. Hollup, S.A., Molden, S., Donnett, J.G., Moser, M.B., and Moser, E.I. (2001). Accumulation of hippocampal place fields at the goal location in an annular watermaze task. J. Neurosci. *21*, 1635–1644.
23. Gauthier, J.L., and Tank, D.W. (2018). A dedicated population for reward coding in the hippocampus. Neuron *99*, 179–193. e7.
24. Butler, W.N., Hardcastle, K., and Giocomo, L.M. (2019). Remembered reward locations restructure entorhinal spatial maps. Science *363*, 1447–1452.
25. Boccara, C.N., Nardin, M., Stella, F., O'Neill, J., and Csicsvari, J. (2019). The entorhinal cognitive map is attracted to goals. Science *363*, 1443–1447.
26. Howe, M.W., Tierney, P.L., Sandberg, S.G., Phillips, P.E.M., and Graybiel, A.M. (2013). Prolonged dopamine signalling in striatum signals proximity and value of distant rewards. Nature *500*, 575–579.

27. van der Meer, M.A.A., Johnson, A., Schmitzer-Torbert, N.C., and Redish, A.D. (2010). Triple dissociation of information processing in dorsal striatum, ventral striatum, and hippocampus on a learned spatial decision task. Neuron *67*, 25–32.

28. van der Meer, M.A.A., and Redish, A.D. (2011). Theta phase precession in rat ventral striatum links place and reward information. J. Neurosci. *31*, 2843–2854.

29. Fiorillo, C.D., Tobler, P.N., and Schultz, W. (2003). Discrete coding of reward probability and uncertainty by dopamine neurons. Science *299*, 1898–1902.

30. Carelli, R.M., and Deadwyler, S.A. (1994). A comparison of nucleus accumbens neuronal firing patterns during cocaine self-administration and water reinforcement in rats. J. Neurosci. *14*, 7735–7746.

31. Lavoie, A.M., and Mizumori, S.J. (1994). Spatial, movement- and reward-sensitive discharge by medial ventral striatum neurons of rats. Brain Res *638*, 157–168.

32. Schultz, W., Apicella, P., Scarnati, E., and Ljungberg, T. (1992). Neuronal activity in monkey ventral striatum related to the expectation of reward. J. Neurosci. *12*, 4595–4610.

33. Tennant, S.A., Fischer, L., Garden, D.L.F., Gerlei, K.Z., Martinez-Gonzalez, C., McClure, C., Wood, E.R., and Nolan, M.F. (2018). Stellate cells in the medial entorhinal cortex are required for spatial learning. Cell Rep *22*, 1313–1324.

34. Domnisoru, C., Kinkhabwala, A.A., and Tank, D.W. (2013). Membrane potential dynamics of grid cells. Nature *495*, 199–204.

35. Harvey, C.D., Collman, F., Dombeck, D.A., and Tank, D.W. (2009). Intracellular dynamics of hippocampal place cells during virtual navigation. Nature *461*, 941–946.

36. Chen, G., King, J.A., Burgess, N., and O'Keefe, J. (2013). How vision and movement combine in the hippocampal place code. Proc. Natl. Acad. Sci. USA *110*, 378–383.

37. Diehl, G.W., Hon, O.J., Leutgeb, S., and Leutgeb, J.K. (2017). Grid and nongrid cells in medial entorhinal cortex represent spatial location and environmental features with complementary coding schemes. Neuron *94*, 83–92. e6.

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

39. Boccara, C.N., Sargolini, F., Thoresen, V.H., Solstad, T., Witter, M.P., Moser, E.I., and Moser, M.-B. (2010). Grid cells in pre- and parasubiculum. Nat. Neurosci. *13*, 987–994.

40. Hardcastle, K., Maheswaranathan, N., Ganguli, S., Giocomo, L.M., Hardcastle, K., Maheswaranathan, N., Ganguli, S., and Giocomo, L.M. (2017). A multiplexed, heterogeneous, and adaptive code for navigation in medial entorhinal cortex. Neuron *94*, 375–387. e7.

41. Zilli, E.A. (2012). Models of grid cell spatial firing published 2005–2011. Front. Neural Circuits *6*, 16.

42. Kinkhabwala, A.A., Gu, Y., Aronov, D., and Tank, D.W. (2020). Visual cue-related activity of cells in the medial entorhinal cortex during navigation in virtual reality. eLife *9*, e43140.

43. Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. (2017). Proximal policy optimization algorithms. Preprint at arXiv. https://doi.org/10.48550/arXiv.1707.06347.

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

45. Kuipers, B., and Kuipers, B. (1999). The Spatial Semantic Hierarchy {(v.2). Artif. Intell. *119*, 191–233.

46. Mehta, M.R., Quirk, M.C., and Wilson, M.A. (2000). Experience-dependent asymmetric shape of hippocampal receptive fields. Neuron *25*, 707–715.

47. Hok, V., Lenck-Santini, P.-P., Roux, S., Save, E., Muller, R.U., and Poucet, B. (2007). Goal-related activity in hippocampal place cells. J. Neurosci. *27*, 472–482.

48. Duvelle, É., Grieves, R.M., Hok, V., Poucet, B., Arleo, A., Jeffery, K.J., and Save, E. (2019). Insensitivity of place cells to the value of spatial goals in a two-choice flexible navigation task. J. Neurosci. *39*, 2522–2541.

49. Takamiya, S., Shiotani, K., Ohnuki, T., Osako, Y., Tanisumi, Y., Yuki, S., Manabe, H., Hirokawa, J., and Sakurai, Y. (2021). Hippocampal CA1 neurons represent positive feedback During the learning process of an associative memory task. Front. Syst. Neurosci. *15*, 718619.

50. Sarel, A., Finkelstein, A., Las, L., and Ulanovsky, N. (2017). Vectorial representation of spatial goals in the hippocampus of bats. Science *355*, 176–180.

51. Egorov, A.V., Hamam, B.N., Fransén, E., Hasselmo, M.E., and Alonso, A.A. (2002). Graded persistent activity in entorhinal cortex neurons. Nature *420*, 173–178.

52. Dhillon, A., and Jones, R.S. (2000). Laminar differences in recurrent excitatory transmission in the rat entorhinal cortex in vitro. Neuroscience *99*, 413–422.

53. Yoshida, M., and Hasselmo, M.E. (2009). Persistent firing supported by an intrinsic cellular mechanism in a component of the head direction system. J. Neurosci. *29*, 4945–4952.

54. Peng, Y., Barreda Tomás, F.J., Klisch, C., Vida, I., and Geiger, J.R.P. (2017). Layer-specific organization of local excitatory and inhibitory synaptic connectivity in the rat presubiculum. Cereb. Cortex *27*, 2435–2452.

55. LaChance, P.A., Todd, T.P., and Taube, J.S. (2019). A sense of space in postrhinal cortex. Science *365*, eaax4192.

56. Banino, A., Barry, C., Uria, B., Blundell, C., Lillicrap, T., Mirowski, P., Pritzel, A., Chadwick, M.J., Degris, T., Modayil, J., et al. (2018). Vector-based navigation using grid-like representations in artificial agents. Nature *557*, 429–433.

57. Sorscher, B., Mel, G.C., Ocko, S.A., Giocomo, L., and Ganguli, S. (2020). A unified theory for the computational and mechanistic origins of grid cells. Preprint at bioRxiv. https://doi.org/10.1101/2020.12.29.424583.

58. Sürmeli, G., Marcu, D.C., McClure, C., Garden, D.L.F., Pastoll, H., and Nolan, M.F. (2015). Molecularly defined circuitry reveals input-output segregation in deep layers of the medial entorhinal cortex. Neuron *88*, 1040–1053.

59. Swanson, L.W., and Köhler, C. (1986). Anatomical evidence for direct projections from the entorhinal area to the entire cortical mantle in the rat. J. Neurosci. *6*, 3010–3023.

60. Insausti, R., Herrero, M.T., and Witter, M.P. (1997). Entorhinal cortex of the rat: cytoarchitectonic subdivisions and the origin and distribution of cortical efferents. Hippocampus *7*, 146–183.

61. Gerlei, K., Passlack, J., Hawes, I., Vandrey, B., Stevens, H., Papastathopoulos, I., and Nolan, M.F. (2020). Grid cells are modulated by local head direction. Nat. Commun. *11*, 4228.

62. Matsumoto, J. (2019). Matsumotoj/tetroplater (first release). https://doi.org/10.5281/zenodo.3334858.

63. Osborne, J.E., and Dudman, J.T. (2014). RIVETS: a mechanical system for in vivo and in vitro electrophysiology and imaging. PLoS One *9*, e89007.

64. Siegle, J.H., López, A.C., Patel, Y.A., Abramov, K., Ohayon, S., and Voigts, J. (2017). Open Ephys: an open-source, plugin-based platform for multi-channel electrophysiology. J. Neural Eng. *14*, 045003.

65. Lopes, G., Bonacchi, N., Frazão, J., Neto, J.P., Atallah, B.V., Soares, S., Moreira, L., Matias, S., Itskov, P.M., Correia, P.A., et al. (2015). Bonsai: an event-based framework for processing and controlling data streams. Front. Neuroinform. *9*, 7.

66. Masís, J., Mankus, D., Wolff, S.B.E., Guitchounts, G., Joesch, M., and Cox, D.D. (2018). A micro-CT-based method for quantitative brain lesion characterization and electrode localization. Sci. Rep. *8*, 5184.

67. Xie, S., Wallace, R.J., and Pankaj, P. (2020). Time-dependent behaviour of demineralised trabecular bone - Experimental investigation and development of a constitutive model. J. Mech. Behav. Biomed. Mater. *109*, 103751.

68. Chrobak, J.J., and Buzsáki, G. (1998). Gamma oscillations in the entorhinal cortex of the freely behaving rat. J. Neurosci. *18*, 388–398.

69. Colgin, L.L., Denninger, T., Fyhn, M., Hafting, T., Bonnevie, T., Jensen, O., Moser, M.B., and Moser, E.I. (2009). Frequency of gamma oscillations routes flow of information in the hippocampus. Nature *462*, 353–357.

70. Kornienko, O., Latuske, P., Bassler, M., Kohler, L., and Allen, K. (2018). Non-rhythmic head-direction cells in the parahippocampal region are not constrained by attractor network dynamics. eLife *7*, e35949.

71. Chrobak, J.J., and Buzsáki, G. (1994). Selective activation of deep layer (V-VI) retrohippocampal cortical neurons during hippocampal sharp waves in the behaving rat. J. Neurosci. *14*, 6160–6170.

72. Chung, J.E., Magland, J.F., Barnett, A.H., Tolosa, V.M., Tooker, A.C., Lee, K.Y., Shah, K.G., Felix, S.H., Frank, L.M., and Greengard, L.F. (2017). A fully automated approach to spike sorting. Neuron *95*, 1381–1394.e6.

73. Virtanen, P., Gommers, R., Oliphant, T.E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., et al. (2020). SciPy 1.0: fundamental algorithms for scientific computing in Python. Nat. Methods *17*, 261–272.

74. Benjamini, Y., and Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. J. Royal Stat. Soc. *57*, 289–300.

75. Kropff, E., Carmichael, J.E., Moser, M.-B., and Moser, E.I. (2015). Speed cells in the medial entorhinal cortex. Nature *523*, 419–424.

76. Tsao, A., Sugar, J., Lu, L., Wang, C., Knierim, J.J., Moser, M.-B., and Moser, E.I. (2018). Integrating time from experience in the lateral entorhinal cortex. Nature *561*, 57–62.

77. Góis, Z.H.T.D., and Tort, A.B.L. (2018). Characterizing speed cells in the rat hippocampus. Cell Rep *25*, 1872–1884.e4.

78. Bates, D., Mächler, M., Bolker, B.M., and Walker, S.C. (2015). Fitting linear mixed-effects models using lme4. J. Stat. Softw. *67*, 1–48.

79. Fox, J., and Weisberg, S. (2018). An R Companion to Applied Regression (SAGE Publications).

80. Clark, M., and Bracci, E. (2018). Dichotomous dopaminergic control of ventral pallidum neurons. Front. Cell. Neurosci. *12*, 260.

81. Ye, J., Witter, M.P., Moser, M.-B., and Moser, E.I. (2018). Entorhinal fast-spiking speed cells project to the hippocampus. Proc. Natl. Acad. Sci. USA *115*, E1627–E1636.

82. Skaggs, W.E., McNaughton, B.L., and Gothard, K.M. (1993). An information-theoretic approach to deciphering the hippocampal code. In Advances in Neural Information Processing Systems, *5* (NIPS 1992), pp. 1030–1037.

83. Elman, J.L. (1990). Finding structure in time. Cogn. Sci. *14*, 179–211.

84. Kingma, D.P., and Ba, J. (2014). Adam: a method for stochastic optimization. Preprint at arXiv. https://doi.org/10.48550/arXiv.1412.6980.

85. Bergstra, J., and Bengio, Y. (2012). Random search for hyper-parameter optimization. J. Mach. Learn. Res. *13*, 281–305.

86. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. (2019). PyTorch: an imperative style, high-performance deep learning library. In Advances in Neural Information Processing Systems (Curran Associates, Inc.), pp. 8026–8037.

87. Kostrikov, I. (2018). Pytorch implementations of reinforcement learning algorithms. https://github.com/ikostrikov/pytorch-a2c-ppo-acktr-gail.

CellPress
OPEN ACCESS

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Deposited data		
Electrophysiology data	This study	https://datashare.ed.ac.uk/handle/10283/777; https://doi.org/10.7488/ds/3515
Histology	This study	Supplemental information section of the manuscript
Experimental models: Organisms/strains		
C57BL/BNCrl mice	Charles River	C57BL/BNCrl
Software and algorithms		
Python	Python Software Foundation	Version 3.5.1
R	The R Project for Statistical Computing	Versions 3.3.1 – 4.2.0.
Analysis code	This study.	https://github.com/MattNolanLab/Ramp_analysis; https://doi.org/10.5281/zenodo.6993158
Simulation code	This study	https://github.com/MattNolanLab/Ramp_simulations; https://doi.org/10.5281/zenodo.6982553
Other		
Electrode interface board	Neuralynx	EIB-16
Head post	Ronal	Rivets
Headstage	Intan technologies	RHD-16
Data acquisition board	OpenEphys	N/A

RESOURCE AVAILABILITY

Lead contact

Further information and requests for resources and reagents should be directed to and will be fulfilled by the lead contact: Matthew Nolan (mattnolan@ed.ac.uk).

Materials availability

This study did not generate new unique reagents.

Data and code availability

- Data is deposited at https://datashare.ed.ac.uk/handle/10283/777 (see key resources table for DOIs).
- All original code used for analysis of the data is available at https://github.com/MattNolanLab (see key resources table for DOIs).
- Any additional information required to reanalyze the data reported in this paper is available from the lead contact upon request.

EXPERIMENTAL MODEL AND SUBJECT DETAILS

All animal experiments were carried out under a project license granted by the UK Home Office, were approved by the Animal Welfare and Ethical Review Board of the University of Edinburgh School of Medicine and Veterinary Medicine, and conformed with the UK Animals (Scientific Procedures) Act 1986 and the European Directive 86/609/EEC on the protection of animals used for experimental purposes.

Experiments used 11 male C57BL/BNCrl wild type mice (Charles River) 7-14 weeks (mean = 12.6, SD = 0.71) old at the time of the implant surgery. Males were used for consistency in their size, which was important for operation of the behavioral apparatus. One week prior to surgery animals were group housed (3-5 mice per cage) in a holding room on a reversed 12 hour on/off light cycle (light

from 7 pm to 7 am). After surgery mice were single housed in cages that contained a running wheel and cardboard igloo for additional enrichment. Prior to surgery and for 2-3 days following surgery mice were given standard laboratory chow and water *ad libitum*.

METHOD DETAILS

Microdrive fabrication

Microdrive fabrication for in vivo electrophysiology recordings was similar to our previous work.[61] 16-channel microdrives consisting of 4 tetrodes were built by threading 90% platinum, 10% iridium wire (18 μm HML-coated, Neuralynx) to an EIB-16 board (Neuralynx) via an inner cannula (21 gauge 9 mm long). The board was then covered in epoxy and once dried a poor lady frame (Axona) was cemented to the side. An outer cannula (17 gauge 7 mm), sanded at a 45 degree angle on one end to fit the curvature of the skull, was placed around the inner cannula. The outer cannula was secured temporarily using vaseline, allowing it to be lowered during the surgery. Tetrodes were then trimmed to ~3 mm using ceramic scissors (Science Tools, Germany) and gold plated (Non-cyanide gold plating solution, Neuralynx) using a custom built tetrode plater[62] until the impedance was between 150 and 200 kΩ.

Surgical procedures

Before surgery, tips of the tetrodes were washed with ethanol and then with sterile saline. General surgical procedures were carried out as described previously.[58] Inhalation anesthesia was induced using 5 % isoflurane / 95 % oxygen, and sustained at 1 – 2 % isoflurane / 98-99 % oxygen with a flow rate of 1 L / minute. An oval incision was made to expose the surface of the skull and a RIVETS head-post[63] attached to the skull and held in place with UV curing resin cement (RelyX Unicem, 3M). For electrical grounding, two small craniotomies were drilled on the left hemisphere ~3.4 mm lateral, and ~1 mm rostral relative to Bregma and the center of the intraparietal plate respectively. M1 x 4 mm screws (AccuGroup SFE-M1-4-A2) were then implanted just below the surface of the skull and cemented in place.

The microdrive was attached to an Omnetics to Mill-Max adaptor (Axona, HSADPT-NN1), which was then attached to the stereotaxic frame. The tetrodes were lowered 1.2 - 1.4 mm into the brain, beginning at 3.4 mm lateral from Bregma (right hemisphere) and along the lambdoid suture, and at an angle of -16° in order to target retrohippocampal cortices. The outer cannula was then lowered and sterile Vaseline carefully placed around the rim to seal the craniotomy. The implant was fixed to the skull with UV curing resin. After the resin hardened, the grounding wires were wrapped around the grounding screws and fixed with silver paint (RS components 101–5621). A layer of resin was applied to cover the grounding screws and any holes were filled with dental cement (Simplex Rapid). Mice recovered on a heat mat for ~20 minutes, and had unlimited access to Vetergesic jelly (0.5 mg / kg of body weight buprenorphine in raspberry jelly) for 12 hours after surgery. Mice were given a minimum of 2 days postoperative recovery before proceeding.

Handling and habituation

Handling and habituation of mice was similar to that described previously.[33,61] Mice were handled twice a day for 5 - 10 minutes for one week following surgery. Animals were then habituated to the virtual reality setup for 10 and 20 minutes respectively over two consecutive days. After each habituation session animals were given access to soy milk in their cages to familiarize them with the reward and were given access to a larger arena for 5 - 10 minutes of free exploration.

Food restriction

To motivate mice to obtain rewards in the location memory task their access to food was restricted so that their body weight was maintained at ~85% of its baseline value.[33] From 4-5 days before starting training an allotted food amount, determined on a daily basis to maintain the animal's weight at its target value, was made available to the mice. The baseline weight of the animal was calculated from its weight prior to restriction and normalized to the expected daily growth for the animal's age.

Experimental design overview

Experimental days involved two recording sessions; one in virtual reality and one in an open arena (Figure 5A). On these days animals were collected from the holding room 30 - 60 minutes before the first recording session. Before recording, animals were handled for 5 - 10 minutes, weighed and placed for 10-20 minutes in a cage filled with objects and a running wheel. Animals were then head fixed in the virtual reality setup and performed the location memory task for approximately 30 minutes. After the recording session, animals were placed back in the object-filled arena for 30 minutes before a second recording session in an open field arena for 30 minutes. They were then returned to their home cage. Tetrodes were typically lowered by 50 - 100 μm after each session. At the end of the experiment animals were sacrificed and the tissue processed for anatomical localization of the recording location (described below). Animals in which the recording location was outside the retrohippocampal areas were excluded from analysis (n = 2).

Location memory task

Mice were trained to obtain rewards at a marked location on a virtual linear track as described previously.[33] During recording sessions animals were head fixed in the virtual reality setup for approximately 30 minutes. Mice were head fixed using a Rivets clamp (Ronal Tool Company)[63] and ran on a custom made cylindrical treadmill fitted with a rotary encoder (Pewatron). A feeding tube placed in front of the animal dispensed soy milk rewards (5 - 10 μl per reward). Virtual tracks were generated using Blender3D (https://blender.com/). The track had length 2 m, with a 60 cm track zone, a 20 cm reward zone, a second 60 cm track zone and a 60 cm

black box to separate successive trials. The walls and floor of the track were labelled with a repeating pattern (Figure 1A). On trials in which the reward zone was visible it was indicated by distinct vertical green and black bars. At any point on the track the visible distance ahead of the mouse was 40 cm.

Electrophysiological signals from the 16 channel microdrive were acquired using an Intan headstage connected via an SPI cable (Intan Technologies, RHD2000 6-ft (1.8 m) standard SPI interface cable) attached to an Open Ephys acquisition board.[64] Signals were filtered using the board's analogue high and low cut filter (2.5 Hz -7603.8 Hz). Behavioral variables including position, trial number and trial type were calculated in Blender3D at 60 Hz as described previously.[33] Using Blender's embedded Python interpreter (Python API), custom scripts wrote each variable (60 Hz, 12-bit precision) to a data acquisition (DAQ) microcontroller (Arduino Due) which in turn sent the signals to the OpenEphys acquisition board via HDMI.

Open arena exploration

The open arena, as described in Gerlei et al.,[61] consisted of a metal box with a square floor area, removable metal walls, and a metal frame (Frame parts from Kanya UK, C01-1, C20-10, A33-12, B49-75, B48-75, A39-31, ALU3).

For motion and head direction tracking, a camera (Logitech B525, 1280 x 720 pixels Webcam, RS components 795-0876) was attached to the frame of the open arena. To record motion alongside head-direction data, red and green polystyrene balls were attached to the sides of the headstage. A custom Bonsai script then tracked the polystyrene balls during recording sessions based on their color.[65] To synchronize the position and electrophysiology data, an LED was attached to the side of the open arena in the field of view of the camera. Pulses were sent to the LED via an Arduino (Arduino Uno) as well as to Open Ephys acquisition board via the I/O board. The pulses were separated by randomly generated intervals (20 to 60 seconds) so that the two datastreams could be reconciled during subsequent analysis.[65]

Electrophysiological signals were recorded using the custom made 16 channel microdrive described above. An Intan recording headstage with an SPI cable (Intan Technologies, RHD2000 6-ft (1.8 m) Ultra Thin SPI interface cable C3216) attached was used to connect the microdrive to a commutator (SPI cable adapter board, Intan Technologies C3430 and 3D printed holder, custom designed by Patrick Spooner). The commutator was connected to an Open Ephys acquisition board[64] that sent the electrophysiology signal to a computer where the signal was filtered between 2.5 Hz - 7603.8 Hz and displayed using the OpenEphys GUI.

Post recording anatomical assessment of tetrode location

To determine tetrode locations we used either Cresyl violet staining (n = 3 mice) or microcomputed tomography (microCT) (n = 7 mice) (Data S1). Two mice that underwent cresyl violet staining had no visible tetrode tracks and are omitted from Data S1. For these mice tetrode location was inferred by the presence of theta modulation of the local field potential and the spiking activity of recorded neurons (Figure S3). For both methods mice were anesthetized with isoflurane and given a lethal dose of sodium pentobarbital (Euthatal, Meridal Animal Health, UK).

For cresyl violet staining, we made an electrolytic lesion using 2 x 2 ms, 20 μA current pulses per tetrode. The mice were fixed by transcardial perfusion of ice cold phosphate buffered saline (PBS; Invitrogen) followed by 4 % paraformaldehyde (PFA; Sigma Aldrich) in 0.1 M phosphate buffer (PB; Sigma Aldrich). Brains were removed, left in 4 % PFA in 0.1 M PB overnight and then in 30 % sucrose in PBS for two nights. A freezing microtome was used to cut 50 μm sagittal slices, which were mounted on polarized slides, dried overnight and then placed in xylene before rehydration in varying concentrations of ethanol (100%, 90% then 70% for 1-2 minutes each). Once washed, sections were placed in Cresyl violet for 2 mins before additional washing and dehydration (70% then 90% for 1.5 minutes each followed by two 30 second washes in 100%). Sections were then placed in xylene and mounted with Dibutylphthalate Polystyrene Xylene (DPX). Imaging of sections was performed with a slide scanner (Zeiss Axioscan.Z1). Sections were then manually inspected for the lesion site.

For microCT imaging, we modified a protocol used previously for imaging rat brains.[66] Animals were perfused with a mixture of PFA and glutaraldehyde and then left in the same solution for 2 nights. Brains were then washed in ddH_20 before being incubated at 4° for two weeks in 2 % osmium tetroxide (2% OsO_4). After osmium staining brains were washed in ddH_20, dehydrated in ethanol and then embedded in resin. Once the resin had cured, we placed the sample in a micro-CT scanner (Skyscan 1172, Bruker, Kontich, Belgium) which generated a 3D image of the sample. Scanning parameters were: source voltage 54 kV, current 185 μA, exposure 885 ms with a 0.5 mm aluminum filter between the X-ray source and the sample.[67] After acquiring the scan the dataset was reconstructed (CTan software, v1.13.5.1) and viewed with DataViewer (Bruker).

For the Cresyl violet and microCT methods, tetrodes were localized relative to landmarks in version 2 of the Allen Reference Atlas for the mouse brain (https://mouse.brain-map.org/static/atlas). Tetrodes were assigned either to the MEC (n = 5 mice), the pre/para-subiculum (n = 2), or when the location was ambiguous between these areas to retrohippocampal cortex (n = 3) (Data S1). Theta frequency field potentials and theta phase locking of spike firing (Figure S3) were consistent with previous reports for these regions.[39,68–70] We did not attempt to differentiate between presubiculum and parasubiculum as in the mouse their borders are not well defined without additional markers and there is disagreement between reference atlases when using landmarks as reference points. Location of all tetrodes in the retrohippocampal region, defined as including the MEC, presubiculum and parasubiculum (e.g., see Chrobak and Buzsáki[71]) was corroborated by electrophysiological recording of a theta field potential and theta modulation of the firing of neurons (cf. Boccara et al.[39]) (see Figure S5 and analyses section below). In contrast, neurons in nearby structures such as the postrhinal cortex do not show theta phase locked firing.[55] The animals from which we were

unable to localize the tetrodes anatomically, but in which field potential recordings and spike firing showed theta modulation, were assigned to the retrohippocampal group.

QUANTIFICATION AND STATISTICAL ANALYSIS

Analyses were carried out using Python (version 3.5.1) and R versions: 3.3.1 to 4.2.0.

Spike sorting

Spikes were isolated from electrophysiological data using an automated pipeline based on MountainSort (v 0.11.5 and dependencies)[72] and similar to that described in Gerlei et al.[61] Data pre-processing steps converted Open Ephys files to mda format, filtered signals between 600 Hz - 6000 Hz and then performed spatial whitening over all channels. Events were detected from peaks > 3 standard deviations above baseline and separated by at least 0.33 ms from other events on the same channel. A 10-dimensional feature space was generated using the first 10 principal components of the detected waveforms and spike sorting using the ISO SPLIT algorithm performed on the feature space.[72] Cluster quality was evaluated using metrics for isolation, noise-overlap, and peak signal to noise ratio.[72] Units that had a firing rate > 0.5 Hz, isolation > 0.9, noise overlap < 0.05, and peak signal to noise ratio > 1 were accepted for further analysis.

Classification of ramp-like activity

To classify ramping neurons (Figures 1 and S4) we generated firing rate maps for each trial by binning the track into 1 cm bins, summing the number of spikes in each bin and dividing by the time the animal spent there. Bins in which the animal's speed was < 3 cm/s were removed. For the purposes of data visualization, firing rates were smoothed by convolution with a Gaussian filter (SD = 2 cm, SciPy Python package[73]) and then averaged across trials. For each neuron, we also generated 1000 shuffled datasets by adding a random time drawn from a uniform distribution between 20-580 seconds and added this to the timestamp of each spike. Random times were drawn independently for each trial and shuffle. Spike times that exceeded the recording length were wrapped around to the start of the session. Spike locations were recomputed from the shuffled spike times and a firing rate map generated as described previously. We then fit linear models to firing rate as a function of position across the regions of the track before (0-60 cm) or after the reward zone (80-140 cm). This was done for each of the 1000 shuffled firing rates and the original dataset for each neuron. The significance (p) values obtained from fitting the linear models were corrected for multiple comparisons using the Benjamini & Hochberg method.[74] We classified neurons as having ramp-like activity with positive (or negative) slopes if the linear model was significant (corrected $p < 0.01$) and the coefficients of the model were outside the 5 - 95 % intervals of the shuffled data. Neurons not meeting these criteria were allocated to the unclassified group. We focused our analyses on six groups of neurons separated according to their slope classifications: positive before the reward zone and after the reward zone (+ +), positive before the reward zone and negative after (+ -), positive before the reward zone and unclassified after (+ un), negative before the reward zone and negative after (- -), negative before the reward zone and positive after (- +) or negative before the reward zone and unclassified after (- un).

Ramp score calculation

As a complementary method to identify ramp-like firing we calculated a ramp score based on the correlation between firing rate and the animal's position (Figure S2). This ramp score is related to the score used previously for identification of speed cells.[75] For each cell we calculated separate ramp scores for the track segments before and after the reward zone (Figure S2A).

We first smoothed the mean firing rate of each cell in each binned position using a Savitzky-Golay filter (window size: 7, polynomial order: 1) and then identified locations before the reward zone at which peaks and troughs in the firing rate were present by the change of the sign of the slope. Using all turning points (peaks and troughs) in the smoothed curve, together with two endpoints (i.e. beginning and end) of the track segment before the reward zone, we construct a set of potential track segments. For example, if there is a single peak at 30 cm, with two endpoints of the zone at 0 cm and 60 cm, then the set of segments will be (0, 30), (30, 60), (0, 60), where each segment is defined by its starting and ending locations.

To avoid bias from location bins with many data points (e.g. when the animal slowed down in the reward zone), we oversample the original unsmoothed firing rate so that all segments have the same number of points. Then in each of the track segments d_i, using the balanced firing rate fr, we calculate the Pearson correlation coefficient r_i between the firing rate and the location loc.

$$r_i = corrcoeff_{d_i}(fr, loc)$$

To calculate the ramp score for each segment we first determine the ratio of the length of the track segment to the whole track,

$$C_i = \frac{d_i}{D}$$

where d_i is the length of the i-th track segment, D is the length of the track region of interest (e.g. the regions before or after the reward zone). The ramp score S_i of this segment is then the harmonic mean of the correlation coefficient and the segment ratio:

$$S_i = \frac{2 \cdot |r_i| \cdot C_i}{|r_i| + C_i}$$

A harmonic mean is used to ensure that a high ramp score will mean that the cell has a high correlation coefficient and a high segment ratio.

To determine the ramp score for each cell, the ramp scores of all possible segments were calculated and the final ramp score S_R of the cell is chosen as the ramp score that has the largest magnitude. The sign of the ramp score indicates whether the slope is positive or negative.

$$S_R = sign(r_{s_{max}}) \cdot max_i(S_i)$$

where $r_{s_{max}}$ is the correlation coefficient of the track segment corresponding to the maximum ramp score.

Generalized linear mixed effect models for position, speed and acceleration

To quantify the influence of position, speed and acceleration on the firing rate of each neuron, we built generalized linear mixed effect models to predict firing rate using position, speed and acceleration as fixed effects and trial number as a random effect. For each neuron, all variables were binned into 100 ms blocks and bins when the animal's speed was < 3 cm/s were removed.[76,77] For the purposes of data visualization, but not model fitting, firing rates and kinematic variables were smoothed by convolution with a Gaussian filter (SD = 200 ms, SciPy Python package[73]). The average speed and acceleration of the animal across the whole recording was calculated, and bins in which speed or acceleration was 3 standard deviations (SDs) above the average for that variable were removed. A Poisson model with the configuration 'Rate ~ Position + Speed + Acceleration + (1 + Position | Trials)' and log linker function was fit using the lme4 package (version 1.1-12) in R.[78] The ANOVA function provided by the car package (version 3.0-9) was used to calculate significance values for each model coefficient.[79] To classify neurons as being modulated by position (P), speed (S), acceleration (A), or a combination thereof, we used a significance threshold of 0.01. For example, a neuron with $p < 0.01$ for position and acceleration and ≥ 0.01 for speed coefficients was classified as a PA neuron. p values obtained from the ANOVA were corrected for multiple comparisons using the Benjamini & Hochberg method.[74]

Classification of trial outcomes

Trials were classified by task performance into hits and misses based on whether the mouse registered a stop (speed < 4.7 cm/s) within the reward zone[33] or not. Miss trials were split into near-hits (tries) and run-throughs (runs), by comparing their average speeds in the reward zone to hit trials. The 95th percentile of average speeds in the reward zone for hit trials in a single session was used to discriminate between tries (< 95th percentile speed) and run trials (> 95th percentile speed).

Calculation of average speed

The average speed as a function of location on the track for each session was calculated by binning speed into 1 cm location bins for each trial and then smoothing by convolution with a Gaussian filter (SD = 2 cm, SciPy Python package[73]). Speed for each location bin was then averaged over trials. The mean speed across all sessions was calculated for each animal for each trial outcome. The mean and standard error across all animals was calculated from the means for each animal.

Calculation of spiking theta index

Theta index was calculated following methods in Kornienko et al.[70] (Figure S3A). Spike times were transformed into instantaneous firing rates. A power spectrum was then estimated using Welch's method. The theta index was defined as

$$T = \frac{\theta - baseline}{\theta + baseline}$$

where θ is the mean theta band power (6-10 Hz) and baseline is the mean power of two adjacent frequency bands (3-5 and 11-13 Hz).

Power spectral analyses of local field potentials

We computed the power spectra across recording sessions for each channel for all raw voltage samples where the mouse was moving (> 3 cm/s). Power was estimated using Welch's method using a 2 second (60000 samples at 30Hz sample rate) Hanning window and averaged across channels.

Spike half width calculation

Spike half width (the width of the spike at half-maximal amplitude) for each neuron was calculated as described in Clark and Bracci.[80] For each neuron 50 spike waveforms were randomly selected and an average waveform generated. Spike amplitude was defined as the difference between the baseline voltage, calculated as the lowest value, and the positive voltage peak, calculated as the highest value. Spike half width was then calculated as the time (ms) the voltage stayed above half the maximal spike amplitude.

Classification of putative inhibitory and excitatory cells

Putative interneurons were defined as neurons with mean firing rate > 10 Hz, similar to previous studies in entorhinal populations[81] (Figure S7C). Excitatory neurons were classified as mean firing rates < 10 Hz. Consistent with this, spike half width in cells classified as putative interneurons were typically shorter than cells classified as excitatory (Figure S7C).

Classification of cell types in the open arena

Neurons were identified across recording sessions in the location memory task and open arena by concatenating the raw voltage channels from each session and then spike sorting. To classify neurons based on their activity in the open arena we followed methods in Diehl et al.[37] to assign cells into one of five hierarchically organized functional cell types: grid > border > other spatial > head direction > non-spatial. Neurons were assigned to the highest cell type in the hierarchy for which their corresponding identity metric was greater than the 99th percentile of the same scores from 1000 shuffled datasets. For example, a cell with grid score above the 99th percentile was classified as a grid cell irrespective of scores in other categories. To generate shuffled spike data, we drew a single value from a uniform distribution between 20-580 seconds and added this to the timestamp of each spike. Spike times that exceeded the recording length were wrapped around to the start of the session. Spike locations were recomputed from the shuffled spike times and spatial scores calculated similar to measured data. Using this method we generated 1000 shuffles per cell and used the 99th percentile of the scores as a threshold for belonging to the five cell types.

We used established scores for grid, border, head direction and spatial stability.[37,38,82] Grid scores were defined as the difference between the minimum correlation coefficient for rate map autocorrelogram rotations of 60 and 120 degrees and the maximum correlation coefficient for autocorrelogram rotations of 30, 90 and 150 degrees.[37,38,82] The firing rate map was calculated by binning spikes into 2.5 cm bins and dividing by the total time occupied in each bin and then smoothed with a Gaussian kernel. Autocorrelograms were calculated by sliding the rate map over all x and y bins and calculating a correlation score. Fields were detected in this autocorrelogram by converting it into a binary array using 20% of the maximal correlations as a threshold. If the binary array had more than 7 local maxima, a grid score was calculated. Correlations between the rotated autocorrelograms were then calculated using only a ring containing the 6 local maxima closest to the center of the binary array and excluding the maximum at the center. The ring was detected based on the average distance of the 6 fields near the center of the autocorrelogram (middle border = 1.25* average distance, outer border = 0.25 * average distance).

To calculate border scores,[37,38,82] firing fields were identified from the firing rate map by collecting neighboring bins with firing rates greater than 30% the maximal firing rate and covering at least 200 cm^2. For each putative field, the mean firing distance d_m was computed as the average distance to the wall for each bin, divided by the mean firing rate of the field. To obtain a value between 0 and 1, d_m was normalized by the shortest distance from the center of the arena to the wall. The maximal coverage c_m of a field was measured as the maximum proportion of field bins occupying any of the four edges of the firing rate map. The border score was then given by

$$b = \frac{c_m - d_m}{c_m + d_m}$$

To calculate head direction scores,[37,38,82] the head direction recorded at times corresponding to firing events were binned into 360 bins between 0 and 2π and was normalized by the duration of time spent occupying each directional bin. This polar histogram was then smoothed with a rolling sum with a window size of 23 degrees. The head direction score is the length of the mean vector of this polar histogram. To obtain this length, the x and y components are computed by first calculating dx and dy in a unit circle (radius = 1) in steps of 1 degree as

$$dx = \frac{cos(angle)}{radius}, dy = \frac{sin(angle)}{radius}$$

The head direction score is then given using the Pythagorean theorem by

$$\sqrt{(\Sigma polarHistogram \cdot dx)^2 + (\Sigma polarHistogram \cdot dy)^2}$$

For assessing spatial stability, we used two separate metrics, spatial information and the within-session spatial correlation to identify spatial cells.[3,37–40] We calculated the spatial information per spike as

$$\sum_{i=1}^{N} P_i \left(\frac{R_i}{R}\right) \cdot log_2 \left(\frac{R_i}{R}\right)$$

where i indexes a position bin in the firing rate map, N is the number of bins, P_i is the occupancy probability, R_i is the firing rate in the bin, R is the mean firing rate. The within-session spatial correlation is calculated by computing the Pearson correlation between the firing rate map computed from the first half session and the second half session. Bins that were not visited in both halves were excluded from the calculation.

Neural network model

We use a reinforcement learning algorithm to train artificial neural networks to solve a task similar to that used for the neuronal recordings. In the artificial task, the track consists of 1010 to 1150 locations (depending on track randomization), each with a corresponding image. The track layout was similar to the experimental track, but differed in that visual distance to the reward zone was increased to make the task more challenging and the visual cues indicating the reward zone location were removed in order to promote learning of cue-independent behavioral strategies. The reinforcement learning agent could either choose to take a step forward or to stop. If the agent chooses to move forwards, the image corresponding to the next location in the track is shown

to the network. The optimal behavior is to keep moving forwards to the hidden reward zone, stop once to acquire the reward, and start moving to the end of the track to proceed to the next trial.

We cropped and re-sampled images used for the track in the recording experiment to 89x89 pixels and used them as the input to the artificial network. Each image frame of the track was first processed by a three-layer convolutional neural network before being fed into a discrete-time recurrent neural network (RNN).[83] A two-unit linear layer was used to read out the stop/go decision from the RNN. The recurrent states were reset after every 10 trials (which make up a training episode). To make the RNN outputs positive and comparable to neural firing rates, we chose the rectified linear unit (ReLU) as the activation function.

To simulate trial-to-trial variability in movement speed observed in the experimental data, the speed at which the agent moves was randomized at the start of each trial. For example, on 'slow' trials, the agent moved forwards 3 locations every time they chose to step forwards, but on 'fast' trials they moved 7 locations. The possible trial speeds were 3, 4, 5, 6, and 7. To simulate impreciseness of motor control, the number of locations moved forwards was also noisy, for example on the fastest trials there is an 80% chance of moving 7 locations forwards, with 15% probability for moving 7 or 9 (± 1) locations forwards, and 5% probability for 10 or 6 (± 2) locations forwards. Each black box zone was randomly padded by 0 to 70 locations. The artificial track was made up of repeated sequences of images, meaning it is not possible to know the true location on the track given a single image. These factors also decouple time from location so that the agent adopts a path integration strategy.

The task given to the agent differed from the task used for experimental recordings in that there was no visual reward zone cue. Training trials were equivalent to 'non-beaconed trials' and included a visual reward indicator in the form of a grey bar that was overlaid at the top of the image input for 5 timesteps if the agent stopped in the reward zone for the first time on a trial. This was designed to be analogous to the immediate feedback the experimental animals receive from a tone and soy milk reward. Probe trials, in which the indicator and reward were absent, were used only for evaluation of the network.

To train the agent to perform the task we used the proximal policy optimization algorithm,[43] which learned an action and value function. Documentation of the algorithm, including pseudocode, is available here: https://spinningup.openai.com/en/latest/algorithms/ppo.html. During training, stopping in the reward zone for the first time on a given trial resulted in a reward of +100, whereas any other stops resulted in a reward of -1. Moving forwards was a neutral action and had a reward of 0. These rewards were totaled for each trial. The network weights were updated using stochastic-gradient descent to maximize the reward.[84] To find the optimal hyperparameters for training (such as learning rate), we performed random parameter searches using the Edinburgh Compute and Data Facility (ECDF),[85] using PyTorch 1.2.0 and cuDNN 10.1.105[86] and adapted an existing Python implementation of the proximal policy optimization algorithm.[87] The hyperparameters used are in Table S1.