Task-anchored grid cell firing is selectively associated with successful path integration-dependent behaviour

Harry Clark[1], Matthew F. Nolan[1]*

[1] Centre for Brain Discovery Sciences, Simons Initiative for the Developing Brain, Hugh Robson Building, University of Edinburgh, Edinburgh, United Kingdom.

* Corresponding author, email: mattnolan@ed.ac.uk

Abstract

Grid firing fields have been proposed as a neural substrate for spatial localisation in general or for path integration in particular. While there are theoretical arguments to support each possibility, it has been challenging to directly test whether and when grid cells contribute to behaviours. Here, we investigate firing of grid and non-grid cells in the medial entorhinal cortex during a location memory task. We find that grid firing can either be anchored to the task environment, or can encode distance travelled independently of the task reference frame. Anchoring varied between and within sessions, while spatial firing of non-grid cells was either coherent with the grid population, or was stably anchored to the task environment. We took advantage of the variability in task-anchoring to evaluate whether and when encoding of location by grid cells might contribute to behaviour. We find that when reward location is indicated by a visual cue performance is similar regardless of whether grid cells are task-anchored or not, arguing against a role for grid representations when location cues are available. By contrast, in the absence of the visual cue, performance was enhanced when grid cells were anchored to the task environment. Our results suggest that anchoring of grid cells to task reference frames selectively enhances performance when path integration is required.

Introduction

The ability to generate and manipulate internal representations of the sensory world is fundamental to cognitive functions of the brain. Grid representations generated by neurons in the medial entorhinal cortex (MEC) are thought to be critical for spatial behaviours and other cognitive functions that require structured representations (Moser et al., 2008; Whittington et al., 2022). However, the range of behaviours to which grid cells contribute is unclear (Ginosar et al., 2023). On the one hand, theoretical arguments that grid cell populations can generate high capacity codes imply that they could in

principle contribute to all spatial behaviours (Fiete et al., 2008; Mathis et al., 2012; Sreenivasan and Fiete, 2011). On the other hand, if the behavioural importance of grid cells follows from their hypothesised ability to generate position representations by integrating self-motion signals (McNaughton et al., 2006), then their behavioural roles may be restricted to tasks that involve path integration strategies.

Experiments that have manipulated circuits containing grid cells support the idea that they contribute to spatial behaviours. Lesions of the MEC disrupt place representations in the hippocampus and impair performance in spatial memory tasks (Brun et al., 2008; Hales et al., 2021, 2018, 2014; Miao et al., 2015; Schlesiger et al., 2015; Steffenach et al., 2005). These impairments are often incomplete in that hippocampal place representations remain, although they are less stable, and some spatial tasks can still be solved. More selective genetic manipulations support similar conclusions. Deletion during postnatal development of NMDA receptors from neurons in the MEC and nearby structures reduces the number of detected grid cells in these areas while having less effect on other spatial cell types (Gil et al., 2017). This manipulation impairs path integration without affecting other spatial behaviours (Gil et al., 2017). In contrast, targeted inactivation of stellate cells in layer 2 of the MEC, which are thought to be a major grid cell population (Gu et al., 2018; Rowland et al., 2018), impairs learning of both path integration-dependent behaviours and cue-based navigation more generally (Qin et al., 2018; Tennant et al., 2018). However, for all of these manipulations it is difficult to establish whether impairments result from deficits in grid firing per se, or from other alterations in the circuit and its potential for plasticity. It is also unclear whether functions that in these experiments appear resistant to perturbations of the MEC could have been restored by adaptive compensatory changes.

These challenges are common to efforts to test hypothesised functional roles for neural codes using perturbation strategies. A complementary approach is to take advantage of variability in the expression of behaviours and candidate neural representations. Thus, hypothesised functions for neural codes can be corroborated by correlations with behavioural outcomes, while dissociations between representations and behaviour may rule out hypothesised roles for a given code. In the case of grid codes, the idea that they provide a general purpose spatial code predicts that they are always available and are stably anchored to the external environment (Figure 1A). This notion is challenged by observations that in circular track environments, grid firing patterns are maintained but are no longer anchored to the environment (Jacob et al., 2019). In this case, the grid representations are informative about distance travelled but not about absolute position (Figure 1B). Spatial representations within populations of MEC neurons that include grid cells can also be unstable with the network spontaneously remapping between different representations of location (Low et al., 2021). Thus, it appears that grid representations are not necessarily stably anchored to the external world, but it is unclear whether this instability impacts performance of spatial tasks. Here, we asked whether similar instability of grid activity manifests in a goal directed task, and if so whether it can dissociate proposed behavioural roles for grid firing patterns.

We address this question by investigating grid and non-grid firing during a task in which mice learn the location of a reward on a virtual track (Figure 1C)(Tennant et al., 2022, 2018). On cued trials, the mice

received rewards for stopping within a 'reward zone' that was marked by a distinct visual cue. In contrast to cue-rich virtual environments often used to study grid cells (e.g. (Campbell et al., 2018; Domnisoru et al., 2013)), this is the only spatially localised cue available after initiation of a trial. Removal of this cue enables testing of whether the mice are able to use a path integration strategy, in which case they should continue to selectively stop in the reward zone, or a cue-based strategy in which case they should no longer stop selectively at the reward zone location (Figure 1C)(see also (Tennant et al., 2022, 2018). We find that during the task, grid cells can either be anchored to the track reference frame ('task-anchored'), or can maintain a periodic firing pattern independent of the track reference frame ('task-independent'). Adoption of these anchoring modes varied both between and within recording sessions. On trials when the reward zone cue was visible, adoption of the task-anchored representation did not predict task performance. By contrast, when the reward zone cue was absent, task-anchored grid firing was associated with successful localisation of the reward zone. Thus, our results suggest that task anchoring of the grid cell network selectively enhances performance of behaviours that require path integration.

Results

We recorded from neurons in the MEC of 9 wild-type mice exploring an open arena and then performing a location memory task in a virtual linear track environment. By comparing the hexagonal symmetry of spatial autocorrelograms of neural activity in the open arena with corresponding shuffled data we identified 103/1881 neurons as grid cells (11.4 ± 17.4 grid cells / mouse, range 0.4 % to 9.5 %, Data S2) (see Methods for classification procedures). These grid cells had field sizes of 7.5 ± 3.3 cm and grid spacing of 72.5 ± 13.7 cm and were found in dorsomedial parts of the MEC. Until indicated otherwise, we report analyses of neurons across all trials of the location memory task regardless of whether the reward zone is indicated by the cue or whether mice stop at the rewarded location.

Grid cells exhibit either task-anchored or task-independent firing

A priori, we envisaged two scenarios for activity of grid cells during the location memory task. Given well established spatial firing of grid cells in open arenas we might expect that grid cell activity is anchored to the task reference frame (Figure 1A, D). Alternatively, given distance encoding but location-independent firing of grid cells in circular tracks (Jacob et al., 2019), we might expect the activity of grid cells to be periodic but independent from the task reference frame (Figure 1B, E).

To distinguish these scenarios we estimated the periodicity of each neuron's activity as a function of distance moved by the mouse using the Lomb-Scargle method (Lomb, 1976; Scargle, 1982; VanderPlas, 2018) (Figure 1, Supplements 1-3). This approach yields periodograms indicating power as a function of oscillation frequency and an associated estimate of the false alarm probability. We validated these estimates for detection of periodicity associated with grid firing using synthetic and

shuffled data (Figure 1, Supplements 4-5). According to the task-anchored firing scheme, firing fields should occur at the same positions on each trial (Figure 1A). In this case, because the virtual track repeated every 200 cm, significant peaks in the periodogram should occur at integer multiples of the spatial frequency of the repeating track (Figure 1D, Figure 1, Figure Supplement 3). By contrast, in the task-independent firing scheme, grid representations do not anchor to the track, but maintain firing fields that are periodic with respect to distance run (Figure 1B). In this case, significant peaks in the periodogram would reflect the distance between firing fields repeating independently of the task reference frame (Figure 1E, Figure 1, Figure Supplement 3). A further possibility is that in the location memory task the activity of grid cells is no longer periodic. In this case, peaks in the periodogram above the false alarm threshold should be absent.

We initially tested these predictions for neural activity across a complete behavioural session (n = 103 grid cells, N = 61 sessions, 233 ± 135 trials / session)(Figure 2). We found that 68 of 103 grid cells had peaks in their periodograms within 5 % of an integer multiple of the spatial frequency of the track repetition, consistent with their activity being anchored to the reference frame provided by the virtual track (Figure 2A, D-E). We refer to these neurons as showing 'session-level task-anchored' grid firing. By contrast, 32 of 103 grid cells had periodograms with peaks at frequencies outside 5 % of an integer multiple of the track length, indicating that their activity was not coupled to the task reference frame (Figure 2B, D-E). We refer to these neurons as showing 'session-level task-independent' grid firing. Grid cells with periodograms lacking peaks above the false alarm threshold, which we will refer to as aperiodic grid cells, were rare (3/103)(Figure 2C-E). When we carried out similar analyses for non-grid cells we found that the proportion of neurons with task-anchored firing was similar, but task-independent periodic firing was rarer, while aperiodic firing neurons were more common (Figure 2E).

To validate the periodogram-based classification we calculated the mean firing of grid cells as a function of track position. Consistent with their classification, task-anchored grid cells showed clear firing rate peaks associated with specific track locations (Figure 2A). By contrast, task-independent and aperiodic grid cells lacked clear firing rate peaks (Figure 2B-C). These differences manifest as substantially higher spatial information scores, but similar average firing rates, for grid cells with session-level task-anchored firing compared with task-independent firing (Figure 2F). Features of the periodogram such as peak power and peak width did not differ between task-anchored grid cells and task-independent grid cells (Figure 2F). The aperiodic grid cells had much lower mean firing rates than the task-anchored grid or task-independent grid cells suggesting that their apparent lack of periodicity reflects inactivity (Figure 2F).

These data suggest that during the location memory task we consider here, grid firing can either be anchored to the track and therefore be directly informative about position relative to the task reference frame, or be independent of the track and therefore may only be directly informative about distance travelled within the behavioural reference frame. Anchoring is consistent with previous reports of grid cell activity on virtual and real world linear tracks (Domnisoru et al., 2013), while task independence is

consistent with encoding of distance but not position by grid cells in real world circular tracks (Jacob et al., 2019).

Grid cells switch between task-anchored and task-independent firing

We next asked if the mode adopted by the grid cells at the level of a whole session, either task-anchored or task-independent, was stable across individual trials within the session (Figure 3A, left), or if cells could switch mode (Figure 3A, right). Visual inspection of firing rate heat maps indicated that for some grid cells their firing pattern was stable across most trials within a session (Figures 3B-C). However, for many grid cells there appeared to be clear changes in anchoring within a session (Figure 3D-E). These switches could not be explained by variation between trials in the availability of cues or rewards, as these were interleaved in blocks that repeated throughout a session (see Methods), whereas periods in which grid cell activity was in a given mode extended across the repeating blocks (e.g. Figures 3D,E, 4A, 5E,F).

To quantify switching between firing modes we evaluated rolling periodograms across each session. We classified each periodogram window as task-anchored, if the periodogram peaks occurred at integer multiples of the spatial frequency at which the track repeats, as task-independent if there were periodogram peaks at other spatial frequencies, or as aperiodic if there were no peaks above the false alarm threshold (see Methods and Figure 3, Figure Supplements 1-3). For 26.2% of grid cells (27/103) the classification was consistently (> 85% of trials) task-anchored or task-independent (Figure 3F), whereas for 73.8% of grid cells (76/103) no single coding scheme accounted for more than 85% of trials (Figure 3F). Grid cells identified as having task-anchored firing in our initial session-level analysis (Figure 2) were biassed towards high proportions of task-anchored trials (Figure 3B, F), whereas grid cells identified as task-independent at the session-level showed a bias towards high proportions of task-independent trials (Figure 3C, F), although in both groups many neurons showed variation between trials indicating that session-level analyses may obscure dynamic changes in task-anchoring. Differences in spatial information and periodogram properties between grid cells classified as task-anchored or task-independent based on their session-level firing patterns (Figure 2F) were nevertheless maintained when we instead compared neurons with firing mode that was consistent across trials within a session (Figure 3, Figure Supplement 4). For non-grid cells, consistent with session-level analyses, we again found that a sub-population showed task anchored firing at a trial-level, while task-independent periodic firing was substantially less common than in the grid cell population (Figure 3, Figure Supplement 5).

To validate the window-based analyses we compared mean firing rate as a function of track position between trials classified as task-anchored and those classified as task-independent. Consistent with the classification scheme correctly separating task-anchored from task-independent activity, spatially localised firing rate fields were present during task-anchored firing but were reduced or abolished when firing was task-independent (e.g. see rate maps Figures 3D-E). This difference was associated with substantially lower spatial information scores for trials with task-independent firing, while mean

firing rates were similar (Figure 3G). As a further test, we compared the durations of periods of activity in a given state to durations generated from shuffled data sets. If classification was by chance then the distributions of durations should be similar. By contrast, among all grid cells periods of task-anchored or task-independent firing extended in blocks across multiple trials, with the distribution of block lengths substantially different to that generated by shuffled data (Figure 3, Figure Supplement 6)(P = 0.0003, KS = 0.066; DF = 1765, 2359; Kolmogorov-Smirnov test). Variation between grid cells in their firing mode also could not be explained by the relative proportion of cued trials in a session (Figure 3, Figure Supplement 7).

Given that populations of grid cells show coordinated dynamics that are consistent with their forming networks with continuous attractor dynamics (Barry et al., 2007; Waaga et al., 2021, 2021; Yoon et al., 2013), changes in anchoring should be coherent across grid cells and any non-grid cells that coordinate with the grid network. To test this, we compared activity patterns of simultaneously recorded grid and non-grid cells (Figure 4). Grid cells typically transitioned between task-anchored and task-independent firing modes at the same time (for examples see Figure 4A-C and Figure 5E-F). This manifested as strong agreement between grid cells in their session-long order of task-anchored and task-independent epochs (Figure 4D-E). By contrast, the activity of non-grid cells was more diverse. Some non-grid cells transitioned between task-anchored and task-independent modes at the same time as simultaneously recorded grid cells, whereas others remained stably anchored throughout the recording session (Figure 4A-C and Figure 5E-F). As a result, session-long agreement scores for pairs of grid and non-grid cells were more variable and on average lower than for pairs of grid cells (Figure 4D-E). Non-grid cells also differed from the grid cell population in that their activity maintained similar spatial information between trials when grid cells were task-anchored versus when they were task-independent, whereas for grid cells task-independence was associated with reduced spatial information (Figure 4F).

Together, these data indicate that grid cells can switch between task-anchored and task-independent firing modes within a behavioural session. This switching happens coherently across grid cell populations, which is consistent with grid cells forming networks with continuous attractor dynamics (Barry et al., 2007; Gardner et al., 2021; Waaga et al., 2021; Yoon et al., 2013). Our data also suggest that non-grid cells within the MEC form multiple populations, with some having activity that is coherent with the grid cell network, whereas others do not show task-independent periodic firing but instead maintain stable spatial representations independently from grid cells.

Task-anchored coding by grid cells is selectively associated with successful path integration-dependent reward localisation

Our analyses indicate that grid cells exhibit either task-anchored or task-independent firing, and that their activity can switch between these modes within a recording session. Since task-anchored firing fields could be read out directly to estimate track location, but location may only be inferred indirectly from task-independent activity, we reasoned that the presence or absence of task-anchoring could be

used to assess whether grid firing contributes to the ongoing behaviour. Because in grid networks the activity of an individual neuron is informative about the network state as a whole (Fiete et al., 2008; Gardner et al., 2021; Waaga et al., 2021), then in principle activity of any grid cell is indicative of whether the grid network as a whole is in a task-anchored or task-independent mode (See also Figure 4). Thus, if anchoring of grid representations to the task environment is critical for localisation of the reward, then task-anchored coding of individual neurons should predict successful trials. On the other hand, if behavioural performance is maintained when the grid representation is task-independent, then it is unlikely that anchored grid representations are necessary for reward localisation.

To distinguish these possibilities we took advantage of variation in behavioural outcomes, which were such that mice either stopped correctly in the reward zone ('hit' trials), slowed down on approach to the reward zone but did not stop ('try' trials) or maintained a high running speed across the reward zone ('run' trials)(Figure 5A-B and Figure 5, Figure Supplements 1-2). We separately evaluated outcomes from trials in which the reward zone cue was visible and a reward available ('beaconed trials'), trials in which the reward zone cue was omitted and a reward was available ('non-beaconed trials'), and trials in which the cue and the reward were both omitted ('probe trials'). Because trials of each type were interleaved into blocks that were repeated across a session (see Methods), while periods of task-anchored and task-independent activity were typically maintained for blocks of many trials, differences in behavioural outcomes could not be explained by association of the grid anchoring mode with particular trials types.

We first compared stopping behaviour when grid cells showed task-anchored firing that was stable within a session to when they showed task-independent firing that was stable within a session. On cued trials the spatial organisation of stopping behaviour (Figures 5C vs 5D, 6A), and the proportion of hit trials (Figure 6B) was similar for both groups. By contrast, on non-beaconed and probe trials stopping behaviour was clearly spatial when grid activity was task-anchored, but spatial organisation was largely absent when grid activity was task-independent (Figures 5C vs 5D, 6A), while the proportion of hit trials was substantially reduced (Figure 6B). These observations are consistent with grid representations being required for path integration-dependent but not cued localisation of the reward zone. This session-level analysis has the advantage that it focuses on large blocks of time (sessions) during which the mode of grid firing was stable, but the disadvantage that it excludes many sessions in which the mode of grid firing switches between task-anchored and task-independent.

We addressed this by using additional trial-level comparisons to evaluate all behavioural sessions, including those when the grid mode was unstable. On cued trials the spatial organisation of stopping behaviour (Figures 5E-F, 6C), and the proportion of hit trials (Figure 6D), was similar irrespective of whether grid cell fields were task-anchored or task-independent. On trials in which the reward zone cue was hidden, the relationship between firing and behavioural outcome was more complex. In many sessions, localisation of the reward occurred almost exclusively when cell firing was task-anchored and not when it was task-independent (Figure 5E-F). In a few sessions, we observed spatial stopping behaviour comparable to cued trials, even when grid firing was predominantly task-independent (Figure 5, Figure Supplement 3). In one session, we also observed grid cells shifting phase between

task-anchored epochs (Figure 5E) similar to findings in Low et al. (2022), although this did not appear to alter task performance. On average, spatial stopping behaviour was reduced on non-beaconed and probe trials during which grid firing was task-independent (Figure 6C), and the proportion of successful trials was reduced compared with when grid firing was task-anchored (Figure 6D). The differences in outcomes between task anchored versus task independent trials were not associated with differences in running speed profiles (Figure 6, Figure Supplement 1) indicating they were not a consequence of a difference in motor behaviour.

Additional analyses of relationships between firing mode and trial outcomes were consistent with these observations. Thus, task-anchored firing occurred on a greater proportion of non-beaconed and probe-trials that were hits compared with task-independent firing, but on a smaller proportion of run trials, in which mice ran through the reward zone without slowing down (Figure 6E). In addition, the odds ratio for receipt of reward on task-anchored versus task-independent trials was close to 1 when cues were available, but was much larger in the absence of cues (Figure 6F). Analyses using an alternative classification method, based on template matching of the task-anchored firing rate profile (Figure 6, Figure Supplement 3A-B), also indicated a similar behavioural performance on cued trials, but impaired performance on uncued trials when grid cells were in a task-independent mode (Figure 6, Figure Supplement 3C-F).

Together, these analyses demonstrate that anchoring of grid codes to track position is not necessary to successfully obtain rewards when visible cues are present. By contrast, when visual cues that indicate the reward zone are absent, task-anchored coding appears to promote successful localisation.

Discussion

Our results support specific roles for task-anchored grid representations in path integration-dependent behaviours while arguing against the idea that grid codes provide a representation of location that is used more generally. Thus, we find that grid representations can either be anchored to position in a task environment, or can provide an environment-independent distance metric (Figure 2), that grid cells can switch between these operating modes within a behavioural session (Figures 3-4), and anchoring of grid firing fields to location is not required for cued reward localisation, but appears to promote path integration-dependent reward localisation (Figures 5-6). By contrast, while some non-grid cells had activity that switched modes coherently with grid cells, many non-grid cells did not show task-independent periodic activity (Figure 3, Figure Supplement 5, Figures 4-5, Figure 5, Figure Supplement 2), suggesting that the MEC may implement multiple, parallel spatial computations.

Task-anchoring of grid cells varies within and between behavioural sessions

A standard view of grid firing is that it provides an 'always on' high capacity representation of current location. In contrast, our results demonstrate that grid representations are not necessarily anchored to the behaviourally relevant environmental reference frame. This extends previous observations that on circular tracks grid cells read out path integrated distance rather than absolute position (Jacob et al., 2019). Our results show that both modes of grid representation can be observed during the same behavioural task, and that the grid network can switch between these operating modes within a session. Consistent with the reported continuous attractor dynamics of grid networks, we find that simultaneously recorded grid cells switch modes coherently (Figures 4-5), although without recording from all grid cells simultaneously we can not rule out the possibility that switching reflects a subset of grid cells that become disconnected from the wider network. The mode switching we observe here appears to differ from previously reported spontaneous remapping of MEC network states (Low et al., 2021) as the previously reported phenomenon was independent of any task contingencies and involved remapping between different task-anchored representations. By contrast, we show here switching between task-anchored representations that may be directly useful for solving the task at hand, and task-independent representations that appear unlikely to contribute to solving the task.

Grid cells are a relatively small proportion of the neurons in the MEC and an important question is the extent to which switches between task-anchored and task-independent modes extend to other cell types. Among non-grid cells, a sub-population appeared to switch between task-anchored and task-independent modes coherently with grid cells. However, task-independent periodic activity was relatively rare among non-grid cells (Figure 3 Figure Supplement 5). A possible explanation for this dissociation is that the MEC might contain multiple sub-networks, with populations of grid cells and coherent non-grid cells sharing continuous attractor dynamics, while other populations of non-grid cells operate independently from the grid network. For example, non-grid cells may be differentially influenced by visual cues (Casali et al., 2018; Kinkhabwala et al., 2020) or could generate location representations through ramping profiles rather than discrete firing fields (Tennant et al., 2022). The presence in the MEC of functionally distinct grid and non-grid networks might also account for why manipulations that perturb grid cells selectively disrupt path integration-dependent behaviours (Gil et al., 2017), whereas manipulations that target layer 2 stellate cells, which include grid and non-grid cells, impair path integration-dependent and cued behaviours (Qin et al., 2018; Tennant et al., 2018).

The causes of switching between task-anchored and task-independent representations may be an important focus for future investigation. It appears unlikely that task-independent coding by grid cells result from failure of upstream circuits to generate appropriate visual representations, for example through a shift in visual attention, as mice performed well on the visually cued version of the location memory task when grid cells were task-independent. A perhaps more promising hypothesis for future investigation is that switching reflects modulatory pathways reducing the impact of visual inputs to the grid system, possibly reflecting top down control mechanisms, shifts in brain state, or uncertainty about whether visual or motor signals indicate the correct environment.

Spatial roles of grid cells may be specific to path integration-dependent behaviours

Our finding that cued identification of a reward location is similar when grid cells are task-anchored or task-independent, suggests that grid representations are not required for cued recall of locations. Thus, when both visual and grid inputs are both available to downstream decision making circuits the grid input appears not to be used. If it was, then the inconsistent positional signals from the task-independent grid codes would impair performance. An implication of this result is that cue rich tracks often used to investigate grid activity patterns may not engage behaviours that require anchored grid firing. Several observations suggest the MEC could nevertheless be required for recall of cued locations. First, inactivation of stellate cells in layer 2 of the MEC causes deficits in the task we use here and in other visually cued tasks (Qin et al., 2018; Tennant et al., 2018). Second, other spatial tasks that involve selections between cued locations appear to require the MEC (e.g. (Gaskin and White, 2013, 2010)). In this case, spatial representations used for the task could be encoded by other functional cell types, for example neurons that encode location through border (Solstad et al., 2008) or ramping firing fields (Tennant et al., 2022), or through cue-responsive cells (Casali et al., 2018; Keene et al., 2016; Kinkhabwala et al., 2020).

On trials in which the reward zone cue is absent, efficient reward localisation relies on path integration from the start of the track (Tennant et al., 2018). We found that on these trials, task-anchored grid firing was associated with a spatially localised stopping strategy and a higher proportion of successful trials compared with when grid firing was task-independent. This corroborates key predictions of hypothesised roles for grid cells in path integration (McNaughton et al., 2006). It is conceivable that behaviourally relevant computations are implemented elsewhere in the brain with grid anchoring to the track an indirect consequence, but explanations of this kind are hard to reconcile with evidence that stellate cells in the MEC are required for the task we used here (Tennant et al., 2018), or with evidence for specific roles of grid firing in path integration based on genetic manipulations that abolish grid firing without affecting other functional cell types (Gil et al., 2017). Nevertheless, our finding of residual localisation performance on task-independent trials suggests that additional neural mechanisms may also support path integration-dependent behaviour. This could reflect additional mechanisms for the implementation of path integration, for example through ramping activity (cf.(Tennant et al., 2022)). Alternatively, mice could in principle estimate track location with a system that utilises information about distance travelled obtained from task-independent grid representations. If multiple mechanisms support path integration then it will be important to establish when each contributes. For example, because grid representations are available on the first entry to an environment they may be important for behaviour in newly experienced locations, whereas for familiar locations complementary representational strategies that emerge with learning may be sufficient.

Ideas and speculation

Our results point to a specific role for grid firing in path integration-dependent behaviour, and demonstrate the importance of anchoring of grid representations to task environments. One

implication of our results is that rather than being an 'always on' tracking system, grid cell networks vary in their engagement with the environment. This may reflect control of the grid network by attentional or other top down mechanisms. Alternatively, as the CA1 region of the hippocampus provides a major input to the MEC, instability of grid anchoring could be an indirect consequence of mechanisms that control the structure and stability of place cell maps (Krishnan et al., 2022; Pettit et al., 2022). In either case, our results motivate a focus on grid cell activity in tasks that require path integration, while an implication of our finding for investigations of grid cell activity using cue-rich environments is that in these experimental settings grid cells may not be influencing behavioural outcomes. Our results also offer a new perspective on interindividual differences in path integration by humans (e.g. (Chrastil et al., 2017; Lakshminarasimhan et al., 2018; Petzschner and Glasauer, 2011)). Thus, rather than resulting from variation in path integration per se, differences between individuals could instead result from variation in the anchoring of grid representations underlying path integration. This could be important as a potential mechanism for deficits in spatial localisation associated with neurological and neurodevelopmental disorders (Kunz et al., 2015; Newton et al., 2023; Noel et al., 2020).

Materials and Methods

All experiments were carried out under a UK Home Office project licence, were approved by the Animal Welfare and Ethical Review Board of the University of Edinburgh College of Medicine and Veterinary Medicine, and conformed with the UK Animals (Scientific Procedures) Act 1986 and the European Directive 86/609/EEC. Three of the nine mice used here were also part of a previous study (Tennant et al., 2022).

Microdrive fabrication and surgical procedures

Microdrive fabrication and surgical procedures were similar to our previous work (Gerlei et al., 2020; Tennant et al., 2022). Microdrives containing 4 tetrodes were built by threading 90% platinum, 10% iridium tetrode wires (18 µm HML-coated, Neuralynx) to an EIB-16 board (Neuralynx) via an inner cannula (21 gauge 9 mm long). The board was covered in epoxy and a poor lady frame (Axona) cemented to the side. An outer cannula (17 gauge 7 mm), placed around the inner cannula, was secured temporarily using vaseline, allowing it to be lowered during the surgery. Tetrodes were trimmed to ~3 mm using ceramic scissors (Science Tools, Germany) and gold plated (Non-cyanide gold plating solution, Neuralynx) to give an impedance between 150 and 200 kΩ at 1kHz.

Before surgery, tips of the tetrodes were washed with ethanol and then sterile saline. Anaesthesia was induced using 5 % isoflurane / 95 % oxygen, and sustained at 1-2 % isoflurane / 98-99 % oxygen. After exposing the surface of the skull a RIVETS head-post (Osborne and Dudman, 2014) was attached to the skull with UV curing resin cement (RelyX™ Unicem, 3M). For electrical grounding, two M1 x 4 mm screws (AccuGroup SFE-M1-4-A2) were implanted through small

craniotomies drilled on the left hemisphere ~3.4 mm lateral, and ~1 mm rostral relative to Bregma and the centre of the intraparietal plate respectively. The microdrive was attached to a stereotaxic frame via an Omnetics to Mill-Max adaptor (Axona, HSADPT-NN1) and the tetrodes lowered 1.2 - 1.4 mm into the right hemisphere of the brain, beginning at 3.4 mm lateral from Bregma and along the lambdoid suture, and at an angle of -15° in the posterior direction. The outer cannula was lowered and sealed with sterile Vaseline, and the implant fixed to the skull with UV curing resin. After the resin hardened, the grounding wires were wrapped around the grounding screws and fixed with silver paint (RS components 101-5621). The grounding screws were covered with resin and any holes filled with dental cement (Simplex Rapid). After the surgery, mice recovered for ~20 minutes on a heat mat, had unlimited access to Vetergesic jelly (0.5 mg / kg of body weight buprenorphine in raspberry jelly) for 12 hours, and before proceeding were given a minimum of 2 days postoperative recovery.

Behavioural and electrophysiological recording

The behavioural set up, training procedures and recording approaches were similar to those described previously (Tennant et al., 2022, 2018). Mice were handled twice a day for 7 days following surgery. They were then habituated to the virtual reality setup for 10 - 20 minutes per day over two consecutive days. After each habituation session the mice were given access to soy milk to familiarise them with the reward and were given access to an open arena for 5 - 10 minutes of free exploration. From 4-5 days before starting training their access to food was restricted so that their body weight was ~85% of its baseline value, calculated from its weight prior to restriction and normalised to the expected daily growth for the animal's age.

Experimental days involved recording from mice in the open arena and then in the virtual location memory task. On a few days this order was reversed without apparent effects on the results obtained. Mice were collected from the holding room 30 - 60 minutes before recording, were handled for 5 - 10 minutes, weighed and placed for 10 - 20 minutes in a cage containing objects and a running wheel. Between recording sessions mice were placed back in the object-filled playground for 30 minutes. Tetrodes were typically lowered by 50 - 100 μm after each session. The open arena consisted of a metal box with a square floor area, removable metal walls, metal frame (Frame parts from Kanya UK, C01-1, C20-10, A33-12, B49-75, B48-75, A39-31, ALU3), and an A4-sized cue card in the middle of one of the metal walls. For the open field exploration session, mice were placed in the open arena while tethered via an ultrathin SPI cable and custom-build commutator and left unprompted for 30 minutes to freely move around. For the location memory task mice were trained to obtain rewards at a location on the virtual linear track. Mice were head-fixed using a RIVETS clamp (Ronal Tool Company, Inc) and ran on a cylindrical treadmill fitted with a rotary encoder (Pewatron). Virtual tracks, generated using Blender3D (blender.com) had length 200 cm, with a 60 cm track zone, a 20 cm reward zone, a second 60 cm track zone and a 60 cm black box to separate successive trials. The distance visible ahead of the mouse was 50 cm. The reward zone was either marked by distinct vertical green and black bars on “beaconed” trials, or was not marked by a visual cue at all on “non-beaconed” or “probe” trials. A feeding tube placed in front of the animal dispensed soy milk rewards (5 - 10 μl per reward) if

the mouse stopped in the reward zone, however was not dispensed on probe trials. A stop was registered in Blender3D if the speed of the mouse dropped below 4.7 cm/s. Speed was calculated on a rolling basis from the previous 100 ms at a rate of 60 Hz.

Trials were delivered in repeating blocks throughout a recording session. For example, 3 beaconed trials (B) followed by 2 non-beaconed trials (N) with the order repeating until the end of the session. To encourage learning and engagement in both beaconed and non-beaconed trials, the first day of training typically used a trial type ratio of 3 beaconed trials to 1 non-beaconed trials. As training progressed we then increased the proportion of non-beaconed trials up to a ratio of 1 beaconed trial to 4 non-beaconed trials. Examples of trial blocks include BBBBN, BBBN, BBN, BN, BNN, BNNN, BBNNN and BBNNNNNNN where each character indicates the trial type of each trial within a block. In some sessions we replaced single non-beaconed trials in trial blocks with a probe trial (P). Examples of trial blocks with probe trials include BBBBNBBBBP, BBBNBBBP, BBNBBP and BBNNNNNNNP. The ratio and order of trial type delivery was not found to affect the results obtained (Figure 3, Figure Supplement 7).

Electrophysiological signals were acquired using an Intan headstage connected via an SPI cable (Intan Technologies, RHD2000 6-ft (1.8 m) standard SPI interface cable) attached to an Open Ephys acquisition board. Signals were filtered (2.5 Hz - 7603.8 Hz). For the location memory task, behavioural variables including position, trial number and trial type were calculated in Blender3D at 60 Hz and sent via a data acquisition (DAQ) microcontroller (Arduino Due) to the OpenEphys acquisition board. In the open arena, motion and head-direction tracking used a camera (Logitech B525, 1280 x 720 pixels Webcam, RS components 795-0876) attached to the frame. Red and green polystyrene balls were attached to the sides of the headstage and were tracked using a custom script written in Bonsai (Lopes and Monteiro, 2021). Synchronisation of position and electrophysiology data used an LED attached to the side of the open arena in the field of view of the camera, with randomly timed trigger pulses sent to the LED via an Arduino board (Arduino Uno) and to the Open Ephys acquisition board via the I/O board.

Following experiments, tetrodes were localised using a microCT scanner (Data S1). Mice were anaesthetised with isoflurane and then a lethal dose of sodium pentobarbital (Euthatal, Meridal Animal Health, UK), were perfused with a mixture of PFA and glutaraldehyde, and the head with the microdrive still intact on the skull left in the same solution for 2 nights. All tissue and bone except that attached to the microdrive was removed before washing the brains in ddH_20 and incubating at 4° for two weeks in 2 % osmium tetroxide (2% OsO_4). Brains were then washed in ddH_20, dehydrated in ethanol and then embedded in resin. After the resin had cured the brains were imaged in a micro-CT scanner (Skyscan 1172, Bruker, Kontich, Belgium). 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. The scan dataset was reconstructed (CTan software, v1.13.5.1) and viewed with DataViewer

(Bruker). 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).

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; Gerlei et al., 2020). Recordings from the open field and virtual reality tasks were concatenated for spike sorting. Pre-processing steps converted Open Ephys files to mda format, filtered signals between 600 Hz - 6000 Hz and performed spatial whitening over all channels. Events were detected from peaks > 3 standard deviations above baseline and separated by at least 0.33 ms from other events on the same channel. The first 10 principal components of the detected waveforms were used as inputs to the ISOSPLIT algorithm. Cluster quality was evaluated using isolation, noise-overlap, and peak signal to noise ratio metrics (Chung et al., 2017). Units with firing rate > 0.2 Hz, isolation > 0.9, noise overlap < 0.05, and peak signal to noise ratio > 1 were used for further analysis. Downstream analyses were carried out using Python (version 3.8.1) and R versions: 4.2.3.

Analysis of neural activity in the open arena

For analysis of neural activity in the open arena, firing rate maps were calculated by binning spikes into 2.5 cm bins, dividing by the total time occupied in each bin and then smoothed with a Gaussian kernel. Autocorrelograms were calculated by sliding the rate map over all x and y bins and calculating a correlation score. Grid scores were defined as the difference between the minimum correlation coefficient for rate map autocorrelogram rotations of 60 and 120 degrees and the maximum correlation coefficient for autocorrelogram rotations of 30, 90 and 150 degrees (see (Sargolini et al., 2006)). Fields were detected in the autocorrelogram by converting it into a binary array using 20% of the maximal correlations as a threshold. If the binary array had more than 7 local maxima, a grid score was calculated. Correlations between the rotated autocorrelograms were then calculated using a ring containing the 6 local maxima closest to the centre of the binary array and excluding the maximum at the centre. The ring was detected based on the average distance of the 6 fields near the centre of the autocorrelogram (middle border = 1.25* average distance, outer border = 0.25 * average distance). To compute the spatial stability of cells in the open arena, the within-session spatial correlation was calculated by computing the Pearson correlation between the firing rate map from the first half session and the second half session. Bins that were not visited in both halves were excluded from the calculation. Neurons were classified as grid cells when their grid score and spatial stability score was in the 99th percentile of the same scores from 1000 shuffled datasets. Shuffled spike data was generated by drawing a single value from a uniform distribution between 20-580 seconds and adding this to the timestamp of each spike. Spike times that exceeded the recording duration were wrapped to the start of the session. Spike locations were recomputed from the shuffled spike times and spatial scores recalculated.

Analysis of behaviour during the location memory task

Plots of running speed as a function of location on the virtual track were generated by first binning speed into 1 cm location bins for each trial and then smoothing by convolution with a Gaussian filter (SD = 2 cm, SciPy Python package).

Trials were classified into hits and misses based on whether the mouse stopped (speed < 4.7 cm/s) within the reward zone. Miss trials were further split by comparing their average speeds in the reward zone to hit trials; the 95th percentile of speeds in the reward zone for hit trials was used to discriminate between try trials (< 95th percentile speed) and run trials (> 95th percentile speed; Figure 5, Figure Supplement 1). Trials in which the mouse's average speed outside of the reward zone was < 10 cm/s were left unclassified.

To compute stop density profiles (e.g. Figure 5C-F), stops were counted within 1 cm location bins and the counts were divided by the number of trials to obtain the number of stops per cm per trial. This was smoothed by convolution with a Gaussian filter (SD = 1 cm). To evaluate the stop density when aggregating sessions or for epochs of trials within a session (e.g. Figure 6A,C), the same procedure was applied to first generate a stop density profile for the trials of interest within a session. We then subtracted an average of density profiles calculated in the same way for shuffled data from the same trials but in which stop locations were randomly drawn from a uniform distribution of track locations. In this way, stop densities below zero in the subtracted profiles can be interpreted as below chance relative to their average shuffled distributions, and stop densities greater than zero can be interpreted as greater than chance. Aggregate stop density profiles were then generated by averaging the individual subtracted profiles. Where average stop density plots were shown for coding epochs (e.g. Figure 6C, Figure 6, Figure Supplement 2E), these were weighted both on the proportion of trials classified in a particular coding scheme (e.g. a code weight of 0.05 when 5% of trials were classified as task-anchored for a single grid cell) and the number of simultaneously recorded grid cells such that single sessions were weighted equally (e.g. a session weight of 0.2 (1/5) when five grid cells are simultaneously recorded in a session). Similar weighted averages were also applied to plots where the proportion of hit, try and run trials differed within a session in a similar vein (Figure 6E, Figure 6, Figure Supplement G).

Analysis of neural activity during the location memory task

Firing rate maps for each trial were generated by dividing the track into 1 cm bins, summing the number of spikes in each bin and dividing by the time the animal spent there. Firing rates were smoothed with a Gaussian filter (SD = 2 cm). Spatial information was calculated in bits per second as

$$\sum_{i=1}^{N} P_i \lambda_i log_2(\frac{\lambda_i}{\lambda})$$

where i indexes a position bin in the firing rate map, N is the number of bins, P_i is the occupancy probability, λ_i is the firing rate in the bin and λ is the mean firing rate. When spatial information scores were generated for epochs within a session, we took the same number of trials for each epoch (e.g.

when comparing task-anchored and task-independent epochs). This was done by randomly subsetting the epoch with the greater number of trials to match the number of trials of the epoch with the smaller number of trials.

To quantify the spatial periodicity of neural firing, the Lomb-Scargle method of least-squares spectral analysis (LSSA) as implemented by the Astropy Python module (The Astropy Collaboration et al., 2022) was used to generate a frequency spectrum in the spatial domain (Lomb, 1976; Scargle, 1982). A periodogram was computed every 10 cm with a sample length equal to 3 track lengths (600 cm). Track locations were normalised between 0 and 1 so that spatial frequencies corresponded to the number of oscillations per trial. Spatial frequencies > 5 were discarded from further analysis as no grid cells were found with grid spacings < 40 cm. Individual periodograms were combined to create an average periodogram across the session (Figure 1, Figure Supplement 1 for an illustrative example).

To distinguish spatially periodic firing from aperiodic firing (Figure 2), we compared the spectral peaks from a cell's averaged periodogram to a false alarm threshold. The threshold was calculated with a bootstrapping method that used 1000 shuffled instances of the cell's firing, with the aim of disrupting spatial periodicity while preserving any local temporal firing. This was inspired by the field shuffle procedure used for grid cells in 2 dimensional environments (Barry and Burgess, 2017). First, firing fields were identified by detecting the peaks and troughs in a smoothed version of the cell's firing rate map (convolution with a Gaussian filter, SD = 4 cm) with a minimum peak distance of 20 cm, smaller peaks were removed until all conditions of the detection were satisfied using the SciPy function scipy.signal.find_peaks. Fields were defined as the region between adjacent troughs. The field positions were reallocated in an unsmoothed rate map to random positions, preserving the spatial organisation of the field, while bins not attributed to a firing field filled the remaining gaps. The shuffled unsmoothed rate map was then smoothed by convolution with a Gaussian filter (SD = 2 cm) and the periodogram calculated. This was repeated 1000 times and the 99th percentile calculated from the distribution of shuffled peak powers used to create the false alarm threshold. Cells with a measured peak power below this threshold were classified as aperiodic (Figure 1, Figure Supplement 2 for an illustrative example).

To establish whether periodic cells had activity anchored to the track we calculated the difference between the peak spatial frequency of their periodogram and the nearest positive-integer value. We identified task-anchored periodic cells as those for which the difference was ≤ 0.05, and task independent period cells as those for which the difference was > 0.05. Using this metric yielded high prediction accuracies when comparing our classification to true labels when we simulated task-anchored or task-independent grid cells (Figure 1, Figure Supplement 5).

To classify activity within a recording session, we calculated average periodograms across a rolling window with a size of 200 individual periodograms, which equated to 10 trials (considering 20 samples per 200 cm track with 10 cm steps between periodogram increments). The peak of the

average periodogram from a single window was identified, and the peak power and the spatial frequency at which this occurred extracted. To determine if the peak reflected a periodic signal, it was compared with an adjusted false alarm threshold (see below). Windows with periodograms containing peaks above the threshold were then classified as task-anchored or task-independent as for session-level periodograms, while other windows were classified as aperiodic. To classify at the level of individual trials, midpoints of each rolling window were extracted and assigned to trial numbers by assessing which trial number the midpoint was closest to. For example, a rolling window with a midpoint of 10.4 would be assigned to trial number 10, while a midpoint of 10.6 (or 10.5 in border cases) would be assigned to trial number 11. All rolling windows with their respective classifications from a single trial would be pooled and counted. The classification with the most counts for that trial number would then be assigned to that trial (Figure 3, Figure Supplement 1 for an illustrative example). When assigning a common classification across a group of cells recorded simultaneously (e.g. Figure 4B), we used the mode of their classification, in cases where the group only contained two cells and there was not 100 % agreement between these two cells, one cell was selected at random to represent the common population code.

To select the optimal rolling window size we considered how the false alarm threshold changes as a function of the number of periodograms used to compute the average periodogram (Figure 3, Figure Supplements 1 and 2). As the number of samples within the rolling window increases, the false alarm threshold decreases (Figure 3, Figure Supplement 3). To account for this, an adjusted false alarm threshold was calculated using the first 200 periodograms from a field shuffle to compute an average periodogram. This was repeated for 1000 field shuffles and the 99th percentile of the peak powers used as the adjusted false alarm threshold. We also found the number of periodograms used to calculate the average periodogram greatly affected the prediction accuracy and bias. We opted for a rolling window size equal to 200 periodograms as this was found to achieve high accuracy with minimal bias (Figure 3, Figure Supplement 2).

To calculate the coding agreement between any two simultaneously recorded cells, the coding schemes were compared across the course of the session. The agreement score for a cell pair was equal to the percentage of trial classifications that agreed. In order to compare the measured agreement against chance for any cell pair, shuffled arrangements of trial classifications were generated for one of the two cells by splicing the trial classification raster where consecutive trials switch coding schemes and then reordering them at random (Figure 4D). This procedure was repeated 10 times for each cell pair and agreement scores calculated accordingly. Where shuffled agreement scores are visualised (Figure 4E), the average of the shuffled agreement scores are shown.

A potential weakness of the Lomb-Scargle method that we used for trial-level classification is that the window over which classification is made has limited resolution, while simulations on artificial data suggested that classification could be biassed towards task-independent firing depending on how

frequent transitions between coding schemes occur. We therefore implemented a second method for trial-level classification (Figure 6, Figure Supplement 2). With this method we first created an average firing rate profile of the task-anchored trials obtained from the Lomb-Scargle method. We then used this average as a template which we correlated with the firing rate profile of each trial. For analyses in Figure 6, Figure Supplement 2, trials with a correlation coefficient ≥ 0.5 were then classified as task-anchored positive (TA+) and task-anchored negative (TA-) otherwise. We note that when using this method we discarded units for which less than 15% of trials were originally classified as task-anchored, as for these units we were unable to generate templates of sufficient quality.

Simulation of firing during the location memory task

To evaluate classification of grid firing as task-anchored or task-independent we first simulated various cell types including grid cells, place cells, ramp cells and aperiodic cells (Figure 1, Figure Supplement 4). For each cell type, we simulated an agent moving with a constant velocity of 10 cm/s across 100 trials of a 200 cm long linear track and logged the locations visited with a sampling rate of 1000 Hz. For each cell-type, the probability of firing at any given location was defined by a probability density function (PDF) with a range of 0 to 1. The average firing rate was set by multiplying this normalised PDF by a scalar variable P_{max}(spike) which by default was set to 0.1. Firing events were then assigned to each sampled location based on the scaled PDF. Firing rate maps and subsequent periodograms were created as described above.

PDFs for grid cells were created by positioning Gaussians kernels at equidistant locations along the track, with kernel standard deviation equal to 0.1 multiplied by a given grid spacing between the kernels. To simulate task-anchored grid codes, the Gaussian kernels were positioned at the same track location on each trial, whereas to simulate task-independent grid codes the kernels were positioned independently from the track with distances equal to the grid spacing. To simulate field jitter, a displacement of the kernel position was drawn from another Gaussian distribution with mean = 0 cm and SD = 0 cm (for no jitter) or SD = 10 cm (for default jitter). A random jitter was drawn for each field and was used to shift fields accordingly. The PDF for the place cell example was made up of a singular Gaussian kernel (mean = 100 cm, SD = 10 cm) positioned at the centre of the track and was repeated every trial. The PDF for the ramp cell example consisted of a linear ramp from the start of the track (0 cm) to the end of the track (200 cm). The PDF for the “random field” cell was created by first generating the PDF for the place cell example and then passing this through the field shuffle as described (see Figure 1, Figure Supplement 2). The PDF for the Gaussian noise example was a uniform distribution.

To generate PDFs of grid cells alternating between task-anchored and task-independent codes, representations of each type were generated and merged (Figure 3, Figure Supplement 2). Considering the rolling classification computes a prediction label across a number of trials, we reasoned the manner and frequency of the alternation would affect prediction. We simulated this alternation in both blocks of trials and at the level of single trials. For the simulations that alternated in

blocks of trials, the initial trial was randomly assigned to either the task-anchored PDF or the task-independent PDF with equal probability. For all subsequent trials there was a 10% chance of alternating to the other PDF (e.g. task-anchored to task-independent or task-independent to task-anchored). For simulations that alternated at the level of single trials, every trial was randomly assigned to either the task-anchored PDF or the task-independent PDF, with equal probability.

To evaluate our classification of periodic firing at the level of individual cells, we simulated 500 task-anchored grid cells and 500 task-independent grid cells with grid spacings uniformly distributed between 40 - 400 cm and compared the true labels of these simulated cells with the predicted classifications (Figure 1, Figure Supplement 5). To determine what spatial frequency tolerance to use in our classification, we classified the simulated dataset across the full range of spatial frequency thresholds. This was repeated using a range of P_{max}(Spike) and jitter SD values. To simplify the analysis, no field shuffle was computed on simulated data and therefore no false alarm threshold was used. This effectively forced our classifier to label cells as task-anchored or task-independent without the possibility of an aperiodic label. The prediction accuracy was calculated as the percentage of true task-anchored and task-independent coding grid cells with a correct prediction label. The prediction bias was calculated as the percentage of actual task-anchored cells minus the percentage of predicted task-anchored cells. As the classification in the simulation analysis had only two valid labels, positive bias represents over classification as task-independent whereas negative bias represents over classification as task-anchored. In plots showing prediction bias, positive and negative bias is relabeled to reflect this.

To evaluate classification of periodic firing at the level of individual trials, we simulated 100 grid cells with grid spacings uniformly distributed between 40 - 400 cm that could alternate between task-anchored and task-anchored task-independent coding task-independent either in blocks of trials or every trial (see above) and compared the true labels of these simulated cells and trials with the predicted classifications (Figure 3. Figure Supplement 2). To determine how many periodograms to average over (or if any), we classified the simulated dataset across a range of rolling window sizes to map at what rolling window size we could maximise prediction accuracy and minimise bias between task-anchored and task-independent classifications. Again, this was repeated using a range of P_{max}(Spike) and jitter SD values and no field shuffles were computed. The prediction accuracy of our classification was calculated as the average percentage of true task-anchored and task-independent coding trials with a correct prediction label. The prediction bias was calculated as the average percentage of actual task-anchored trials minus the percentage of predicted task-anchored trials across cells.

Statistical analyses

Group comparisons used linear mixed effect models (Figures 2F, 3G, 4E-F, Figure 3 Figure Supplement 4) or generalised linear mixed effect models (Figures 4E, 6B, D-F) implemented using

lme4 (Bates et al., 2015), lmerTest (Kuznetsova et al., 2017) and glmmTMB (Brooks et al., 2017) packages within R (v4.3.2; R Core Team 2023), with model comparisons using ANOVA and post-fitting pairwise comparisons (Searle et al., 1980) using estimated marginal means (https://github.com/rvlenth/emmeans). Fits of firing rate, spatial information and peak width in Figures 2F, 3G and Figure 3, Figure Supplemental A-B, D used log transformed data. Fits in 4E (comparisons of shuffled data) used a beta family function, and in Figure 6 used a binomial family function with logit linker. Random effects had a nested structure to account for animals and sessions (all models), and where appropriate neuron identity (Figures 4E, 6). For analyses in Figure 6 trials classified as 'aperiodic' were removed from the dataset to facilitate direct comparison of trials with task-anchored and task-independent aperiodic grid firing. To estimate the effect size of the relative influence of task-anchored firing versus task-independent firing on task performance (Figure 6F), odds ratios and confidence intervals were extracted from the full model using sjPlot (https://CRAN.R-project.org/package=sjPlot).

Data and code availability

Data will be made available at https://datashare.ed.ac.uk/handle/10283/777 and code at https://github.com/orgs/MattNolanLab/.

Acknowledgements

We thank Caswell Barry for helpful discussions. This work was supported by the Wellcome Trust (200855/Z/16/Z to MFN), MRC Precision Medicine PhD programme (HC) and the Simons Initiative for the Developing Brain (to MFN). This work made use of resources provided by the Edinburgh Compute and Data Facility. For the purpose of open access, the author has applied a CC BY public copyright license to any Author Accepted Manuscript version arising from this submission.

References

Barry C, Burgess N. 2017. To be a Grid Cell: Shuffling procedures for determining "Gridness." doi:10.1101/230250

Barry C, Hayman R, Burgess N, Jeffery KJ. 2007. Experience-dependent rescaling of entorhinal grids. *Nat Neurosci* **10**:682–684. doi:10.1038/nn1905

Bates D, Mächler M, Bolker BM, Walker SC. 2015. Fitting linear mixed-effects models using lme4. *J Stat Softw* **67**:1–48.

Brooks ME, Kristensen K, van Benthem KJ, Magnusson A, Berg CW, Nielsen A, Skaug HJ, Machler M, Bolker BM. 2017. glmmTMB balances speed and flexibility among packages for zero-inflated generalized linear mixed modeling. *R J* **9**:378–400. doi:10.3929/ethz-b-000240890

Brun VH, Leutgeb S, Wu H-Q, Schwarcz R, Witter MP, Moser EI, Moser M-B. 2008. Impaired spatial representation in CA1 after lesion of direct input from entorhinal cortex. *Neuron* **57**:290–302. doi:10.1016/j.neuron.2007.11.034

Campbell MG, Ocko SA, Mallory CS, Low IIC, Ganguli S, Giocomo LM. 2018. Principles

governing the integration of landmark and self-motion cues in entorhinal cortical codes for navigation. *Nat Neurosci* **21**:1096–1106. doi:10.1038/s41593-018-0189-y
Casali G, Shipley S, Dowell C, Hayman R, Barry C. 2018. Entorhinal Neurons Exhibit Cue Locking in Rodent VR. *Front Cell Neurosci* **12**:512. doi:10.3389/fncel.2018.00512
Chrastil ER, Sherrill KR, Aselcioglu I, Hasselmo ME, Stern CE. 2017. Individual Differences in Human Path Integration Abilities Correlate with Gray Matter Volume in Retrosplenial Cortex, Hippocampus, and Medial Prefrontal Cortex. *eNeuro* **4**:ENEURO.0346-16.2017. doi:10.1523/ENEURO.0346-16.2017
Chung JE, Magland JF, Barnett AH, Tolosa VM, Tooker AC, Lee KY, Shah KG, Felix SH, Frank LM, Greengard LF. 2017. A Fully Automated Approach to Spike Sorting. *Neuron* **95**:1381-1394.e6. doi:10.1016/j.neuron.2017.08.030
Domnisoru C, Kinkhabwala AA, Tank DW. 2013. Membrane potential dynamics of grid cells. *Nature* **495**:199–204. doi:10.1038/nature11973
Fiete IR, Burak Y, Brookings T. 2008. What grid cells convey about rat location. *J Neurosci* **28**:6858–6871. doi:10.1523/JNEUROSCI.5684-07.2008
Gardner RJ, Hermansen E, Pachitariu M, Burak Y, Baas NA, Dunn BJ, Moser M-B, Moser EI. 2021. Toroidal topology of population activity in grid cells. *Cold Spring Harb Lab*. doi:10.1101/2021.02.25.432776
Gaskin S, White NM. 2013. Parallel processing of information about location in the amygdala, entorhinal cortex and hippocampus. *Hippocampus* **23**:1075–1083. doi:10.1002/hipo.22179
Gaskin S, White NM. 2010. Temporary inactivation of the dorsal entorhinal cortex impairs acquisition and retrieval of spatial information. *Neurobiol Learn Mem* **93**:203–207. doi:10.1016/j.nlm.2009.09.012
Gerlei K, Passlack J, Hawes I, Vandrey B, Stevens H, Papastathopoulos I, Nolan MF. 2020. Grid cells are modulated by local head direction. *Nat Commun* **11**:4228. doi:10.1038/s41467-020-17500-1
Gil M, Ancau M, Schlesiger MI, Neitz A, Allen K, De Marco RJ, Monyer H. 2017. Impaired path integration in mice with disrupted grid cell firing. *Nat Neurosci*. doi:10.1038/s41593-017-0039-3
Ginosar G, Aljadeff J, Las L, Derdikman D, Ulanovsky N. 2023. Are grid cells used for navigation? On local metrics, subjective spaces, and black holes. *Neuron* **0**. doi:10.1016/j.neuron.2023.03.027
Gu Y, Lewallen S, Kinkhabwala AA, Domnisoru C, Yoon K, Gauthier JL, Fiete IR, Tank DW. 2018. A Map-like Micro-Organization of Grid Cells in the Medial Entorhinal Cortex. *Cell* **175**:736-750.e30. doi:10.1016/j.cell.2018.08.066
Hales JB, Reitz NT, Vincze JL, Ocampo AC, Leutgeb S, Clark RE. 2021. A role for medial entorhinal cortex in spatial and nonspatial forms of memory in rats. *Behav Brain Res* 113259. doi:10.1016/j.bbr.2021.113259
Hales JB, Schlesiger MI, Leutgeb JK, Squire LR, Leutgeb S, Clark RE. 2014. Medial Entorhinal Cortex Lesions Only Partially Disrupt Hippocampal Place Cells and Hippocampus-Dependent Place Memory. *Cell Rep* 1–9. doi:10.1016/j.celrep.2014.10.009
Hales JB, Vincze JL, Reitz NT, Ocampo AC, Leutgeb S, Clark RE. 2018. Recent and remote retrograde memory deficit in rats with medial entorhinal cortex lesions. *Neurobiol Learn Mem*. doi:10.1016/j.nlm.2018.07.013
Jacob P-Y, Capitano F, Poucet B, Save E, Sargolini F. 2019. Path integration maintains spatial periodicity of grid cell firing in a 1D circular track. *Nat Commun* **10**:840. doi:10.1038/s41467-019-08795-w
Keene CS, Bladon J, McKenzie S, Liu CD, O'Keefe J, Eichenbaum H. 2016. Complementary Functional Organization of Neuronal Activity Patterns in the Perirhinal, Lateral Entorhinal, and Medial Entorhinal Cortices. *J Neurosci* **36**:3660–3675. doi:10.1523/JNEUROSCI.4368-15.2016
Kinkhabwala AA, Gu Y, Aronov D, Tank DW. 2020. Visual cue-related activity of cells in the medial entorhinal cortex during navigation in virtual reality. *Elife* **9**.

doi:10.7554/eLife.43140
Krishnan S, Heer C, Cherian C, Sheffield MEJ. 2022. Reward expectation extinction restructures and degrades CA1 spatial maps through loss of a dopaminergic reward proximity signal. *Nat Commun* **13**:6662. doi:10.1038/s41467-022-34465-5
Kunz L, Schröder TN, Lee H, Montag C, Lachmann B, Sariyska R, Reuter M, Stirnberg R, Stöcker T, Messing-Floeter PC, Fell J, Doeller CF, Axmacher N. 2015. Reduced grid-cell–like representations in adults at genetic risk for Alzheimer's disease. *Science* **350**:430–433. doi:10.1126/science.aac8128
Kuznetsova A, Brockhoff PB, Christensen RHB. 2017. lmerTest Package: Tests in Linear Mixed Effects Models. *J Stat Softw* **82**.
Lakshminarasimhan KJ, Petsalis M, Park H, DeAngelis GC, Pitkow X, Angelaki DE. 2018. A Dynamic Bayesian Observer Model Reveals Origins of Bias in Visual Path Integration. *Neuron* **99**:194-206.e5. doi:10.1016/j.neuron.2018.05.040
Lomb NR. 1976. Least-squares frequency analysis of unequally spaced data. *Astrophys Space Sci* **39**:447–462. doi:10.1007/BF00648343
Lopes G, Monteiro P. 2021. New Open-Source Tools: Using Bonsai for Behavioral Tracking and Closed-Loop Experiments. *Front Behav Neurosci* **15**.
Low IIC, Williams AH, Campbell MG, Linderman SW, Giocomo LM. 2021. Dynamic and reversible remapping of network representations in an unchanging environment. *Neuron*. doi:10.1016/j.neuron.2021.07.005
Mathis A, Herz A, Stemmler M. 2012. Resolution of Nested Neuronal Representations Can Be Exponential in the Number of Neurons. *Phys Rev Lett* **109**:1–5. doi:10.1103/PhysRevLett.109.018103
McNaughton BL, Battaglia FP, Jensen O, Moser EI, Moser M-B. 2006. Path integration and the neural basis of the "cognitive map." *Nat Rev Neurosci* **7**:663–678. doi:10.1038/nrn1932
Miao C, Cao Q, Ito HT, Yamahachi H, Witter MP, Moser M-B, Moser EI. 2015. Hippocampal Remapping after Partial Inactivation of the Medial Entorhinal Cortex. *Neuron* **88**:590–603. doi:10.1016/j.neuron.2015.09.051
Moser EI, Kropff E, Moser M-B. 2008. Place cells, grid cells, and the brain's spatial representation system. *Annu Rev Neurosci* **31**:69–89. doi:10.1146/annurev.neuro.31.061307.090723
Newton C, Pope M, Rua C, Henson R, Ji Z, Burgess N, Rodgers CT, Stangl M, Dounavi M-E, Castegnaro A, Koychev I, Malhotra P, Wolbers T, Ritchie K, Ritchie C, O'Brien J, Su L, Chan D. 2023. Path integration selectively predicts midlife risk of Alzheimer's disease. doi:10.1101/2023.01.31.526473
Noel J-P, Lakshminarasimhan KJ, Park H, Angelaki DE. 2020. Increased variability but intact integration during visual navigation in Autism Spectrum Disorder. *Proc Natl Acad Sci U A* **117**:11158–11166. doi:10.1073/pnas.2000216117
Osborne JE, Dudman JT. 2014. RIVETS: a mechanical system for in vivo and in vitro electrophysiology and imaging. *PLoS One* **9**:e89007. doi:10.1371/journal.pone.0089007
Pettit NL, Yuan XC, Harvey CD. 2022. Hippocampal place codes are gated by behavioral engagement. *Nat Neurosci* 1–6. doi:10.1038/s41593-022-01050-4
Petzschner FH, Glasauer S. 2011. Iterative Bayesian Estimation as an Explanation for Range and Regression Effects: A Study on Human Path Integration. *J Neurosci* **31**:17220–17229. doi:10.1523/JNEUROSCI.2028-11.2011
Qin H, Fu L, Hu B, Liao X, Lu J, He W, Liang S, Zhang K, Li R, Yao J, Yan J, Chen H, Jia H, Zott B, Konnerth A, Chen X. 2018. A Visual-Cue-Dependent Memory Circuit for Place Navigation. *Neuron* **99**:47-55.e4. doi:10.1016/j.neuron.2018.05.021
Rowland DC, Obenhaus HA, Skytøen ER, Zhang Q, Kentros CG, Moser EI, Moser M-B. 2018. Functional properties of stellate cells in medial entorhinal cortex layer II. *Elife* **7**. doi:10.7554/eLife.36664
Sargolini F, Fyhn M, Hafting T, McNaughton BL, Witter MP, Moser M-B, Moser EI. 2006. Conjunctive representation of position, direction, and velocity in entorhinal cortex.

Science **312**:758–762. doi:10.1126/science.1125572
Scargle JD. 1982. Studies in astronomical time series analysis. II - Statistical aspects of spectral analysis of unevenly spaced data. *Astrophys J* **263**:835. doi:10.1086/160554
Schlesiger MI, Cannova CC, Boublil BL, Hales JB, Mankin E a, Brandon MP, Leutgeb JK, Leibold C, Leutgeb S. 2015. The medial entorhinal cortex is necessary for temporal organization of hippocampal neuronal activity. *Nat Neurosci*. doi:10.1038/nn.4056
Searle SR, Speed FM, Milliken GA. 1980. Population Marginal Means in the Linear Model: An Alternative to Least Squares Means. *Am Stat* **34**:216–221. doi:10.1080/00031305.1980.10483031
Skaggs WE, McNaughton BL. 1996. Theta phase precession in hippocampal. *Hippocampus* **6**:149–172.
Solstad T, Boccara CN, Kropff E, Moser M-B, Moser EI. 2008. Representation of geometric borders in the entorhinal cortex. *Science* **322**:1865–1868. doi:10.1126/science.1166466
Sreenivasan S, Fiete I. 2011. Grid cells generate an analog error-correcting code for singularly precise neural computation. *Nat Neurosci*. doi:10.1038/nn.2901
Steffenach HA, Witter M, Moser MB, Moser EI. 2005. Spatial memory in the rat requires the dorsolateral band of the entorhinal cortex. *Neuron* **45**:301–313. doi:10.1016/j.neuron.2004.12.044
Tennant SA, Clark H, Hawes I, Tam WK, Hua J, Yang W, Gerlei KZ, Wood ER, Nolan MF. 2022. Spatial representation by ramping activity of neurons in the retrohippocampal cortex. *Curr Biol*. doi:10.1016/j.cub.2022.08.050
Tennant SA, Fischer L, Garden DLF, Gerlei KZ, Martinez-Gonzalez C, McClure C, Wood ER, Nolan MF. 2018. Stellate Cells in the Medial Entorhinal Cortex Are Required for Spatial Learning. *Cell Rep* **22**:1313–1324. doi:10.1016/j.celrep.2018.01.005
The Astropy Collaboration, Price-Whelan AM, Lim PL, Earl N, Starkman N, Bradley L, Shupe DL, Patil AA, Corrales L, Brasseur CE, Nöthe M, Donath A, Tollerud E, Morris BM, Ginsburg A, Vaher E, Weaver BA, Tocknell J, Jamieson W, van Kerkwijk MH, Robitaille TP, Merry B, Bachetti M, Günther HM, Aldcroft TL, Alvarado-Montes JA, Archibald AM, Bódi A, Bapat S, Barentsen G, Bazán J, Biswas M, Boquien M, Burke DJ, Cara D, Cara M, Conroy KE, Conseil S, Craig MW, Cross RM, Cruz KL, D'Eugenio F, Dencheva N, Devillepoix HAR, Dietrich JP, Eigenbrot AD, Erben T, Ferreira L, Foreman-Mackey D, Fox R, Freij N, Garg S, Geda R, Glattly L, Gondhalekar Y, Gordon KD, Grant D, Greenfield P, Groener AM, Guest S, Gurovich S, Handberg R, Hart A, Hatfield-Dodds Z, Homeier D, Hosseinzadeh G, Jenness T, Jones CK, Joseph P, Kalmbach JB, Karamehmetoglu E, Kałuszyński M, Kelley MSP, Kern N, Kerzendorf WE, Koch EW, Kulumani S, Lee A, Ly C, Ma Z, MacBride C, Maljaars JM, Muna D, Murphy NA, Norman H, O'Steen R, Oman KA, Pacifici C, Pascual S, Pascual-Granado J, Patil RR, Perren GI, Pickering TE, Rastogi T, Roulston BR, Ryan DF, Rykoff ES, Sabater J, Sakurikar P, Salgado J, Sanghi A, Saunders N, Savchenko V, Schwardt L, Seifert-Eckert M, Shih AY, Jain AS, Shukla G, Sick J, Simpson C, Singanamalla S, Singer LP, Singhal J, Sinha M, Sipőcz BM, Spitler LR, Stansby D, Streicher O, Šumak J, Swinbank JD, Taranu DS, Tewary N, Tremblay GR, de Val-Borro M, Van Kooten SJ, Vasović Z, Verma S, Cardoso JV de M, Williams PKG, Wilson TJ, Winkel B, Wood-Vasey WM, Xue R, Yoachim P, ZHANG C, Zonca A. 2022. The Astropy Project: Sustaining and Growing a Community-oriented Open-source Project and the Latest Major Release (v5.0) of the Core Package. doi:10.3847/1538-4357/ac7c74
VanderPlas JT. 2018. Understanding the Lomb–Scargle Periodogram. *Astrophys J Suppl Ser* **236**:16. doi:10.3847/1538-4365/aab766
Waaga T, Agmon H, Normand VA, Nagelhus A, Gardner RJ, Moser M-B, Moser EI, Burak Y. 2021. Grid-cell modules remain coordinated when neural activity is dissociated from external sensory cues. *bioRxiv*. doi:10.1101/2021.08.29.458100
Whittington JCR, McCaffary D, Bakermans JJW, Behrens TEJ. 2022. How to build a cognitive map. *Nat Neurosci* **25**:1257–1272. doi:10.1038/s41593-022-01153-y

Yoon K, Buice MA, Barry C, Hayman R, Burgess N, Fiete IR. 2013. Specific evidence of low-dimensional continuous attractor dynamics in grid cells. *Nat Neurosci* **16**:1077–1084. doi:10.1038/nn.3450

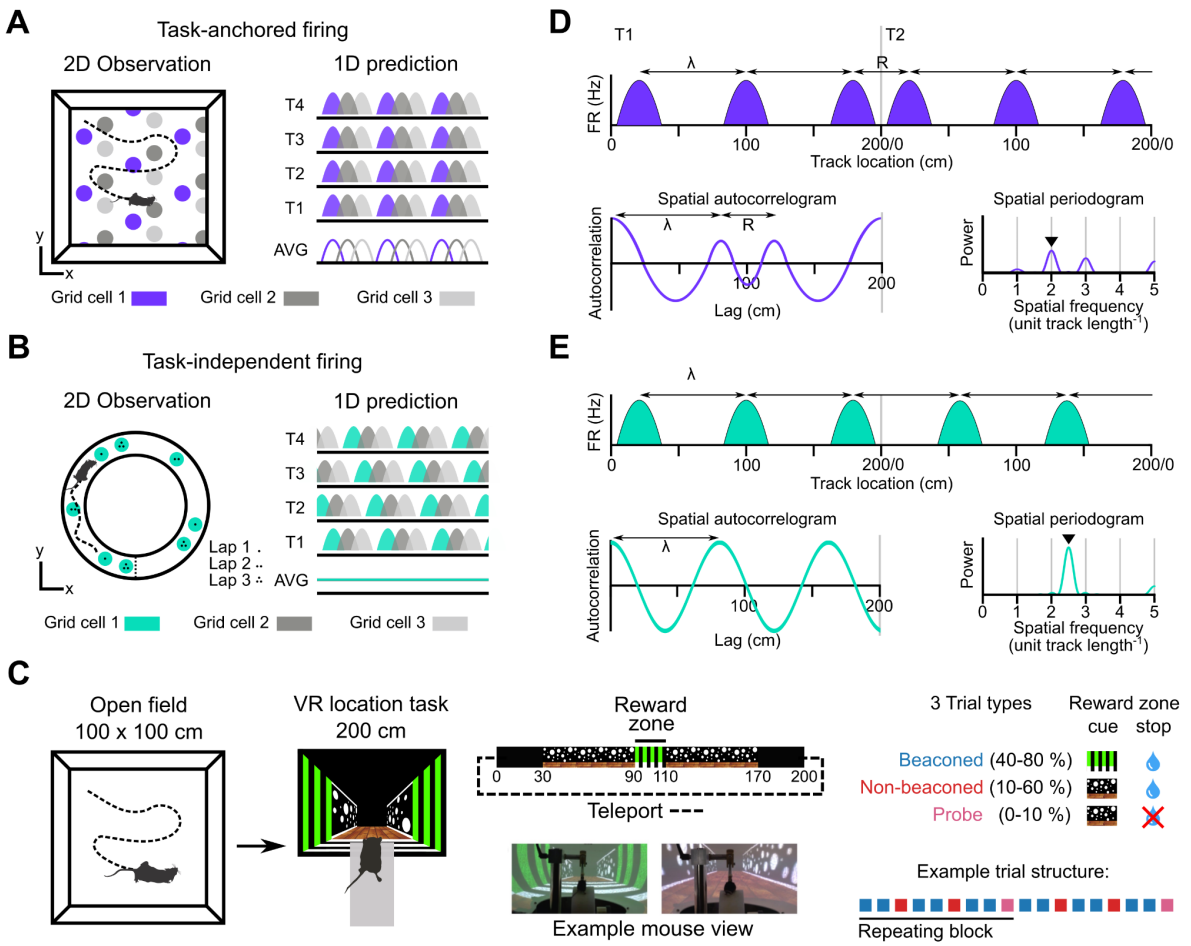

Figure 1. Models for grid representation and experiment design

(A-B) Predicted task-anchored (A) and task-independent (B) firing of grid cells in a 1D environment (right) given firing patterns of grid cells previously observed in square (A) and circular (B) 2D arenas (left). T1-T4 indicate consecutive trials in the 1D environment, AVG indicates expected average across many trials.

(C) Neurons were recorded in an open arena and then in a location memory task. Trials were configured with a reward for stopping in a visually cued zone (beaconed), or a reward for stopping in the same zone but with the cue absent (non-beaconed), or without the visual cue or the reward (probe). Trial percentages indicate the proportion of trial types experienced in a single session; in any given session this proportion was fixed and trials were interleaved in a fixed repeating pattern (see Methods).

(D) In the task-anchored coding scheme, a grid cell fires with field spacing λ and resets its firing every trial by anchoring its fields to the same track location, with a realignment lag R observed in the spatial autocorrelogram. Fields locations remain constant on each trial and thus peaks in the periodogram occur at integer spatial frequencies relative to the track repetition (see Figure 1, Figure Supplement 3)

(E) In a task-independent coding scheme, a grid cell fires with field spacing λ and continues to fire at regular intervals without anchoring to the track. Unless field spacing and the track length are integer divisible, the location of fields varies across trials, and thus the peak of the periodogram is not constrained to an integer spatial frequency (see Figure 1, Figure Supplement 3).

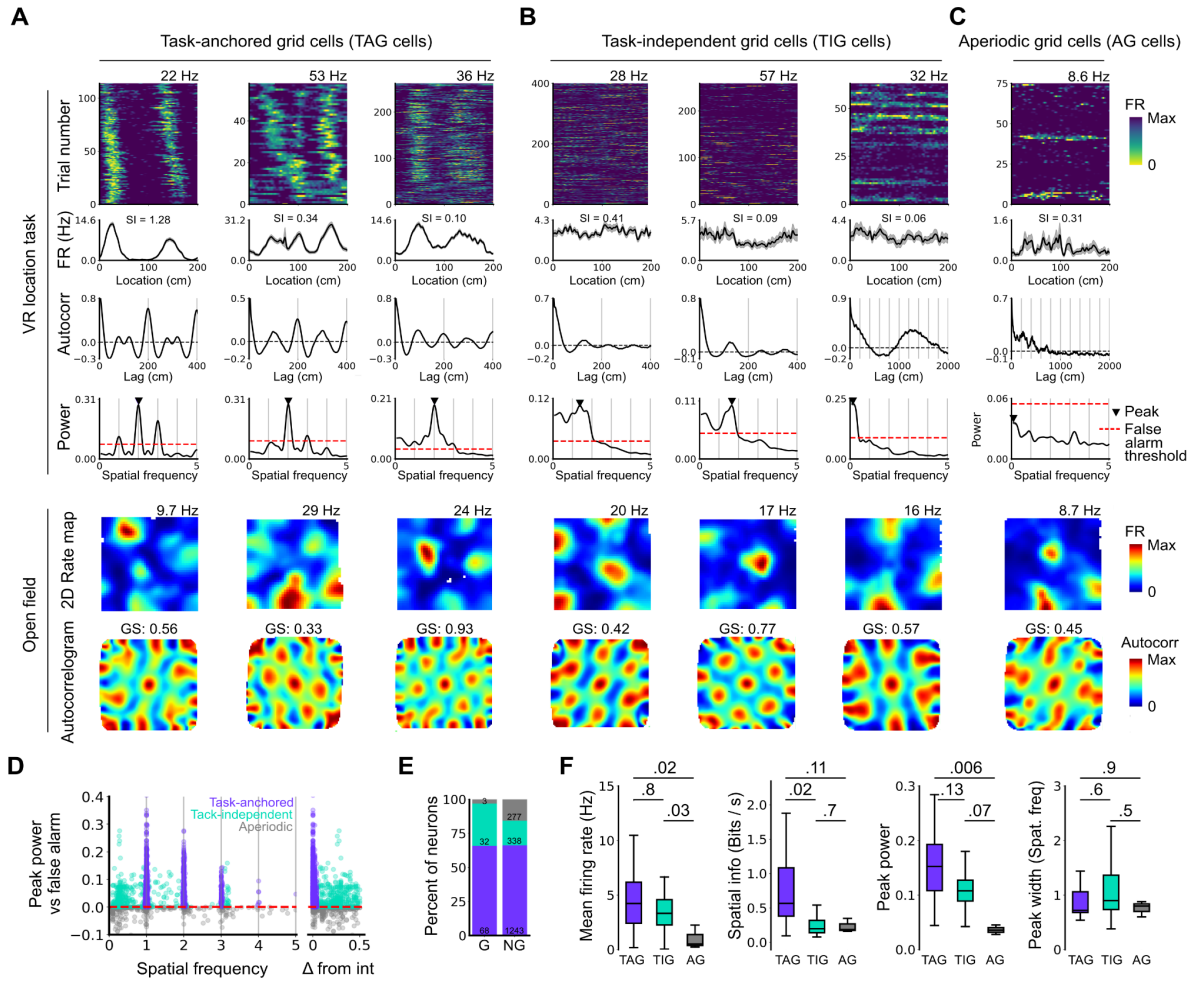

Figure 2. Grid cells operate in task-anchored or task-independent firing modes

(A-C) Examples of grid cells with activity during the location memory task classified as task-anchored (A), task-independent (B), or aperiodic (C) at the session-level. Examples are ordered in their respective groups by their spatial information on the track. From top to bottom plots show: heap map of firing rate as a function of track position, spatial autocorrelation of the track firing rate, periodogram of the track firing, open field firing rate map and open field spatial autocorrelogram. The red line indicates the false alarm threshold estimated from shuffled data and significant peaks are labelled with a triangle. X-axis scales are adjusted on the virtual reality spatial autocorrelation to better illustrate the long-range periodic signal.

(D) Peak power as a function of the spatial frequency at which the peak occurs for all recorded cells. The red dashed line indicates the false alarm threshold generated from shuffled data.

(E) Percentage of grid (G) and non-grid (NG) cells classified to task-anchored, task-independent and aperiodic groups.

(F) Comparison between task-anchored (TAG), task-independent (TIG) and aperiodic (AG) grid cells of mean firing rate (ANOVA: DF = 2, P = 0.006, X^2 = 10.215; pairwise comparisons: TAG vs TIG, DF = 4.12, P = 0.8, T = 0.562; TAG vs AG, DF = 42.07, P = 0.02, T = 2.906; TIG vs AG, DF = 67.29, P = 0.03, T = 2.588), spatial information scores (ANOVA: DF = 2, P = 0.008, X^2 = 9.54; pairwise comparisons: TAG vs TIG, DF = 73.6, P = 0.02, T = 2.815; TAG vs AG, DF = 95.5, P = 0.11, T = 2.036; TIG vs AG, DF = 96.6, P = 0.7, T = 0.783), peak

power (ANOVA: DF = 2, P = 0.001, X^2 = 13.792; pairwise comparisons: TAG vs TIG, DF = 19.0, P = 0.13, T = 2.033; TA vs A, P = 0.006, T = 3.239, DF = 54.9; TI vs A, P = 0.07, T = 2.274, DF = 80.4), and peak width (ANOVA: DF = 2, P = 0.15, X^2 = 3.7963; pairwise comparisons: TAG vs TIG, P = 0.61, T = -1.029, DF = 3.27; TA vs A, P = 0.88, T = 0.472, DF = 45.76; TI vs A, P = 0.55, T = 1.053, DF = 71.22).

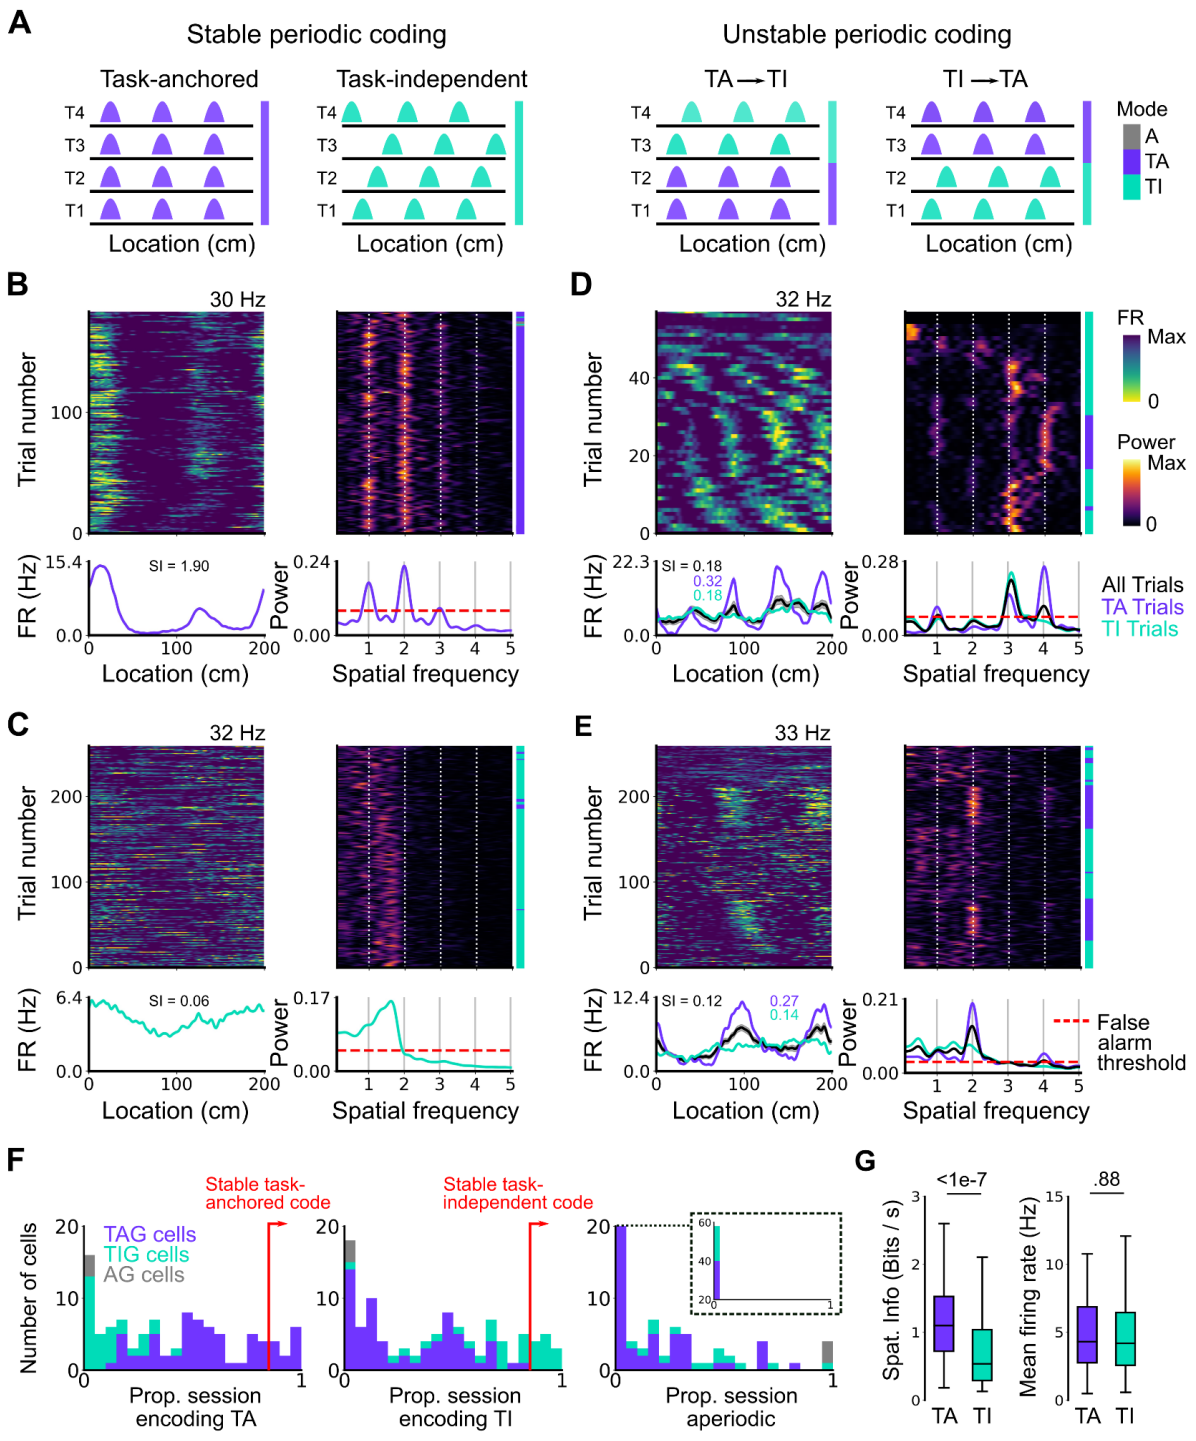

Figure 3. Coding schemes switch within behavioural sessions

(A) In a 'stable coding' scenario grid cells remain either task-anchored or task-independent throughout the recording session (left), whereas with 'unstable coding' the grid activity switches between task-anchored and task-independent modes (right).
(B-E) Example trial-by-trial firing rate heat maps (upper left), corresponding rolling periodogram heat maps (upper right), mean rate maps (lower left) and mean periodograms (lower right) for neurons exhibiting stable task-anchored coding (B), stable task-independent

coding (C) and unstable coding in which representations switch between task-anchored and task-independent (D-E).

(F) Distribution across all recorded grid cells of task-anchored trials (left), task-independent trials (centre) and aperiodic trials (right). Session-level cell task-anchored grid (TAG), task-independent grid (TIG) and aperiodic grid (AG) cell classifications are differently coloured.

(G) Spatial information was higher for trials on which grid cells were task-anchored compared to when they were task-independent (left)(ANOVA: $P < 1e\text{-}7$, $X^2 = 30.98$, DF = 1), whereas the average firing rate was similar between task-anchored and task-independent trials (right)(ANOVA: $P = 0.88$, $X^2 = 0.022$, DF = 1).

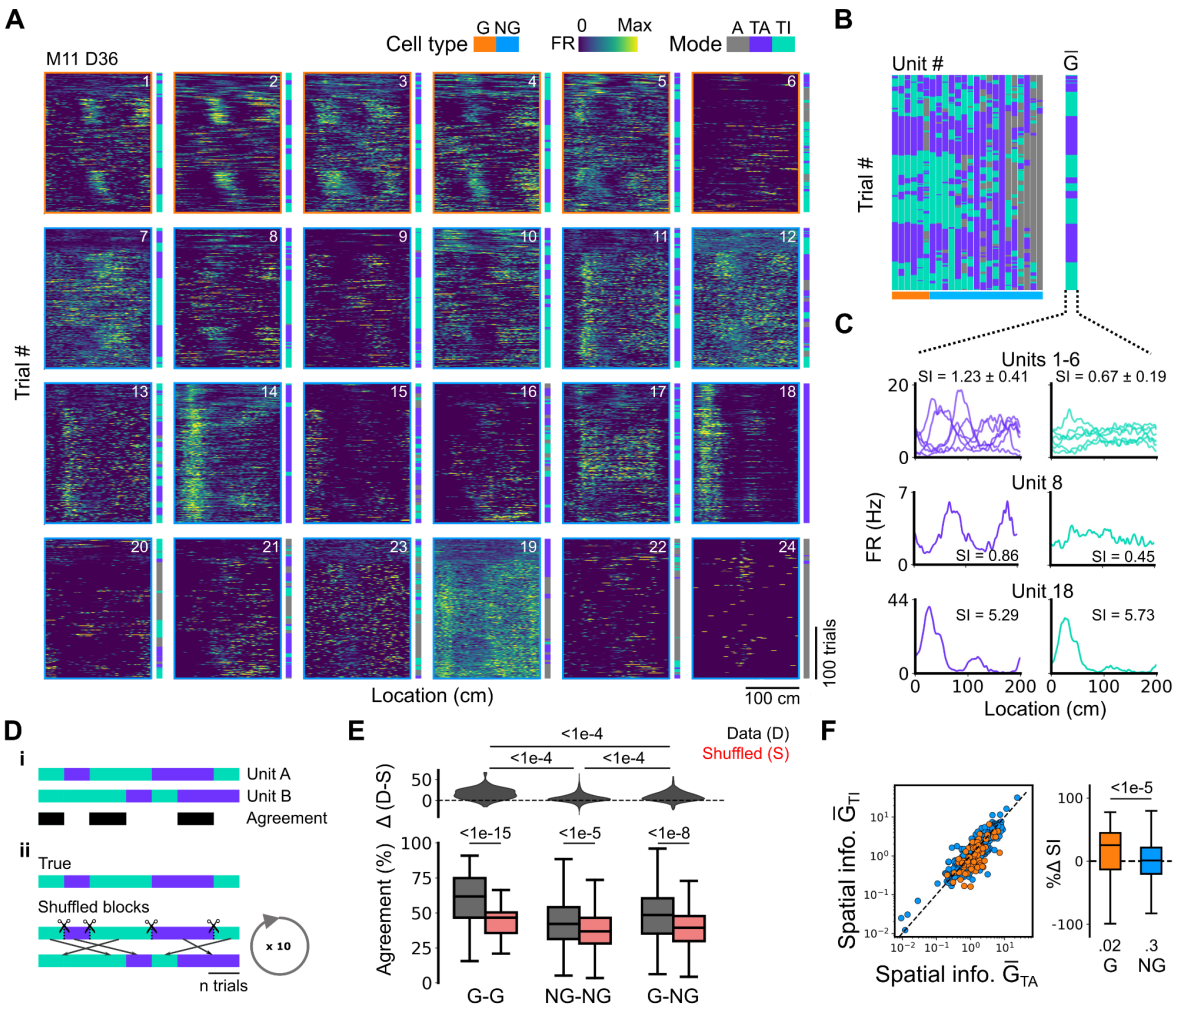

Figure 4. Grid cells and non-grid cells switch between coding schemes coherently

(A) Joint activity of 6 simultaneously recorded grid cells (orange frames) and 18 non-grid cells (blue frames) from a single session. For each cell, the firing rate map across trials (left) is shown next to the trial classification (right).

(B) Classifications for all grid cells (GC) and non-grid cells (NG) as shown in (A), ordered by their agreement to the most common classification within the recorded network at any particular trial. The common classification for recorded grid cells is shown as $\overline{G}$.

(C) Mean firing rate as a function of position for exemplar units from (A) when $\overline{G}$ was task-anchored (left) and task-independent (right).

(D) Strategy for assessing agreement between cells in their firing mode (i) and for generating shuffled datasets (ii).

(E) Agreement in the firing mode between each combination of grid (G) and non-grid (NG) cells and corresponding scores for the shuffled data (lower), and the difference between the shuffled and actual scores (upper). Agreement was greater between grid cell pairs than between pairs involving non-grid cells (ANOVA: G-G, $P < 1e\text{-}15$, $X^2 = 65.72$, DF = 1, NG-NG, $P < 1e\text{-}5$, $X^2 = 21.22$, DF = 1, G-NG, $P < 1e\text{-}8$, $X^2 = 33.26$, DF = 1; pairwise comparisons: G-G vs NG-NG, $P < 1e\text{-}4$, T = -10.768, DF = 10720, G-G vs G-NG, $P < 1e\text{-}4$, T = -8.578, DF = 10710, NG-NG vs G-NG, $P < 1e\text{-}4$, T = -7.373, DF = 10366).

(F) Spatial information of individual cells during trials in which $\overline{G}$ is task-independent as a function of spatial information during trials in which $\overline{G}$ is task-anchored (left). The difference

in spatial information between sessions classed as task-anchored or task-independent on the basis of grid cell activity was greater for grid than non-grid cells (right, ANOVA: P < 1e-5, X^2 = 21.1, DF = 1; G vs zero, P = 0.018, T = 2.723, DF = 11.87; NG vs zero, P = 0.3, T = -1.278, DF = 2.58). The percentage change in spatial information was calculated as $100 \cdot (SI_{\overline{G=TA}} - SI_{\overline{G=TI}})/SI_{\overline{G=TA}}$.

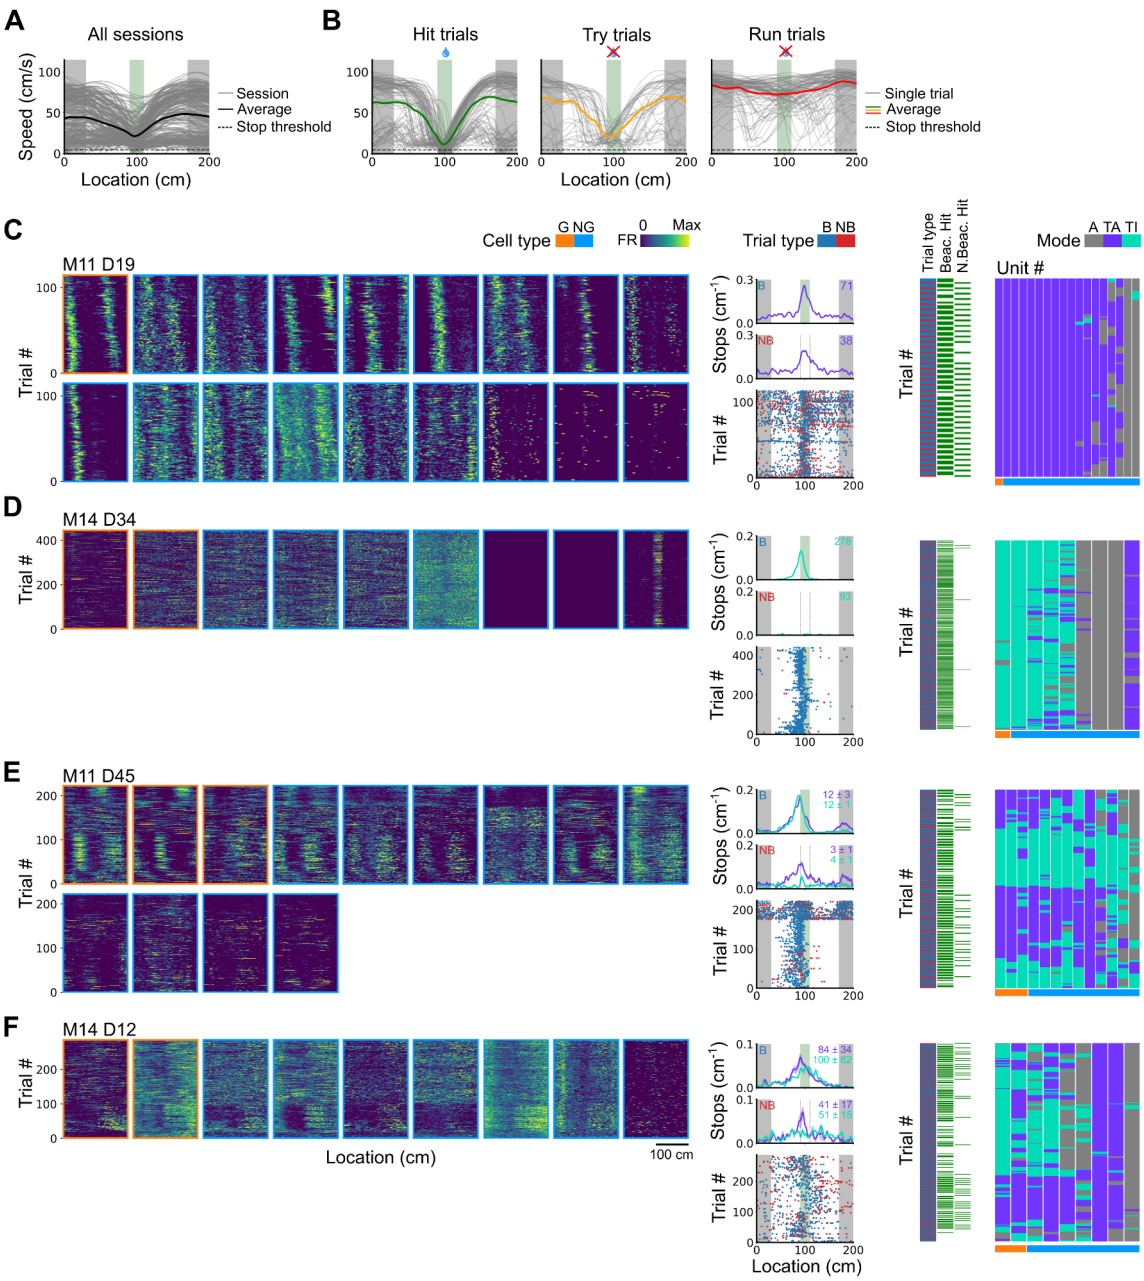

Figure 5. Spatial behaviour during task-anchored and task-independent grid modes

(A) Averages across each behavioural session of running speeds as a function of track position for all sessions

(B) Running speed as a function of track position for trial outcomes classified as hit, try or run for an example session.

(C-F) Examples of variation in the behaviour-related activity of grid and non-grid cells recorded on the location memory task, illustrating firing patterns that are stable and task-anchored (C) or task-independent (D) firing, and unstable firing where cells switch between task-anchored and task-independent modes (E-F). Plots show all simultaneously recorded cells' firing rate maps in each session (left), stop rasters (lower centre), stop density on beaconed (B) and non-beaconed (NB) trials coloured according to whether grid cells were task-anchored or task-independent (upper centre) and a summary of raster of behaviour and cell classifications (right). Shaded regions in stop density plots represent

standard error of the mean measured across epochs. The number of trials classified in a particular coding scheme is also provided with the stop density plot. Grid cells and non-grid cells are colour coded by bounding boxes around the firing rate map.

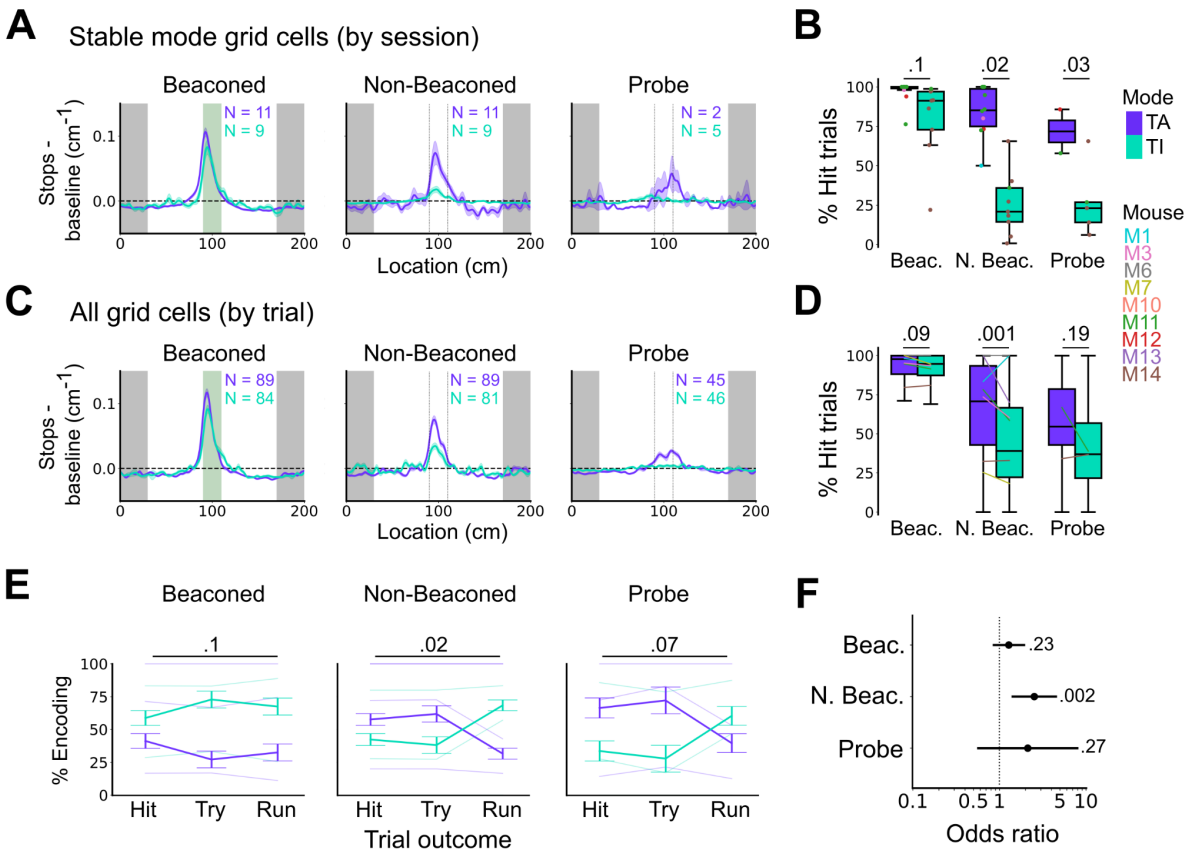

Figure 6. Anchoring of grid firing to the task reference frame promotes localisation by path integration but is not required for cued localisation

(A, C) Stopping probability relative to baseline as a function of position for sessions with stable codes (A) and for epochs within sessions containing both stable and unstable codes. (C). Shaded regions in A and C represent standard error of the mean measured across sessions and epochs respectively. N represents the number of sessions with stable codes (A) and the number of grid cells with coding epochs (C). Stable sessions were defined as a session with at least one grid cell for which >85 % of trials were in a single coding mode.

(B, D) Percentage of hits on beaconed, non-beaconed and probe trials when the code is task-anchored or task-independent for (B) sessions with stable codes (ANOVA: Beaconed, P = 0.1, X^2 = 2.70, DF = 1; Non-beaconed, P = 0.022, X^2 = 5.24, DF = 1; Probe, P = 0.033, X^2 = 4.55, DF = 1) and when epochs within sessions are task-anchored or task-independent (D)(ANOVA: Beaconed, P = 0.09, X^2 = 2.84, DF = 1; Non-beaconed, P = 0.001, X^2 = 10.14, DF = 1; Probe, P = 0.19, X^2 = 1.70, DF = 1, see Methods).

(E) Percentage of trials with hit, try and run outcomes in which grid cell firing is task-anchored (TA) or task-independent (TI) (ANOVA: Beaconed, P = 0.11, X^2 = 4.49, DF = 2; Non-beaconed, P = 0.02, X^2 = 7.75, DF = 2; Probe, P = 0.07, X^2 = 5.40, DF = 2). Error bars denote standard error of the mean measured across the mean values for each animal. Faded lines show percentage values for individual mice that contained hit, try and run trials for a given trial type.

(F) Odds ratio between receipt of reward for trials on which epochs contain task-anchored firing relative to trials on which epochs contain task-independent firing (Beaconed-G, P = 0.23; Non-beaconed-G, P = 0.001; Probe-G, P = 0.27).

Supplemental Figures

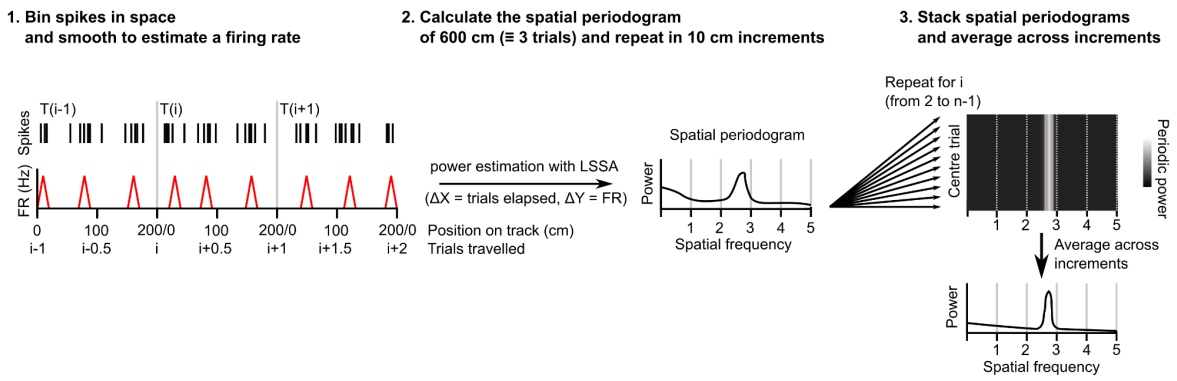

Figure 1, Figure Supplement 1. Procedure for extracting a session-level periodogram from a set of spike timestamps

A single cell's spikes were binned in space (1 cm bin size), divided by the time spent in each bin then smoothed using a Gaussian kernel (SD = 2 cm). Least squares spectral analysis (LSSA) was computed on a signal equivalent to 3 track lengths using the Lomb-Scargle implementation in the Astropy python module (Δx = distance elapsed in trials, Δy = firing rate)(The Astropy Collaboration et al., 2022). The signal was advanced 10 cm and repeated for the whole session. All individual periodograms generated were then used to generate the average periodogram.

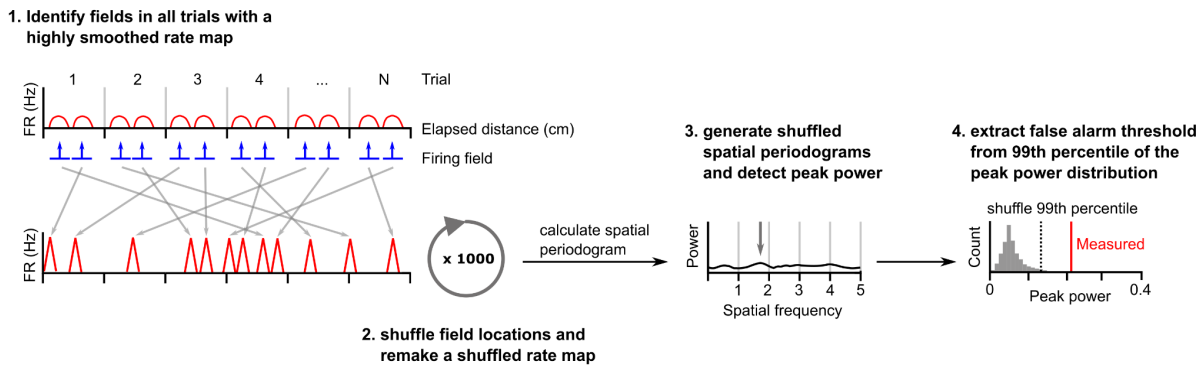

Figure 1, Figure Supplement 2. Procedure for estimating a false alarm threshold for a given cell

To generate shuffled datasets a cell's spikes were first binned in space (1 cm bin size), divided by the time spent in each bin then smoothed using a Gaussian kernel (SD = 4 cm). Firing fields were detected and the original field bins reallocated to random locations on a new unsmoothed rate map. Bins not attributed to a field were allocated to the remaining gaps in the new map. The unsmoothed shuffled rate map was then smoothed with a Gaussian kernel (SD = 2 cm). An average periodogram was calculated and peak power detected. This was repeated 1000 times. The false alarm threshold was set equal to the 99th percentile of the peak power of periodograms generated from the shuffled dataset.

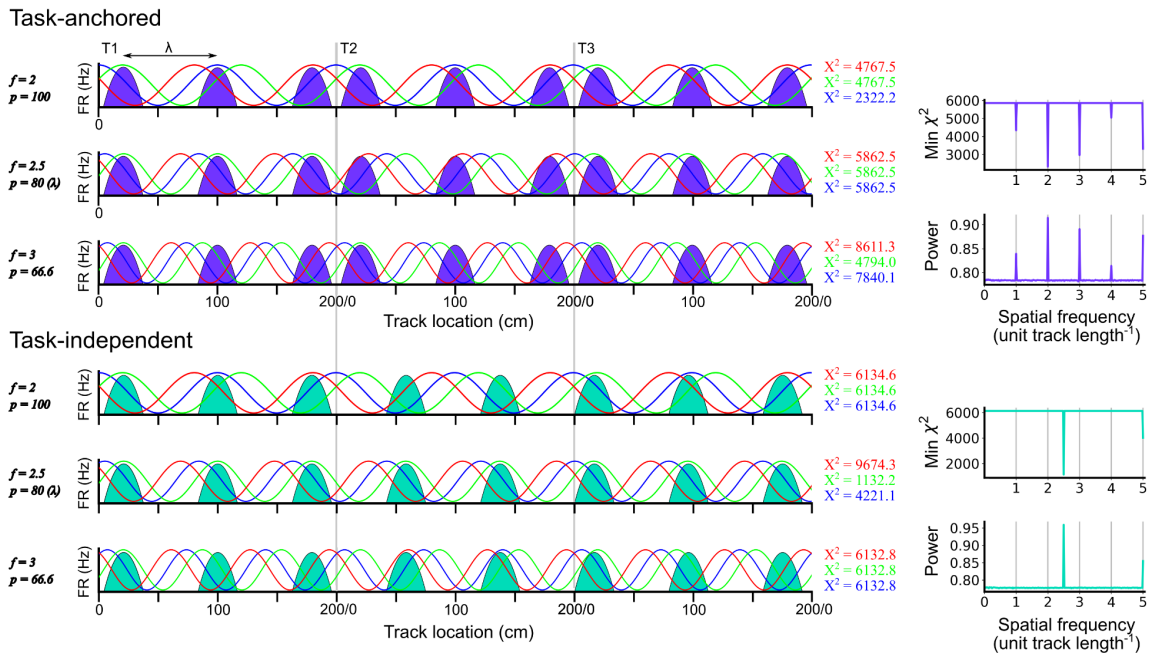

Figure 1, Figure Supplement 3. Identification of task-anchored and task-independent periodic firing

Examples of least squares spectral analysis (LSSA) used to estimate the sinusoidal components that best describe the firing rate profile of a task-anchored grid code (upper) and a task-independent code (lower). In each plot the schematised firing fields (solid colours) are shown as a function of track position. Each row shows examples of different sinusoidal components (f = frequency, p = period) at different phases. LSSA minimises the chi-squared error by searching over all phases (three example phases are shown in red, green and blue per row) and amplitudes (amplitude is ignored here for simplicity) for each spatial frequency. For task-anchored grid representations, the chi-squared error (min χ^2) is minimal at positive-integer spatial frequencies that correspond to sinusoidal waves with periods that are integer divisible by the track length. For task-independent grid representations the chi-squared error is minimal at a spatial frequency that corresponds to the underlying periodic firing rate profile. The standard normalised periodogram is created by normalising the chi-squared errors of the periodic model around a non-varying reference model. The resulting power is a dimensionless quantity that lies within the range 0 to 1.

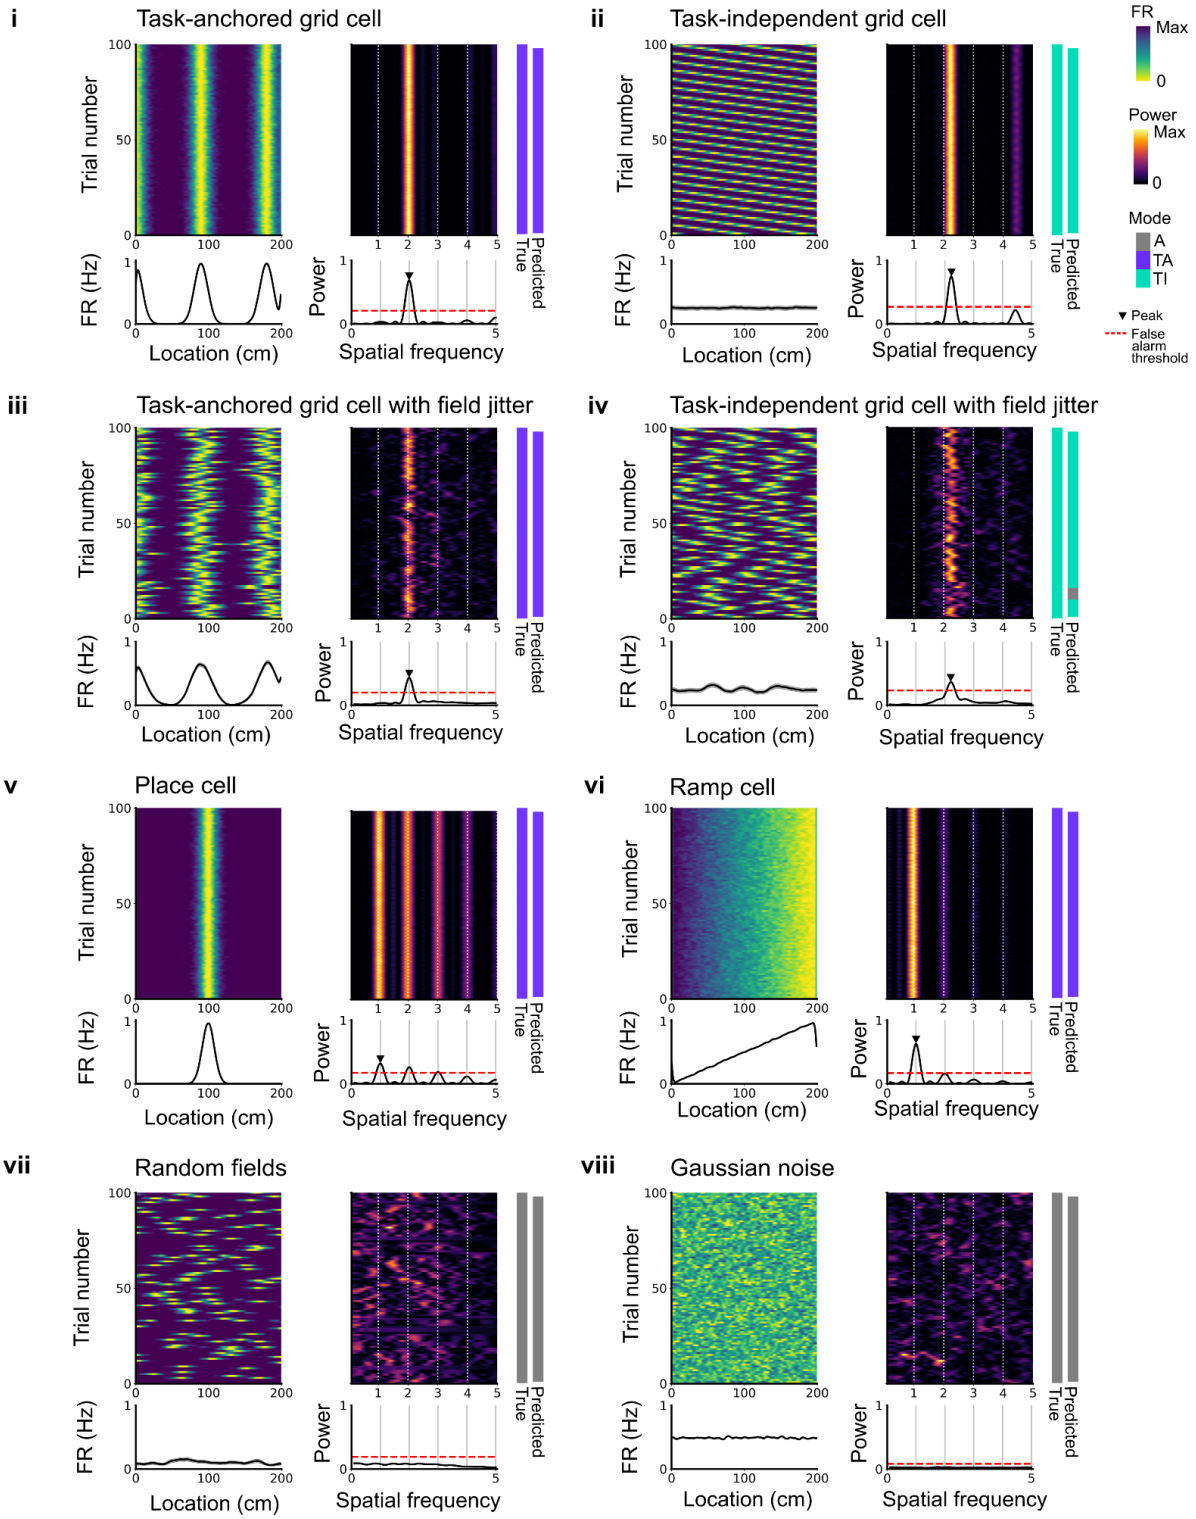

Figure 1, Figure Supplement 4. Expected spatial periodicity for different functional cell types

Plots show simulations of track firing for (i) a task-anchored grid cell, (ii) a task-independent grid cell, (iii) a task-anchored grid cell with field jitter, (iv) a task-independent grid cell with field jitter, (v) a place cell, (vi) a ramp-like cell, (vii) a cell with randomly positioned fields and (viii) a cell with Gaussian noise-like activity. Each group of panels shows firing rate heat map

by trial (upper left) and corresponding average firing rate (lower left), rolling periodogram (upper centre) and corresponding average periodogram (lower centre), the true and predicted trial classification (upper right). The predicted trial classification is established using analyses described in Figure 1, Figure Supplements 1-3.

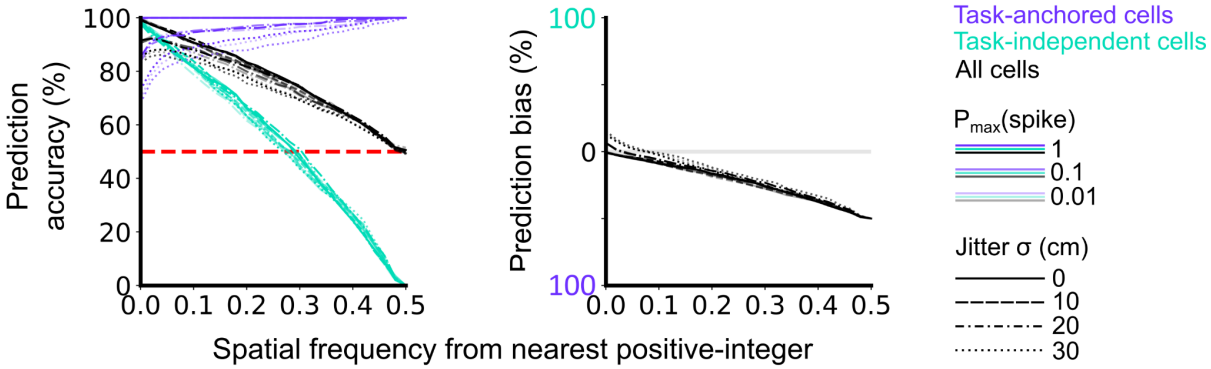

Figure 1. Figure Supplement 5. Validation of accuracy and bias for classification of task-anchored and task-independent modes

500 task-anchored grid cells and 500 task-independent grid cells with grid spacings uniformly distributed between 40 - 400 cm were simulated to identify the optimal threshold to classify based on spatial frequency. Simulations were repeated using different spike rate probabilities and different field jitters (defined by the standard deviation of the field locations from their preassigned locations). Roughly, P_{max}(spike) values of 0.01, 0.1 and 1 covered average firing rates on the order of magnitudes of 10^{-1}, 10^{0}, 10^{1} Hz respectively whereas jitter SD values of 0, 10, 20 and 30 cm covered progressively greater deviations from perfect periodic firing. Classifications were made on each set of simulations in the same way as for analysis of experimental data (see Figure 1, Figure Supplements 1-4). Prediction accuracy (left) and bias (right) are shown as a function of the spatial frequency from the nearest positive-integer used as a threshold to distinguish task-anchored and task-independent modes. Chance level accuracy is denoted with a red dashed line.

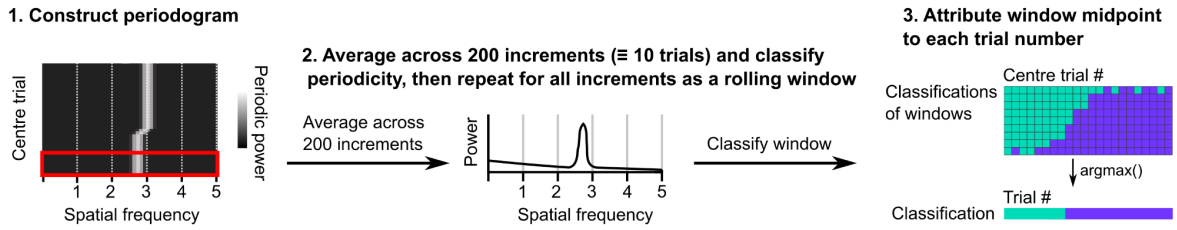

Figure 3. Figure Supplement 1. Procedure for classifying periodicity on a rolling basis. Periodograms were generated as in Figure 1, Figure Supplement 1. Sliding windows containing 200 consecutive periodograms were then extracted and averaged. A classification was made on the averaged periodogram using the peak power, spatial frequency and an adjusted false alarm threshold (see Methods). The classification and midpoint location of the window were stored and the procedure repeated for the full session. Windows were assigned to trials based on the location of their middle points and then the classification occurring most often was assigned to that trial.

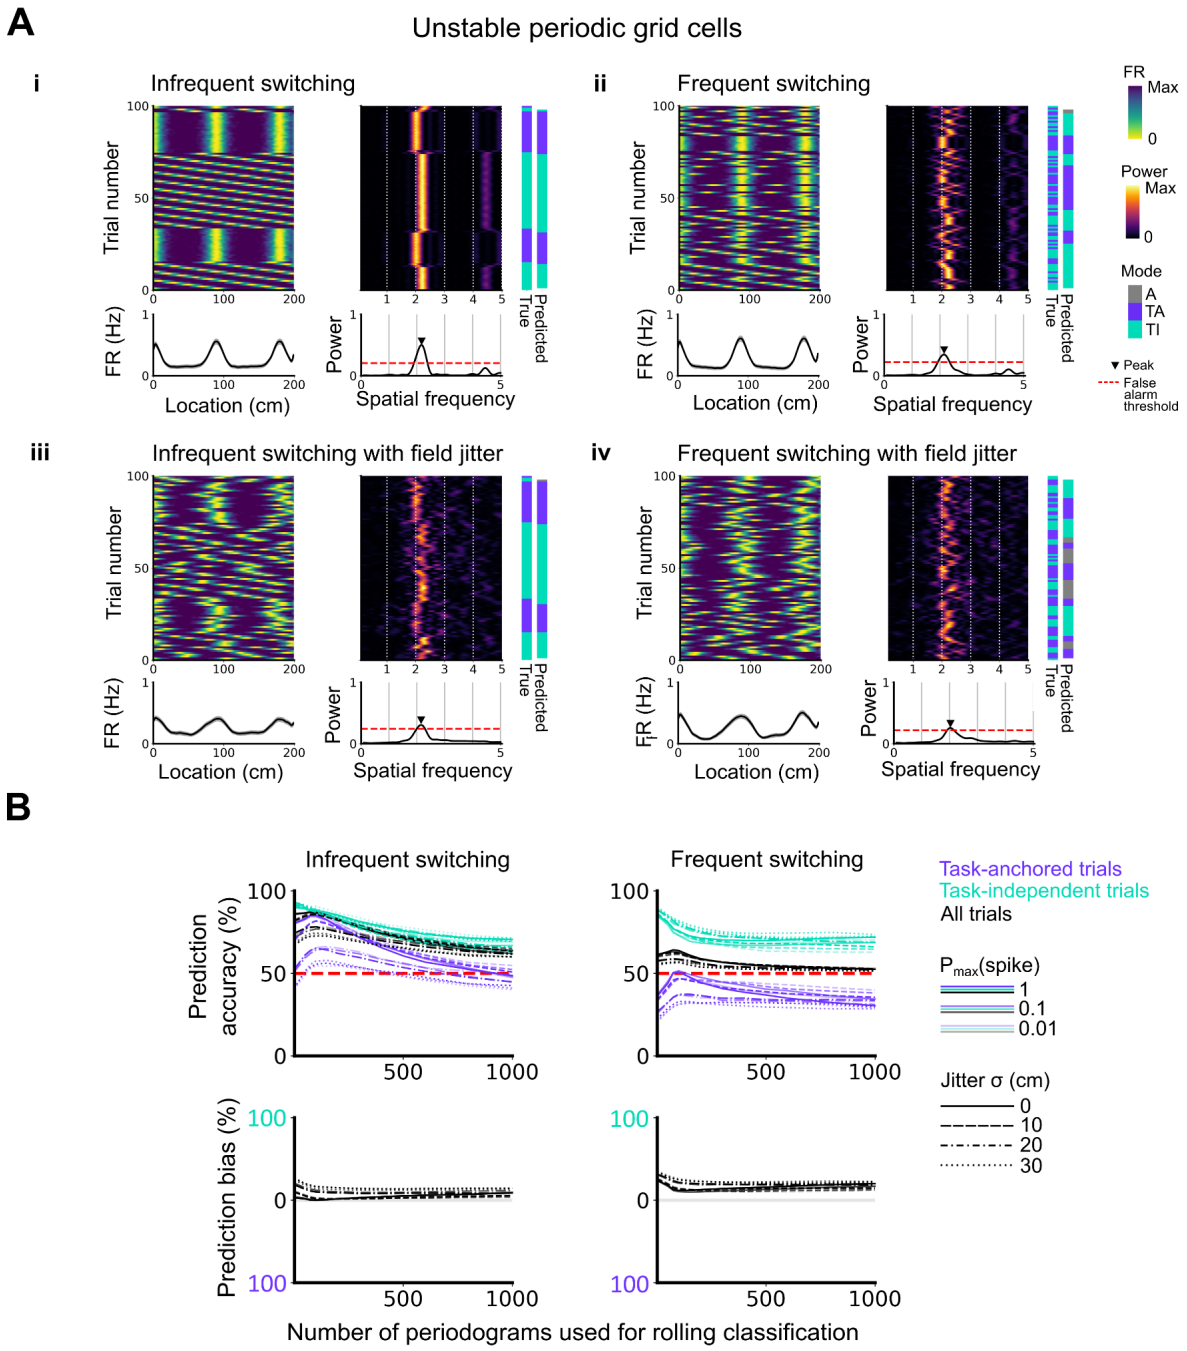

Figure 3. Figure Supplement 2. Evaluation of classification accuracy and bias as a function of the number of periodograms used for rolling classification

(A) Examples of the activity of simulated grid cells that switch between task-anchored and task-independent encoding modes of operation, either infrequently (i, iii) or relatively frequently (ii, iv), and with (iii, iv) and without (i, ii) jitter. Each example shows the firing rate heatmap by trials (upper left), the average firing rate map (lower left), the rolling periodogram (upper right), the average periodogram (lower right) and the classification labels (true and predicted; far right).

(B) To assess classification accuracy under different rolling windows, we calculated the prediction accuracy and bias (see Methods). Chance level is denoted with a red dashed line. Plots on the left show the prediction accuracy (top) and the prediction bias (bottom) when grid cells alternate between task-anchored and task-independent firing fields in blocks of

trials (as in i, iii above), whereas the right plots shows the same information for a higher switching frequency (as in ii, iv above).

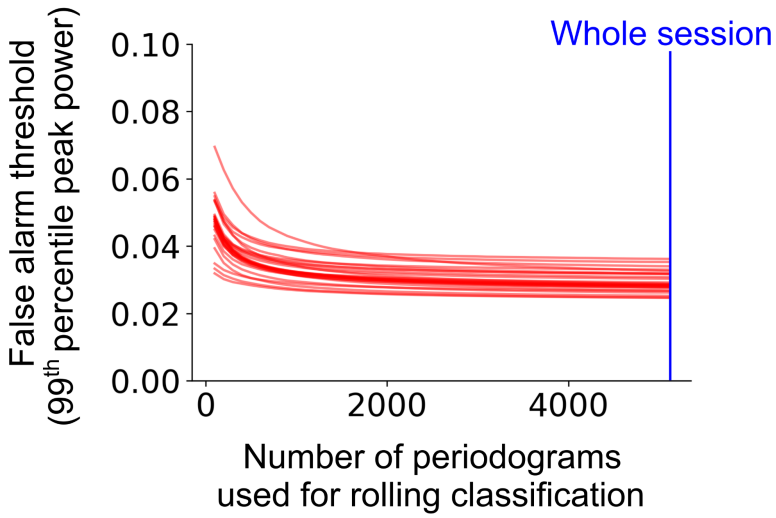

Figure 3. Figure Supplement 3. Classification accuracy and bias on the level of trial as a function of rolling window size

The false alarm threshold was defined by the 99th percentile of the distribution of the peak power of the average periodogram for 1000 shuffled instances of a cell's firing rate profile across the full experimental session. The false alarm threshold decreases asymptotically as the number of periodograms used for the rolling window classification is increased. Each red line represents a single cell within a representative session. The blue line represents the maximum number of periodograms in this session.

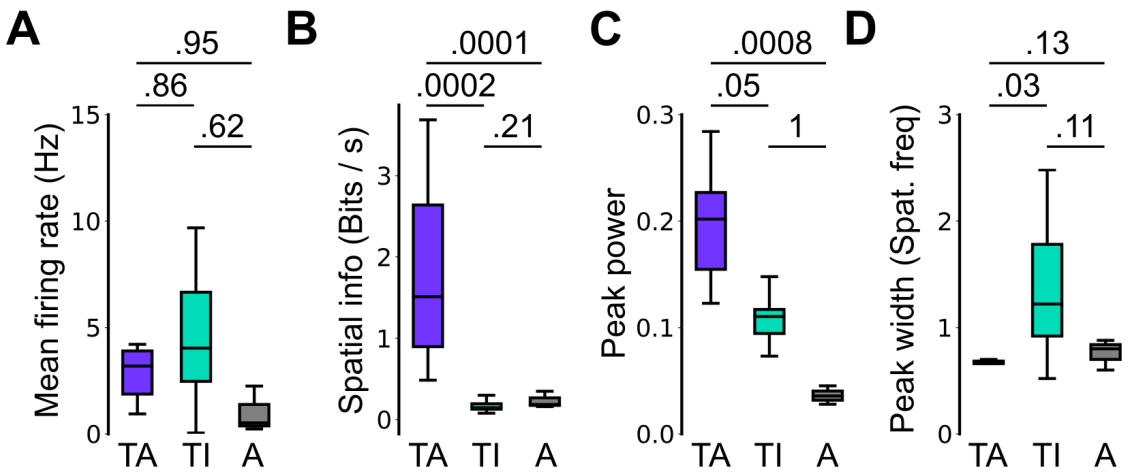

Figure 3. Figure Supplement 4. Spatial firing properties of stable task-anchored and task-independent grid cells. This is a repeat of the analysis shown in Figure 2F, but categorising cells according to the proportion of trials on which their firing is a given mode. Thus, whereas in 2F task-anchored, task-independent and periodic firing classifications are assigned on the basis of average periodograms across the behavioural session, here we group cells according to whether they show stable (> 85% of trials) task-anchored, task-independent or periodic firing. Conclusions from this analysis are similar to those of Figure 2F.

(A) Comparison of mean firing rate (ANOVA: P = 0.49 , X^2 = 1.42, DF = 2; pairwise comparisons: TA vs TI, P = 0.86, T= -0.532, DF = 21.4; TA vs A, P = 0.95, T = -0.292, DF = 96.7; TI vs A, P = 0.62, T = -0.937, DF = 23.1).

(B) Comparison of spatial information scores (ANOVA: P < 1e-4, X^2 = 21.47, DF = 2; pairwise comparisons: TA vs TI, P = 0.0002, T = 4.229, DF = 89.9; TA vs A, P = 0.0001, T = -4.412, DF = 94.2; TI vs A, P = 0.21, T = 1.703, DF = 85.2).

(C) Comparison of peak power (ANOVA: P = 0.0007, X^2 = 14.42, DF = 2; pairwise comparisons: TA vs TI, P = 0.05, T= 2.383, DF = 88.4; TA vs A, P = 0.0008, T = -3.781, DF = 98.9; TI vs A, P = 0.996, T = -0.080, DF = 86.5).

(D) Comparison of peak width (ANOVA: P = 0.002, X^2 = 12.78, DF = 2; pairwise comparisons: TA vs TI, P = 0.03, T = -2.840, DF = 16.2; TA vs A, P = 0.13, T = 1.949, DF = 95.3; TI vs A, P = 0.11, T = -2.162, DF = 16.0).

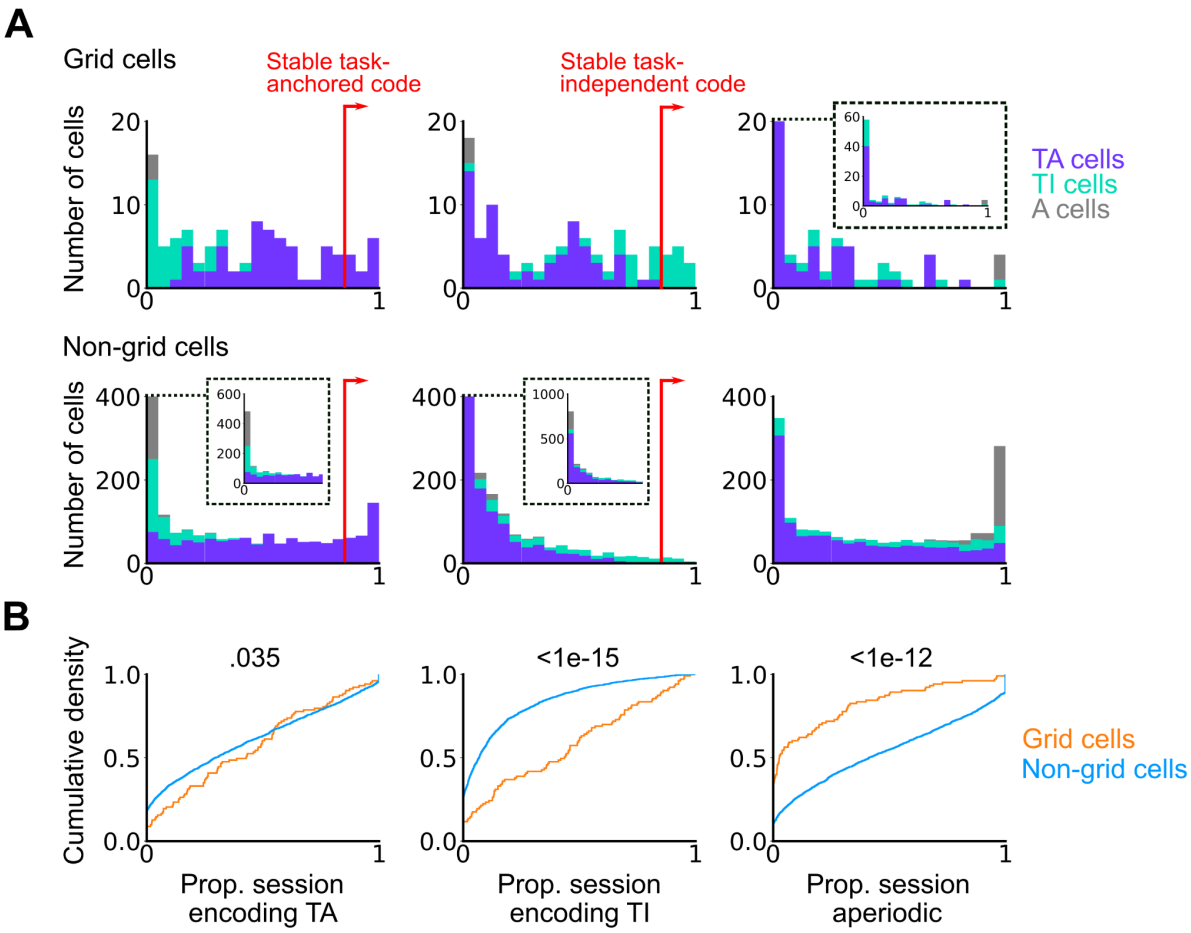

Figure 3. Figure Supplement 5. Grid cells exhibit task-independent periodic codes more frequently than non-grid cells. Extension of Figure 3F.

(A) Distribution of all recorded grid cells (top) and non-grid cells (bottom) encoding task-anchored (left), task-independent (centre) and aperiodic (right) epochs within a session. Session-level cell classifications are represented by the colour of the corresponding coding group.

(B) Cumulative density for grid and non-grid cells encoding task-anchored, task-independent and aperiodic epochs within a recording session. (Task-anchored: P = 0.035, KS = 0.14; task-independent: P < 1e-15, KS = 0.42; aperiodic: P < 1e-12, KS = 0.38, $N_{grid\ cells}$ = 103, $N_{non\text{-}grid\ cells}$ = 1778, Kolmogorov-Smirnov test).

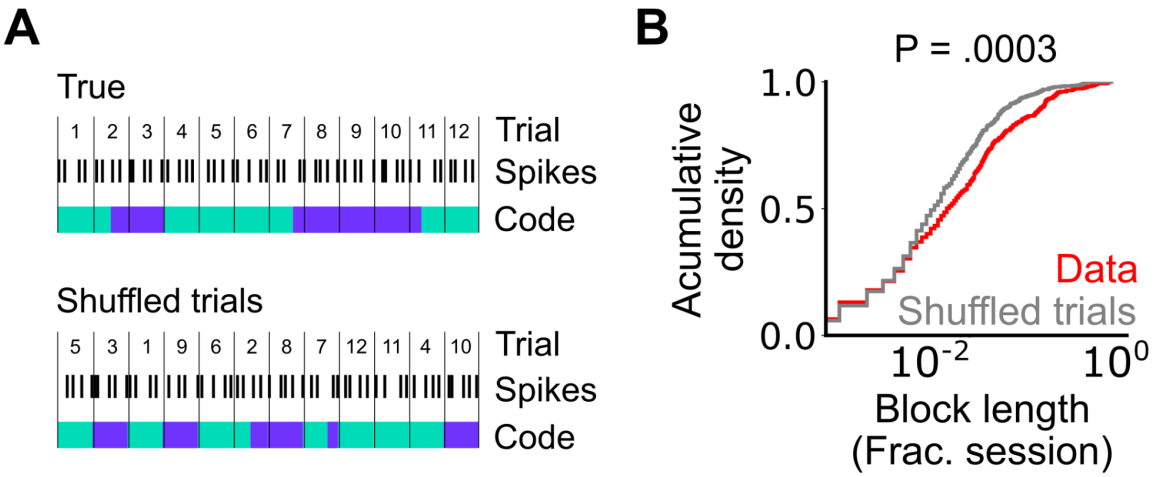

Figure 3. Figure Supplement 6. Assessment of the length of periodic coding blocks
(A) We compared the length of coding blocks - sequences of trials in which grid cells operate in the same mode - in the observed data with the length in datasets shuffled at the level of trials. To generate shuffled data spike locations were reallocated to a new position in the firing profile by randomly shuffling the trials (see Methods). The rolling classification was then recomputed and length of coding blocks measured.
(B) Cumulative histogram of block lengths for the observed grid cell population (red) and the same grid cell population with the trial order shuffled and blocks recalculated (grey). (P = 0.0003; KS = 0.066; $N_{blocks(data)}$ = 1765 $N_{blocks(shuffled)}$ = 2359; Kolmogorov-Smirnov test)

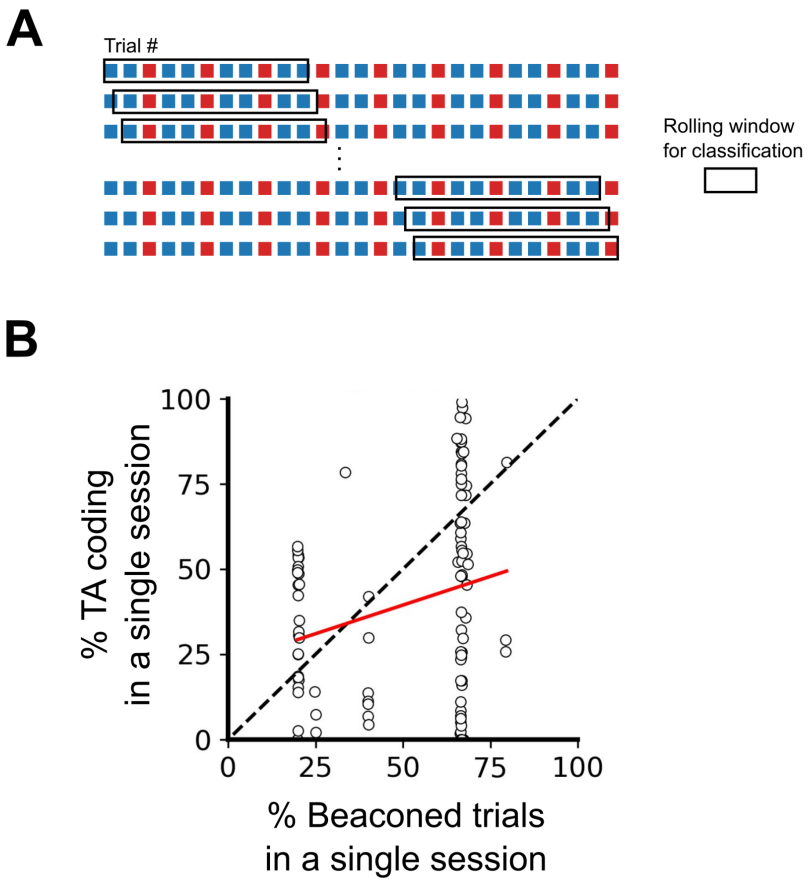

Figure 3, Figure Supplement 7. Differences in order and ratio of trial types do not explain the variability in task-anchoring.

(A) Example trial structure with repeating beaconed and non-beaconed trials in a ratio of 2 beaconed trials to 1 non-beaconed trial. The rolling window over which the rolling classification of periodicity is calculated is larger than the repeating trial block. Therefore attribution of a periodic classification to a single trial cannot be explained by the trial type or order.

(B) The percentage of trials in which each grid cell showed task anchored coding as a function of the percentage of beaconed trials experienced during the sessionThe percentage of beaconed trials did not have a significant effect on the occurrence of task-anchored codding ($P = 0.95$, $X^2 = 0.032$, $DF = 1$, ANOVA comparison of binomial family GLMMs with and without % beaconed trials as a fixed effect, with mouse ID, session ID and neuronal ID as random effects and a logit linker function).

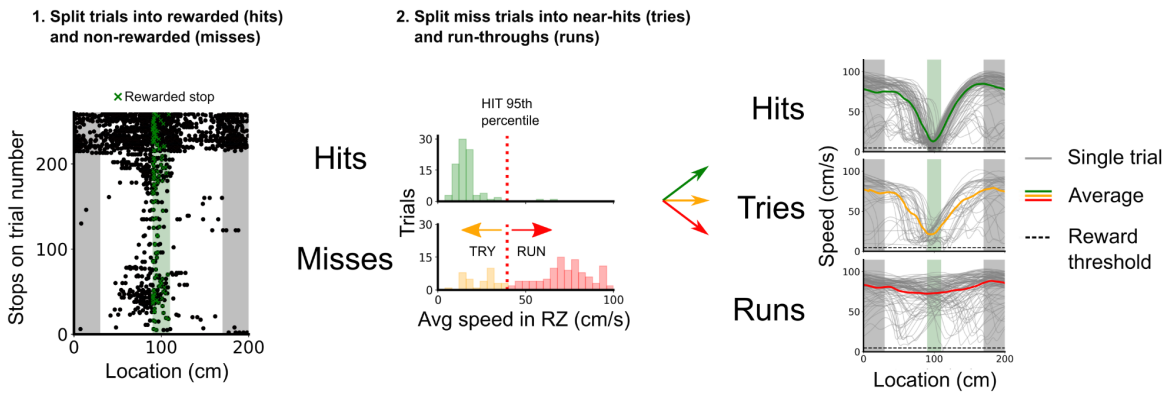

Figure 5, Figure Supplement 1. Classification of trial outcomes

Trials were classified into hits, tries and runs, based on stopping locations and the speed profile on a given trial. First, trials were classified into hits and misses based on whether a stop was registered in the reward zone or not. Next, a distribution of the average speeds in the reward zone is drawn for hit and miss trials, the 95th percentile of the hit average speeds in the reward zone was used to split the miss trials into near hits (tries) and run-throughs (runs). Finally, trials in which the mouse's average speed outside of the reward zone was < 10 cm/s were removed to better discriminate trial outcomes based on the mouse's engagement.

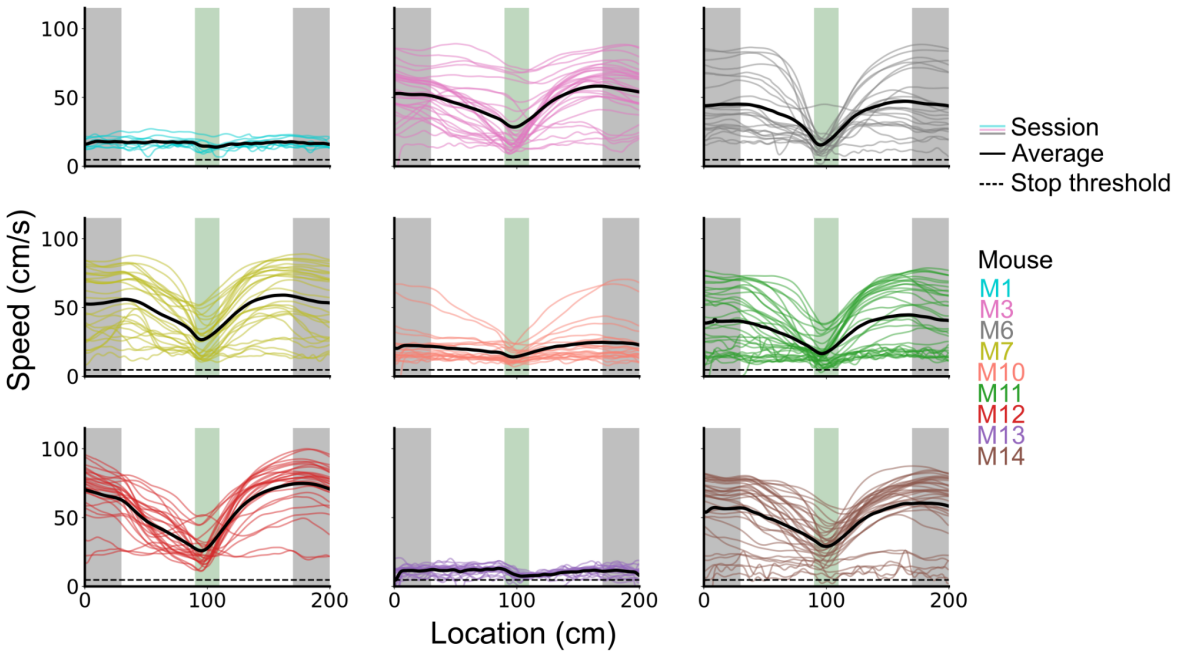

Figure 5, Figure Supplement 2. Speed profiles across behavioural sessions for each mouse.

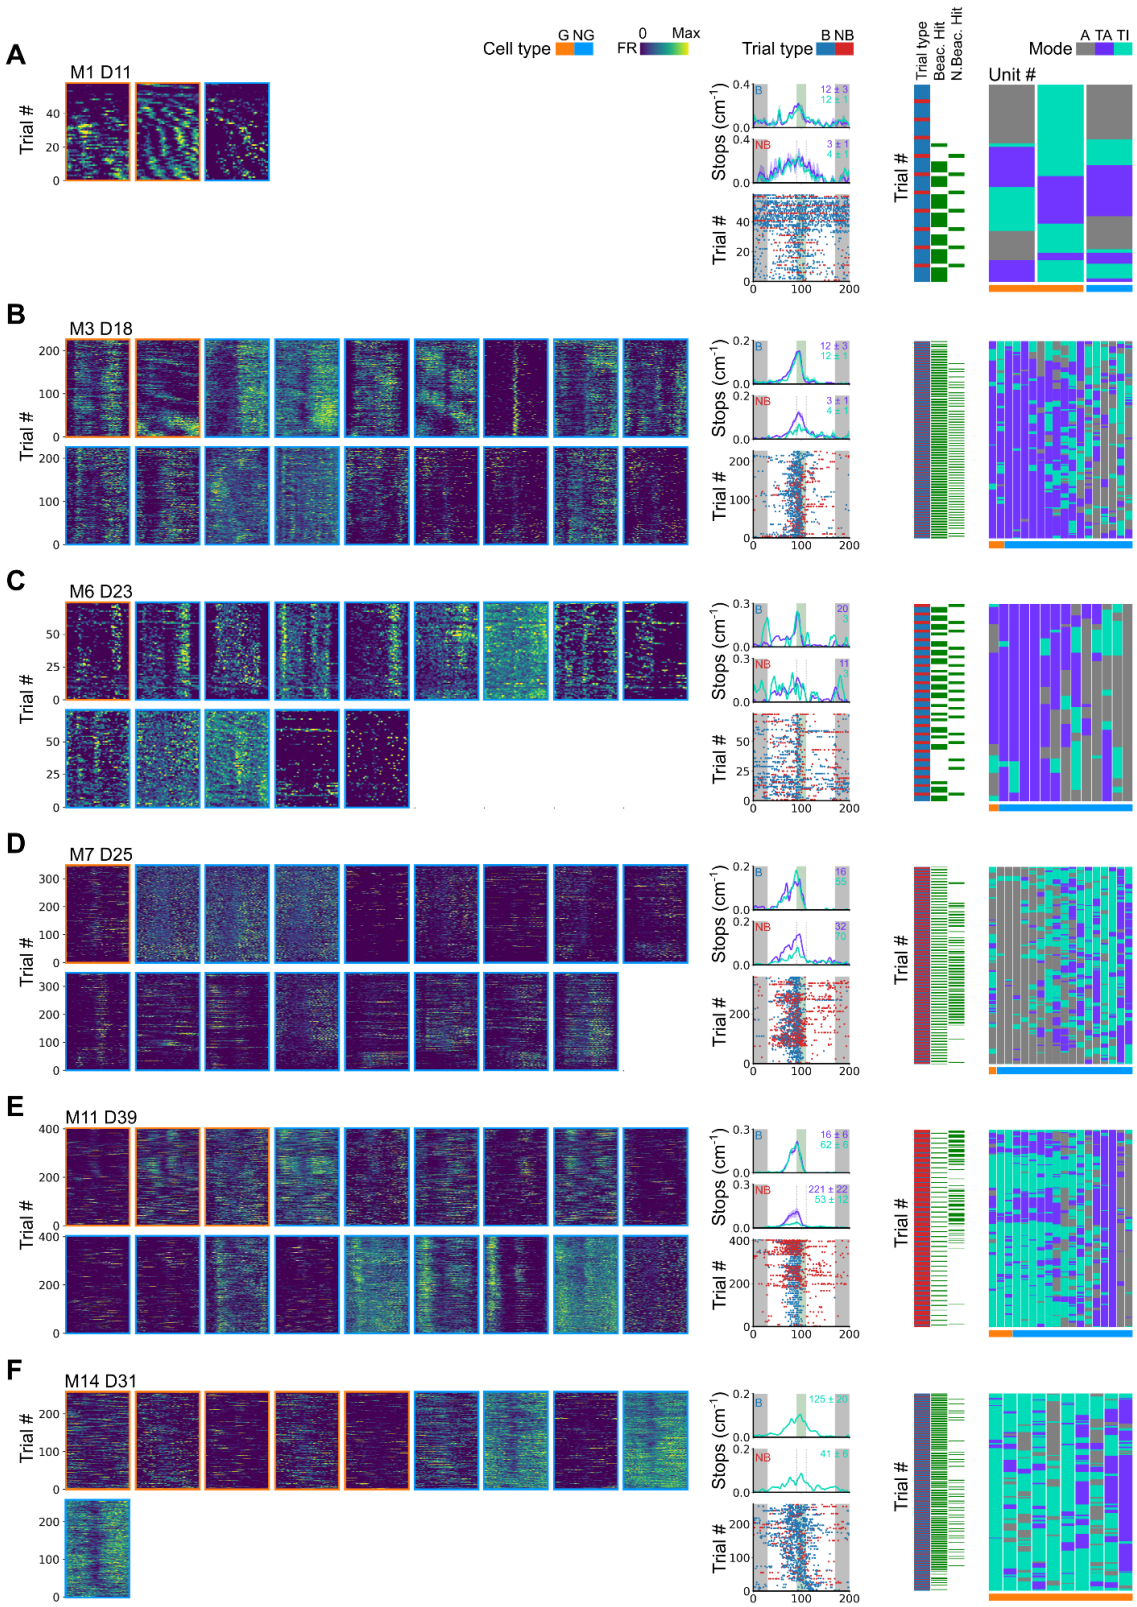

Figure 5, Figure Supplement 3. Further examples of spatial behaviour during task-anchored and task-independent grid modes

(A-F) Examples of variation in the behaviour-related activity of grid and non-grid cells recorded on the location memory task. Plots show all simultaneously recorded cells' firing

rate maps in each session (left), stop rasters (lower centre), stop density on beaconed (B) and non-beaconed (NB) trials coloured according to whether grid cells were task-anchored or task-independent (upper centre) and a summary of raster of behaviour and cell classifications (right). Shaded regions in stop density plots represent standard error of the mean measured across epochs. The number of trials classified in a particular coding scheme is also provided with the stop density plot. Grid cells and non-grid cells are colour coded by bounding boxes around the firing rate map.

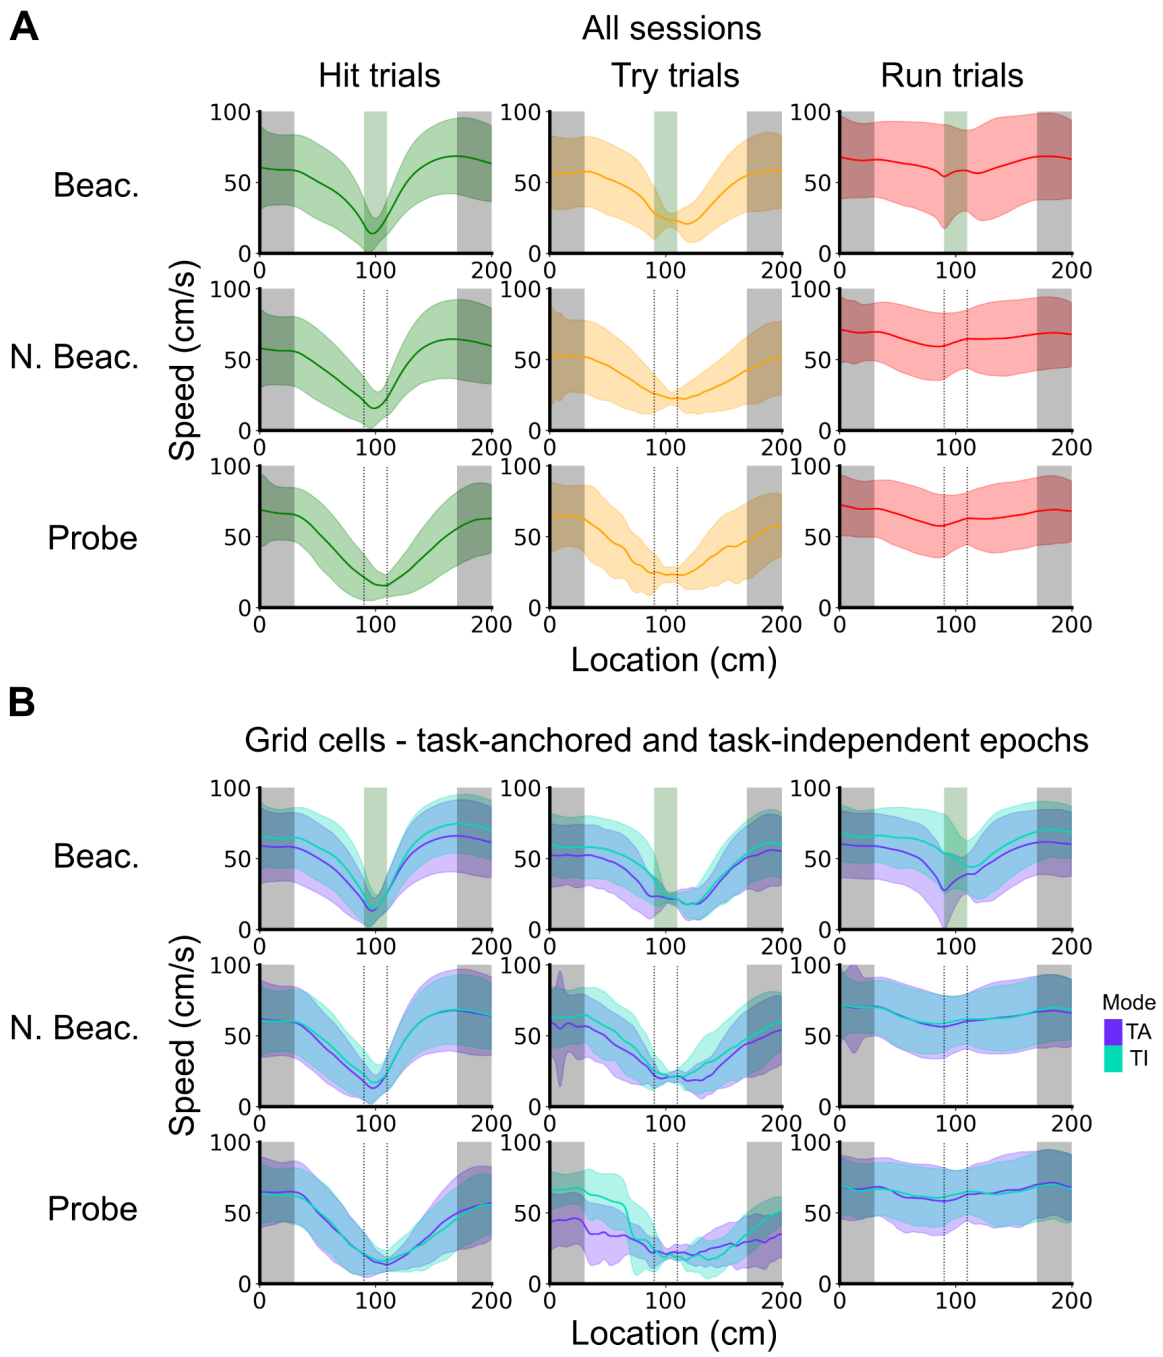

Figure 6, Figure Supplement 1. Speed profiles across different trial outcomes

(A) Running speed as a function of track location for all combinations of trial types and trial outcomes. Shaded regions represent standard deviations measured across trials.

(B) Same as in A, however trials were subdivided further based on the coding scheme of grid cells. In cases where multiple grid cells were recorded simultaneously, the same trial was used multiple times to calculate the average speed profile.

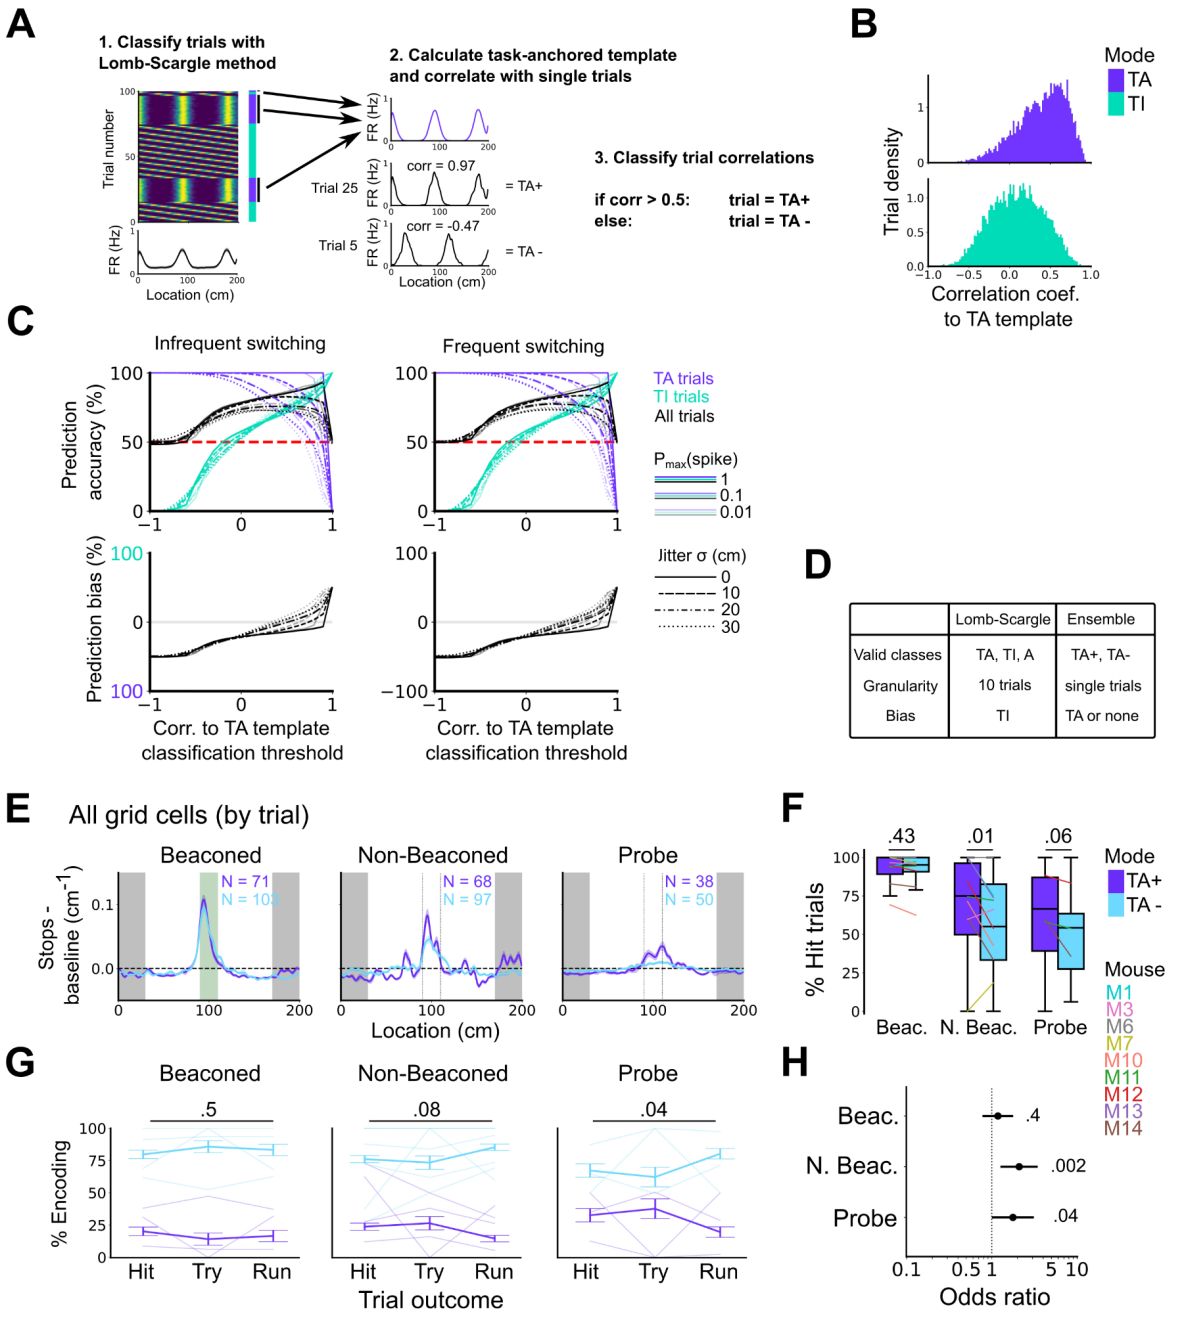

Figure 6, Figure Supplement 2. High correlation to the task-anchoring template of grid firing promotes localisation by path integration but is not required for cued localisation.

(A) Procedure for classifying trials using an ensemble method (Lomb-Scargle + template correlation).

(B) Trial density of all trial template correlations for trials classified as task-anchored (upper) or task-independent (bottom) using the Lomb-Scargle method. The same trial may be represented multiple times based on the number of simultaneously recorded cells.

(C) To assess classification accuracy using the ensemble method, we calculated the prediction accuracy and bias of simulated grid cells (see Methods). Chance level is denoted with a red dashed line. Plots on the left show the prediction accuracy (top) and the prediction

bias (bottom) when grid cells alternate between task-anchored and task-independent firing fields in blocks of trials (as in i, iii above), whereas the right plots shows the same information for a higher switching frequency (as in ii, iv above).

(D) Summary table detailing the pros and cons of each method.

(E) Stopping probability relative to baseline as a function of position for epochs within sessions containing both stable and unstable codes. (E). Shaded regions represent standard error of the mean measured across epochs. N represents the number of grid cells with coding epochs.

(F) Percentage of hits on beaconed, non-beaconed and probe trials when the code is task-anchored or task-independent for epochs within sessions are task-anchored or task-independent (ANOVA: Beaconed, P = 0.43, X^2 = 0.62, DF = 1; Non-beaconed, P = 0.01, X^2 = 6.43, DF = 1; Probe, P = 0.06, X^2 = 3.52, DF = 1).

(G) Percentage of trials with hit, try and run outcomes in which grid cell firing is highly correlated with task-anchored firing (ANOVA: Beaconed, P = 0.51, X^2 = 1.35, DF = 2; Non-beaconed, P = 0.08 , X^2 = 5.14, DF = 2; Probe, P = 0.04, X^2 = 6.67, DF = 2). Error bars denote standard error of the mean measured across the mean values for each animal. Faded lines show percentage values for individual mice that contained hit, try and run trials for a given trial type.

(H) Odds ratio between receipt of reward for trials on which epochs contain task-anchored firing relative to trials on which epochs contain task-independent firing (Beaconed, P = 0.4; Non-beaconed, P = 0.002; Probe, P = 0.04).

Supplemental Data

M1: MEC

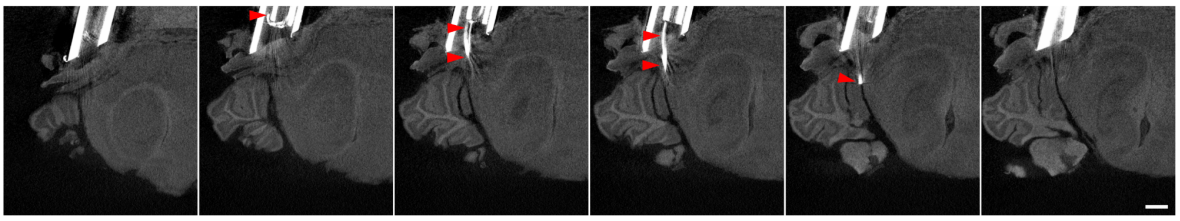

M2: Deep MEC

M3: MEC

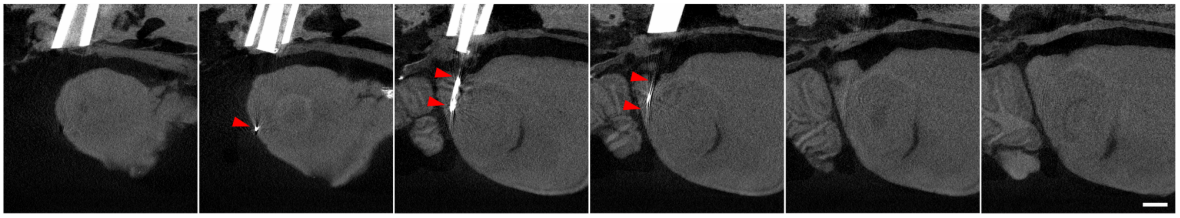

M6: MEC

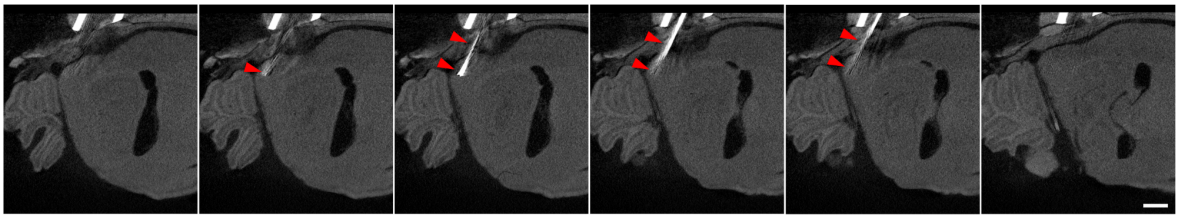

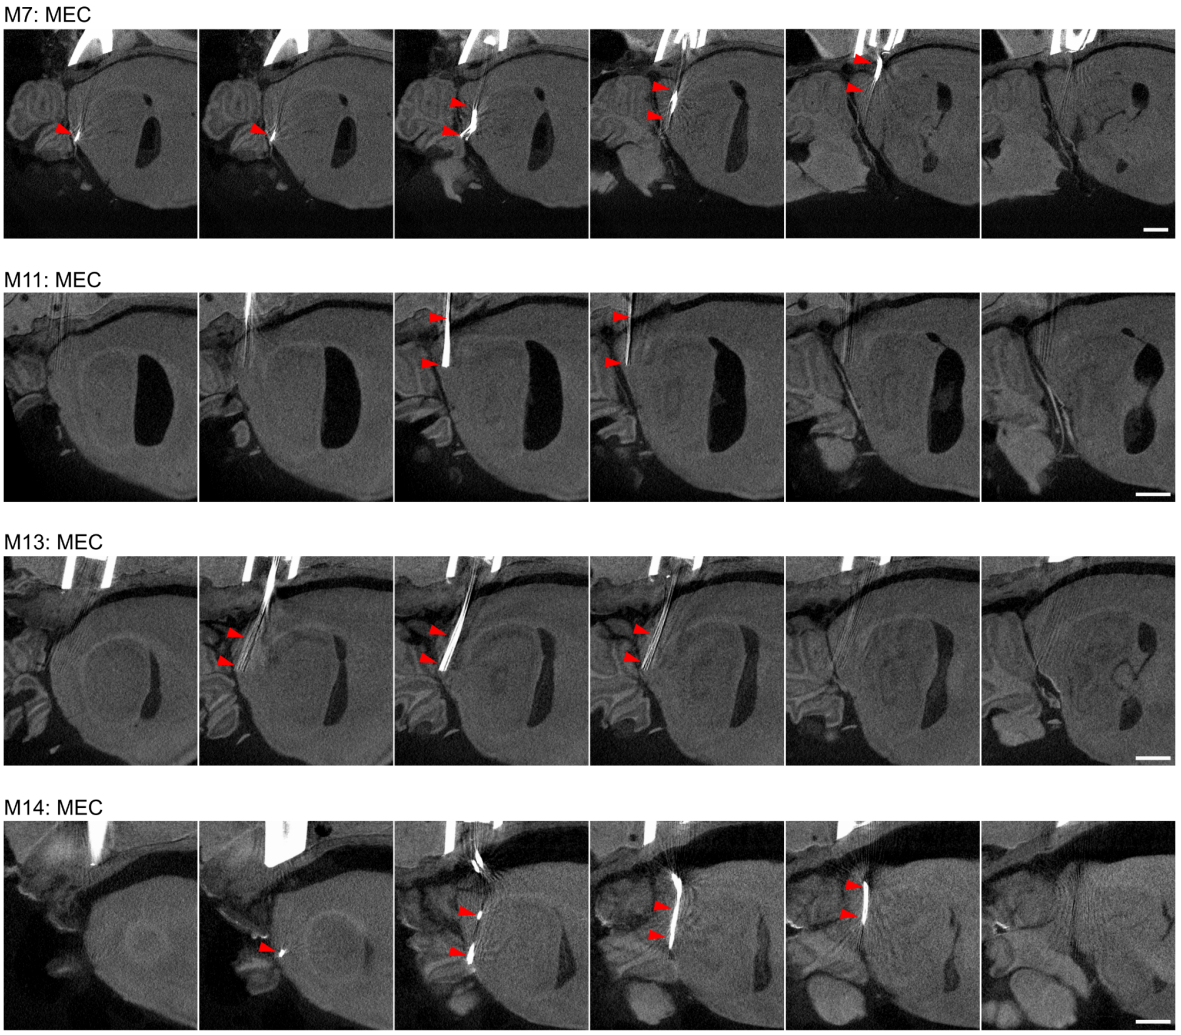

Supplemental Data 1. Tetrode localisation.
Micro-CT images used for assessment of tetrode locations in eight mice. For each animal, sagittal slices are presented lateral to medial from left to right and the classification of the tetrodes target is shown at the top left. This classification is based on the terminal location of the tetrode and the distance travelled during the experiment (see Methods). M10 sustained significant damage during preparation for microCT imaging and was not imaged. In all images red triangles point to the tetrode tracks. Scale bar denotes 1 mm.

Mouse	N cells	N sessions	N trials / session	N grid cells (%)	Tetrode location
M1	52	11	45.5 ± 33.6	3 (5.8)	MEC
M3	243	26	270.3 ± 93.5	7 (2.9)	MEC
M6	229	21	172.3 ± 99.8	1 (0.4)	MEC
M7	191	23	268.1 ± 137.8	4 (2.1)	MEC
M10	66	16	74.3 ± 34.8	1 (1.5)	Unclassified
M11	557	35	173.1 ± 126.1	53 (9.5)	MEC
M12	105	21	244.4 ± 68.5	1 (0.95)	Unclassified
M13	160	25	47.5 ± 25.1	7 (4.3)	MEC
M14	278	37	316.4 ± 148.2	26 (9.4)	MEC
Average	209 ± 152.5	23.9 ± 8.3	179.1 ± 103.4	11.4 ± 17.4 (4.5 ± 3.5)	
Total	1881	215		103	

Supplemental Data 2. Summary table of recorded cells and estimated tetrode locations. Tetrode locations were estimated from microCT images as Supplemental Data 2. For average measurements, standard deviations are provided. Mice used in Tennant et al. 2022 were M3, M6 and M7.