Analogue representation of a spatial memory by ramp-like neural activity in retrohippocampal cortex

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

[1]Centre for Discovery Brain Sciences, [2]Simons Initiative for the Developing Brain, [3]Centre for Statistics, University of Edinburgh, Edinburgh, United Kingdom, [4]Current address: Center for Neural Science, New York University, New York, United States. [5]Equal contributing authors.

Abstract

When solving navigational problems, remembering that a critical location is approaching can enable appropriate behavioural choices without waiting for sensory signals. Whereas multiple types of neuron in the hippocampus and retrohippocampal cortex represent locations using codes based on discrete spatial firing fields, analogue neural representations may be advantageous for efficiently recalling proximity to locations of behavioral importance. Here, we identify retrohippocampal neurons that use analogue ramp-like changes in firing rate to represent location as mice navigate a virtual environment in which they have learned the location of a reward. Ramp-like firing patterns had positive or negative slopes and could switch polarity or reset their rate at the reward location. These transitions were maintained when location cues were removed, indicating that path integration is sufficient to drive recall of the track structure expressed through ramp-like codes. We suggest that analogue ramp-like representations encode learned models for goal-directed navigation.

Introduction

Keeping track of a current location relative to remembered locations is important for solving navigation problems. For example, are we approaching the coffee shop, or have we passed it? In solving problems of this kind, neurons with receptive fields that are selective for location, such as place cells and grid cells, play key roles (Hafting et al., 2005; Moser et al., 2017; O'Keefe and Dostrovsky, 1971). These neurons are central to theories of navigation and memory (e.g. (McClelland et al., 1995; O'keefe and Nadel, 1978; O'Reilly and Rudy, 2001; Whittington et al., 2019)). Because they have discrete receptive fields (Figure 1A), their representations can effectively constitute a labelled line code for space, analogous to labelled line codes in sensory systems, with their sequential activity encoding routes, for example to and from the coffee shop. In this framework, memories are formed by associations between labels. However, labelled line codes are relatively inefficient in their use of neurons, and additional parameters useful for navigation, such as whether we are moving closer to or further from the coffee shop, or any other target, are challenging to read out directly for distances

greater than the width of individual firing fields. Theoretical models suggest that continuously graded neural representations of location using ramp-like analogue codes may be particularly useful here (Figure 1B)(Burak and Fiete, 2009; Bush et al., 2015; Foster et al., 2000; Gershman, 2014; Gustafson and Daw, 2011; Stachenfeld et al., 2017; Stemmler et al., 2015). However, the extent to which recall of learned environments engages analogue representational strategies is unclear.

The retrohippocampal cortices, which include the entorhinal cortex, presubiculum and parasubiculum are critical for computations that underlie spatial and contextual memory (Buckmaster et al., 2004; Hales et al., 2014; Jacob et al., 2017; Kitamura et al., 2015; Qin et al., 2018; Sauvage et al., 2010; Steffenach et al., 2005; Suh et al., 2011; Tennant et al., 2018). A constellation of neurons in the retrohippocampal cortices employ discrete firing fields to represent navigational variables (Bicanski and Burgess, 2020; Hafting et al., 2005; Moser et al., 2017; O'Keefe, 1976) and provide input to hippocampal place cells. In addition to grid cells, other retrohippocampal neurons use discrete firing fields to encode the location of environmental borders, visual cues, vectorial relationships to objects and head direction (Campbell et al., 2018; Hafting et al., 2005; Høydal et al., 2019; Kinkhabwala et al., 2020; Lever et al., 2009; O'Keefe, 1976; Sargolini et al., 2006; Solstad et al., 2008; Taube et al., 1990). Discrete representations similar to those found in retrohippocampal cortices also emerge in artificial neural networks trained to predict location and head direction, suggesting that discrete codes are part of a general neural solution to estimation and tracking of location (Banino et al., 2018; Cueva and Wei, 2018; Sorscher et al., 2019). In contrast to these discrete representations of location, encoding of time and movement speed by retrohippocampal neurons is achieved through continuous changes in firing rate (Kropff et al., 2015; Tsao et al., 2018). The abundance in retrohippocampal cortices of neurons with discrete spatial receptive fields also contrasts with reward and decision-related circuits in which ramp-like analogue representations of learned locations relative to rewards have been found (Carelli and Deadwyler, 1994; Engelhard et al., 2019; Fiorillo et al., 2003; Howe et al., 2013; Lavoie and Mizumori, 1994; van der Meer and Redish, 2011; van der Meer et al., 2010; Schultz et al., 1992). The presence in the postrhinal cortex, which is a major input to the retrohippocampal cortex, of neurons that encode continuous representations of distance to the centre of an environment suggests that analogue signals may reach the retrohippocampal cortices (LaChance et al., 2019). However, in conditions in which distance encoding neurons were found in the postrhinal cortex, similar analogue representations were absent from adjacent retrohippocampal regions (LaChance et al., 2019). Nevertheless, it is possible that analogue representations of location have not been observed in retrohippocampal cortices because they are prominent only when they contribute solutions to behavioural problems that engage recall of learned environments.

Here, we investigate representations of location during behavioural tasks that are efficiently solved by recalling the structure of an environment in which behaviourally salient locations are separated by areas lacking unique cues. We recorded from neurons in the MEC and pre/parasubiculum of mice that

were rewarded for stopping on a virtual linear track at a specific location, which we will call the reward zone (Figure 1C)(Tennant et al., 2018). The reward zone is marked by a prominent visual cue but otherwise the track lacks unique spatial cues. With training in this task, mice efficiently obtain rewards by running directly to the reward zone (Figure 1D). This spatial behaviour is maintained when the visual cue indicating the location of the reward zone is removed, indicating that mice can solve the task using a path integration strategy as well as a cue-based strategy (Tennant et al., 2018). Thus, in this task a memory for location on the track is critical to efficiently obtaining rewards, and as this memory is expressed behaviourally, even in the absence of location-specific cues, it can be used to evaluate cue-driven and path integration-dependent representations.

We demonstrate here that neurons in retrohippocampal cortices encode an animal's current position in the task, with respect to behaviourally relevant boundaries, using ramp-like continuous increases or decreases in their firing rate. These analogue representations were primarily driven by output from a path integrator, but in the region of the reward zone their firing rate slope could switch direction and their firing rate could reset. These location-dependent changes were maintained in the absence of reward and location cues, indicating that the structure of ramp-like activity is determined by recall of a memory of the environment. Recurrent neural networks trained to solve tasks with a similar structure also generate ramp-like representations suggesting that analogue codes may provide a general mechanism for estimation of location in learned environments. Together, our results demonstrate that the retrohippocampal cortices encode learned models of the world which serve as neural substrates for tracking current location relative to behaviourally salient environmental features.

Results

We trained mice to stop at the reward zone on the virtual track (Figure 1C-D and Tennant et al., 2018). During training, 'beaconed trials', in which a visible cue indicated the location of the reward zone, were interleaved with 'non-beaconed trials', in which the cue was hidden but a reward was still available. When mice reached a performance criterion of more than 75% correct trials on two consecutive days (see Methods, Supplemental Figure 1), alternate non-beaconed trials were replaced with 'probe trials' in which reward was not available and the cue was again hidden (see also Tennant et al., 2018). We recorded neural activity using tetrodes that were targeted to the retrohippocampal cortices (Figure 1E and Supplemental Figure 2), and were typically advanced by approximately 100 µm after each recording session. Following automated spike sorting we obtained 649 isolated single units recorded during sessions that took place after the training criterion was passed and that contained > 30 correct trials (105 sessions, 7 mice).

Continuous ramp-like changes in firing rate encode location

We focus our initial analysis on beaconed trials in which animals stopped in the reward zone (see Methods). During these trials, we observed many neurons with spiking activity that continuously increased or decreased as a function of position on the track (Figure 1F). This contrasts with discrete firing fields that are typical of place and grid cell fields recorded previously on linear virtual tracks (Chen et al., 2013; Domnisoru et al., 2013; Harvey et al., 2009).

To more systematically evaluate these putative ramp-like representations, we further restrict our initial analysis to neural activity recorded when mice moved through the region of the track between the start box and the reward zone. We refer to this region as the 'outbound' zone (Figure 2A). We identified neurons with firing rates that varied significantly (linear model fit, $p < 0.01$ after correction for multiple comparisons, see Methods) as a function of position within the outbound zone (Figure 2A-B). With this analysis, we classified 23.0 % of neurons (n = 149/649) as having positive slopes and 33.6 % of neurons (n = 218/649) as having negative slopes (Figure 2B-C, F). The remaining neurons were categorised as unclassified (n = 282/649). The proportion of neurons with positive (25.2 ± 2.8%, range: 18.5 - 35.6 %) or negative slopes (34.8 ± 4.6, range: 22.2 - 51.0 %) was similar across all animals. By contrast, in shuffled datasets 99.4% (n = 644106 / 648351) of neurons were categorised as unclassified (Figure 2B). As a second approach to identify neurons with ramp-like activity we calculated for each neuron a 'ramp score' (Figure 2D); this score is based on the correlation between firing rate and position and is equal to +/- 1 for firing that varies linearly with positive/negative slopes, and peaks and troughs at the start and end of the outbound zone (Supplemental Figure 3A). Ramps scores were largely consistent with the classification based on linear fits of neuronal firing (Figure 2F and Supplemental Figure 3B). For neurons classified as having positive or negative slopes, the absolute ramp scores (+ve slopes: 0.39 ± 0.22; -ve slopes: -0.42 ± 0.20) were greater than for unclassified neurons (-0.03 ± 0.31; vs +ve slopes: $p = 2.22 \times 10^{-16}$, $t = 14.93$; vs -ve slopes: $p = 2.22 \times 10^{-16}$, $t = -16.2$; Welch Two sample t-test). Moreover, the maxima and minima in the firing rate of neurons identified as having negative slopes, were clustered towards the start and end of the outbound zone respectively, and vice-versa for neurons with positive slopes (Figure 2E). In contrast, for shuffled data and for unclassified neurons the maxima and minima were evenly distributed across the track (Supplemental Figure 3C). Thus, these data indicate that ramp-like changes in firing rate are a common feature of neural activity in retrohippocampal cortices during the linear location estimation task.

To evaluate whether ramp-like firing depends on the location of the recorded neurons within retrohippocampal cortex, we classified neurons based on the location of their recording tetrodes as being in either the pre- or parasubiculum, the MEC, or where tetrode localisation was ambiguous as simply being in the retrohippocampal area (see Methods, Supplemental Figure 2). Theta frequency field potentials and associated phase locking of spike firing (Supplemental Figure 4A-C) were consistent with previous reports for these regions (Boccara et al., 2010; Chrobak and Buzsáki, 1998;

Colgin et al., 2009; Kornienko et al., 2018)(Supplemental Figure 4D-E). We found neurons with ramp-like activity in all areas and in general the properties of ramp-like changes in firing rate were similar (Figure 2F and Supplemental Figure 5). Thus, ramp-like analogue representations appear to be a general feature of neural activity in retrohippocampal cortex during the location estimation task.

Because speed and acceleration, as well as position, vary during the task, we asked whether these variables alone or in combination with position were better able to predict each neuron's firing rate. When we fit mixed models, that included position, speed and acceleration as fixed effects and trial number as a random effect, to the firing rate of neurons classified as having positive or negative slopes (Figure 3A), the resulting standardized coefficients for position were on average greater than for speed or acceleration (Figure 3B). By assessing the significance of each predictor variable (Wald test, $p < 0.01$), we further classified the ramping neurons according to whether their firing rate was best predicted by each variable alone or in combination (Figure 3C). A majority of neurons were classified as representing either position alone (26/149 +ve slope neurons, 35/218 -ve slope neurons)(Figure 3D-E), position in combination with speed (37/149 +ve slope neurons, 65/218 -ve slope neurons)(Figure 3F-G), or position in combination with speed and acceleration (31/149 +ve slope neurons, 39/218 -ve slope neurons)(Supplemental Figure 6A,B). In the latter cell populations, the standardised coefficients for position were on average greater than for speed or acceleration (Figure 3F, Supplemental Figure 6A,B). Ramping neurons classified with firing rate solely modulated by acceleration (8/149 +ve slope neurons, 9/218 -ve slope neurons) or speed (12/149 +ve slope neurons, 18/218 -ve slope neurons) were relatively rare (Figure 3C and Supplemental Figure 6C, D). However, neurons modulated solely by speed were more common among neurons that did not have a positive or negative slope (88/282 unclassified neurons), indicating that encoding of speed is in general not able to account for ramp-like activity (Figure 3C). We also asked if ramp-like activity could be explained by dependence of neuronal firing on time rather than position (cf. (Tsao et al., 2018)). We found that for neurons with firing on the outbound zone classified in either positive or negative slope groups, the relationship between firing rate and time was steeper for short versus long duration trials, and the firing rate was usually better fit with models in which location rather than time was the independent variable (Supplemental Figure 7). Thus, ramp-like neural activity primarily reflects encoding of location on the track, and in some neurons may be conjunctive with speed or acceleration.

Together, these analyses indicate that during a learned behaviour in which location is critical for receiving rewards, a substantial proportion of neurons in the retrohippocampal cortices generate ramp-like representations of an animal's position.

Path integration drives ramp-like activity

What signals drive ramp-like activity? Ramp-like encoding of location could be driven, either directly or indirectly, by the output of a neural path-integrator. Alternatively, it could be a response to the upcoming visual cue indicating the reward zone. To distinguish these possibilities, we compared neuronal firing on beaconed trials with firing on non-beaconed and probe trials (Figure 4A). On the latter trials, because the visual cue indicating the reward zone was absent, ramp-like activity driven by the visual cue should be abolished, whereas activity driven by a path integrator should be unaffected. We restricted our initial analyses to neurons that on beaconed trials demonstrated position encoding ramp-like activity, which we identify on the basis of their classification as having positive or negative firing rate slopes (Figure 2) and the dependence of their activity on position alone or in combination with speed and/or acceleration (Figure 3).

We found that for most position encoding ramping neurons the slope on the non-beaconed and probe trials had the same classification as on beaconed trials (n = 177/257)(Figure 4B-E). Following fitting of mixed effects models with position, speed and acceleration as independent variables, standardised coefficients for position were consistently larger than for speed or acceleration for these neurons (ANOVA $p = 1.36 \times 10^{-8}$, f=18.55; P vs A, adj $p = 4.42 \times 10^{-6}$; P vs S, adj $p = 7.63 \times 10^{-8}$; Tukey's test)(Figure 4C). These data are consistent with a path integration mechanism driving the ramp-like activation of these neurons. A smaller proportion of the position encoding ramp neurons were classified differently on non-beaconed and probe compared with beaconed trials (n = 61/257). The difference in the activity of these neurons was localised to the region of the outbound track zone in which the reward zone was visible to the mouse (Figure 4F - G). Thus, these data indicate that for most neurons showing ramp-like positional activity a path integration mechanism is sufficient to drive changes in firing rate across the outbound track zone, while for a minority of neurons the visual reward zone cue drives ramp-like activity.

Ramp-like activity is reset by the reward zone cue

Classic neural integrator models generate ramp-like changes in firing through integration of a continuous input (Brody et al., 2003; Carpenter, 2011; Major and Tank, 2004). This can enable location signals to be computed by integrating speed inputs. In this framework, landmarks might reset ramp-like activity, but for a given neuron the slope of the ramp as a function of position should be stable. Alternatively, ramp-like activity might emerge downstream from grid cell networks that perform path integration (e.g. (Bush et al., 2015)), or as trajectories in the state space of a recurrently connected neural network, with entrance to a specific trajectory depending on local landmarks (e.g. (Liu and Buonomano, 2009; Sussillo and Abbott, 2009; Sussillo et al., 2015)). In these scenarios, initialisation and resetting of activity by landmarks could modify the slope of ramp-like activity as well as resetting the firing rate.

To distinguish these possibilities, we compared ramp-like activity between the outbound and homebound regions of the track (Figure 5A). We focus on activity during beaconed trials of the neurons with ramp-like positional firing driven by a path integration-like mechanism (n = 177/649 neurons). We compared fits of firing rate as a function of position for the outbound and homebound track zones (Figure 5B). In contrast to classic integrator models, we find many neurons for which the ramp-like activity changes direction between outbound and homebound zones of the track (n = 50/177)(Figure 5C-E). We refer to these neurons as +- and -+ , using + and - to indicate the slope on the outbound and homebound regions of the track respectively. For firing rates averaged across these populations, the change in slope was apparent in the region between the outbound and homebound zones (Figure 5E). We also found neurons with ramping activity present only on the outbound region of the track (51/177)(Supplemental Figure 8A, D). Moreover, even among neurons for which the direction of ramp-like activity was similar between outbound and homebound zones, the slope of the ramping activity could nevertheless differ substantially (Figure 5C). These observations argue against neural integrator models in which firing rate is a fixed function of a velocity encoding input.

Among neurons for which the direction of ramp-like activity remained the same on the homebound and outbound track regions (++ and -- groups), their firing rate could nevertheless show abrupt offsets between the end of the outbound and start of the homebound regions of the track. These offsets were clearly apparent in the activity of individual neurons (Figures 1F and 5F) and in the averaged activity of ++ and -- neurons (Figure 5G). To quantify this further, we compared activity of the neurons in the first 5 cm of the homebound region, with predicted activity at the same location extrapolated from the fits of their firing in the outbound region (Figure 5H). We find that for the majority of ++ and -- neurons their activity in this region was outside of the 99% confidence interval of the prediction (n = 57/76) and for almost all neurons the resulting offset was towards the firing rate at the start of the outbound zone, indicating that the firing tends to reset at the start of the homebound zone (Figure 5I).

In summary, these data show that ramp-like activity is present on outbound and homebound segments of the track, but the properties of ramps may differ substantially between these regions. Neurons can switch polarity (+- and -+ neurons), or have similar polarity but tend to offset their firing rate at the transition between outbound and homebound regions of the track (++ and --neurons). These observations are not predicted by standard neural integrator models that perform path integration when presented with velocity inputs, but may be consistent with activity patterns generated downstream from grid cell networks that perform path integration or by more complex recurrent neural network models.

Visual cues are not required to reset ramp-like activity

Transitions in slope direction and offsets in the firing rate at the reward zone could be triggered by sensory input from the visual cue or by receipt of the reward. Alternatively, they could reflect read out of internally generated estimates of location from a learned model of the track. To distinguish these possibilities we evaluated neural activity on probe trials, in which the visual cue and reward are absent (Figure 6A). We focussed again on neurons with path integration-driven, position-dependent activity.

Strikingly, we found that for many neurons switching of ramp direction and offsets in firing rate at the reward zone were maintained on probe trials, indicating that these location-dependent changes in activity pattern did not require the visual cue or the reward (Figure 6). The slope of firing rate as a function of position for the homebound track zone was similar for beaconed and probe trials ($r^2 = 0.46$, $p = 1.21 \times 10^{-23}$)(Figure 6B) and 78/177 neurons had firing rate profiles with a similar slope classification on beaconed and probe trials (Figure 6C). Among neurons that on beaconed trials were classified as having slope that switched direction between the outbound and homebound track zones, 7/25 +- neurons and 14/25 -+ neurons had the same slope classification (Figure 6C) and similar firing rate profiles on probe trials (Figure 6D-E). For neurons no longer classified as switching slope, this was because on probe trials the slope on the homebound region of the track no longer reached the threshold for significance. Among neurons classified on beaconed trials as similar slopes in outbound and homebound track regions, 10/26 ++ neurons and 20/50 -- neurons had the same classification on probe trials (Figure 6C). These neurons had similar firing rate profiles on beaconed and probe trials (Figure 6F, G) and again showed a strong bias towards resetting their activity during the transition between the outbound and homebound regions of the track (Figure 6H, I). For neurons no longer classified as having the same outbound and homebound slopes, this was again because on probe trials the slope on the homebound region of the track no longer reached the threshold for significance after correcting for multiple comparisons (Supplemental Figure 8C, F). Together, these data indicate that for many neurons with ramp-like activity that is driven by path integration, the visual cue and reward are not required for offsets in firing rate and changes in the slope direction that take place between the outbound and homebound regions of the track, indicating that ramp-like activity of these neurons reflects readout of an internally generated learned model of the track.

Dependence of ramp-like firing on behavioural performance

To what extent is ramp-like activity associated with task success? If ramp-like activity provides 'always-on' positional representations then it should be maintained on trials when mice do not stop in the reward zone. On the other hand, if ramp-like activity serves as a decision signal for stopping, or if it relies on attentional signals, then it may be absent on trials in which mice run through the reward zone without stopping. To distinguish these possibilities we took advantage of beaconed trials on which the mice failed to stop in the reward zone (Figure 7A-B). We grouped trial outcomes into those in which the mice successfully stopped within the reward zone, which we refer to as a 'hit trials', those in which mice slow down but do not stop in the reward zone, which we refer to as 'try trials', and those

in which mice maintain a high running speed throughout the reward zone, which we refer to as 'run trials'. We reasoned that separation of try and run trials should dissociate effects of running speed as well as stopping-related versus positional codes.

We find that for many positional ramping neurons, offsets in their firing rate and transitions in ramp direction at the reward zone were maintained on miss trials. Among neurons with path integration-driven position firing for which ramping activity switched direction between the outbound and homebound regions of the track (-+ and +- groups), the mean activity profiles were similar on hit, run and try trials (Figure 7C). For neurons with path integration-driven ramp-like activity that had a similar direction on the outbound and homebound track regions (++ and -- groups), the mean activity profiles were similar on hit and try trials, but on average showed less prominent activity resetting in the region of the reward zone during run trials (Figure 7C). However, when we focussed only on neurons with ramping activity that was correlated with position and not also conjunctively modulated by speed or acceleration (group P in Figure 3), the difference in firing profile on run trials was no longer apparent (Figure 7D), suggesting that in the larger population of neurons (groups P, PA, PS and PSA from Figure 3) it reflected modulation of activity by speed or acceleration. Consistent with this, neurons with ramping activity that was correlated with running speed but not position (groups S and SA in Figure 3), had firing rate profiles that were similar on hit and try trials but differed on run trials (Figure 7E). Because activity on run trials was not included in our initial assessment of these neurons as speed modulated, this difference corroborates their initial classification. When we quantified the effect of trial outcome on offsets in the activity of neuronal firing, we found modest effects of trial outcome when considering all neurons with path integration driven and position modulated activity (P, PS, PA and PSA groups)(Figure 7F), no effect on neurons with activity modulated by position but not acceleration or speed (P group)(Figure 7G) and substantial effects on offsets for neurons with activity modulated by running speed (S and SA groups)(Figure 7H).

Together, these analyses demonstrate that changes in the ramp slope or offset in ramp firing represent transitions between the outbound and homebound part of the track independently of trial outcome. This is consistent with positional ramp-like firing contributing to 'always-on' representations of location. These analyses also further validate our classification of ramp-like activity as primarily driven by position, rather than speed or acceleration, as the latter variables differed on missed trials but ramp-like activity and its transitions between track zones was maintained.

Recurrent neural networks learn to estimate location using ramp-like representations

If ramp-like representations reflect a general neural solution to distance estimation then they may emerge in artificial neural networks trained to solve similar location estimation tasks. To investigate this possibility, we trained a network model that received visual input corresponding to the virtual track and that generated stop or go signals to control movement along the track. The model was organised

such that visual signals were processed by a three layer convolutional neural network with the output fed into a recurrent neural network that generated movement control signals (Figure 8A). Networks trained with a reinforcement learning algorithm (Schulman et al., 2017) successfully learned the task (Figure 8B). Before training, the agent controlled by the network stopped frequently on the track with no particular spatial stopping pattern (Figure 8C). After training, the agent stopped at the reward zone on non-beaconed (Figure 8D) and probe trials (Figure 8E) indicating that the network solves the task using a distance-based strategy. Experimental mice exhibit similar stopping patterns during learning (Tennant et al., 2018).

When we analysed activity of the units in the network in the same way as for the experimental data, we found units with ramp-like activity in the outbound and homebound regions of the track (Figure 8F-J). Just as for the experimental data, the ramp-like activity could have either a positive or a negative slope Figure 8F, J), was clearly distinguishable from activity patterns generated by shuffled datasets (Figure 8G-H) and could show changes in slope and resets in firing rate at the reward zone (Figure 8F, I). Similar to the experimental data (e.g. Figure 2C), ramps generated by the trained recurrent network more frequently had a negative rather than a positive slope, although in contrast to the experimental data negative slopes were typically steeper than positive slopes (Figure 8I). Thus, these data show that recurrent neural networks trained on a distance estimation task generate ramp-like activity with properties similar to ramp-like firing of neurons in the retrohippocampal cortex.

Discussion

We show that retrohippocampal neurons generate structured analogue representations of a learned environment using ramp-like changes in firing rate that are bounded by salient locations. These ramp-like changes in firing rate appear to be driven by a neural path integrator rather than by location-specific cues. This coding scheme is complementary to but differs fundamentally from the discrete firing fields of place and grid cells that are critical components of current theories for spatial cognition, memory and abstract conceptual thinking (Behrens et al., 2018; Bellmund et al., 2018; O'keefe and Nadel, 1978). Ramp-like codes may be well suited to represent internal models for control of learned goal-directed behaviours (e.g. (Foster et al., 2000; Gershman, 2014)).

Analogue representation of a spatial memory

A striking feature of ramp-like firing is that it is consistent with readout of a learned model in which the environment is broken into compartments bounded by salient cues. In this framework, offsets in the frequency and changes in the slope of ramping activity mark transitions between compartments (Figure 5), while position within a compartment is represented by absolute firing rate (Figures 2-3). Because changes in slope and firing rate offsets were maintained in the absence of rewards and after removal of cues that indicate the reward zone (Figure 6), they are unlikely to be a direct response to

external location cues, but instead are consistent with recall from memory of the structure of the environment. This separation of the environment into compartments differs from the global metric provided by grid cells (Hafting et al., 2005; McNaughton et al., 2006), but is reminiscent of robotic control strategies that segment the environment (Kuipers and Kuipers, 1999). As the slope and direction of ramp-like firing differs between compartments, representations of position may be specific to individual compartments, potentially enabling ramp-like codes to combinatorially represent location and local context, for example whether the current location is in the homebound or outbound part of the track.

Analogue representation of location through ramp-like changes in firing rate contrasts with well studied firing patterns of neurons in the hippocampal formation such as place and grid cells. The ramp-like changes in firing rate of retrohippocampall neurons is more similar to ramp-like spatial firing of neurons in the ventral striatum but differs in several key respects (cf.(Carelli and Deadwyler, 1994; Fiorillo et al., 2003; Howe et al., 2013; Lavoie and Mizumori, 1994; van der Meer and Redish, 2011; van der Meer et al., 2010; Schultz et al., 1992)). First, activity of neurons in retrohippocampal cortices either ramps up or down on approach to the rewarded area, whereas in the striatum ramps generally involve an increase in firing rate. Second, ramps in the retrohippocampal cortices were observed after as well as before the reward zone, whereas ramping activity of striatal neurons appears to be primarily in anticipation of a rewarded location. Third, ramping activity of retrohippocampal neurons was similar on probe trials when there was no reward obtainable and on beaconed trials when a reward was obtained, whereas striatal neurons show elevated activity after a reward is received. The ramp-like activity of retrohippocampal neurons also has some similarities to centre-distance activity of neurons recorded in the postrhinal cortex during foraging behaviours in an open arena (LaChance et al., 2019). However, whereas +/- and -/+ cells that we describe are reminiscent of centre distance firing patterns, activity that resembles +/+ and -/- patterns, in particular their offsets in firing rate in the region of the reward zone, has not been reported for neurons in postrhinal cortex. A further key difference is that under conditions in which centre-distance cells were recorded in the postrhinal cortex, similar activity patterns were absent from the retrohippocampal cortices (LaChance et al., 2019). This may reflect differences in the structure of the environment and in the task contingencies, which in the recordings we report here involved a spatial learning task rather than unconstrained foraging. The distinct features of ramp-like firing of retrohippocampal neurons may be particularly well suited to contribute representations of the structure of the learned environment and location relative to salient landmarks.

Mechanisms for ramp-like positional representation

How are ramp-like representations of location generated? Classic neural integrator models based on either recurrent network dynamics or intrinsically generated neural activity generate ramp-like

changes in firing rate in response to a continuous input (Fransén et al., 2006; Major and Tank, 2004). Consistent with these models retrohippocampal circuits also contain abundant recurrent connections and neurons that fire persistently in response to brief input (Dhillon and Jones, 2000; Egorov et al., 2002; Peng et al., 2017; Yoshida and Hasselmo, 2009), and receive speed signals that could be used to drive path integration (Carvalho et al., 2020; Justus et al., 2016). Nevertheless, these models do not easily account for changes in the slope of ramp-like activity between outbound and homebound regions of the track (Figures 5 and 6), although it is possible that the integrator mechanism, or its speed inputs, are reconfigured on each part of the track. Alternatively, movement through each compartment of the track may be encoded through a distinct trajectory in the state space of recurrently connected networks. This scenario differs from the classic neural integrators in that representations for each location would be learned from experience (Liu and Buonomano, 2009; Sussillo and Abbott, 2009; Sussillo et al., 2015). Our simulations of abstracted recurrent networks trained to solve the location estimation task appear consistent with such a mechanism. While retrohippocampal circuits contain recurrent connectivity that may support computations of this kind (Dhillon and Jones, 2000; Peng et al., 2017; Simonnet and Fricker, 2018), they could also take place in upstream structures (LaChance et al., 2019). A further possibility is that ramp-like activity is generated by integration of inputs from other retrohippocampal neurons that encode location through discrete firing fields (Burak and Fiete, 2009; Bush et al., 2015; Stachenfeld et al., 2017; Stemmler et al., 2015). Such representations are abundant locally within the MEC and parasubiculum, and include the activity of grid cells, which may mediate path integration (Gil et al., 2017; McNaughton et al., 2006; Tennant et al., 2018), and neurons with activity driven by visual inputs (Casali et al., 2018; Kinkhabwala et al., 2020; Pérez-Escobar et al., 2016).

What is the relationship between ramp-like activity, visual inputs and memory? Sharp changes in firing between the black box and the start and end of the track (e.g. Figure 1) are consistent with strong visual control of the activity of retrohippocampal neurons (e.g. (Casali et al., 2018; Chen et al., 2016; Kinkhabwala et al., 2020; Pérez-Escobar et al., 2016)). Thus, it is likely that visual signals are critical for registering ramp-like activity to the track. In contrast, while changes in ramp slope and offsets in firing rate occurring between the outbound and homebound compartments of the track must be learned from their sensory cues, which could be visual or reward-related, the maintenance of transitions on probe trials indicates that they can be recalled from memory. In our simulations, recall of the transition location is triggered by ramping activity reaching a threshold that then moves the network into a new state encoding the subsequent track compartments (Figure 8). In these simulations the motion signals derived from visual input are sufficient to update ramping activity. Alternatively, if ramping activity is generated by integration of signals from upstream neurons with discrete firing fields (cf. (Burak and Fiete, 2009; Bush et al., 2015; Stachenfeld et al., 2017; Stemmler et al., 2015)), then recall of the transitions triggered following sequential activation of grid cells or

hippocampal place cells could also be maintained in the absence of location-specific cues (Fischler et al., 2019).

Implications for learned behaviours

How might ramp-like representations in retrohippocampal cortices contribute to spatial behaviours? 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, the correct stopping location is read out from the identity of the activated place cells (Bush et al., 2015; Fiete et al., 2008). However, readout of metric representations of distance to the stopping position from discrete representations is challenging without continuous intermediate representations (Bush et al., 2015; Fiete et al., 2008; Stemmler et al., 2015). The ramp-like firing properties of retrohippocampal neurons that we describe here are consistent with predictions about the form these intermediate representations should take (Fiete et al., 2008; Stemmler et al., 2015). Thus, using ramp cell activity, distance to a remembered location of interest can be read out from changes in firing rate, while reconfiguration of ramp-like activity at salient landmarks would provide information about the local context. An important future question will be to establish the extent to which ramping activity requires input from place or grid cells, or whether processed sensory input is sufficient for ramping activity. In either scenario, the structure of the environment could be stored in the weights of synaptic inputs from these neurons. Given that deep layers of entorhinal cortex have extensive projections throughout the telencephalon (Insausti et al., 1997; Sürmeli et al., 2015; Swanson and Köhler, 1986), ramping activity could then 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 path integration-like mechanism drives analogue changes in firing rate that are configured to represent position relative to salient landmarks and that may be suitable for control of behaviours by downstream brain structures.

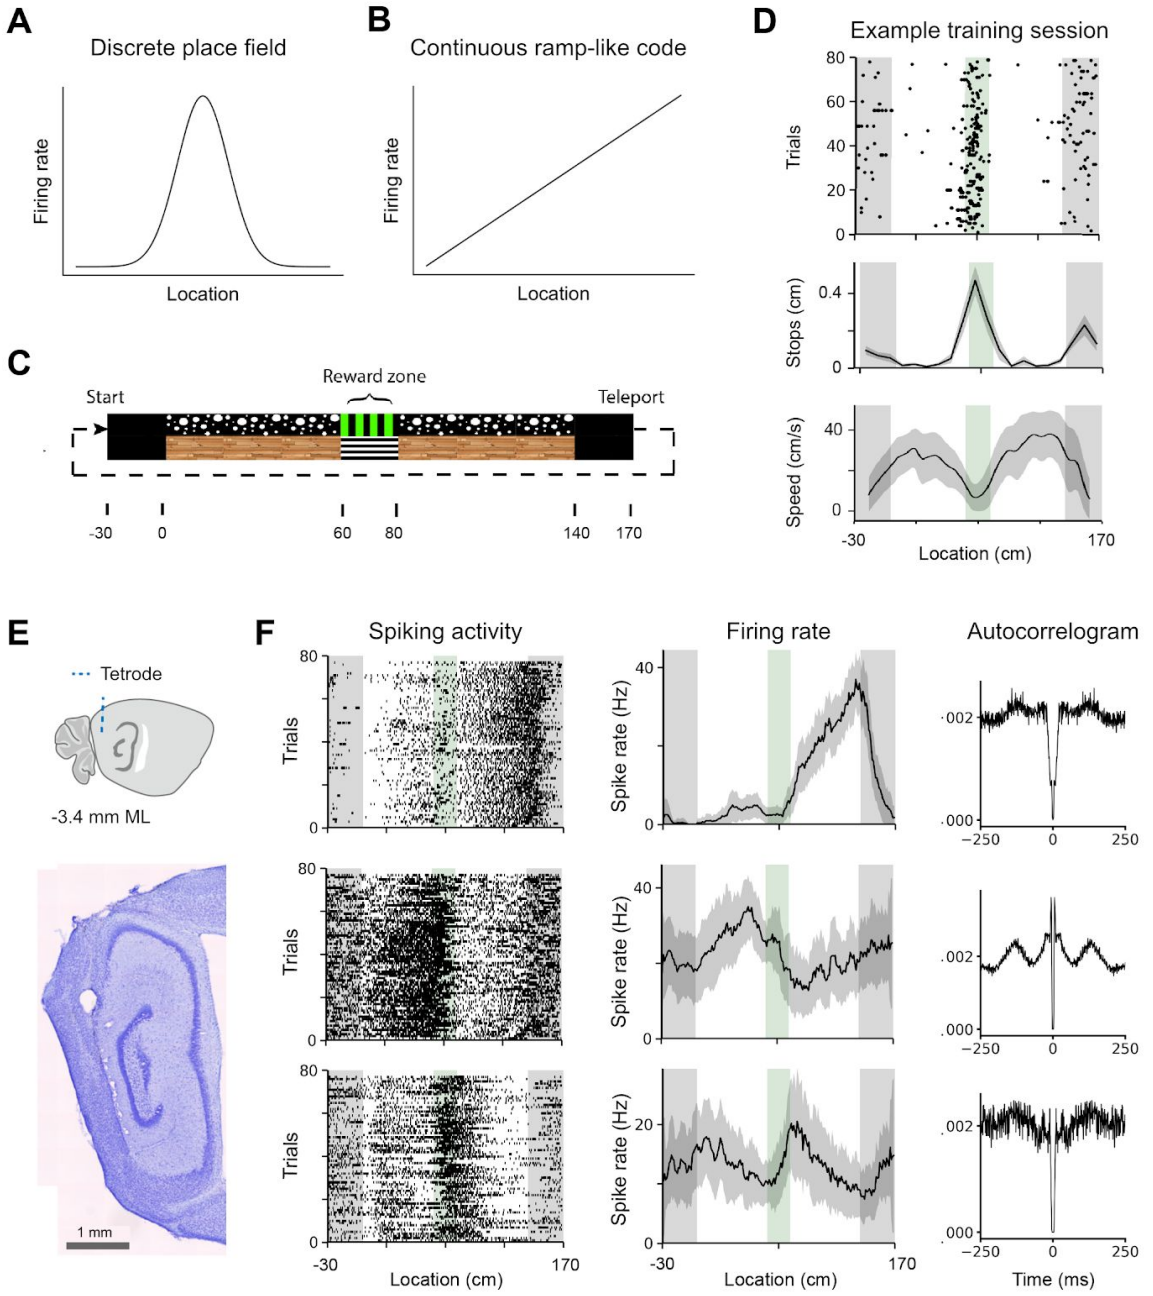

Figure 1. Ramp-like neuronal activity during recall of a learned location

(A, B) Schematic 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 behaviour 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. Standard error of the mean is in grey. The vertical green bar indicates the location of the reward zone.

(E) Tetrodes were targeted towards the dorsal retrohippocampal region (upper). In this experiment the tetrodes were classified as terminating in the retrohippocampal area (lower). Additional sections are shown in Supplemental Figure 2.

(F) Example rasters of spike location (left column), the corresponding mean firing rate (middle column) and spike autocorrelograms (right column) from 3 neurons recorded simultaneously during beaconed trials. Standard error of the firing rate is in grey. The neurons are recorded from the animal in (E) on the training day shown in (D).

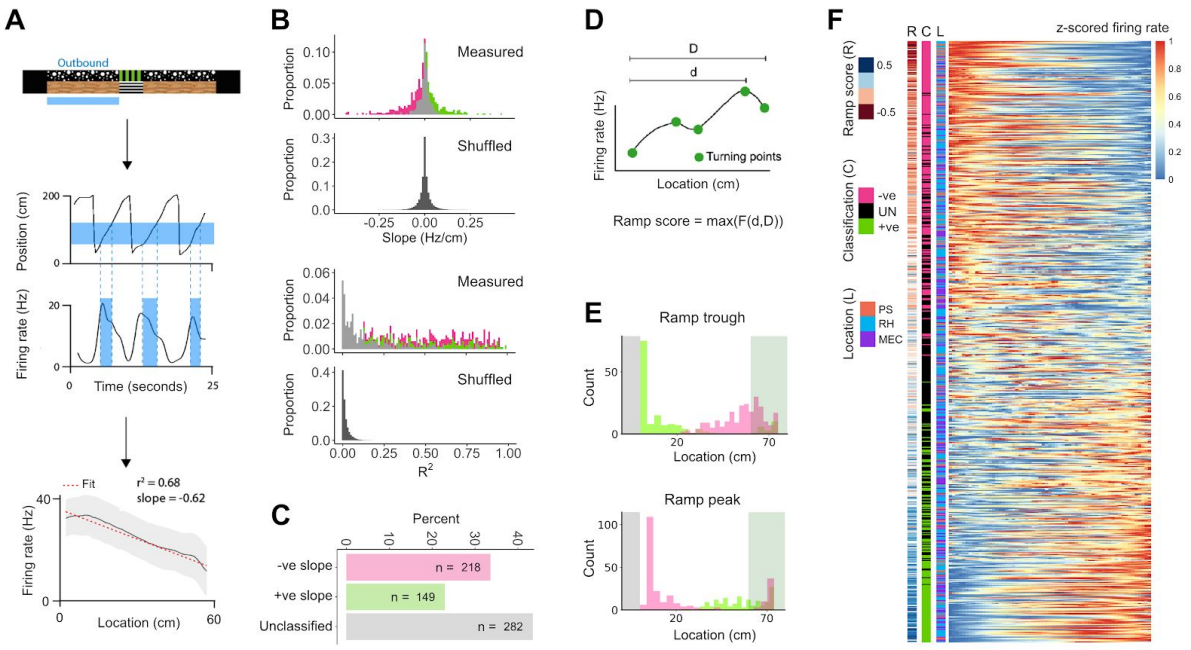

Figure 2. Quantification of ramp-like activity

(A) Procedure for fitting firing rate as a function of location on the outbound region of the track (0-60 cm, blue bar).

(B) The distribution of slope (upper) and r^2 (lower) of the model fit for each neuron. Neurons are colour-coded as having a positive slope (green), negative slope (pink) or unclassified (grey). Distribution of r^2 of the model fit for each shuffled dataset (n=627372) is plotted in black below the corresponding experimental data.

(C) Proportions of neurons classified in each group.

(D) Calculation of ramp score. D is the length of the track segment, d is the length of a section of the track bounded by extrema (i.e. a peak or a trough) in the relationship between firing rate and location. For all combinations of peaks and troughs, correlations between the smoothed firing rate and position are calculated and normalised to the segment length (see Methods). The ramp score for a particular neuron is then taken as the maximal normalised correlation score for that neuron.

(E) Location of peaks and troughs in firing rate for sections used to calculate the ramp score for each ramp-like neuron. Colours are as for (B) and (C).

(F) Normalised firing rates in the outbound region of the track plotted for all neurons as a function of location. Neurons are ordered along the y axis ordered according to the slope of their firing rate. Columns to the left indicate the ramp score (R), the classification of activity based on fit of the linear model (C) and the location within the retrohippocampal region (L).

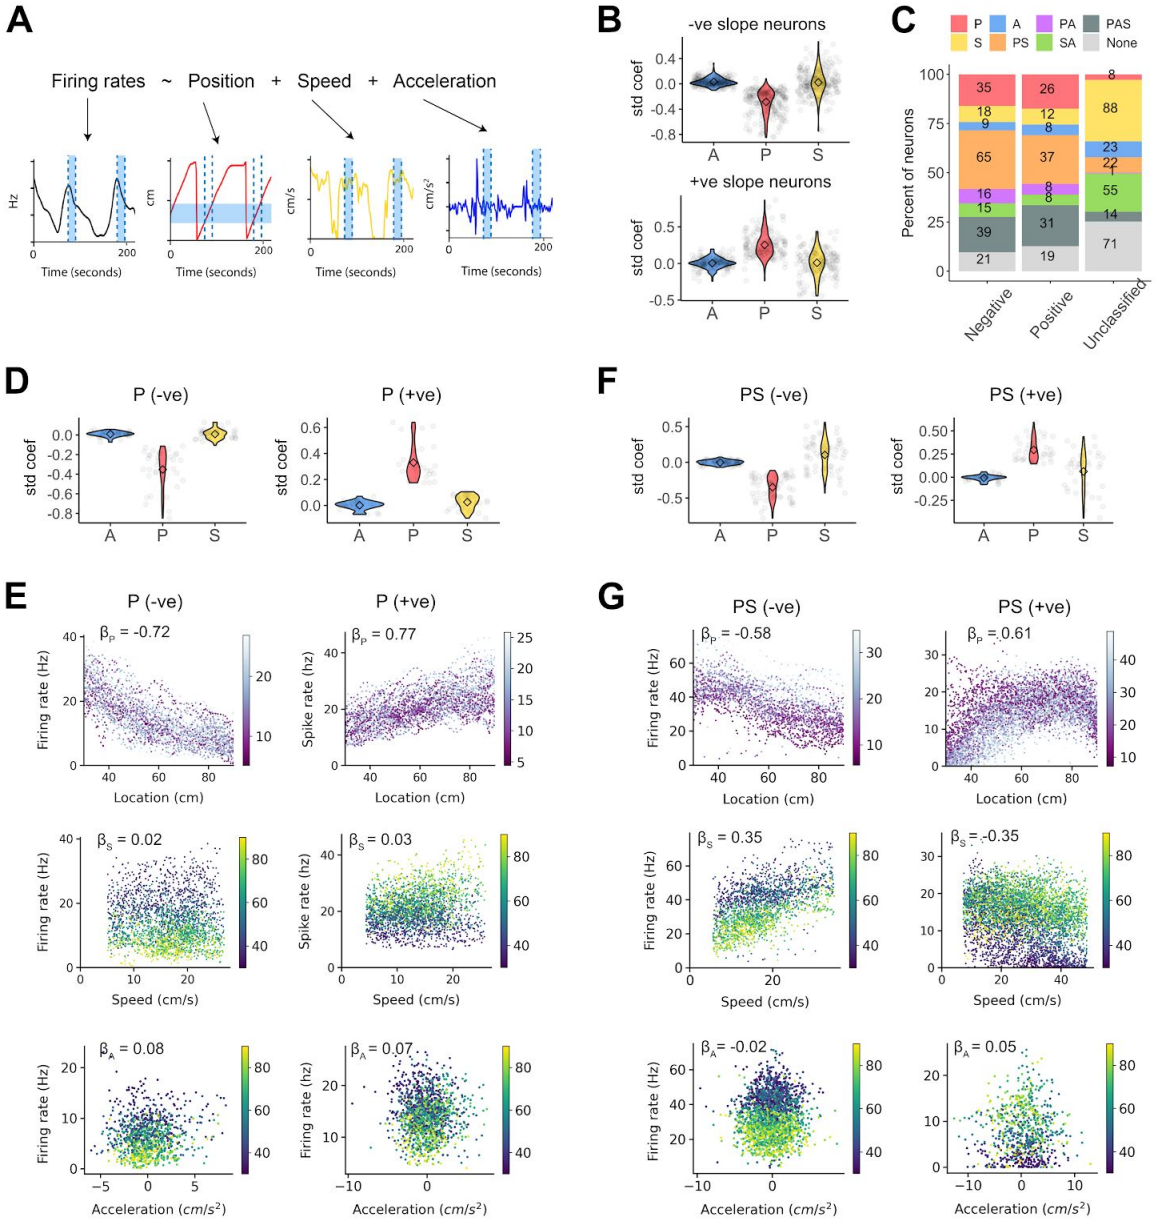

Figure 3. Differential influence of position, speed and acceleration on ramp-like firing

(A) Models that include speed (S), acceleration (A) and position (P) as fixed effects were fit to the firing rate in the outbound track region. Trial number was included as a random effect.

(B) Standardised coefficients obtained from the model fits for negative (upper) and positive (lower) slope neurons. Coefficients for speed, acceleration and position differed for neurons with negative and positive slopes (+ve slope: $p < 2 \times 10^{-16}$, f = 133; -ve slope: $p < 2 \times 10^{-16}$, f = 257.5; ANOVA). For neurons in negative and positive slope groups the position coefficients were larger than those for acceleration (+ve slope: adj p = $p < 1 \times 10^{-7}$; negative slope: adj $p < 1 \times 10^{-7}$; Tukey's test) and speed (+ve slope: adj p = $p < 1 \times 10^{-7}$; negative slope: adj $p < 1 \times 10^{-7}$ Tukey's 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 as unclassified.

(D, F) Distribution of standardised coefficients for neurons classified as having negative slopes (left) or positive slopes (right) and encoding either position alone (D) or a combination of position and speed (F).

(E, G) Examples of firing rate as a function of position (upper), running speed (middle) and acceleration (lower) for neurons corresponding to the groups in (D and F).

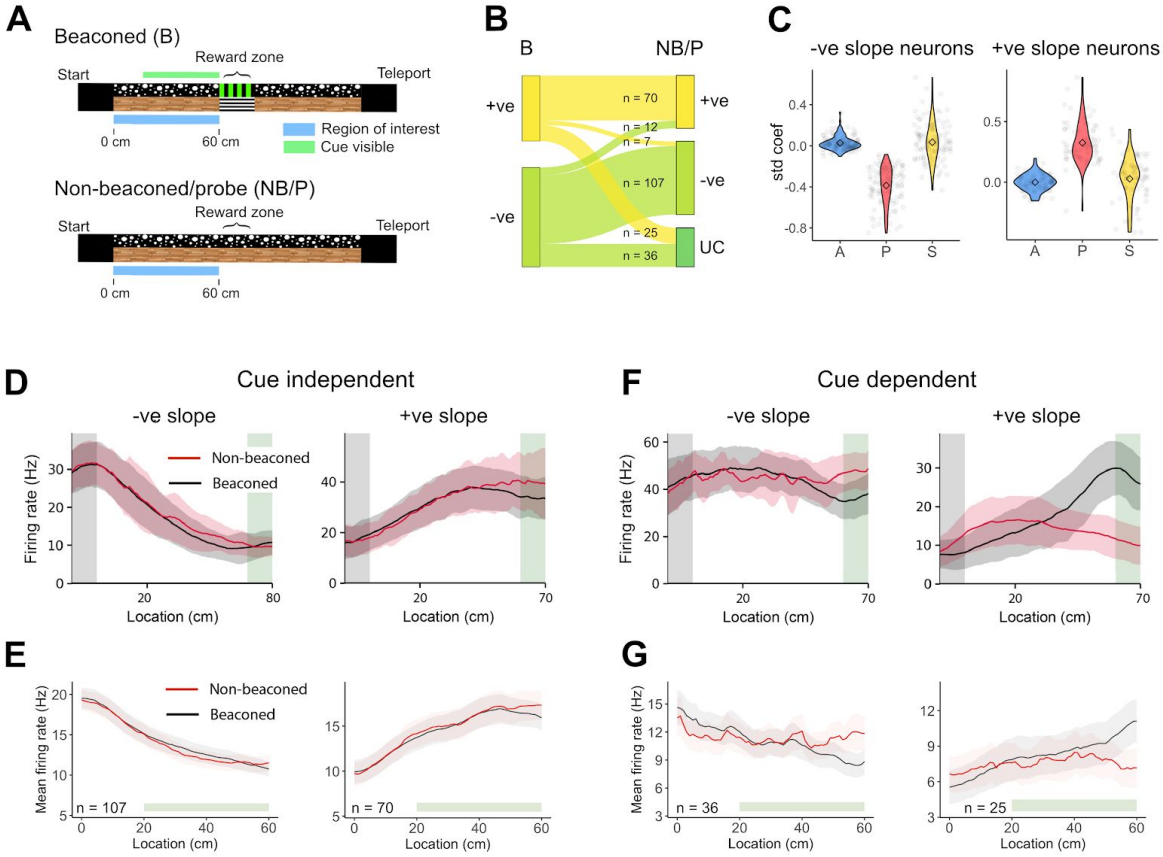

Figure 4. Ramp-like activity does not require visible cues that mark the reward zone

(A) Schematic of the track for beaconed and non-beaconed/probe trials. The green bar indicates the region of the track on which the upcoming reward zone cue is visible.

(B) Classification of the slope of neurons on beaconed (left) and non-beaconed/probe trials (right).

(C) Standardised coefficients on non-beaconed/probe trials calculated as in Figure 3 for neurons that were classified as having negative (left) or positive (right) slopes on beaconed and non-beaconed/probe trials.

(D, F) Mean firing rate as a function of position for example neurons with classification that was independent from (D) or that was modified by the presence of the reward zone cue (F).

(E, G) Population average firing rate as a function of location for neurons with slope classification that was independent from (E) or that was modified by the presence of the reward zone cue (G), and grouped according to whether the slope on beaconed trials was negative (left) or positive (right). In (D-F) the SEM is indicated in grey.

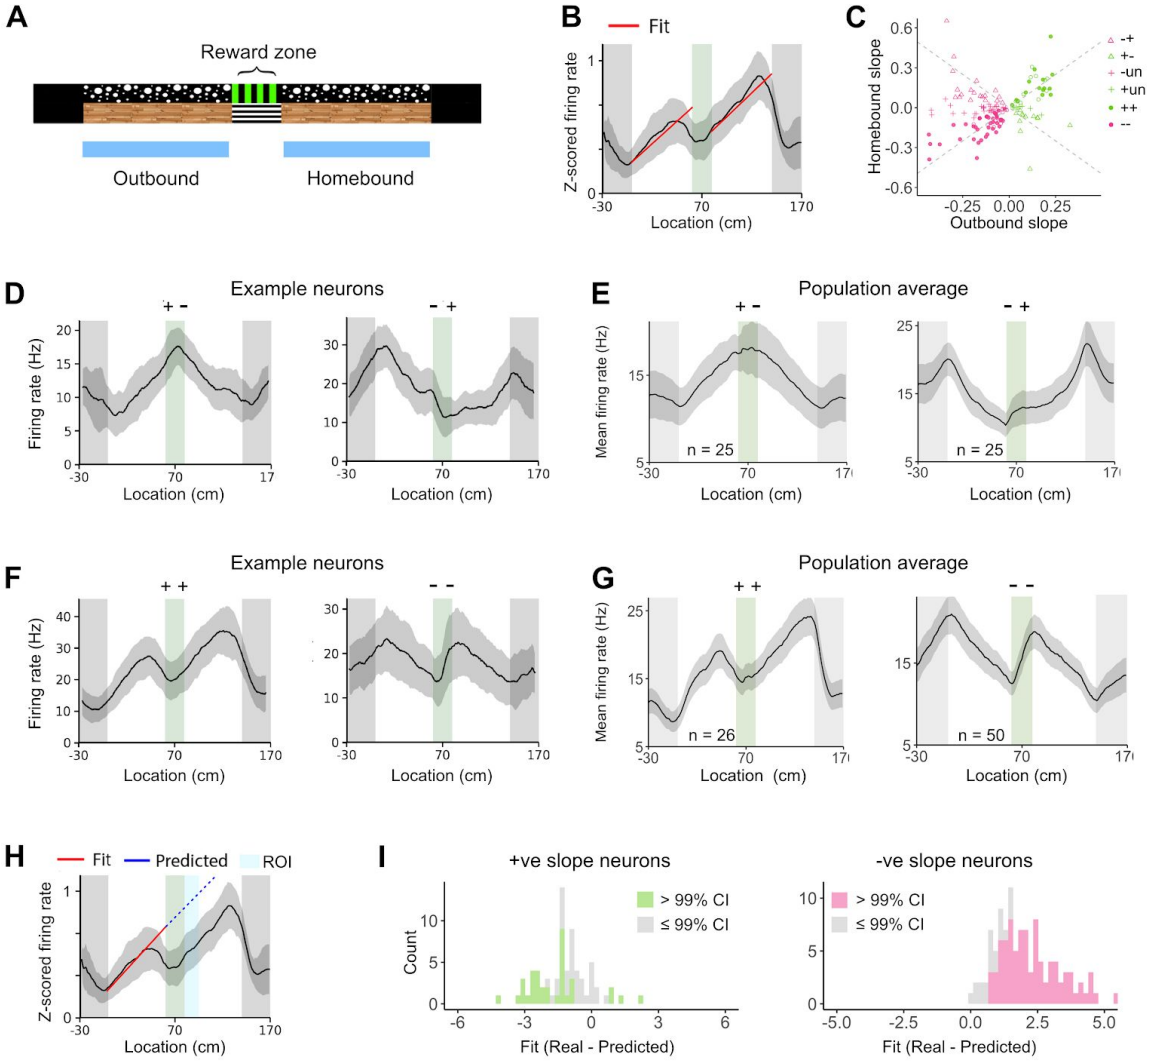

Figure 5. Ramp-like activity is modified between outbound and homebound regions of the track

(A) Schematic indicating outbound and homebound regions of the track used for analyses.

(B) Spiking activity as a function of location was fit separately for outbound and homebound regions of the track.

(C) Slope of the fit to the firing rate on the outbound track region as a function of slope on the homebound region. Dashed lines indicate points for which the slope would be identical on the homebound and outbound track segments. Symbols indicate classification of individual neurons.

(D-E) Examples (D) and population averages (E) for neurons with slope that switches to the opposite direction between the outbound and homebound regions of the track. Solid lines indicate the mean and grey shaded areas indicate SEM.

(F-G) Examples (F) and population averages (G) of firing rate as a function of position for neurons classified as having negative (left) or positive (right) slopes on both outbound and homebound regions of the track.

(H) Offsets in firing rate were calculated by subtracting the actual firing rate in the first 5 cm of the homebound track region with the firing rate predicted by extrapolating the outbound fit.

(I) Distribution of offsets for neurons classified on beaconed trials as having ++ (left) or -- (right) activity patterns.

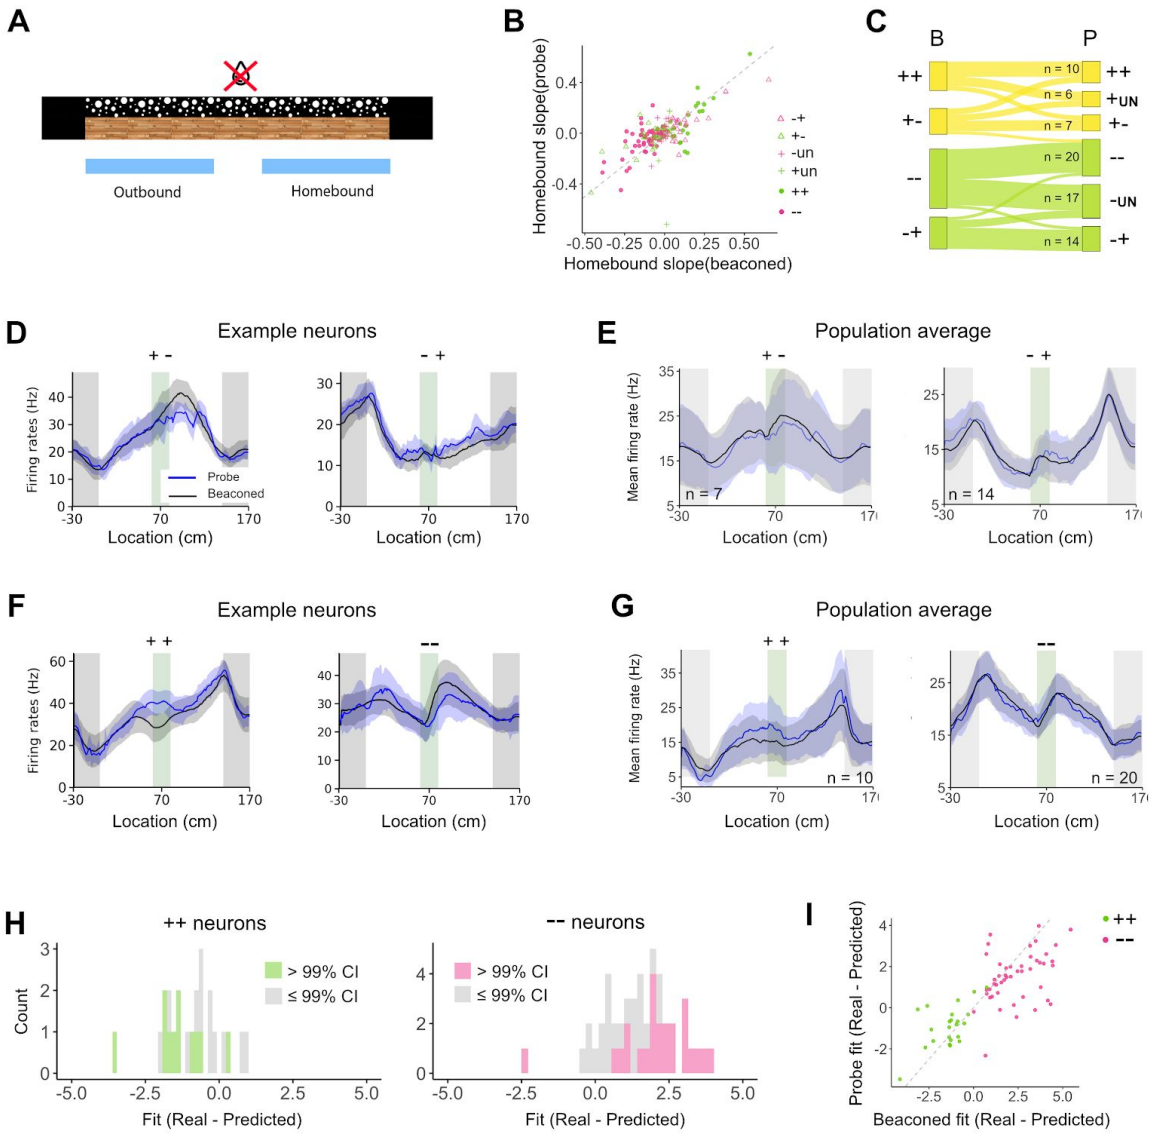

Figure 6. Location cues are not required for changes in ramp properties between outbound and homebound track segments

(A) Schematic of track configuration. These analyses consider only probe trials; visual cues indicating the reward zone are absent and no reward is delivered.

(B) Slope of the fit to the firing rate on the homebound track region in beaconed trials versus probe trials. The dashed line indicates points for which the slope would be identical on the homebound and outbound track segments. Slopes in the homebound region were correlated between beaconed and probe trials ($r^2 = 0.46$, $p < 2 \times 10^{-16}$).

(C) Classification of neurons based on the slope of their activity on homebound and outbound track segments on beaconed (B) and probe trials (P). N number is displayed for groups contiaining 5 or more cells.

(D-E) Individual examples (D) and population averages (E) of firing rate as a function of track position for neurons with slope that changes direction between the outbound and homebound regions of the track on beaconed (black) and probe trials (blue). Solid lines indicate the mean and shaded areas indicate the SEM.

(F-G) Examples (F) and population averages (G) of firing rate as a function of track position for neurons classified as ++ (left) or -- right on beaconed (black) and probe trials (blue).

(H) Firing rate offsets in the first 5 cm of the homebound track region, compared with predicted firing rate from fits to activity in the outbound region of the track, for path integration-driven neurons classified on beaconed trials as having ++ or -- firing profiles.

(I) Firing rate offsets in the first 5 cm of the homebound track region on probe trials as a function of the same offset on beaconed trials for the same neurons as in (H). Offsets were correlated between beaconed and probe trials ($r^2 = 0.56$, $p = 4.53 \times 10^{-14}$)

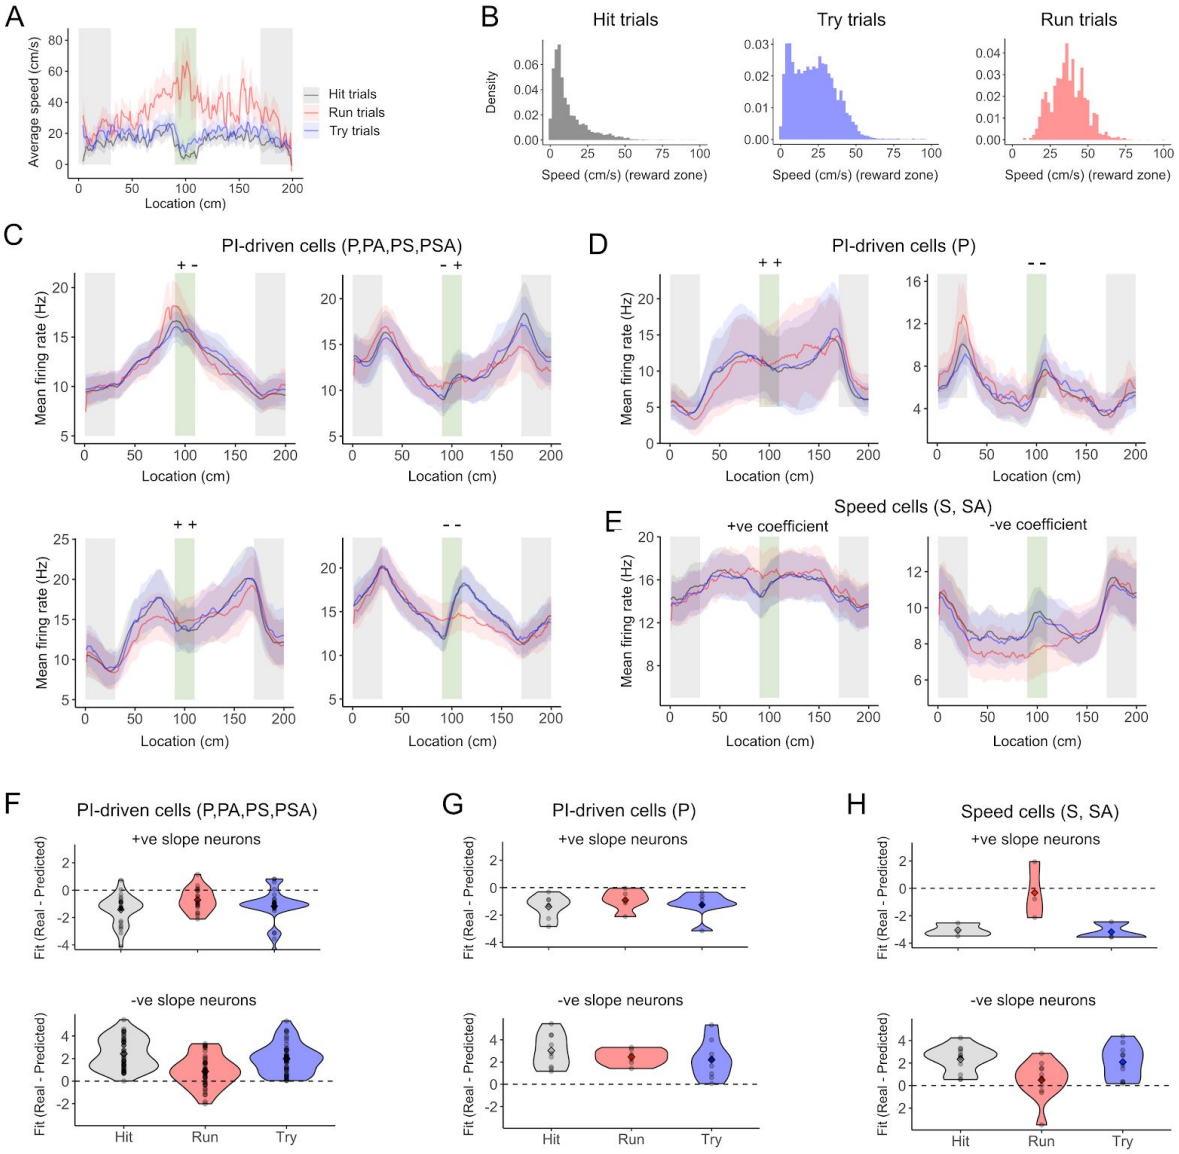

Figure 7. Relationship between task performance and ramp-like activity

(A) Population average of mean running speed as a function of position for each trial outcome across all animals (n=7). Shaded regions correspond to the standard error of the mean across animals.

(B) Distribution of trial level mean running speeds within the reward zone for hit, try and run trial outcomes. The histograms are calculated from all cells (n = 649) and all trials with a given outcome (hit: n=35989; try: n=42276; run: n=10173).

(C) Population averaged firing rate of position modulated neurons (P, PA, PS, PSA groups) as a function of position for +-, -+, ++ and -- groups on hit, try and run through trials. Shaded regions correspond to the standard error of the mean across cells.

(D) Population averaged firing rates plotted as in (C) but only for ++ and -- neurons modulated solely by position and not in combination with speed or acceleration (P group).

(E) Population averaged firing rate as a function of position on hit, try and run through trials for speed modulated neurons (S, SA groups). Data includes neurons identified neurons all neurons from the S and SA groups irrespective of their slope.

(F-H) Firing rate offsets in the first 5 cm of the homebound track region on run, try and hit trials for all position modulated neurons (P, PS, PA, PSA groups)(F), neurons modulated solely by position (P group)(G) and speed modulated neurons (S and SA groups)(H). Offsets for hit, run and try trials did not differ for position neurons (P group) with negative or positive slopes (+ve slope: $p = 0.65$, $f = 0.45$; -ve slope: $p = 0.48$, $f = 0.76$; ANOVA). Offsets for hit, run and try trials for position neurons (P, PS, PA, PSA groups) were significantly different for negative but not positive slopes (+ve slope: $p = 0.11$, $f = 2.3$; -ve slope: $p = 5.48 \times 10^{-7}$, $f = 16.7$; ANOVA) similar to neurons modulated by speed (S and SA groups) (+ve slope: $p = 0.05$, $f = 4.73$; -ve slope: $p = 0.01$, $f = 5.29$; ANOVA).

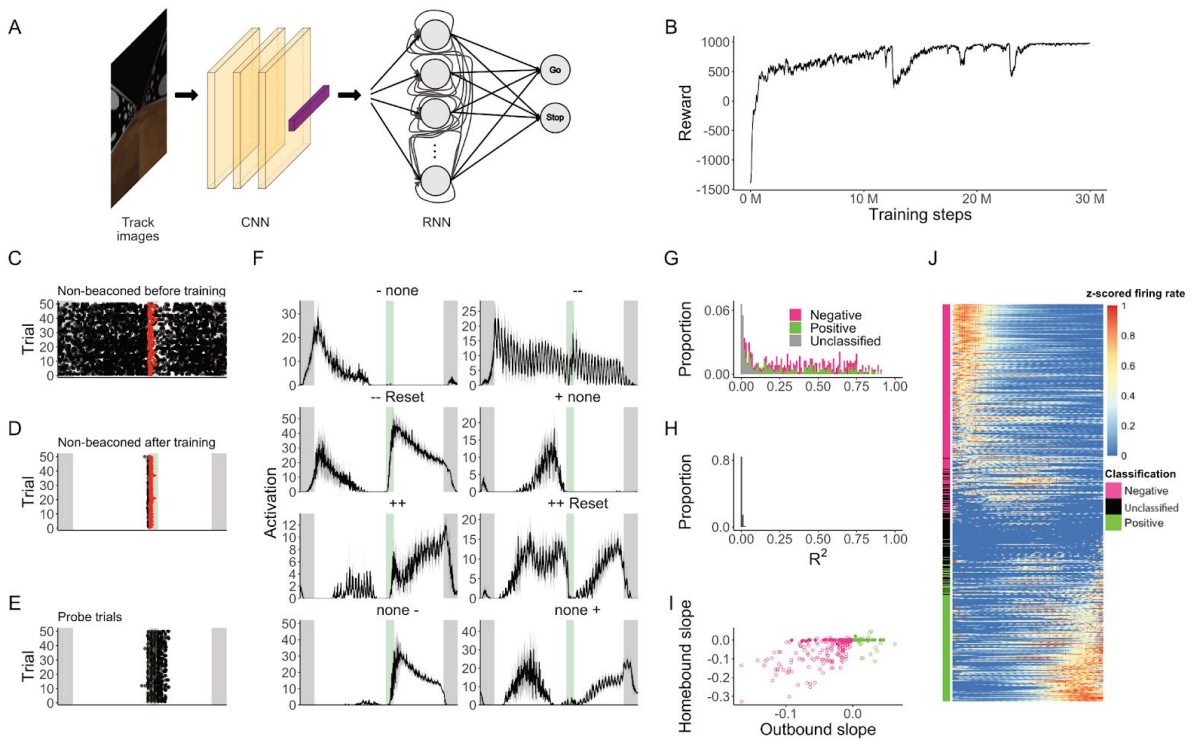

Figure 8. Recurrent networks trained to estimate location generate ramp-like representations

(A) Schematic of the network. Images of the track are played into a convolutional neural network (CNN) that sends outputs to a recurrently connected network (RNN) that sets the activity of movement control (Go / Stop) neurons. 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. A typical trial involves ~ 250 steps.

(C-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 outbound track zone. Units are classified as in Figure 2.

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

(I) Slopes of the firing rate of RNN units in the homebound region as a function of slopes in the outbound region.

(J) Normalised firing rate as a function of position on the outbound track region for all units in the RNN.

Materials and methods

Experimental models

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. Male and female mice, aged 7–16 weeks, were used for all experiments.

Seven male mcos wild type mice 7-14 weeks (mean = 12.8, SD = 4.7) old at the time of the implant surgery were used. 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 singly 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*.

Experimental methods

Microdrive fabrication

Microdrive fabrication for in vivo electrophysiology recordings was similar to Gerlei et al (2020). 16-channel microdrives consisting of 4 tetrodes were built by threading 90% platinum 10% iridium wire (18µm 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 (https://github.com/MatsumotoJ/Tetroplater) 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 (Sürmeli et al., 2015). Inhalation anaesthesia 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 (Osborne and Dudman, 2014) 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.

To implant the microdrive it was attached to an Omnetics to Mill-Max adaptor (Axona, HSADPT-NN1), which was then attached to the stereotaxic frame. We lowered the tetrodes, at an angle of -16° in order to target retrohippocampal cortices, 1.2 - 1.4 mm into the brain, beginning at 3.4 mm lateral from Bregma (right hemisphere) and along the lambdoid suture. The outer cannula was then lowered and sterile Vaseline carefully placed around the rim to provide a protective seal around the craniotomy. The implant was then fixed to the skull with UV curing resin. After the cement hardened, the grounding wires were wrapped around the grounding screws and fixed with silver paint (RS components 101-5621). A layer of resin cement was applied to cover the grounding screws and any holes were filled with dental cement (Simplex Rapid). Mice recovered on a heat mat for approximately 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 in the holding room before handling and habituation to the virtual reality set up.

Handling and habituation

Handling and habituation of mice was similar to that described in Tennant et al (2018). 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 familiarise 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 behavioural task their access to food was restricted so that their body weight was maintained at approximately 85% of its baseline value (Tennant et al., 2018). 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 normalised to the expected daily growth for the animal's age.

Virtual reality task

Mice were trained to locate a rewarded location on a virtual linear track as described previously (Tennant et al., 2018). On experimental days animals were collected from the holding room 30 - 60 minutes before the 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. During recording sessions animals were head fixed in the virtual reality setup for approximately 45 minutes.

Mice were head fixed in the virtual reality setup using a Rivets clamp (Ronal Tool Company, Inc)(Osborne and Dudman, 2014) 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 (blender.com). The track had length 2 m, with a 60 cm outbound zone, a 20 cm reward zone, a 60 cm homebound 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 via 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 (Siegle et al., 2017). Signals were filtered using the board's analogue high and low cut filter (2.5 Hz -7603.8 Hz). Behavioural variables including position, trial number and trial type were calculated in Blender3D at 60 Hz as described previously (Tennant et al., 2018). 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.

After the recording session, animals were placed back in the object filled arena and then a second open arena before being returned to their home cage. Tetrodes were typically lowered by 100 µm after each session.

Post recording assessment of tetrode location

To determine tetrode locations we used either Cresyl violet staining (n = 4 mice) or microcomputed tomography (microCT) (n = 3 mice) (Figure S2). One mouse that underwent cresyl violet staining had no visible tetrode tracks and is omitted from Figure S2. For both methods mice were anaesthetised 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 0.2 M phosphate buffer (PB; Sigma Aldrich). Brains were removed, left in 4 % PFA in 0.2 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 polarised slides, dried overnight and then placed in xylene before rehydration in varying concentrations of ethanol. Once washed, sections were placed in Cresyl violet for 2 mins before additional washing and dehydration. 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 (Masís et al., 2018). 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 aluminium filter between the X-ray source and the sample (Xie et al., 2020). 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 localised 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 = 2 mice), the pre/parasubiculum (n = 2), or when the location was ambiguous between these areas to retrohippocampal cortex (n = 3)(Supplemental Figure 2). 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, 1994)) was corroborated by electrophysiological recording of a theta field potential and theta modulation of the firing of neurons (cf. (Boccara et al., 2010))(see Figure Supplemental Figure 4 and analyses section below). In contrast, neurons in nearby structures such as the postrhinal cortex do not show theta phase locked firing (LaChance et al., 2019). In one animal, initial recording sessions that did not show a clear theta-modulated local field potential or theta phase locked spike firing were omitted. The animal from which we were unable to localise the tetrodes anatomically, but in which field potential recordings and spike firing showed theta modulation, was assigned to the retrohippocampal group.

Quantification and statistical analyses

Analyses were carried out using Python (version 3.5.1) and R version: 3.3.1 (2016-06-21). All code will be made available at https://github.com/MattNolanLab.

Spike sorting

Spikes were isolated from electrophysiological data using an automated pipeline based on MountainSort (v 0.11.5 and dependencies)(Chung et al., 2017) and similar to that described in Gerlei et al (2020). 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 to remove correlated noise. Events were detected from peaks > three standard deviations above baseline and at least 0.33 ms away from other events on the same channel. A 10-dimensional feature space was then generated using the first 10 principal components of the detected waveforms and spike sorting using the ISO SPLIT algorithm performed on that feature space (Chung et al., 2017). Cluster quality was evaluated using metrics for isolation, noise-overlap, and peak signal to noise ratio (Chung et al., 2017). 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 neurons based on the slope of their mean firing rate (Figure 2) we generated firing rate maps 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. Firing rates were then smoothed with a gaussian filter (2 cm width, 2 SD, Scipy Python package (Virtanen et al., 2020)). For each neuron, we also generated 999 shuffled datasets in which average firing rates were randomly allocated to track bins. We then fit linear models to firing rate as a function of position across the region of the track corresponding to either the outbound or the homebound zone. The p values obtained from fitting the linear models were corrected for multiple comparisons using the Benjamini & Hochberg method (Benjamini and Hochberg, 1995). We classified neurons as having ramp-like activity with positive (or negative) slopes if the linear model was significant (corrected $p < 0.01$), the coefficients of the model were outside the 5 - 95 % intervals of the shuffled data, and the slope was positive (or negative). Neurons not meeting these criteria were allocated to the unclassified group.

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. This ramp score is related to the score used previously for identification of speed cells (Kropff et al., 2015). For each cell we calculated separate ramp scores for the homebound and outbound track zones.

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 within the outbound track zone at which peaks and troughs in the firing rate are present. Using all turning points (peaks and troughs) in the smoothed curve, together with two endpoints (i.e. beginning and end) of the outbound regions, 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 length of the track region of interest (e.g. the homebound or outbound 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.

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 used mixed models incorporating 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 (Góis and Tort, 2018; Tsao et al., 2018). Firing rate and speed was then smoothed with a gaussian filter (2 cm width, 2 std, Scipy Python package (Virtanen et al., 2020)). 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 model with the configuration 'Rate ~ Position + Speed + Acceleration + (1 + Position | Trials) + error' was fit using the lme4 package (version 1.1-12) in R (Bates et al., 2015). The ANOVA function provided by the car package (version 3.0-9) was used to calculate significance values for each model coefficient (Fox and Weisberg, 2018). To classify neurons as being modulated by position, speed, acceleration, 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.

Mixed effect models for evaluating trial outcome

To evaluate possible effects of trial outcome, we used mixed effect models that included position and trial outcome as fixed effects and trial number as a random effect. Possible outcomes were: 'hit' trials in which animals stopped in the reward zone; 'try' trials if the animal did not stop in the reward zone but its speed in this region was within the 5-95% confidence intervals of hit trials; and 'miss' trials when the animal did not stop in the reward zone and its speed was outside the 5-95% confidence intervals for hit trials.

Similar to analyses of position, all spikes that occurred when the animal's speed was < 3 cm/s were removed and bins in which speed or acceleration were > 3 standard deviations (SDs) above the mean removed from the recording. A model with the configuration 'Rates ~ Position * Trial_outcome + (1 + Position | Trials) + error' was fit for each neuron using the lme4 package in R (Bates et al., 2015). Significance values for each model coefficient were then calculated using the ANOVA function provided by the car package (Fox and Weisberg, 2018). We used a significance threshold of 0.01 for the interaction between position and trial outcome to classify neurons ramp-like activity that depended on the trial outcome. Sessions in which mice had less than 3 trials classified as either run or try were not included in the analyses.

Average speed calculation

For Figure 7a, 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 with a gaussian filter (2 cm width, 3 std, Scipy Python package (Virtanen et al., 2020)). 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. (2018). Spike times were transformed into instantaneous firing rates. A power spectrum was then estimated using Welch's method. The theta index was defined as θ-baseline/θ+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.

Neural network model

We trained an artificial neural network to use a reinforcement learning algorithm 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 randomisation), each with a corresponding image. The reinforcement learning agent can 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 behaviour 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 track 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) (Elman, 1990). 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 randomised 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 optimisation algorithm (Schulman et al., 2017), which learned an action and value function. 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 totalled for each trial. The network weights were updated using stochastic-gradient descent to maximise the reward (Kingma and Ba, 2014). To find the optimal hyperparameters for training (such as learning rate), we performed random parameter searches using the Edinburgh Compute and Data Facility (ECDF) (Bergstra and Bengio, 2012), using PyTorch 1.2.0 and cuDNN 10.1.105 (Paszke et al., 2019) and adapted an existing Python implementation of the proximal policy optimisation algorithm (Kostrikov, 2018). The hyperparameters used are in Supplemental Table 1.

Acknowledgements

We thank Dr 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) and the BBSRC (BB/ L010496/1); by grants from the Simons Initiative for the Developing Brain to M.F.N; by a College of Medicine and Veterinary Medicine PhD Studentship, funded by the Thomas Work Fellowship, to K.G. and by the Wellcome Trust (108890/Z/15/Z) Translational Neuroscience PhD programme to I.H. This work made use of resources provided by the Edinburgh Compute and Data Facility.

Supplemental figures

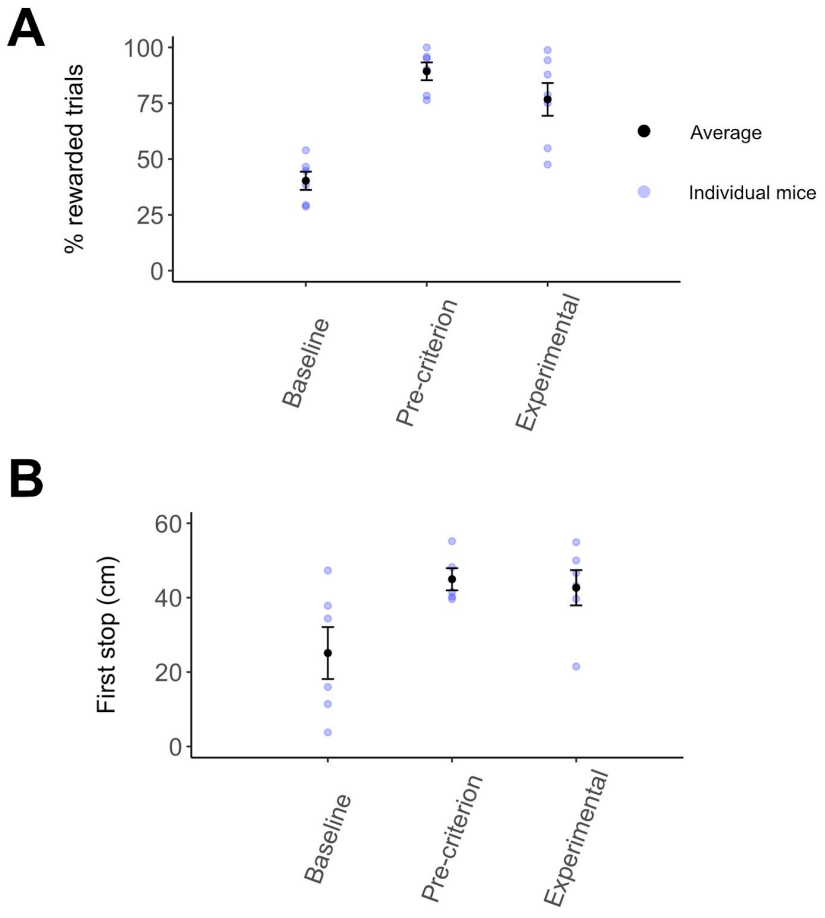

Supplemental Figure 1. Behavioural training

(A) The proportion of rewarded trials for each mouse (blue) and averaged across mice (black, ± SEM) for the first training day (baseline), the two consecutive days in which mice reached the criterion for introduction of probe trials (Pre-criterion) and the days after the criterion was passed (Experimental). Mice took 13.3 ± 7.1 days to reach the performance criterion.

(B) The mean location at which mice first stopped on each trial for each mouse (blue) and averaged across all mice (black, ± SEM). Categories are as in (A).

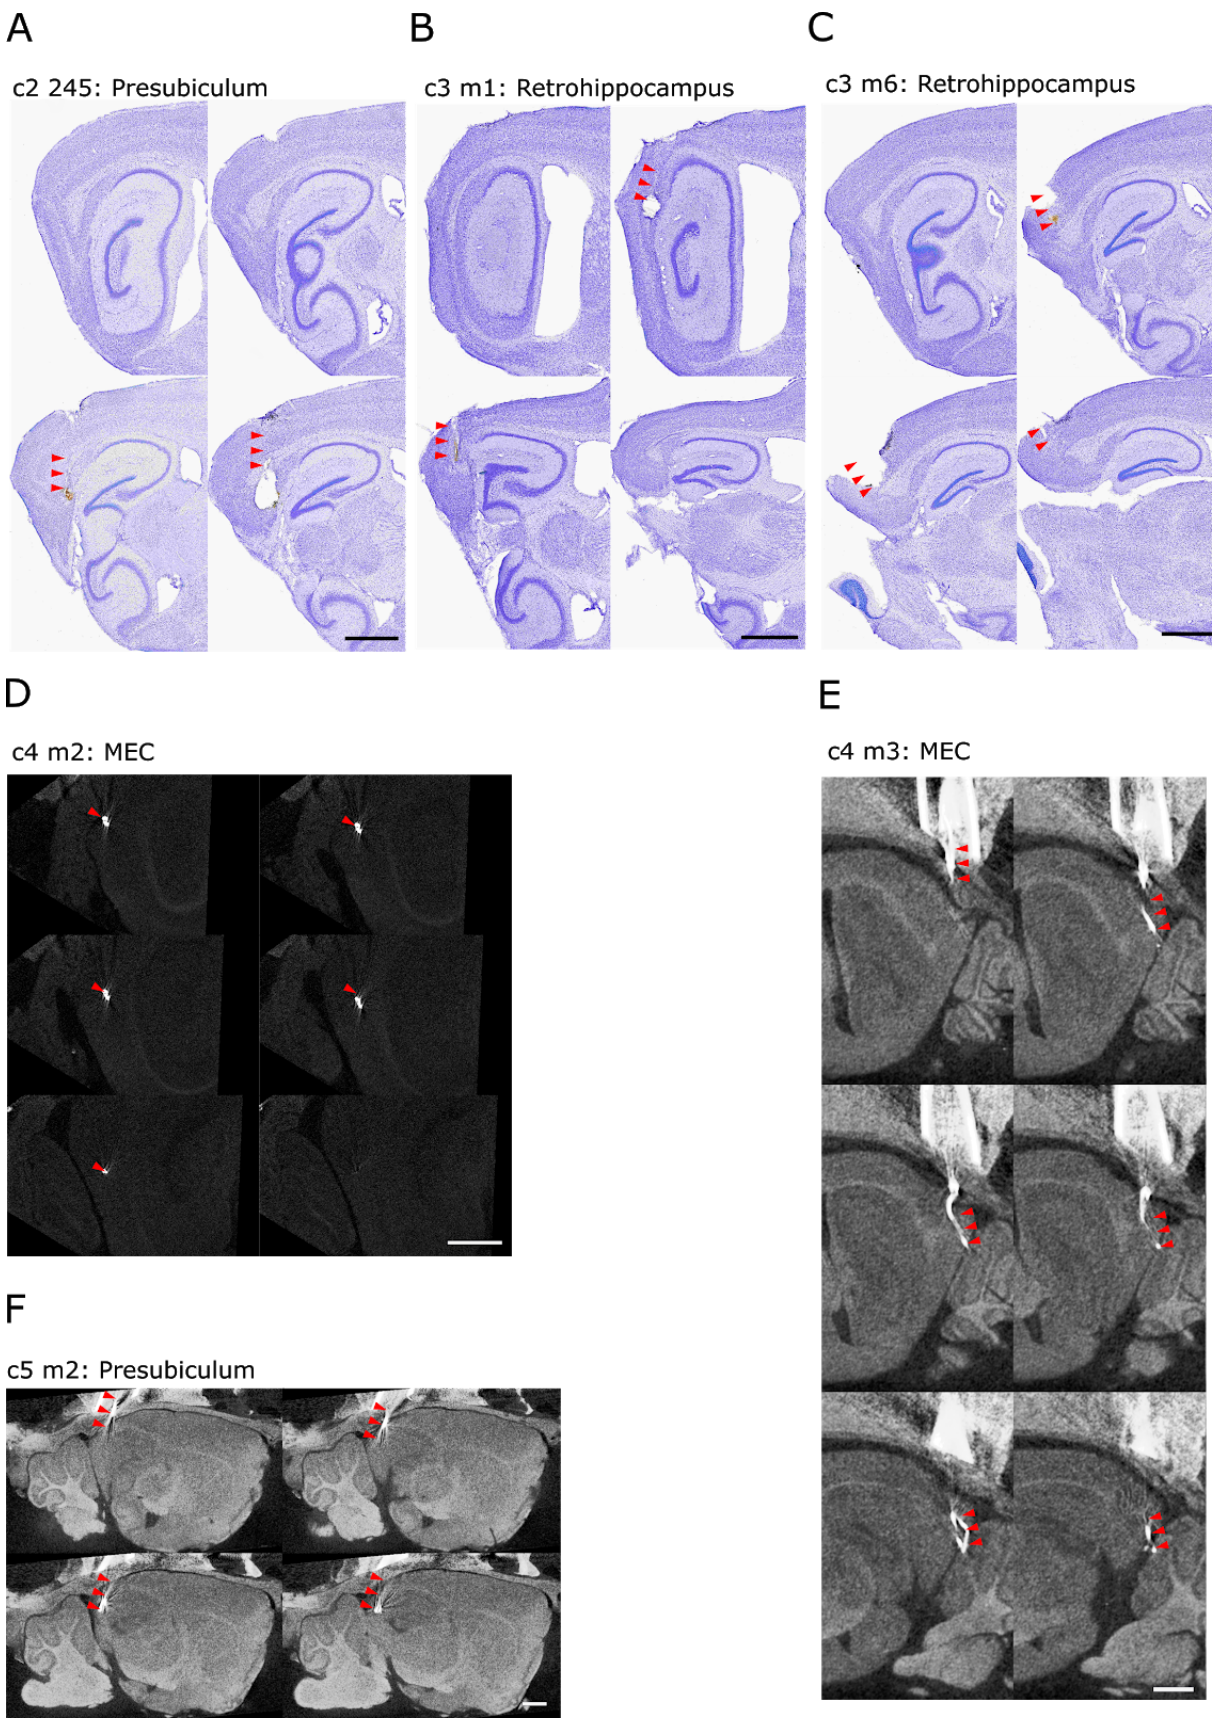

Supplemental Figure 2. Tetrode localisation

(A - C) Cresyl violet stained brain sections used for assessment of tetrode locations in three mice.
(D - F) Micro-CT images used for assessment of tetrode locations in three mice. For each animal, sagittal slices are presented lateral to medial from left to right and the classification of the tetrodes target is shown at the top left. This classification is based on the terminal location of the tetrode and the distance travelled during the experiment (see Methods). One mouse had no visible tetrode tracks in any slice and is not shown. In all images red triangles point to the putative tetrode tracks.

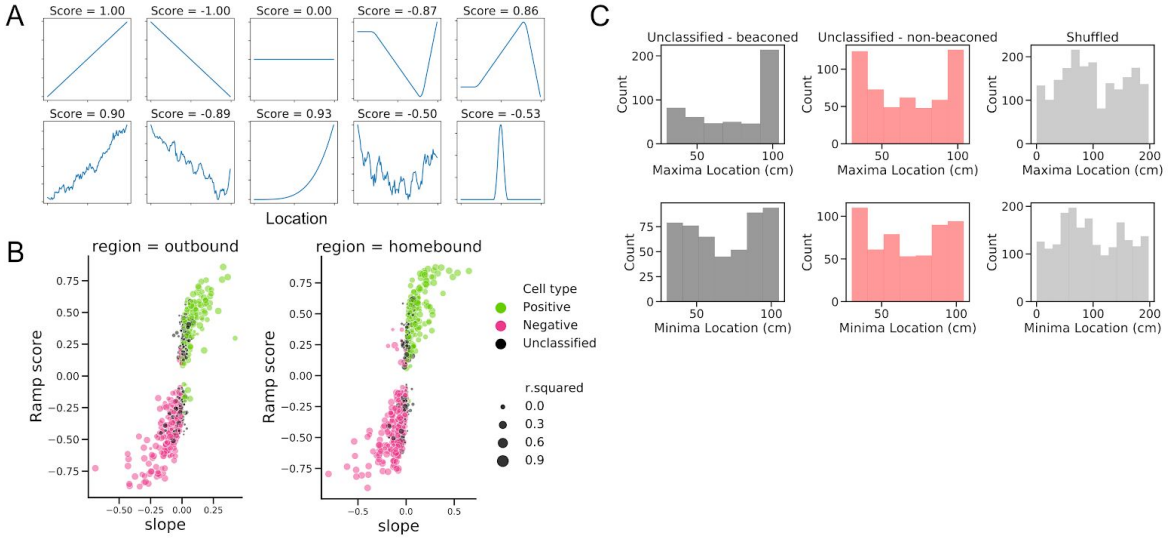

Supplemental Figure 3. Evaluation of ramp scores

(A) Ramp scores for reference distributions.

(B) Ramp scores plotted as a function of the slope of linear fits of firing rate to location within the outbound track zone for all neurons.

(C) Locations of maxima and minima for unclassified neurons. Corresponding plots or neurons classified as having positive or negative slopes are shown in Figure 2 B.

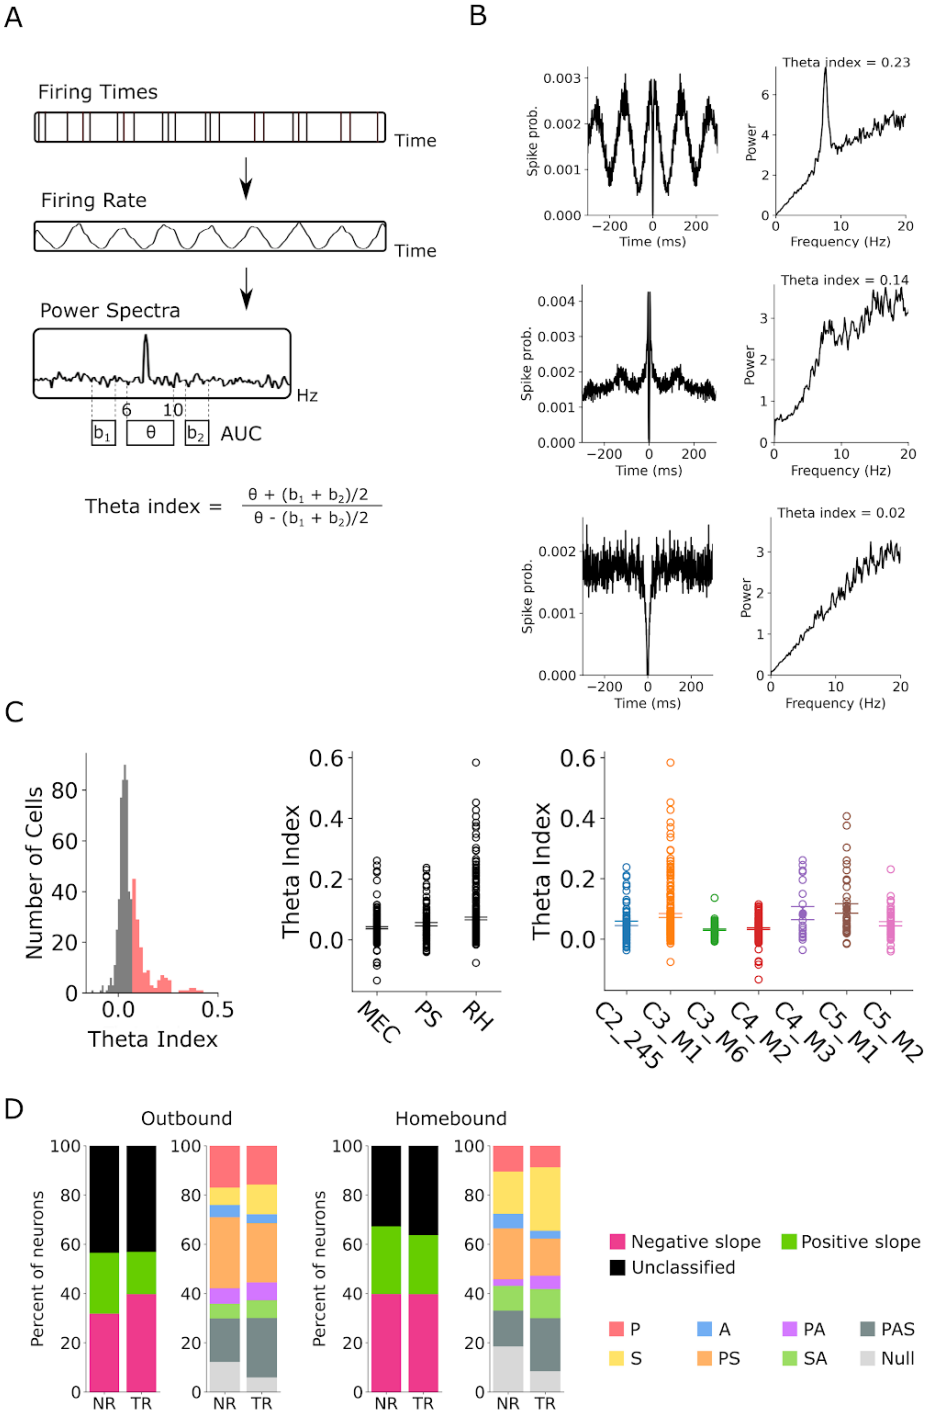

Supplemental Figure 4. Theta modulation of recorded neurons

(A) Schematic illustrating the procedure for calculation of the theta index (see (Kornienko et al., 2018)). Firing indices were transformed into instantaneous firing rates from which a power spectrum was estimated using Welch's method. The theta index was defined as θ-baseline/θ+baseline, where θ is the mean power in the theta band (6-10 Hz) and baseline is the mean power of two adjacent frequency bands (3-5 and 11-13 Hz).

(B) Autocorrelograms (left) and power spectra (right) for three simultaneously recorded neurons that exemplify strong (top), moderate (middle) and low (bottom) theta indices.

(C) Histogram of the theta index across all cells (left) and scatter plots showing the theta index grouped by tetrode location (middle) and mouse (right). Cells were classified as having theta rhythmic firing (red in the histogram) when their theta index was > 0.07 (see (Kornienko et al., 2018)).

(D) Proportions of non-rhythmic (NR) cells and theta rhythmic (TR) cells classified as having positive or negative slopes (see Figure 2) and among these neurons the proportions classified as having firing that correlates with position (P), speed (S) and acceleration (A) (see Figure 3).

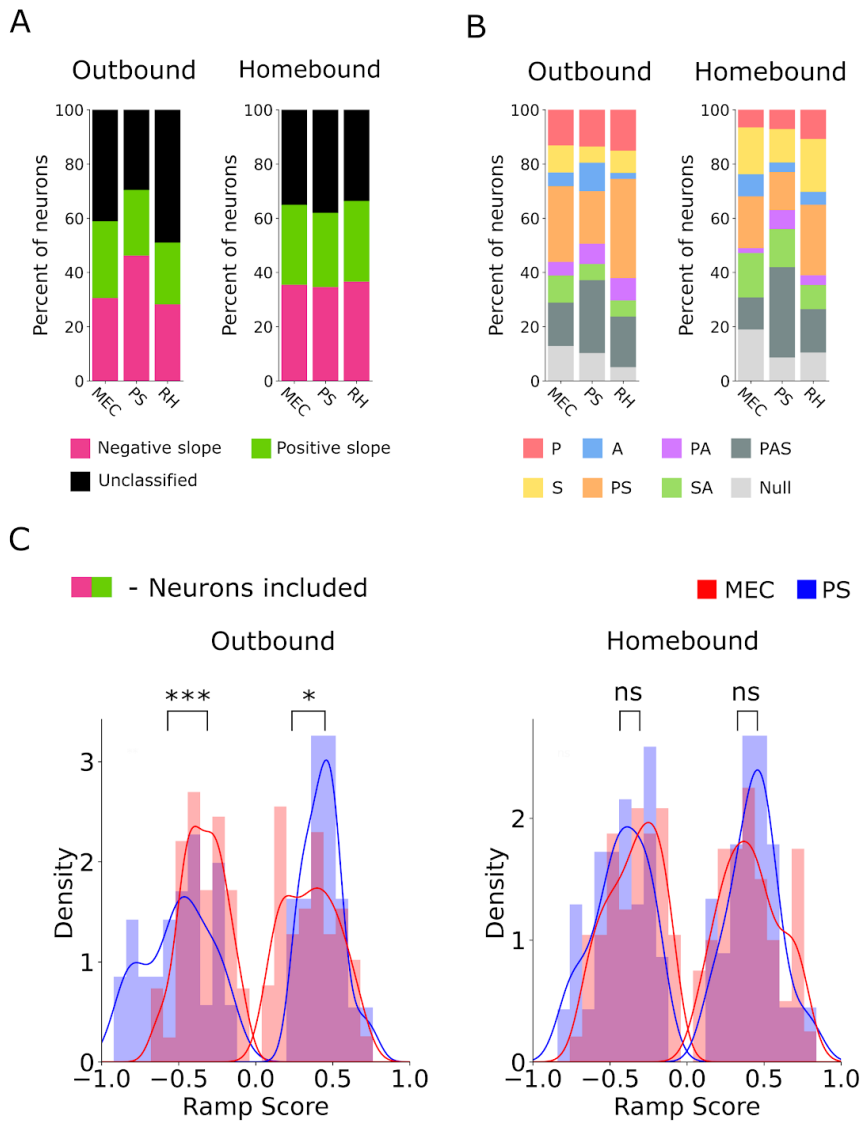

Supplemental Figure 5. Comparison of ramp scores and slopes in medial entorhinal cortex vs pre/parasubiculum

(A) Distribution of neurons with either positive or negative slopes (see Figure 2) on the outbound (left) and homebound (right) track segments according to their locations in the medial entorhinal cortex (MEC), pre/parasubiculum (PS) or retrohippocampal cortex (RH). RH is used for neurons for which anatomical location in either MEC or PS could not be accurately determined.

(B) Distribution among neurons from each subregion with position or negative slopes of firing that correlates with position (P), acceleration (A) or speed (S) after fitting with the linear mixed effect models (see Figure 3).

(C) Distribution of ramp scores for neurons from MEC and pre/parasubiculum that were classified as having positive or negative sopes. Ramp score distributions in MEC and pre/parasubiculum were compared separately for neurons with negative and positive slopes for outbound (-ve neurons p=0.0003, +ve neurons p=0.04, Kolmogorov–Smirnov test) and homebound regions (-ve neurons p=0.4, +ve neurons p=0.4, Kolmogorov–Smirnov test). Smoothed distributions obtained using kernel density estimation are overlaid.

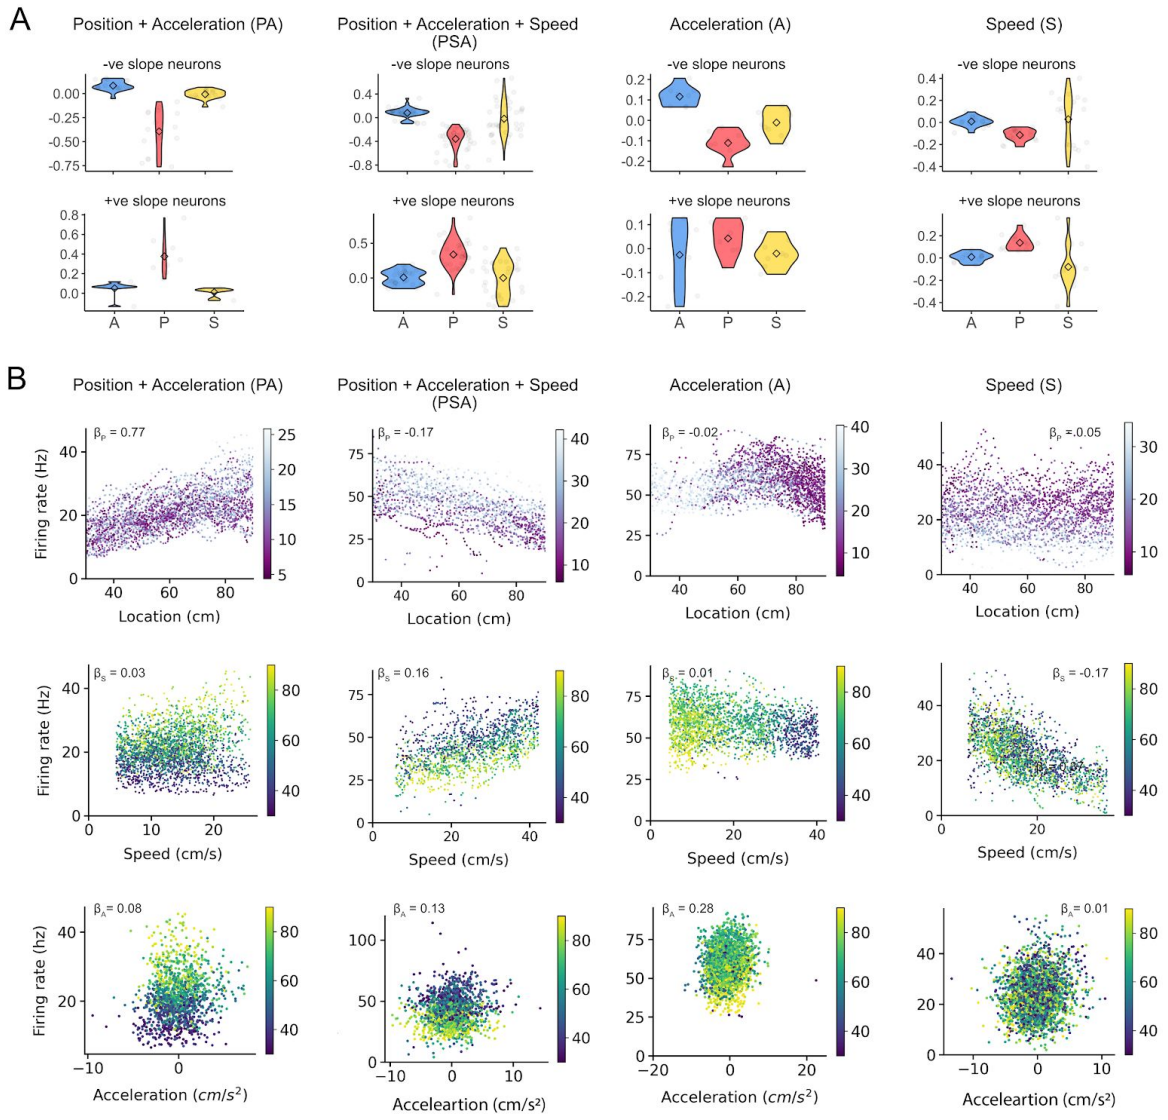

Supplemental figure 6. Firing rate dynamics of neurons with activity modulated by speed and acceleration

(A) Distribution of standardised coefficients for neurons classified as having negative slopes or positive slopes and encoding position and speed (PA), position, speed and acceleration (PSA), acceleration alone (A) or speed alone (S).

(B) Examples of firing rate as a function of position (upper), running speed (middle) and acceleration (lower) for neurons corresponding to the groups in (A). For each neuron, the standardised coefficient is shown in the top right or left hand corner. Fit coefficients for the parameter on the x-axis are shown on each plot.

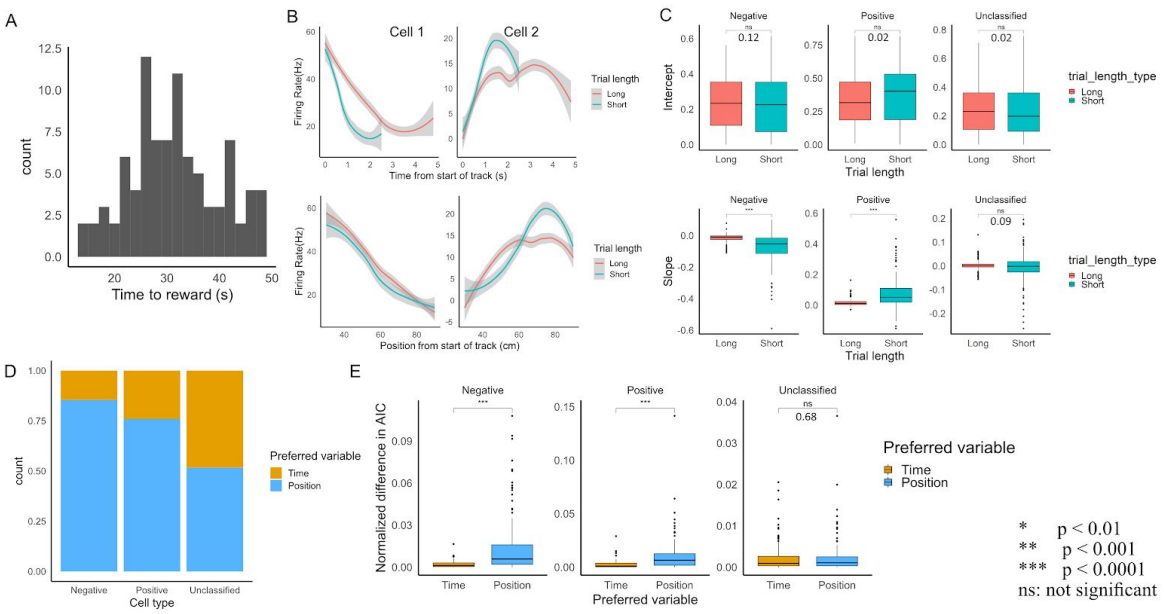

Supplemental figure 7. Evaluation of time versus position encoding

(A) Distribution of time to reach reward in a representative session.

(B) Firing rate of two ramp-like neurons across time for short and long trials plotted as a function of time (upper) or position on the track (lower), from the same behavior session as (A). Short trials are trials that are shorter than the 25th percentile of all the trials in a session, while long trials are those longer than the 75th percentile. Time encoding neurons would be expected to have similar slopes on short and long trials, whereas these neurons have more similar slopes on short and long trials, which is consistent with their encoding position.

(C) Comparison of the intercept (upper) and slope (lower) on short and long trials, for cells classified as having positive or negative slopes, when firing rate on the outbound track zone is fit as a function of time. Significant differences between the slopes of short versus long trials are inconsistent with neurons encoding time (paired Wilcoxon signed-rank test). Values for significance estimates are shown in the figure. Test statistic top row, left to right: V=13371, V=4383, V=23232, bottom row: V=21704, V=1305, V=22255

(D) Percentage of cells encoding time or position based on comparison of Akaike Information Information Criterion (AIC) values obtained from mixed effect models fit to firing rate with either time or position as the fixed effect and trial number as the random effect. The model with the more negative AIC value was assumed to provide a better account of the data. In general the activity of neurons classified as having positive or negative slopes in the outbound track zone were better predicted by models in which position rather than time is a fixed effect.

(E) Comparison of the normalized difference in AIC between the position model and time model for different cell types. The difference is normalized by the sum of the AIC of position and time model. A larger difference is indicative of a model providing a better explanation for the data. Whereas position coding models were often strongly favoured, time encoding models were not (independent Wilcoxon

signed-rank test, p-value for significance estimates are shown in the figure. Test statistic left to right: W=1351, W=1024, W=9616).

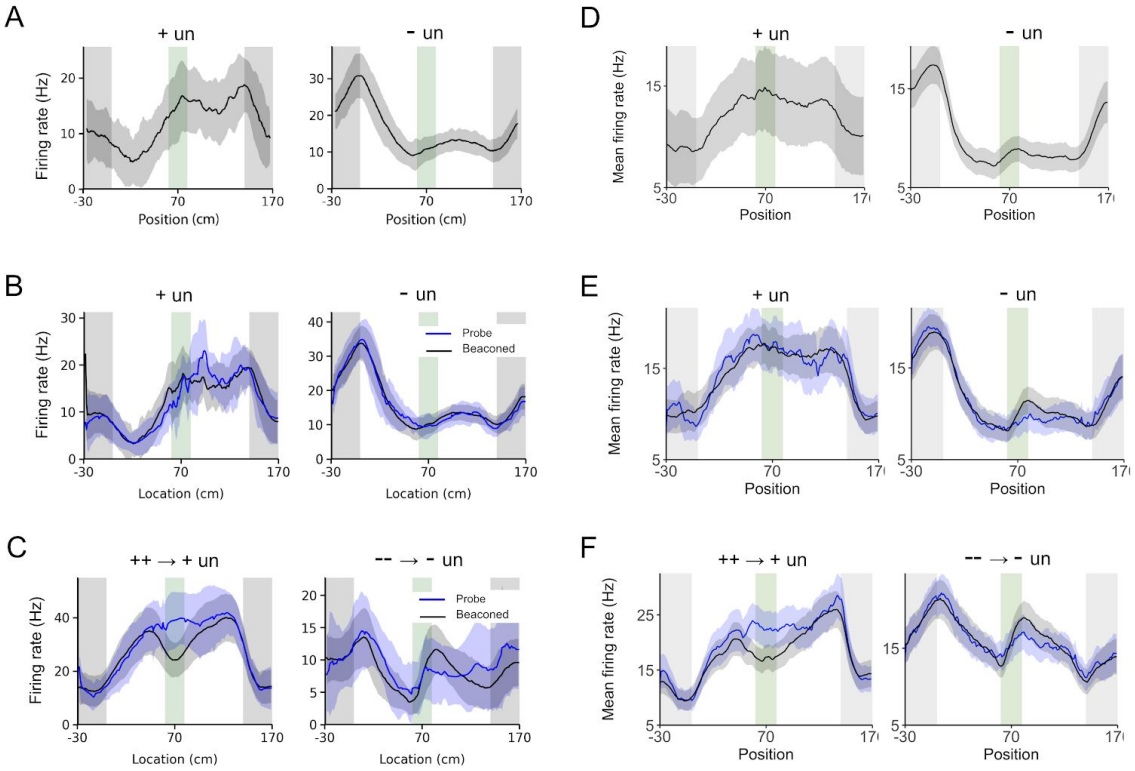

Supplemental Figure 8. Neurons with ramp-like activity on outbound but not homebound regions of the track

Examples (A, B. C) and population averages (D, E, F) on beaconed trials (A, D) and probe trials (B-F) of firing rate as function of location for neurons classified as having positive (left panels) or negative (right panels) firing rate slopes on the outbound track region and no ramping activity on the homebound region in both beaconed and probe trials (B, E) or probe trials alone (C, F).

Supplemental Table 1: Model hyperparameters

Hyperparameters used in the recurrent neural network model

Parameter	Value
Learning rate	1.15e-04
Entropy coefficient	0.01
Epsilon	1e-05
Number of minibatches	4
Number of steps	5000
Number of epochs	14
Weight decay	5e-07
Gamma	0.99
Number of neurons in RNN layer	512

References

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.

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.

Behrens, T.E.J., Muller, T.H., Whittington, J.C.R., Mark, S., Baram, A.B., Stachenfeld, K.L., and Kurth-Nelson, Z. (2018). What Is a Cognitive Map? Organizing Knowledge for Flexible Behavior. Neuron *100*, 490–509.

Bellmund, J.L.S., Gärdenfors, P., Moser, E.I., and Doeller, C.F. (2018). Navigating cognition: Spatial codes for human thinking. Science *362*.

Benjamini, Y., and Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. Journal of the Royal Statisitical Society, Series B *57*, 289–300.

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

Bicanski, A., and Burgess, N. (2020). Neuronal vector coding in spatial cognition. Nat. Rev. Neurosci.

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.

Brody, C.D., Romo, R., and Kepecs, A. (2003). Basic mechanisms for graded persistent activity: discrete attractors, continuous attractors, and dynamic representations. Curr. Opin. Neurobiol. *13*, 204–211.

Buckmaster, C.A., Eichenbaum, H., Amaral, D.G., Suzuki, W.A., and Rapp, P.R. (2004). Entorhinal cortex lesions disrupt the relational organization of memory in monkeys. J. Neurosci. *24*, 9811–9825.

Burak, Y., and Fiete, I.R. (2009). Accurate path integration in continuous attractor network models of grid cells. PLoS Comput. Biol. *5*, e1000291–e1000291.

Bush, D., Barry, C., Manson, D., and Burgess, N. (2015). Using Grid Cells for Navigation. Neuron *87*, 507–520.

Campbell, M.G., Ocko, S.A., Mallory, C.S., Low, I.I.C., Ganguli, S., and Giocomo, L.M. (2018). Principles governing the integration of landmark and self-motion cues in entorhinal cortical codes for navigation. Nat. Neurosci. *21*, 1096–1106.

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.

Carpenter, R.H.S. (2011). What Sherrington missed: the ubiquity of the neural integrator. Ann. N. Y. Acad. Sci. *1233*, 208–213.

Carvalho, M.M., Tanke, N., Kropff, E., Witter, M.P., Moser, M.-B., and Moser, E.I. (2020). A Brainstem Locomotor Circuit Drives the Activity of Speed Cells in the Medial Entorhinal Cortex. Cell Rep. *32*, 108123.

Casali, G., Shipley, S., Dowell, C., Hayman, R., and Barry, C. (2018). Entorhinal Neurons Exhibit Cue Locking in Rodent VR. Front. Cell. Neurosci. *12*, 512.

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. U. S. A. *110*, 378–383.

Chen, G., Manson, D., Cacucci, F., Chen, G., Manson, D., Cacucci, F., and Wills, T.J. (2016). Absence of Visual Input Results in the Disruption of Grid Cell Firing in the Mouse Report Absence of Visual Input Results in the Disruption of Grid Cell Firing in the Mouse. Curr. Biol. 1–8.

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.

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

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.

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.

Cueva, C.J., and Wei, X.-X. (2018). Emergence of grid-like representations by training recurrent neural networks to perform spatial localization.

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

Domnisoru, C., Kinkhabwala, A. a., and Tank, D.W. (2013). Membrane potential dynamics of grid cells. Nature 1–6.

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.

Elman, J.L. (1990). Finding Structure in Time. Cogn. Sci. *14*, 179–211.

Engelhard, B., Finkelstein, J., Cox, J., Fleming, W., Jang, H.J., Ornelas, S., Koay, S.A., Thiberge, S.Y., Daw, N.D., Tank, D.W., et al. (2019). Specialized coding of sensory, motor and cognitive variables in VTA dopamine neurons. Nature.

Fiete, I.R., Burak, Y., and Brookings, T. (2008). What grid cells convey about rat location. J. Neurosci. *28*, 6858–6871.

Fiorillo, C.D., Tobler, P.N., and Schultz, W. (2003). Discrete Codingof Reward Probabilityand Uncertaintyby Dopamine Neurons. Science *299*.

Fischler, W.M., Joshi, N., Devi-Chou, V., Kitch, L., Schnitzer, M., Abbott, L.F., and Axel, R. (2019). Olfactory Landmarks and Path Integration Converge to Form a Cognitive Spatial Map.

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.

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

Fransén, E., Tahvildari, B., Egorov, A.V., Hasselmo, M.E., and Alonso, A.A. (2006). Mechanism of graded persistent cellular activity of entorhinal cortex layer v neurons. Neuron *49*, 735–746.

Gershman, S.J. (2014). Dopamine ramps are a consequence of reward prediction errors. Neural Comput. *26*, 467–471.

Gil, M., Ancau, M., Schlesiger, M.I., Neitz, A., Allen, K., De Marco, R.J., and Monyer, H. (2017). Impaired path integration in mice with disrupted grid cell firing. Nat. Neurosci.

Góis, Z.H.T.D., and Tort, A.B.L. (2018). Characterizing Speed Cells in the Rat Hippocampus. Cell Rep. *25*, 1872–1884.e4.

Gustafson, N.J., and Daw, N.D. (2011). Grid Cells, Place Cells, and Geodesic Generalization for Spatial Reinforcement Learning. PLoS Comput. Biol. *7*, e1002235–e1002235.

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.

Hales, J.B., Schlesiger, M.I., Leutgeb, J.K., Squire, L.R., Leutgeb, S., and Clark, R.E. (2014). Medial Entorhinal Cortex Lesions Only Partially Disrupt Hippocampal Place Cells and Hippocampus-Dependent Place Memory. Cell Rep. 1–9.

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.

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.

Høydal, Ø.A., Skytøen, E.R., Andersson, S.O., Moser, M.-B., and Moser, E.I. (2019). Object-vector coding in the medial entorhinal cortex. Nature.

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.

Jacob, P.-Y., Gordillo-Salas, M., Facchini, J., Poucet, B., Save, E., and Sargolini, F. (2017). Medial entorhinal cortex and medial septum contribute to self-motion-based linear distance estimation. Brain Struct. Funct. *0*, 0–0.

Justus, D., Dalügge, D., Bothe, S., Fuhrmann, F., Hannes, C., Kaneko, H., Friedrichs, D., Sosulina, L., Schwarz, I., Elliott, D.A., et al. (2016). Glutamatergic synaptic integration of locomotion speed via septoentorhinal projections. Nature Publishing Group.

Kingma, D.P., and Ba, J. (2014). Adam: A Method for Stochastic Optimization.

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*.

Kitamura, T., Sun, C., Martin, J., Kitch, L.J., Schnitzer, M.J., and Tonegawa, S. (2015). Entorhinal Cortical Ocean Cells Encode Specific Contexts and Drive Context-Specific Fear Memory. Neuron *87*, 1317–1331.

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*.

Kostrikov, I. (2018). Pytorch implementations of reinforcement learning algorithms.

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

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

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

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.

Lever, C., Burton, S., Jeewajee, A., O'Keefe, J., and Burgess, N. (2009). Boundary vector cells in the subiculum of the hippocampal formation. J. Neurosci. *29*, 9771–9777.

Liu, J.K., and Buonomano, D.V. (2009). Embedding multiple trajectories in simulated recurrent neural networks in a self-organizing manner. J. Neurosci. *29*, 13172–13181.

Major, G., and Tank, D. (2004). Persistent neural activity: prevalence and mechanisms. Curr. Opin. Neurobiol. *14*, 675–684.

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.

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.

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.

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.

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.

Moser, E.I., Moser, M.-B., and McNaughton, B.L. (2017). Spatial representation in the hippocampal formation: a history. Nat. Neurosci. *20*, 1448–1464.

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

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.

O'keefe, J., and Nadel, L. (1978). The hippocampus as a cognitive map (Oxford: Clarendon Press).

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.

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.

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, H. Wallach, H. Larochelle, A. Beygelzimer, F. d\textquotesingle Alché-Buc, E. Fox, and R. Garnett, eds. (Curran Associates, Inc.), pp. 8026–8037.

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.

Pérez-Escobar, J.A., Kornienko, O., Latuske, P., Kohler, L., and Allen, K. (2016). Visual landmarks sharpen grid cell metric and confer context specificity to neurons of the medial entorhinal cortex. Elife *5*, 1–21.

Qin, H., Fu, L., Hu, B., Liao, X., Lu, J., He, W., Liang, S., Zhang, K., Li, R., Yao, J., et al. (2018). A Visual-Cue-Dependent Memory Circuit for Place Navigation. Neuron *99*, 47–55.e4.

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.

Sauvage, M.M., Beer, Z., Ekovich, M., Ho, L., and Eichenbaum, H. (2010). The caudal medial

entorhinal cortex: a selective role in recollection-based recognition memory. J. Neurosci. *30*, 15695–15699.

Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. (2017). Proximal Policy Optimization Algorithms.

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.

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 multichannel electrophysiology. J. Neural Eng. *14*, 045003.

Simonnet, J., and Fricker, D. (2018). Cellular components and circuitry of the presubiculum and its functional role in the head direction system. Cell Tissue Res. *373*, 541–556.

Solstad, T., Boccara, C.N., Kropff, E., Moser, M.-B., and Moser, E.I. (2008). Representation of geometric borders in the entorhinal cortex. Science *322*, 1865–1868.

Sorscher, B., Mel, G., Ganguli, S., and Ocko, S. (2019). A unified theory for the origin of grid cells through the lens of pattern formation. In Advances in Neural Information Processing Systems 32, H. Wallach, H. Larochelle, A. Beygelzimer, F. d\textquotesingle Alché-Buc, E. Fox, and R. Garnett, eds. (Curran Associates, Inc.), pp. 10003–10013.

Stachenfeld, K.L., Botvinick, M.M., and Gershman, S.J. (2017). The hippocampus as a predictive map. Nat. Neurosci. *20*, 1643–1653.

Steffenach, H.A., Witter, M., Moser, M.B., and Moser, E.I. (2005). Spatial memory in the rat requires the dorsolateral band of the entorhinal cortex. Neuron *45*, 301–313.

Stemmler, M., Mathis, A., and Herz, A.V.M. (2015). Connecting multiple spatial scales to decode the population activity of grid cells. Science Advances *1*, e1500816–e1500816.

Suh, J., Rivest, A.J., Nakashiba, T., Tominaga, T., and Tonegawa, S. (2011). Entorhinal Cortex Layer III Input to the Hippocampus Is Crucial for Temporal Association Memory. Science *1415*.

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.

Sussillo, D., and Abbott, L.F. (2009). Generating coherent patterns of activity from chaotic neural networks. Neuron *63*, 544–557.

Sussillo, D., Churchland, M.M., Kaufman, M.T., and Shenoy, K.V. (2015). A neural network that finds a naturalistic solution for the production of muscle activity. Nat. Neurosci. *18*, 1025–1033.

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.

Taube, J.S., Muller, R.U., and Ranck, J.B. (1990). Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. J. Neurosci. *10*, 420–435.

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.

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.

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.

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.

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.

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.