


Distance-tuned neurons drive specialized path integration calculations in medial entorhinal cortex

Graphical abstract



Authors

Malcolm G. Campbell, Alexander Attinger, Samuel A. Ocko, Surya Ganguli, Lisa M. Giocomo

Correspondence

giocomo@stanford.edu

In brief

Campbell et al. use Neuropixels recordings in mice navigating a VR environment to show that MEC neurons are more influenced by path integration than V1 and RSC neurons. These differences are driven by a subset of MEC neurons that exhibit modular distance tuning in darkness, reminiscent of grid cells.

Highlights

- Path integration influences neural activity in MEC more than in V1 or RSC
- Position is encoded retrospectively in V1 and RSC, but prospectively in MEC
- A subset of MEC cells is tuned to distance run in the dark with modular structure
- Distance-tuned neurons drive many of MEC's unique cue integration properties



Article

Distance-tuned neurons drive specialized path integration calculations in medial entorhinal cortex

Malcolm G. Campbell,^{1,3,4} Alexander Attinger,^{1,3} Samuel A. Ocko,² Surya Ganguli,^{1,2} and Lisa M. Giocomo^{1,5,*}

¹Department of Neurobiology, Stanford University, Stanford, CA 94305, USA

²Department of Applied Physics, Stanford University, Stanford, CA 94305, USA

³These authors contributed equally

⁴Present address: Department of Molecular and Cellular Biology, Harvard University, Cambridge, MA 02139, USA

⁵Lead contact

*Correspondence: giocomo@stanford.edu

<https://doi.org/10.1016/j.celrep.2021.109669>

SUMMARY

During navigation, animals estimate their position using path integration and landmarks, engaging many brain areas. Whether these areas follow specialized or universal cue integration principles remains incompletely understood. We combine electrophysiology with virtual reality to quantify cue integration across thousands of neurons in three navigation-relevant areas: primary visual cortex (V1), retrosplenial cortex (RSC), and medial entorhinal cortex (MEC). Compared with V1 and RSC, path integration influences position estimates more in MEC, and conflicts between path integration and landmarks trigger remapping more readily. Whereas MEC codes position prospectively, V1 codes position retrospectively, and RSC is intermediate between the two. Lowered visual contrast increases the influence of path integration on position estimates only in MEC. These properties are most pronounced in a population of MEC neurons, overlapping with grid cells, tuned to distance run in darkness. These results demonstrate the specialized role that path integration plays in MEC compared with other navigation-relevant cortical areas.

INTRODUCTION


In mammals, navigation engages brain regions ranging from primary sensory to higher-order associative areas. During navigation, this range supports the integration of the multiple cues encountered that allow animals to generate an estimate of their position in space. As animals navigate through the world, they must integrate two sources of information, external input from landmark cues and internally generated path integration predictions based on self-motion cues, to compute their position. However, although recent work points to commonalities in the representation of sensory and behavioral variables across cortical areas (Allen et al., 2019; Clancy et al., 2019; Minderer et al., 2019; Musall et al., 2019; Pinto et al., 2019; Stringer et al., 2019), the degree to which cortical areas involved in navigation follow common or specialized algorithms for computing position remains incompletely understood.

To address this question, we considered how the representation of an animal's position is computed in three interconnected brain regions hypothesized to provide complementary computations for visually guided navigation in mice: the medial entorhinal cortex (MEC), primary visual cortex (V1), and retrosplenial cortex (RSC) (Miller and Vogt, 1984; Sugar et al., 2011). Common across these regions are neural representations for an animal's position during navigation and modulation of neural activity by

locomotion (Alexander and Nitz, 2015; Chen et al., 2019; Clancy et al., 2019; Diehl et al., 2017; Fischer et al., 2020; Fiser et al., 2016; Fournier et al., 2020; Hafting et al., 2005; Hardcastle et al., 2017; Keller et al., 2012; Kropff et al., 2015; Mao et al., 2017; Niell and Stryker, 2010; Saleem et al., 2013, 2018). Moreover, MEC, V1, and RSC have all been proposed to implement algorithms for reconciling sensory input with internal models of predicted sensory feedback or spatial position (Alexander and Nitz, 2015; Campbell et al., 2018; Keller et al., 2012). However, MEC, V1, and RSC also show specialized functional coding properties. MEC neurons encode an animal's position and orientation in a world-centered reference frame (Diehl et al., 2017; Hafting et al., 2005; Høydal et al., 2019; Sargolini et al., 2006; Solstad et al., 2008). V1 neurons respond to specific visual stimuli and predicted visual feedback (Hubel and Wiesel, 1962; Keller et al., 2012). RSC neurons conjunctively represent the animal's position in both allocentric and egocentric reference frames (Alexander and Nitz, 2015, 2017). These studies imply some degree of specialization across MEC, RSC, and V1, but the diversity of task and training protocols makes it difficult to directly compare them.

Here, we addressed this gap by using Neuropixels probes to simultaneously record from tens to hundreds of MEC, V1, or RSC neurons in exactly the same task and training conditions as mice navigated virtual reality (VR) environments in which we





(legend on next page)

systematically perturbed the visual cues. The data revealed that path integration influences MEC more than V1 or RSC, and that these calculations are supported by a unique population of MEC neurons tuned to distance run in darkness.

RESULTS

Differential impact of visual cues on spiking activity in MEC, V1, and RSC during navigation

We trained head-fixed mice to traverse a VR linear track for water rewards (Figures 1A and 1B). After training (~2 weeks, >1,000 trials), mice developed stereotyped behavior, slowing down and licking prior to the reward tower (Figure 1C). Mice then performed this task after acute implantation of a silicon probe, which allowed us to record from thousands of single cells (i.e., units) from MEC, V1, and region A29 of RSC (MEC cell: $n = 14,799$, session $n = 90$, mouse $n = 20$; V1 cell: $n = 1,524$, session $n = 24$, mouse $n = 9$; RSC cell: $n = 1,757$, session $n = 18$, mouse $n = 9$) (Figures 1D and 1E; Figure S1A). The cell yield per contact did not differ between brain regions (one-way ANOVA, $F(2,133) = 2.5$, $p > 0.085$), but as a result of the correspondence between the geometry of the probe and MEC anatomy, the number of units per recording session was highest for MEC (mean $\pm$ SEM; MEC: 166 ± 11 cells/session, V1: 65 ± 7 cells/session, RSC: 73 ± 13 cells/session, one-way ANOVA, $F(2,133) = 18.79$, $p < 6.7e-8$, MEC different from V1: $p < 6.9e-7$, MEC different from RSC: $p < 5.8e-6$, Mann-Whitney U test) (Figure 1F). Consistent with a previous report (Hernández-Pérez et al., 2020), population activity traveled from dorsal to ventral MEC in theta-paced waves, whereas we did not observe such activity in RSC or V1 (Figure S1B).

In MEC, V1, and RSC, many cells fired at consistent VR track positions (Figure 1G). We defined such cells as “spatially stable” using a cross-correlation metric (Figure 1H; STAR Methods).


Averaging the firing rates of all spatially stable cells across the VR track revealed clear population-level firing pattern differences between the three regions. In V1 and RSC, firing rates peaked ~20 cm before each visual landmark (Figures 1I and 1J), with the receptive field location influencing spatial firing rate maps in V1 (Figure S1C). In MEC, the average firing rate was relatively constant over the VR track (Figures 1I and 1J). Even so, the position of the mouse on the VR track could be decoded with high accuracy from the neural activity of any of the three regions (Figure 1K). However, we observed more error in decoded position estimates from V1 neural activity. This decoded error emerged from misclassified tower locations, possibly as a result of similar population-level firing patterns in V1 near visual landmarks, leading to the increased prevalence of 80- and 160-cm errors (Figure 1K, inset; Figure S1D). Consistent with this, state space trajectories appeared more convoluted for populations of V1 neurons (Figure 1L). Together, these analyses reveal that neurons in MEC, V1, and RSC all carry information regarding an animal’s spatial position, but that visual landmark cues influence these neural representations in V1 and RSC more than MEC. Nevertheless, visual landmarks were necessary, but not sufficient, to drive spatial firing patterns in MEC (Figures S1E–S1H). Specifically, passive viewing significantly reduced activity and spatial stability in MEC, but spatial patterns were also lost when the monitors were turned off, consistent with dark recordings in open field environments (Chen et al., 2016; Pérez-Escobar et al., 2016).

Conflicts between landmarks and path integration change neural activity patterns in MEC more than in V1 or RSC

We next examined the degree to which visual landmarks versus path integration determined neural activity patterns in MEC, V1, and RSC. We created visuomotor cue conflicts by reducing the

Figure 1. Differential impact of visual landmarks on spiking activity in MEC, V1, and RSC

- (A) VR setup.
(B) Top: schematic of VR track. Bottom: view from behind a mouse.
(C) Average running speed (top) and lick rate (bottom) for all 32 mice (dashed lines = tower locations). Mice slowed down and started licking when approaching the reward tower, indicating they had learned the task.
(D) Example histology images showing probe tracks in MEC, V1, and RSC.
(E) Schematic of dorsal brain surface with location of probe insertion on the cortical surface for all 137 recording sessions in 32 mice. Color of dot indicates the targeted brain region. Dashed lines indicate borders of V1, as well as the borders of A29 and A30 of RSC. Note that we focused on A29 of RSC in all subsequent figures and analyses.
(F) Number of well-isolated units in each recording session. Only units within the target brain area were included.
(G) Spatial firing rate maps across blocks of 16 baseline trials for six example units each in MEC, V1, and RSC (dashed lines = tower locations). Numbers above each rate map: peak firing rate and stability score for the given block of trials.
(H) Histogram of stability scores across blocks of 16 baseline trials. Stability (Stab) scores were higher in V1 and RSC than in MEC (one-way ANOVA, $F(2,96) = 5.15$, $p < 0.0075$, $\text{Stab}_{\text{MEC}} < \text{Stab}_{\text{V1}}$: $p < 0.016$, $\text{Stab}_{\text{MEC}} < \text{Stab}_{\text{RSC}}$: $p < 0.002$, one-sided Mann-Whitney U test). In (H)–(J), values were averaged within session first, then across sessions, and shaded area indicates SEM over sessions, not cells.
(I) Average firing rate of all spatial cells in MEC (green), V1 (red), and RSC (blue) with respect to track location (dashed lines = tower locations). RSC firing rates (FRs) were significantly higher overall (one-way ANOVA, $F(2,94) = 12.17$, $\text{FR}_{\text{RSC}} > \text{FR}_{\text{V1}}$: $p < 1.4 \times 10^{-4}$, $\text{FR}_{\text{RSC}} > \text{FR}_{\text{MEC}}$: $p < 0.005$, one-sided Mann-Whitney U test).
(J) Average spatial autocorrelation for all spatial cells. V1 and RSC, and to a lesser extent MEC, show peaks at lags corresponding to distance between adjacent towers.
(K) Position decoding error as a function of the number of included spatial neurons. Inset shows the distribution of errors for a decoder trained on 20 neurons. Black bar at bottom indicates $p < 0.05$ for $\text{Error}_{\text{V1}} > \text{Error}_{\text{MEC}}$ and $\text{Error}_{\text{V1}} > \text{Error}_{\text{RSC}}$, Mann-Whitney U test. Shading indicates SEM across folds.
(L) Population-level firing rate dynamics during baseline projected onto two dimensions using Uniform Manifold Approximation and Projection (UMAP) for three example sessions (MEC: left, V1: center, RSC: right). Consistent with more frequent misclassification of towers (K) (Figure S1C) and more structured spatial autocorrelation (J), trajectories of population of V1 neurons appeared more convoluted (center), compared with ring-like structures in populations of MEC and RSC neurons.



(legend on next page)

gain between the rotation of the running wheel and the mouse's progression along the VR by a factor of 0.8 for blocks of four trials (number of blocks per session = 1–4, median 2) (Figure 2A) (Campbell et al., 2018). This gain manipulation had minimal effects on the animals' behavior (Figure 2B) and no effect on average firing rates (Figure S2A). If the location of spikes is determined solely by the visually determined position of the mouse in the virtual environment, neural activity patterns should not shift positions between baseline and gain trials. Conversely, if position estimates are influenced by actual distance run on the treadmill due to self-motion, then neural activity patterns will undergo negative (leftward) shifts when plotted in VR coordinates, because the VR now advances less for a given number of steps taken. Alternatively, gain changes could trigger global remapping, which was previously shown to occur in MEC grid cells for large enough conflicts between path integration and landmarks (Campbell et al., 2018).

For these analyses, we considered only spatially stable cells (average peak trial-trial firing rate correlation ≥ 0.5 in the six baseline trials preceding the gain change). In MEC, cells responded to the change in gain either by coherently shifting the spatial firing patterns or remapping, consistent with previous work (Campbell et al., 2018). When cells remapped in response to the gain change, they either returned to the original baseline map following the gain change or they did not (Figure 2C). In contrast, firing patterns in V1 and RSC remained remarkably stable both during and after the gain change (Figure 2C), consistent with the idea that visual cues strongly influenced firing patterns in these brain regions (Figure 1I).

To quantify the effect of the gain change on neural firing patterns, we computed similarity matrices by correlating firing patterns on pairs of trials (Figure 2D). We averaged the similarity matrix over cells for a gain change block and considered each block of trials separately, because remapping patterns and the number of spatially stable cells varied across blocks (MEC: trial

block $n = 123$, session $n = 51$, mouse $n = 18$, 51 ± 3 [SEM] cells per block; V1: trial block $n = 29$, session $n = 15$, mouse $n = 6$, 20 ± 2 [SEM] cells per block; RSC: trial block $n = 25$, session $n = 13$, mouse $n = 7$, 55 ± 8 [SEM] cells per block) (Figure 2E; Figures S2B and S2C). In MEC, we observed variability in the similarity matrices (Figure 2E, left), whereas in V1 and RSC these matrices almost always indicated that neural representations remained stable between baseline, gain change, and post-gain change trials (Figure 2E, right; Figures S2D–S2G). This was the case even for cells with lower baseline stability (Figure S2D). Overall, the gain change altered spatial representations in MEC more than in V1 or RSC, even after accounting for the lower baseline stability in MEC (Figure 2D, bottom).

To identify distinct types of population-level gain change responses, we employed principal-component analysis (PCA) (Figure 2F) and classified population responses as either "Stable" (cluster 1) or "Remapping" (cluster 2, Figures 2F and 2G). In MEC, many trial blocks were assigned to cluster 2, indicating cue conflict often led to remapping of spatial firing patterns (blocks in cluster/total blocks: MEC, cluster 1 = 24/123, cluster 2 = 99/123). This gain-change-induced remapping was remarkably coordinated across co-recorded neurons in MEC (Figures 2H–2J), with remapping patterns more similar within block than between blocks (Figure 2J). We also commonly observed spontaneous remapping in MEC in the absence of any gain change (blocks in cluster/total baseline blocks: cluster 1 = 42/96, cluster 2 = 54/96) (Figure 2G; Figures S2H–S2J) (Low et al., 2020). However, the frequency of remapping was significantly higher during gain change blocks than baseline blocks (Fisher test, $p = 0.00018$). In contrast with the variability in MEC responses, in V1 and RSC, nearly all trial blocks were assigned to cluster 1 (blocks in cluster/total blocks: V1, cluster 1 = 26/29, cluster 2 = 3/29; RSC, cluster 1 = 24/25, cluster 2 = 1/25). This suggests that in V1 and RSC, visual landmarks drive spatial firing patterns, even when visual landmark and path

Figure 2. Conflicts between landmarks and path integration change neural activity patterns in MEC more than in V1 or RSC

- (A) Trial structure for gain change experiments. Baseline trials (black) were interleaved with blocks of four gain change trials (cyan).
- (B) Average running speed along VR track during baseline (black) and gain change trials (cyan). Dotted lines indicate the position of landmarks. Average running speed did not differ between the two conditions (average speed $\pm$ SEM: baseline = 47.0 ± 9.7 cm/s, gain change = 47.4 ± 9.2 cm/s, $p = 0.167$, 216 gain change blocks, Wilcoxon test).
- (C) Spatial firing rate maps for 16 trials surrounding gain change blocks, for example, MEC, V1, and RSC neurons (black bars: baseline trials, cyan bars: gain change trials). V1 and RSC firing patterns were mostly unchanged during gain manipulations (right), while MEC neurons often remapped (left).
- (D) Top: trial-trial similarity matrices for each brain region, averaged over gain change blocks. Bottom: average of the first six rows of the similarity matrices above, excluding the diagonal (error bars indicate SEM over trial blocks).
- (E) As in (D) but each matrix represents an individual gain change block. The average similarity matrices in (D) obscured substantial variability in the degree and pattern of MEC remapping (left), whereas spatial firing patterns of V1 and RSC neurons were more stable and rarely remapped (right).
- (F) Projection onto the first principal component for each gain change block. A threshold of PC1 score = 0.4, defined as the location of the minimum between the two peaks of the score distribution (left), divided gain change blocks into two groups: "Stable" (cluster 1) and "Remap" (cluster 2).
- (G) Top: average similarity matrices for the Stable and Remap clusters identified in (F). Bottom: number of gain change blocks assigned to each cluster by brain region. For comparison, we ran the same analysis on MEC baseline data in which no gain change occurred. Remapping could occur spontaneously in MEC (see also Figure S2J). However, remapping was more frequent in response to a gain change (Fisher test, $p = 1.8e-4$).
- (H) Similarity matrix of an example cluster 2 (Remap) MEC gain change block. Neurons remap in gain change trial 4 (arrow), leading to low pairwise similarity between trials before and after remapping.
- (I) Spatial firing rate maps of four example MEC neurons during the gain change block shown in (H). During the first three gain change trials, the spatial maps are shifted relative to the preceding baseline map. The neurons remap in gain change trial 4 (arrows). Spatial maps in postbaseline trials (trials 11–16) are different from the spatial maps in baseline trials.
- (J) For each MEC trial block in cluster 2 (99 trial blocks total), we assessed coordination of remapping by comparing the correlation of spatial similarity maps for simultaneously recorded spatial cells ("within block") to all other similarity maps ("across block"). Similarity maps were more similar within than across blocks, consistent with population-wide remapping ($p = 8.2e-16$, Wilcoxon signed rank test).
- GC, gain change.




Figure 3. The influence of path integration on position estimates is greater in MEC than in V1 or RSC

(A) Number of spatial cells per gain change block identified as "stable" (cluster 1) in Figure 2G. Bars indicate mean. (B) Example spatial firing rate maps for cells from cluster 1 blocks in MEC ($n = 4$), V1 ($n = 2$), and RSC ($n = 2$). Note the larger map shifts during gain change trials in MEC compared with V1 and RSC. (C) Average trial-trial shift matrices for MEC, V1, and RSC (black: baseline, cyan: gain change). Shifts were computed as the location of the peak spatial cross-correlation for each cell, then averaged over cells within gain change blocks. Red rectangles indicate samples used for estimating average map shift. (D) Average map shifts between gain change trials and baseline trials (outlined in red in C) in MEC, V1, and RSC. Each point represents a gain change block. Map shifts were larger in MEC than in V1 or RSC. Bars indicate mean. *** $p < 0.001$ for map shift significantly different from zero, Wilcoxon signed rank tests.

integration cues are in conflict (Fournier et al., 2020; Saleem et al., 2013, 2018).

Even in the absence of remapping, path integration influences position estimates in MEC more than in V1 or RSC

Although we observed variability in how the MEC neural population responded to the gain change, we focused on "stable" trial blocks (cluster 1) in order to compare cue integration mechanisms across brain regions (MEC: 24 cluster 1 blocks, [mean $\pm$ SEM] 65 ± 8 spatial cells per block; V1: 26 cluster 1 blocks, 21 ± 3 cells per block; RSC: 24 cluster 1 blocks, 55 ± 9 cells per block; Figure 3A). In all three brain regions, we observed that even in stable blocks, firing rate maps shifted coherently in the direction predicted by path integration on gain change trials (Figure 3B). We used spatial cross-correlation to quantify the shift between pairs of trials (Figure 3C). Map shifts on gain change trials were significantly different from zero in all three brain regions (mean map shift $\pm$ SEM: MEC = -4.1 ± 0.5 cm, $p = 2.1 \times 10^{-5}$; V1 = -1.8 ± 0.5 cm, $p = 0.00022$; RSC = -1.4 ± 0.4 cm, $p = 0.0047$, Wilcoxon tests), but these shifts were largest in MEC (one-way ANOVA, $F(2,71) = 8.53$, $p = 0.00048$; shifts in MEC larger than V1, $p = 0.00057$, and RSC, $p = 0.00093$, Mann-Whitney U test; Figure 3D; Figure S3A), indicating a stronger influence of path integration on position estimates in MEC compared with V1 or RSC. Moreover, we found that in MEC, but not V1 and RSC, the map shifts of individual cells systematically deviated from the population average (Figures S3C and S3D). This indicated

that in MEC, cells consistently fell on a spectrum from more landmark driven to more path integration driven, whereas in RSC and V1, the relative influence of path integration versus landmarks was constant across cells.

Position is encoded retrospectively in V1 and RSC but prospectively in MEC

Although the gain-change-induced map shifts we observed could reflect the influence of path integration, map shifts could also be caused by temporal delays between the encoded position and the time of spikes (Figure 4A). We therefore considered the extent to which V1, RSC, and MEC encode exact current position versus past or future position (retrospective/prospective coding) by examining the effect of running speed on spike position in baseline trials (Figure 4B).

We considered blocks of baseline trials and sorted these trials by the running speed of the mouse around the location of peak firing rate for each cell separately. In V1 and RSC, the VR location at which spikes occurred often shifted positively with running speed (Figure 4C), consistent with retrospective coding. In contrast, spikes of MEC cells often shifted in the opposite direction, consistent with prospective coding (Figure 4C). We computed a temporal delay factor for each neuron by maximizing the average trial-trial similarity (Figure 4D). In V1, and to a lesser extent RSC, the average delay factor per session was significantly greater than zero (retrospective coding; mean delay $\pm$ SEM: V1 = 0.12 ± 0.02 s, $p < 4.2 \times 10^{-4}$, RSC = 0.06 ± 0.02 s, $p = 0.016$; Figure 4E). In contrast, the average delay




Figure 4. V1 and RSC code position retrospectively, whereas MEC codes position prospectively

(A) Cartoon illustrating retrospective coding (spikes encode past position, left) versus prospective coding (spikes encode future position, right). Prospective coding makes use of path integration calculations to predict future position. Retrospective coding can reflect delays in sensory processing. Δt , time delay between the mouse being located at the encoded position and the occurrence of spikes (sign convention: $\Delta t > 0$ for retrospective coding and vice versa); v , running speed; Δx , spatial offset between encoded position and the position at which spikes occur.

(B) Schematic showing the predicted influence of running speed on spike position and encoded position, but in opposite directions for retrospective versus prospective coding. Faster running speeds lead to larger shifts in spike position and encoded position, but in opposite directions for retrospective versus prospective coding.

(C) Example spike raster plots for two MEC, V1, and RSC units for 16 baseline trials, around the location of the peak average firing rate, indicated by the red line. Trials are ordered by running speed during the 10 cm preceding the peak. Running speed for each trial is shown on the right of each raster plot.

(D) Temporal delay correction for an example V1 neuron. Top: raster plot with trials ordered by running speed. We searched for the factor that maximized the trial-by-trial correlation of firing rate maps (+110 ms for this unit) and accounted for this delay by changing the time of each spike (equation in middle).

(E) Average estimated temporal delay for each recording, split by brain region. Each dot represents the average temporal delay in a recording session (MEC: 56 sessions, V1: 18 sessions, RSC: 14 sessions). Estimated temporal delay differed from zero in all regions (MEC, $p \leq 0.0068$, V1: $p \leq 4.2 \times 10^{-4}$, RSC: $p \leq 0.016$, Wilcoxon signed rank test). The estimated delay factors were significantly larger in V1 compared with RSC or MEC (one-way ANOVA, $F(2,82) = 13.25$, $p < 1.3 \times 10^{-5}$, Delay_{V1} > Delay_{MEC}, $p \leq 1.3 \times 10^{-7}$, Delay_{V1} > Delay_{RSC}, $p \leq 3.1 \times 10^{-3}$, Mann-Whitney U test). Bars indicate mean.

(F) Change in map shift after spike times was corrected by the optimal temporal delay factor for each cell. Map shifts increased in MEC ($p \leq 0.0027$, Wilcoxon signed rank test) and decreased in V1 and RSC (V1: $p \leq 1.7 \times 10^{-5}$, RSC: $p \leq 0.0015$, Wilcoxon signed rank test).

(G) Residual map shifts after correcting for temporal delay, isolating the influence of path integration on position estimates. Residual map shifts in V1 and RSC were non-zero (V1: $p \leq 0.036$; RSC: $p \leq 0.032$, Wilcoxon signed rank tests).

factor was negative in MEC, indicating prospective coding (mean delay $\pm$ SEM, MEC = -0.01 ± 0.006 s, $p = 0.0068$; Figure 4E) (Battaglia et al., 2004; De Almeida et al., 2012; Kropff et al., 2015). A decoding analysis gave qualitatively similar results (Figure S3F).

Consistent with presence of retrospective coding in V1/RSC and prospective coding in MEC, correcting the spike times of individual cells by their delay factor significantly reduced map shifts in V1 and RSC but increased map shifts in MEC (Figures 4F and 4G). This supports the idea that, during gain change,




Figure 5. Map shifts are influenced by visual contrast in MEC, but not V1 or RSC

(A) Experimental design to test the influence of visual contrast on gain change responses.

(B) Percent of cells with stable gain change responses in high contrast (black), low contrast (gray), or both (green), included in the following analyses (MEC, 400/3,490, 11%; V1, 65/362, 18%; RSC, 137/937, 15%).

(C) Example neuron firing rates in MEC (left), V1 (middle), and RSC (right) during gain changes in high ($c = 100$, top row) and low ($c = 10$, bottom row) contrast. Left panels: black/gray: average rate map during baseline trials; cyan: firing rate map during a single example gain change trial. Right panels: cross-correlation between gain change and baseline maps. The location of the cross-correlation peak, indicated by a red circle, defined the map shift. For each cell, shifts were averaged across trials.

(D) Change in gain-change-induced map shifts between high and low contrast for all cells identified in (B). Negative values indicate increased leftward shifts in low contrast. Only MEC showed increased shifts in low contrast (change in map shift: mean $\pm$ SEM: MEC = -1.5 ± 0.5 cm, $p = 2.7 \times 10^{-5}$; V1 = 0.0 ± 0.3 cm, $p = 0.52$; RSC = -0.5 ± 0.5 cm, $p = 0.39$; Wilcoxon signed rank tests). As a population, shift changes were larger for MEC than for V1 and RSC (MEC versus V1, $p = 0.010$; MEC versus RSC, $p = 0.037$; V1 versus RSC, $p = 0.35$; Mann-Whitney U tests). p values were not corrected for multiple comparisons.


(E) Same data as (D), but with each dot representing the median shift change within each session that had at least five spatially stable cells in both high and low contrast. MEC, but not V1 or RSC, showed significantly larger (more negative) map shifts in low contrast (mean shift change $\pm$ SEM over sessions, MEC = -2.0 ± 0.7 , $n = 14$ sessions, $p = 0.025$; V1 = 0.4 ± 0.4 cm, $n = 4$ sessions, $p = 0.25$; RSC = -0.8 ± 1.4 cm, $n = 6$ sessions, $p = 0.31$; Wilcoxon signed rank tests). Filled circles indicate sessions with significant within-session shift changes ($p < 0.05$, Wilcoxon signed rank tests). In MEC, 5/14 sessions showed significant within-session shift changes (larger leftward shift). In V1, 0/4 sessions had significant changes. In RSC, two sessions showed significantly larger leftward shifts (negative values), whereas another showed significantly smaller leftward shifts (positive values), and three showed no change.

(F) Average normalized firing rate during gain change trials, aligned to the location of peak firing rate in the baseline map. Black (gray) indicate contrast = 100 (10). Shaded area indicates SEM over cells. As in (D) and (E), there were larger leftward map shifts in low contrast in MEC, but not V1 or RSC. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

temporal delays do not drive the map shifts observed in MEC but do contribute to map shifts in V1 and, to a lesser extent, RSC. Moreover, the prospective coding, combined with large map shifts observed on gain change trials in MEC, raises the possibility that MEC uses path integration to predict current position before the arrival of delayed sensory input.

In MEC, but not V1 or RSC, gain-change-induced map shifts increase at low visual contrast

To examine how the quality of visual information influences cue integration, we reduced the contrast of the visual cues on the VR (Figure 5A; Figure S4). This manipulation should reduce the strength (or certainty) of landmark input. If cells combine



(legend on next page)

landmarks and path integration according to their relative certainty (Ernst and Banks, 2002), the decrease in landmark input should increase map shifts. We chose contrast = 10% because it reduced but did not eliminate MEC map stability (Figure S4). We identified cells with stable maps in both high- and low-contrast baseline trials (mean baseline peak correlation ≥ 0.5) and selected gain change trials that were similar to baseline (peak correlation to baseline ≥ 0.5 ; Figure 5B). Lowering the contrast increased map shifts in MEC. In contrast, map shifts in V1 and RSC did not change (Figures 5C–5F). Lowering the contrast to 20% did not influence MEC map shifts (Figure S4H), indicating that the sensitivity of spatial maps to visual contrast is sharp, possibly due to the visual system's ability to adapt to low-contrast conditions (Histed et al., 2012). These results suggest that MEC combines path integration and landmarks according to their relative certainties, whereas V1 and RSC are dominated by visual input during the $g = 0.8$ condition.

A population of putative MEC grid cells tuned to distance run in the dark drives many of MEC's unique cue integration properties

Given the strong influence of path integration on spatial maps in MEC relative to V1 and RSC, we next aimed to examine if MEC

contained pure path integration signals by allowing mice to run on the treadmill in near-total darkness for ~30 min (Figure 6A). Many MEC cells showed striking periodicity of their firing rate as a function of distance run in the dark (Figure 6B), and we used an autocorrelation-based metric to find distance-tuned neurons (Figures 6C and 6D). Using these criteria, 11% of all MEC cells were distance tuned (1,603/14,172), but percentages for individual sessions ranged from 0% to 52% (maximum of 111 co-recorded distance cells; Figures 6E and 6F). In contrast, we did not observe similar distance tuning in V1 or RSC (Figure S5A). Plotting firing rate correlations for all distance cells in individual sessions revealed structured distance tuning that was strikingly reminiscent of grid cells, as preferred distances increased from dorsal to ventral MEC in discrete jumps by a factor of approximately 1.55 (Figures 6E and 6G–6I; Figure S5B) (Fiete et al., 2008; Krupic et al., 2015; Stemmler et al., 2015; Stensola et al., 2012; Towse et al., 2013). In addition, we analyzed previously published data and found a substantial overlap between identified grid cells and distance-tuned cells (Campbell et al., 2018) (Figure S6; grid cells: 8/24 distance tuned; non-grid cells: 5/134 distance tuned; Fisher test, $p = 7.1 \times 10^{-5}$).

Because previous work demonstrated that grid cells respond more to VR gain changes than border cells (Campbell et al.,

Figure 6. A population of putative MEC grid cells tuned to distance run in the dark drives many of MEC's unique cue integration properties

- (A) In dark experiments, mice ran freely in near-complete darkness with no rewards for ~30 min.
- (B) Example firing rate of six MEC cells versus distance run in darkness. A subset of MEC cells showed periodic activation as a function of distance run (blue), whereas others did not (black).
- (C) Spatial autocorrelation of the entire firing rate trace for the six neurons shown in (B). To identify distance cells, we compared autocorrelation peaks with within-cell shuffled distributions.
- (D) Peak prominence versus peak Z score (normalized to shuffled distribution) for all 14,172 MEC neurons recorded in dark sessions. Distance cells (blue) were defined as those with peak prominence > 0.1 and shuffle $p < 0.01$.
- (E) Heatmaps of the firing rate autocorrelation for all distance cells in four example sessions, with cells sorted from dorsal (top) to ventral (bottom).
- (F) Number and percentage of neurons identified as distance tuned in all 57 sessions, colored by mouse (number of distance cells, mean $\pm$ SEM = 28 ± 3 , range = 0–111; percentage, mean $\pm$ SEM over sessions = $13\% \pm 2\%$, range = 0%–52%).
- (G) Distance to MEC border versus preferred spatial period for distance cells (left) and fraction of distance cells relative to MEC border (right). Most distance cells were found in the dorsal 1,000 μm of MEC. A substantial fraction of distance cells was found up to 400 μm above the border of MEC, but the fraction dropped sharply after that, indicating a possible ventral bias in the identification of the MEC boundary.
- (H) Histogram of distance cell periods.
- (I) Same as (H), separated by mouse.
- (J) Average map shift for distance versus non-distance cells during gain = 0.8 cluster 1 (stable) blocks, as in Figure 3. To be included in this analysis, cells had to pass the same spatial stability criteria used in the rest of the manuscript (stability score ≥ 0.5 in 6 baseline trials). Map shifts were significantly larger (more negative) for distance compared with non-distance cells (mean map shift $\pm$ SEM, distance cells = -9.2 ± 0.5 cm, $n = 163$ cells from 8 sessions in 5 mice; non-distance cells = -5.2 ± 0.2 cm, $n = 599$ cells from 12 sessions in 8 mice; $p = 3.8 \times 10^{-16}$, Mann-Whitney U test).
- (K) Difference in average map shift between distance and non-distance cells for each gain change block that had at least five of each (Δ map shift = -2.7 ± 0.6 cm, $n = 8$ gain change blocks from 6 sessions in 4 mice, $p = 0.016$, Wilcoxon signed rank test). Colors indicate mice.
- (L) Trial-to-trial correlation matrices for distance and non-distance cells in remapping blocks, averaged across cells (Figures 2F and 2G).
- (M) Single trial peak correlation to the average baseline firing rate map across the gain change block, separated by distance and non-distance cells. Remapping was more pronounced in distance compared with non-distance cells (dots indicates $p < 0.001$, no dots indicates not significant, Mann-Whitney U test on distance cells versus non-distance cells for each trial). Only remapping blocks were included in this analysis. Error bars indicate SEM over trial blocks.
- (N) Optimal delay factors, computed as in Figure 4D, for distance and non-distance cells separately. Distance cells' delay factors were significantly more negative (prospective) than for non-distance cells (delay factor, mean $\pm$ SEM, distance cells = -49.3 ± 3.8 ms, $n = 549$; non-distance cells = -13.2 ± 2.0 ms, $n = 2,740$; $p = 6.2 \times 10^{-14}$, Mann-Whitney U test).
- (O) Same as N, averaged by session (Δ temporal delay, distance cells – non-distance cells, mean $\pm$ SEM over sessions = -36.5 ± 6.2 ms, $n = 23$ sessions from 7 mice, $p = 3.5 \times 10^{-5}$, Wilcoxon signed rank test). Colors indicate mice.
- (P) Change in map shift in low (10% and 20%) compared with high (100%) contrast. Changes in map shift were larger (more negative) for distance compared with non-distance cells (mean $\pm$ SEM, distance cells, $c = 10$: -1.8 ± 1.8 cm, $n = 52$; $c = 20$: -2.3 ± 0.9 cm, $n = 97$; non-distance cells, $c = 10$: -0.5 ± 0.8 , $n = 138$; $c = 20$: -0.2 ± 0.5 cm, $n = 317$; $p = 0.034$ for difference between distance and non-distance cells, $p = 0.90$ for difference between $c = 10$ and $c = 20$, n-way ANOVA). Boxes indicate the 25th, 50th, and 75th percentiles; whiskers indicate the range of data points not determined to be outliers; and circles indicate outlier data points.
- (Q) Cumulative percent variance explained by principal components of activity of distance and non-distance cells. Error bars indicate SEM over sessions. Small dots above the plot indicate significance of Wilcoxon signed rank tests between distance and non-distance cells (three dots, $p < 0.001$; two, $p < 0.01$; one, $p < 0.05$; no correction for multiple comparisons). A higher variance for distance cells was explained by fewer principal components, indicating a lower-dimensional representation (Yoon et al., 2013).




Figure 7. Large gain changes reveal an influence of non-visual cues in RSC

(A) Trial structure. Blocks of four trials with VR gain = 0.5 were surrounded by blocks of six trials with gain = 1 (same as $g = 0.8$ experiment; Figure 2A). 1–4 blocks per session, median: 2.

(B) Average running speed versus track position (solid lines) $\pm$ SEM (shaded area) over blocks ($n = 208$ blocks from 30 mice) during baseline (black) and $g = 0.5$ (magenta) trials. Average running speed did not differ between the conditions ($p = 0.33$, Wilcoxon signed rank test). Dashed lines indicate tower locations.

(legend continued on next page)

2018; Chen et al., 2019), we asked whether VR gain changes impacted distance cells (putative grid cells) differently than other MEC neurons. In cluster 1 blocks (no remapping, Figure 2), distance cells' spatial maps shifted by more than non-distance cells (Figures 6J and 6K), consistent with a larger influence of path integration. These results demonstrate that a distribution of path integration responsiveness exists in MEC, with putative grid cells occupying the tail of this distribution. The fact that certain neurons within MEC were more influenced by path integration explains why MEC neurons' map shifts were correlated across multiple gain change repetitions (Figures S3C and S3D).

We next asked whether distance cells responded differently than other MEC cells during gain-change-induced remapping (cluster 2, Figure 2). Indeed, although distance cells had higher baseline stability than non-distance cells, the peak correlation to baseline trials following gain change onset dropped by more in distance cells than non-distance cells (Figures 6L and 6M). This indicates that putative grid cells remapped more in response to gain changes than other MEC neurons, consistent with previous work (Campbell et al., 2018).

We hypothesized that MEC uses path integration calculations to predict future position, prior to the arrival of sensory input (Figure 4E). If distance cells drive path integration calculations in MEC, then they should show an even more pronounced prospective bias relative to other MEC cells. Indeed, the optimal temporal delay factor was -49.3 ± 3.8 ms for distance cells versus -13.2 ± 2.0 ms for non-distance cells ($p = 6.2 \times 10^{-14}$, Mann-Whitney U test; Figures 6N and 6O). Additionally, changes in map shift in low contrast were larger for distance cells (Figure 6P). This is consistent with the idea that grid cell attractor dynamics implement Bayesian cue integration, where in conditions of high sensory noise (low contrast), path integration is weighed more strongly, leading to larger shifts (see Discussion).

Finally, attractor dynamics in MEC have been suggested to force grid cell activity to occupy a low-dimensional manifold (Yoon et al., 2013). Consistent with this, we found that in distance cells, a greater percentage of variance was explained by low dimensional PC spaces for distance cells compared with non-distance cells (Figure 6Q), even when cells were matched for mean firing rate or depth within the MEC (Figure S5C). These results support the hypothesis that the activity of putative grid cells occupies a low-dimensional manifold. The ability to record many such cells simultaneously (maximum 111 here) will open avenues

for investigation into the geometry of this low-dimensional activity space (Gardner et al., 2021).

Large gain changes reveal stronger influence of non-visual signals in RSC compared with V1

We next considered how large conflicts between landmarks and path integration cues in the current VR track influence MEC, V1, and RSC (Figures 7A and 7B), previously shown to reliably trigger remapping in MEC (Campbell et al., 2018). The gain manipulation did not affect firing rates (Figure S7A). As expected, the gain = 0.5 condition reliably caused remapping in MEC (Figure 7C, left), whereas V1 neurons remained locked to visual cues (Figure 7C, middle). However, unlike for gain = 0.8, in the gain = 0.5 condition, the spatial firing fields of RSC neurons also showed signatures of remapping and often shifted, appeared, disappeared, or changed in firing rate (Figure 7C, right).

These cell-level observations were confirmed by population-level analyses (Figure 7D–7F). These trial-trial similarity matrices revealed large changes in spatial maps in MEC, intermediate changes in RSC, and almost no change in V1 (Figure 7D; Figures S7B–S7D). Examining average trial-trial similarity matrices for individual trial blocks revealed that MEC populations almost exclusively remapped in response to $g = 0.5$, whereas V1 populations remained locked to landmarks, and RSC showed varying degrees of remapping (Figures 7E and 7F). A complementary decoding analysis gave similar results (Figure 7G). Thus, creating larger conflicts between path integration and landmark cues revealed an influence of non-visual inputs on RSC maps that was not apparent during smaller cue conflicts (Figure 2). However, remapping in RSC was qualitatively different than remapping in MEC. Specifically, RSC maps always returned to their pre-gain change state in the baseline trials following the gain change, whereas this was not the case in MEC (Figures 7D and 7E; Figures S7B and S7D).

Finally, we sought to determine the temporal dynamics of remapping in MEC. Remapping occurred after gain change onset (Figure S7E; median: 3.9 s; 25th percentile: 2.5 s; 75th percentile: 6.0 s). Prior to remapping, the decoding error tended to steadily increase (Figure S7F). In a previously described coupled oscillator attractor model of gain-change-induced remapping (Campbell et al., 2018; Ocko et al., 2018), the phase shift between the landmark input and path integration attractor steadily increases after the onset of gain change until a decoherence threshold is reached (Figure S7G). Thus, the temporal dynamics of MEC remapping were qualitatively consistent with this coupled oscillator model.

(C) Spatial firing rate maps of example neurons in MEC (left), V1 (middle), and RSC (right) during $g = 0.5$ blocks (black bars: baseline trials, magenta bars: $g = 0.5$ trials). Most MEC neurons remapped, with some exceptions (e.g., MEC cell on the bottom right). Almost no V1 neurons remapped. Unlike for the $g = 0.8$ condition (Figure 2C), RSC neurons often remapped in response to $g = 0.5$.

(D) Top: average similarity matrix by brain region for all $g = 0.5$ blocks (MEC: 96 blocks, 56 sessions, 18 mice, median 2 [max 4] gain change blocks per session, 41 ± 4 [SEM] spatial cells per block; V1: 31 blocks, 18 sessions, 6 mice, median 2 [max 2] gain change blocks per session, 22 ± 2 [SEM] spatial cells per block; RSC: 23 blocks, 12 sessions, 7 mice, median 2 [max 2] gain change blocks per session, 50 ± 7 [SEM] spatial cells per block). Similarity decreased at the onset of gain change for MEC, and to a lesser extent RSC, but remained high in V1. Bottom: average of the first six rows of the similarity matrices above, excluding the diagonal (error bars indicate SEM over trial blocks). GC, gain change.

(E) Average similarity matrices as in (D), but for individual $g = 0.5$ blocks.

(F) Projection onto the principal component defined in Figure 2F. V1 blocks were highly separable from MEC blocks, and RSC blocks were intermediate between the two.

(G) Average absolute decoder error in baseline versus $g = 0.5$ trials in MEC, V1, and RSC. The decoder was trained on baseline trials, resulting in higher errors in $g = 0.5$ trials. Points near the y axis indicate remapping (low decoder error in baseline, high decoder error during $g = 0.5$) and can be seen most prominently in MEC but also in RSC. Baseline errors were larger in V1 as a result of tower misclassifications (Figure 1K; Figure S1D).

DISCUSSION

Here, we examined the degree to which neurons in V1, RSC, and MEC followed unique versus shared coding principles during navigation. Although recent work has revealed distributed and heterogeneous coding across cortical regions involved in navigation (Allen et al., 2019; Clancy et al., 2019; Hardcastle et al., 2017; Minderer et al., 2019; Musall et al., 2019; Pinto et al., 2019; Stringer et al., 2019), we report that MEC implements distinct cue integration algorithms compared with V1 and RSC via a sub-population of distance-tuned cells. These MEC distance-tuned cells resulted in a larger influence of path integration on position estimates in MEC compared with V1 and RSC, which could allow MEC to predict an animal's current or future position before the arrival of delayed visual input. Consistent with this idea, MEC cells coded position with a small prospective bias (tens of milliseconds in distance-tuned cells). In contrast, V1 and RSC were similarly strongly influenced by visual landmark features, with delays in sensory processing modulating spike positions in both areas (retrospective coding). However, large conflicts between path integration and landmarks revealed an influence of path integration in RSC (gain-change-induced remapping) that was not present in V1. Together, this work suggests that during navigation, cortical regions can implement distinct algorithms for combining path integration with landmarks, which may allow specialization in how cortical brain regions optimally support different navigational strategies.

Other works have emphasized the influence of non-visual signals, including path integration and visuomotor predictions, on activity in mouse V1 (Fiser et al., 2016; Fournier et al., 2020; Guitchounts et al., 2020; Keller et al., 2012; Pakan et al., 2018; Saleem et al., 2018). Consistent with this work, our data reveal an influence of path integration in V1, in that maps shifted during gain changes by an amount that could not be fully explained by delays in sensory processing. However, the influence of path integration was much stronger in MEC than in V1, and we could not induce remapping in V1 even for large gain changes that induced remapping in RSC. Thus, although path integration can exert a small influence on V1 spatial representations, V1 maps remain locked to visual cues in our task even when RSC and MEC remap. It is possible that the influence of path integration on V1 activity could increase in the total absence of visual landmarks (Pakan et al., 2018), but our results show that the influence of path integration in V1 is small compared with MEC.

In RSC, like V1, we found that visual landmarks strongly influenced neural firing patterns, consistent with work demonstrating that many RSC cells encode visual landmarks (Fischer et al., 2020). In contrast, path integration had less influence on RSC compared with MEC in the current VR task. This seemingly contrasts with work demonstrating that RSC shows degraded neural responses when animals are passively shown the visual stimuli associated with a VR environment or when vision and locomotion are decoupled (Fischer et al., 2020; Mao et al., 2020). However, task engagement and locomotion-driven changes in brain state can alter RSC neural representations (Fischer et al., 2020) and correlations with other brain regions (Clancy et al., 2019). Indeed, in the current work, large visuomotor conflicts revealed

an influence of non-visual inputs on RSC spatial maps. One possibility is that contextual changes such as our large gain change could alter RSC's functional connectivity, upweighting inputs from the hippocampal formation and triggering remapping. In contrast, during small gain changes, when path integration and landmarks do not disagree by much, input from visual cortex could dominate over feedback from the hippocampal formation. This flexibility could allow the mouse to make use of different types of spatial representations in different contexts. Thus, although the current VR task offers a window into comparing the default cue integration algorithms across MEC, RSC, and V1, future work should examine the rigidity versus flexibility of these cue integration algorithms under changing task demands or behavioral states.

A growing body of work has used virtual and augmented reality systems to examine the contributions of path integration and landmarks to place cell activity in the hippocampus (Chen et al., 2013, 2019; Jayakumar et al., 2019). In general, these works found a larger influence of visual input in CA1 compared with MEC. Given this discrepancy, it is likely that CA1 receives information regarding visual features from an additional source, for example, the lateral entorhinal cortex (Knierim et al., 2013).

Despite the advantages afforded by head fixation and VR, it remains challenging to study MEC coding in head-fixed mice, because 1D grid cell firing patterns are difficult to distinguish from those of non-grid MEC cells (Campbell et al., 2018; Domnisoru et al., 2013; Kinkhabwala et al., 2020). However, we observed a population of MEC neurons tuned to distance run in the dark (~11% of MEC cells; range, 0%–52%) in which MEC's unique cue integration properties were most pronounced. These neurons were not observed in V1 or RSC in our dataset, although sequences of neurons spanning fixed distances have been observed in the hippocampus (Villette et al., 2015). The striking modularity of distance tuning in this cell population supports the idea that these neurons correspond to MEC grid cells. Consistent with this, previous reports showed that grid cells retained some distance tuning but lost spatial firing patterns in dark real-world arenas (Chen et al., 2016; Pérez-Escobar et al., 2016; however, see Hafting et al., 2005). Other properties of the distance-tuned cells were also consistent with grid cells: compared with other MEC neurons they were more influenced by path integration (Campbell et al., 2018), showed increased prospective temporal delay (De Almeida et al., 2012; Kropff et al., 2015), and their population activity showed a lower-dimensional structure (Yoon et al., 2013). Indeed, distance neurons significantly overlapped with grid cells in a previously published dataset that included open field recordings (Figure S6). The possibility that grid cells can be identified in 1D head-fixed VR environments by allowing mice to freely run in the dark would improve the field's ability to bring head-fixed techniques to bear to dissect MEC neural circuits.

Our gain change experiments in low-contrast conditions revealed that MEC map shifts were larger when landmark inputs were less certain, although the contrast sensitivity appeared to be sharp. This finding is consistent with the general Bayesian principle that inputs should be weighted according to their certainty (Ernst and Banks, 2002). This result raises two questions: what specific algorithm does MEC use to dynamically weigh

inputs according to their certainty, and what is the neural circuit implementation of this algorithm? One possibility for the algorithm is a Kalman filter, which optimally combines the predictions of an internal model (i.e., path integration estimates) with external inputs (i.e., landmark cues) when dynamics are linear and noise is Gaussian (Denève et al., 2007; Kalman, 1960; Wilson and Finkel, 2009). In terms of neural circuit implementation, our previous work proposed a coupled oscillator attractor model to explain MEC grid cell responses to gain manipulations, specifically the sharp transition between coherent shifts and remapping (Campbell et al., 2018; Ocko et al., 2018) (Figures S8A and S8B). Intriguingly, for small cue conflicts (non-remapping regimen), the attractor model's position estimates match the output of a Kalman filter (Figure S8C; STAR Methods). However, there are some inconsistencies between these models. First, the coupled oscillator attractor model as it stands has no representation of uncertainty, which is an essential ingredient of a Kalman filter. It is possible that uncertainty could be represented by the amplitude of the attractor bump or jitter in its position over time. Second, although the Kalman model can be modified to represent predicted future position by adding an additional path integration-based calculation, the attractor model represents exact current position, which is inconsistent with the prospective bias we observed in putative MEC grid cells. Of note, however, a Kalman filter fails to capture the transition between coherent shifts and remapping we observed in our experimental data. Thus, neither model fully captures our data, and future work should bring them closer together by exploring how attractor models can be altered to incorporate representations of uncertainty (Kutschireiter et al., 2021) and make short timescale predictions of future position.

Further work is also needed to understand the relationship (if any) between the short timescale prospective coding observed here (tens of milliseconds) and the longer timescale predictions that could create successor-like representations in the hippocampus and MEC to support reinforcement learning (Stachenfeld et al., 2017). Although these two types of predictions operate on different timescales, it is possible that mechanisms for making short timescale predictions based on path integration could be co-opted to make longer timescale predictions about the occupancy of non-local states.

STAR★METHODS

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

- **KEY RESOURCES TABLE**
- **RESOURCE AVAILABILITY**
 - Lead contact
 - Materials availability
 - Data and code availability
- **EXPERIMENTAL MODEL AND SUBJECT DETAILS**
 - Animals
- **METHOD DETAILS**
 - Surgery
 - Virtual reality setup
 - Behavioral training
 - Dark sessions in V1 and RSC

- Electrophysiological recording
- Dark and passive playback recordings
- Histology
- **QUANTIFICATION AND STATISTICAL ANALYSIS**
 - Spike sorting and synchronization
 - Spatial firing rate maps
 - Running speed and lick rate
 - Spatial stability and similarity matrix
 - Map shift
 - Position decoding
 - Classification of gain change responses as “stable” or “remapping”
 - Estimating coordination of remapping
 - Estimating temporal delay factors
 - Map shift after temporal delay correction
 - Spatial firing rates in low contrast
 - Spatial autocorrelation and distance tuning
 - Estimating dimensionality of distance and non-distance cell activity using principal components analysis
 - Receptive field mapping in V1
 - Dimensionality reduction of spatial firing rates using UMAP
 - Kalman filter model

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at <https://doi.org/10.1016/j.celrep.2021.109669>.

ACKNOWLEDGMENTS

We would like to thank Bob Schneeveis for engineering assistance, Adriana Diaz for histology assistance, and members of the Giacomo laboratory for discussions. This work was supported by funding from an NSF Graduate Research Fellowship and a Baxter Fellowship (awarded to M.G.C.); funding from the Swiss National Science Foundation (P2BSP3_181743 and P400PB_191076 to A.A.); funding from the James S McDonnell Foundation, the Simons Foundation, and an NSF Career Award (to S.G.); and funding from the Office of Naval Research (N00141812690), Simons Foundation (SCGB 542987SPI), NIMH (MH106475), the James S McDonnell Foundation, and the Vallee Foundation (to L.M.G.).

AUTHOR CONTRIBUTIONS

M.G.C., A.A., and L.M.G. conceived the project and designed the experiments. M.G.C. and A.A. collected and analyzed data. S.A.O. performed modeling and theoretical analysis under supervision of S.G. M.G.C., A.A., and L.M.G. wrote the paper with feedback from S.A.O. and S.G.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: December 15, 2020

Revised: May 25, 2021

Accepted: August 13, 2021

Published: September 7, 2021

REFERENCES

Alexander, A.S., and Nitz, D.A. (2015). Retrosplenial cortex maps the conjunction of internal and external spaces. *Nat. Neurosci.* 18, 1143–1151.

- Alexander, A.S., and Nitz, D.A. (2017). Spatially periodic activation patterns of retrosplenial cortex encode route sub-spaces and distance traveled. *Curr. Biol.* 27, 1551–1560.e4.
- Allen, W.E., Chen, M.Z., Pichamoorthy, N., Tien, R.H., Pachitariu, M., Luo, L., and Deisseroth, K. (2019). Thirst regulates motivated behavior through modulation of brainwide neural population dynamics. *Science* 364, 253.
- Battaglia, F.P., Sutherland, G.R., and McNaughton, B.L. (2004). Local sensory cues and place cell directionality: additional evidence of prospective coding in the hippocampus. *J. Neurosci.* 24, 4541–4550.
- Campbell, M.G., Ocko, S.A., Mallory, C.S., Low, I.I.C., Ganguli, S., and Giocomo, L.M. (2018). Principles governing the integration of landmark and self-motion cues in entorhinal cortical codes for navigation. *Nat. Neurosci.* 21, 1096–1106.
- Chen, G., King, J.A., Burgess, N., and O'Keefe, J. (2013). How vision and movement combine in the hippocampal place code. *Proc. Natl. Acad. Sci. USA* 110, 378–383.
- Chen, G., Manson, D., Cacucci, F., and Wills, T.J. (2016). Absence of visual input results in the disruption of grid cell firing in the mouse. *Curr. Biol.* 26, 2335–2342.
- Chen, G., Lu, Y., King, J.A., Cacucci, F., and Burgess, N. (2019). Differential influences of environment and self-motion on place and grid cell firing. *Nat. Commun.* 10, 630.
- Clancy, K.B., Orsolic, I., and Mrosovsky, T.D. (2019). Locomotion-dependent remapping of distributed cortical networks. *Nat. Neurosci.* 22, 778–786.
- De Almeida, L., Idiart, M., Villavicencio, A., and Lisman, J. (2012). Alternating predictive and short-term memory modes of entorhinal grid cells. *Hippocampus* 22, 1647–1651.
- Denève, S., Duhamel, J.R., and Pouget, A. (2007). Optimal sensorimotor integration in recurrent cortical networks: a neural implementation of Kalman filters. *J. Neurosci.* 27, 5744–5756.
- Diehl, G.W., Hon, O.J., Leutgeb, S., and Leutgeb, J.K. (2017). Grid and non-grid cells in medial entorhinal cortex represent spatial location and environmental features with complementary coding schemes. *Neuron* 94, 83–92.e6.
- Domnisoru, C., Kinkhabwala, A.A., and Tank, D.W. (2013). Membrane potential dynamics of grid cells. *Nature* 495, 199–204.
- Ernst, M.O., and Banks, M.S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature* 415, 429–433.
- Fiete, I.R., Burak, Y., and Brookings, T. (2008). What grid cells convey about rat location. *J. Neurosci.* 28, 6858–6871.
- Fischer, L.F., Mojica Soto-Albors, R., Buck, F., and Harnett, M.T. (2020). Representation of visual landmarks in retrosplenial cortex. *eLife* 9, e51458.
- Fiser, A., Mahringer, D., Oyibo, H.K., Petersen, A.V., Leinweber, M., and Keller, G.B. (2016). Experience-dependent spatial expectations in mouse visual cortex. *Nat. Neurosci.* 19, 1658–1664.
- Fournier, J., Saleem, A.B., Diamanti, E.M., Wells, M.J., Harris, K.D., and Carandini, M. (2020). Mouse visual cortex is modulated by distance traveled and by theta oscillations. *Curr. Biol.* 30, 3811–3817.e6.
- Gardner, R.J., Hermansen, E., Pachitariu, M., Burak, Y., Baas, N.A., Dunn, B.A., Moser, M.-B., and Moser, E.I. (2021). Toroidal topology of population activity in grid cells. *bioRxiv*. <https://doi.org/10.1101/2021.02.25.432776>.
- Guitchounts, G., Masis, J., Wolff, S.B.E., and Cox, D. (2020). Encoding of 3D Head Orienting Movements in the Primary Visual Cortex. *Neuron* 108, 512–525.e4.
- Hafting, T., Fyhn, M., Molden, S., Moser, M.B., and Moser, E.I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature* 436, 801–806.
- Hardcastle, K., Maheswaranathan, N., Ganguli, S., and Giocomo, L.M. (2017). A multiplexed, heterogeneous, and adaptive code for navigation in medial entorhinal cortex. *Neuron* 94, 375–387.e7.
- Hernández-Pérez, J.J., Cooper, K.W., and Newman, E.L. (2020). Medial entorhinal cortex activates in a traveling wave in the rat. *eLife* 9, e52289.
- Histed, M.H., Carvalho, L.A., and Maunsell, J.H.R. (2012). Psychophysical measurement of contrast sensitivity in the behaving mouse. *J. Neurophysiol.* 107, 758–765.
- Høydal, Ø.A., Skytøen, E.R., Andersson, S.O., Moser, M.-B., and Moser, E.I. (2019). Object-vector coding in the medial entorhinal cortex. *Nature* 568, 400–404.
- Hubel, D.H., and Wiesel, T.N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol.* 160, 106–154.
- Jayakumar, R.P., Madhav, M.S., Savelli, F., Blair, H.T., Cowan, N.J., and Knierim, J.J. (2019). Recalibration of path integration in hippocampal place cells. *Nature* 566, 533–537.
- Jun, J.J., Steinmetz, N.A., Siegle, J.H., Denman, D.J., Bauza, M., Barbarits, B., Lee, A.K., Anastassiou, C.A., Andrei, A., Aydın, Ç., et al. (2017). Fully integrated silicon probes for high-density recording of neural activity. *Nature* 551, 232–236.
- Kalman, R.E. (1960). A new approach to linear filtering and prediction problems. *J. Basic Eng.* 82, 35–45.
- Keller, G.B., Bonhoeffer, T., and Hübner, M. (2012). Sensorimotor mismatch signals in primary visual cortex of the behaving mouse. *Neuron* 74, 809–815.
- Kinkhabwala, A.A., Gu, Y., Aronov, D., and Tank, D.W. (2020). Visual cue-related activity of cells in the medial entorhinal cortex during navigation in virtual reality. *eLife* 9, e43140.
- Knierim, J.J., Neunuebel, J.P., and Deshmukh, S.S. (2013). Functional correlates of the lateral and medial entorhinal cortex: objects, path integration and local-global reference frames. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 369, 20130369.
- Kropff, E., Carmichael, J.E., Moser, M.B., and Moser, E.I. (2015). Speed cells in the medial entorhinal cortex. *Nature* 523, 419–424.
- Krupic, J., Bauza, M., Burton, S., Barry, C., and O'Keefe, J. (2015). Grid cell symmetry is shaped by environmental geometry. *Nature* 518, 232–235.
- Kutschireiter, A., Rast, L., and Drugowitsch, J. (2021). Angular Path Integration by Projection Filtering with Increment Observations. *arXiv*, 2102.09650. <https://arxiv.org/abs/2102.09650>.
- Lopez, C.M., Mitra, S., Putzeys, J., Raducanu, B., Ballini, M., Andrei, A., Severi, S., Welkenhuysen, M., Hoof, C.V., Musa, S., et al. (2016). 22.7 A 966-electrode neural probe with 384 configurable channels in 0.13µm SOI CMOS. In 2016 IEEE International Solid-State Circuits Conference (ISSCC) (IEEE), pp. 392–393.
- Low, I.I.C., Williams, A.H., Campbell, M.G., Linderman, S.W., and Giocomo, L.M. (2020). Dynamic and reversible remapping of network representations in an unchanging environment. *bioRxiv*. <https://doi.org/10.1101/2020.10.05.326942>.
- Mao, D., Kandler, S., McNaughton, B.L., and Bonin, V. (2017). Sparse orthogonal population representation of spatial context in the retrosplenial cortex. *Nat. Commun.* 8, 243.
- Mao, D., Molina, L.A., Bonin, V., and McNaughton, B.L. (2020). Vision and Locomotion Combine to Drive Path Integration Sequences in Mouse Retrosplenial Cortex. *Curr. Biol.* 30, 1680–1688.e4.
- McInnes, L., Healy, J., and Melville, J. (2018). UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. *arXiv*, 1802.034126. <https://arxiv.org/abs/1802.034126>.
- Miller, M.W., and Vogt, B.A. (1984). Direct connections of rat visual cortex with sensory, motor, and association cortices. *J. Comp. Neurol.* 226, 184–202.
- Minderer, M., Brown, K.D., and Harvey, C.D. (2019). The spatial structure of neural encoding in mouse posterior cortex during navigation. *Neuron* 102, 232–248.e11.
- Musall, S., Kaufman, M.T., Juavinett, A.L., Gluf, S., and Churchland, A.K. (2019). Single-trial neural dynamics are dominated by richly varied movements. *Nat. Neurosci.* 22, 1677–1686.
- Niell, C.M., and Stryker, M.P. (2010). Modulation of visual responses by behavioral state in mouse visual cortex. *Neuron* 65, 472–479.

- Ocko, S.A., Hardcastle, K., Giocomo, L.M., and Ganguli, S. (2018). Emergent elasticity in the neural code for space. *Proc. Natl Acad. Sci. USA* 115, E11798–E11806.
- Pachitariu, M., Steinmetz, N., Kadir, S., Carandini, M., and Kenneth, D.H. (2016). Kilosort: realtime spike-sorting for extracellular electrophysiology with hundreds of channels. *bioRxiv*. <https://doi.org/10.1101/061481>.
- Pakan, J.M.P., Currie, S.P., Fischer, L., and Rochefort, N.L. (2018). The Impact of Visual Cues, Reward, and Motor Feedback on the Representation of Behaviorally Relevant Spatial Locations in Primary Visual Cortex. *Cell Rep.* 24, 2521–2528.
- Pérez-Escobar, J.A., Kornienko, O., Latuske, P., Kohler, L., and Allen, K. (2016). Visual landmarks sharpen grid cell metric and confer context specificity to neurons of the medial entorhinal cortex. *eLife* 5, e16937.
- Pinto, L., Rajan, K., DePasquale, B., Thiberge, S.Y., Tank, D.W., and Brody, C.D. (2019). Task-dependent changes in the large-scale dynamics and necessity of cortical regions. *Neuron* 104, 810–824.e9.
- Saleem, A.B., Ayaz, A., Jeffery, K.J., Harris, K.D., and Carandini, M. (2013). Integration of visual motion and locomotion in mouse visual cortex. *Nat. Neurosci.* 16, 1864–1869.
- Saleem, A.B., Diamanti, E.M., Fournier, J., Harris, K.D., and Carandini, M. (2018). Coherent encoding of subjective spatial position in visual cortex and hippocampus. *Nature* 562, 124–127.
- Sargolini, F., Fyhn, M., Hafting, T., McNaughton, B.L., Witter, M.P., Moser, M.B., and Moser, E.I. (2006). Conjunctive representation of position, direction, and velocity in entorhinal cortex. *Science* 312, 758–762.
- Shamash, P., Carandini, M., Harris, K., and Steinmetz, N. (2018). A tool for analyzing electrode tracks from slice histology. *bioRxiv*. <https://doi.org/10.1101/447995>.
- Solstad, T., Boccara, C.N., Kropff, E., Moser, M.B., and Moser, E.I. (2008). Representation of geometric borders in the entorhinal cortex. *Science* 322, 1865–1868.
- Stachenfeld, K.L., Botvinick, M.M., and Gershman, S.J. (2017). The hippocampus as a predictive map. *Nat. Neurosci.* 20, 1643–1653.
- Stemmler, M., Mathis, A., and Herz, A.V. (2015). Connecting multiple spatial scales to decode the population activity of grid cells. *Sci. Adv.* 1, e1500816.
- Stensola, H., Stensola, T., Solstad, T., Frøland, K., Moser, M.B., and Moser, E.I. (2012). The entorhinal grid map is discretized. *Nature* 492, 72–78.
- Stringer, C., Pachitariu, M., Steinmetz, N., Reddy, C.B., Carandini, M., and Harris, K.D. (2019). Spontaneous behaviors drive multidimensional, brainwide activity. *Science* 364, 255.
- Sugar, J., Witter, M.P., van Strien, N.M., and Cappaert, N.L.M. (2011). The retrosplenial cortex: intrinsic connectivity and connections with the (para)hippocampal region in the rat. An interactive connectome. *Front. Neuroinform.* 5, 7.
- Towse, B.W., Barry, C., Bush, D., and Burgess, N. (2013). Optimal configurations of spatial scale for grid cell firing under noise and uncertainty. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 369, 20130290.
- Villette, V., Malvache, A., Tressard, T., Dupuy, N., and Cossart, R. (2015). Internally Recurring Hippocampal Sequences as a Population Template of Spatio-temporal Information. *Neuron* 88, 357–366.
- Virtanen, P., Gommers, R., Oliphant, T.E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., et al.; SciPy 1.0 Contributors (2020). SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat. Methods* 17, 261–272.
- Wilson, R.C., and Finkel, L.H. (2009). A neural implementation of the Kalman Filter. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 2728–2736.
- Yoon, K., Buice, M.A., Barry, C., Hayman, R., Burgess, N., and Fiete, I.R. (2013). Specific evidence of low-dimensional continuous attractor dynamics in grid cells. *Nat. Neurosci.* 16, 1077–1084.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Chemicals, peptides, and recombinant proteins		
Dil	Thermo-Fisher	CAT#: V22889
DiO	Thermo-Fisher	CAT#: V22889
DiD	Thermo-Fisher	CAT#: V22889
Deposited data		
Raw data (Neuropixels)	This paper	FigsharePlus: https://www.doi.org/10.25452/figshare.plus.15041316
Raw data (tetrode)	Campbell et al., 2018	FigsharePlus: https://www.doi.org/10.25452/figshare.plus.15026043
Processed data (V1 receptive field locations)	This paper	FigsharePlus: https://www.doi.org/10.25452/figshare.plus.15050289
Experimental models: Organisms/strains		
Mouse: C57BL/6	The Jackson Laboratories	000664
Software and algorithms		
Custom analysis code	This paper	Zenodo: https://www.doi.org/10.5281/zenodo.5138030
Allen Common Coordinate Framework	Allen Institute	https://atlas.brain-map.org/
Allen CCF MATLAB interface	Allen Institute	https://github.com/cortex-lab/allenCCF
SciPy ecosystem of open-source Python libraries (numpy, matplotlib, scipy, etc.)	(Virtanen et al., 2020)	https://www.scipy.org/
UMAP 0.5.1	McInnes et al., 2018	https://github.com/lmcinnes/umap
MATLAB	MathWorks	https://www.mathworks.com/products/matlab.html
Kilosort	Pachitariu et al., 2016	https://github.com/MouseLand/Kilosort
Kilosort2	Stringer et al., 2019, Science	https://github.com/MouseLand/Kilosort
Phy		https://github.com/cortex-lab/phy
Spikes Phy postprocessing tool		https://github.com/cortex-lab/spikes
Other		
Sensapex uMp micromanipulator	Sensapex	https://www.sensapex.com
Phase 3A Option 4 Neuropixels silicon probes and data acquisition system	Jun et al., 2017	https://www.neuropixels.org/probe
Phase 3B Neuropixels 1.0 silicon probes and data acquisition system	Jun et al., 2017	https://www.neuropixels.org/probe
H3 Cambridge Neurotech Probes and data acquisition system	Cambridge Neurotech	https://www.cambridgeneurotech.com/

RESOURCE AVAILABILITY

Lead contact

Further information and requests for resources and reagents should be directed to and will be fulfilled by the Lead Contact, Lisa Giocomo (giocomo@stanford.edu).

Materials availability

This study did not generate new unique reagents.

Data and code availability

Processed electrophysiological and behavioral data have been deposited at Figshare and are publicly available as of the date of publication. DOIs are listed in the [Key resources table](#).

All original code has been deposited at Zenodo and is publicly available as of the date of publication. DOIs are listed in the [Key resources table](#).

Any additional information required to reanalyze the data reported in this paper is available from the lead contact upon request.

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Animals

All techniques were approved by the Institutional Animal Care and Use Committee at Stanford University School of Medicine. Recordings were made from 32 female C57BL/6 mice aged 12 weeks to 6 months at the time of surgery (18 - 26 g, [Table S1](#)). Mice were housed singly or with littermates in transparent cages on a 12-hour light-dark cycle and experiments were performed during the light phase.

METHOD DETAILS

Surgery

Anesthesia was induced with isoflurane (4%; maintained at 1.75%) followed by injection of buprenorphine (0.1 mg/kg). Fiducial marks were made on the skull $\pm$ 3.3 mm lateral from the midline and approximately 4 mm posterior from Bregma, and $\pm$ 2.5 mm lateral from the midline and approximately 3 mm posterior from Bregma for targeting V1 and RSC. A ground screw was affixed to the skull approximately 2.5 mm anterior and 1.5 mm lateral from Bregma on the left side. A stainless steel headbar was attached to the skull using Metabond and the rest of the exposed skull was covered with Metabond. Transparent Metabond was used to allow visualization of the fiducial marks, which would later guide craniotomies and probe placement. After training, which typically lasted 2-3 weeks, mice were again anesthetized with isoflurane and the Metabond and skull were shaved down with a dental drill, posterior to the fiducial mark on both sides. For MEC recordings, bilateral craniotomies, roughly 500 μ m in diameter, were made posterior to the fiducial mark, exposing the transverse sinus. Plastic rings cut from pipette tips ($\sim$ 4 mm diameter) were affixed to the skull around each craniotomy using Metabond. These rings served as wells to hold saline and a layer of silicone oil during recording, which prevented the craniotomy from drying out. Craniotomies were covered with KwikCast and the mouse recovered overnight before the first recording session. A maximum of three recording sessions were performed per hemisphere, each on a different day, for a maximum of six recording sessions on six days total for all mice but one (in which seven recordings sessions were performed across seven days). Durotomies were performed before recording from a craniotomy for the first time, either the same day as the craniotomy (side 1, typically the right side) or three days later (side 2, typically the left side). For recordings targeting V1 and RSC, we followed a similar strategy, except that only one craniotomy, combined with a durotomy, was performed at a time. We then performed 2 recordings per craniotomy before targeting a new location with another brain region. As for MEC recordings, a maximum of 6 recordings per mouse was performed.

Virtual reality setup

The VR setup was the same as in ([Campbell et al., 2018](#)), except that the VR scene was displayed on three 24-inch monitors surrounding the mouse instead of a 14-inch hemispherical projection screen, and the mouse ran on a 6-inch diameter foam cylinder instead of a 20-cm diameter Styrofoam ball. Briefly, rotation of the treadmill was recorded using a rotary encoder (Yumo, 1024 P/R), and data were processed and transmitted to the virtual reality software via a microcontroller (Arduino UNO) using custom software. The virtual reality environment was created and displayed using custom software in Unity 3D (Unity Technologies). The training environment consisted of a floor with checkerboard texture and evenly spaced, identical visual landmarks on both sides. The test environment had the same floor as the training environment, and in addition to the reward tower (which remained the same), contained 4 additional visual landmarks (towers). These five visual landmarks (towers) were identical on the left and right sides and were evenly spaced at 80 cm intervals, covering a total of 400 cm. Upon reaching the tower at 400 cm, mice were 'teleported' to the beginning of the track for the next trial without any visual discontinuity, giving the impression of running along an infinite hallway. Water rewards were delivered via Tygon tubing attached to a metal lick spout mounted in front of the mouse. Delivery was triggered via a solenoid valve, which produced an audible clicking sound upon reward delivery. Licking was detected with a custom-built infrared light barrier. The silicon probes were mounted either on a manual manipulator on a custom rotatable mount (World Precision Instruments) or a motorized micromanipulator (UMP Micromanipulator, Sensapex). Probe holders were placed behind the mouse to minimize visual interference.

Behavioral training

After headbar implantation, mice recovered for three days and were given Baytril (10 mg/kg) and Rimadyl (5 mg/kg) daily. After three days, they were taken off Baytril and Rimadyl and put on water deprivation. They received 0.8 mL of water each day and their weights were monitored to ensure that they remained above 85% of baseline. Training progressed in three stages. In stage one, they were head-fixed on the VR rig and trained to receive water from the lickspout. Water delivery was associated with an audible click of the solenoid. Mice quickly learned this association and began licking upon hearing the click. After stage one, which typically lasted 2 to 4 days, mice progressed to stage two, in which they ran on a training track to receive water rewards (see *Virtual Reality Setup* for description of training track). To encourage mice to run, water rewards ($\sim$ 2 μ L) were automatically delivered whenever they passed

a tower. The tower spacing started at 40 cm and was increased daily to encourage running, up to a maximum of 200 cm. The reward tower on the habituation track was visually identical to the reward tower on the track used for recording. Once mice ran consistently on the habituation track (average running speed > 10 cm/s), they progressed to stage three, in which they ran on the same track that would eventually be used for recording. They ran 50–200 trials per day for 4–20 days (median 14 days) for a total of 800–2500 trials (median 2000 trials) before the first recording session. During this final phase of training, mice developed stereotyped running and licking patterns in which they slowed down and licked prior to the reward tower. Some training sessions were performed on a training rig, but mice were always trained on the recording rig for several days prior to the first recording day to familiarize them with the setup. Mice were deemed fully trained and ready to record when they completed 200 trials within 1 hour for two consecutive days.

Dark sessions in V1 and RSC

For this separate set of experiments, we followed the same procedure as above for head bar implantation, water deprivation, and initial stages of training. Mice were trained on the habituation track, which was progressively elongated as described above. Once mice ran 200 trials within 1 hr on the 400cm long training track, mice were switched to a visual flow only track, where random checkerboard patterns provided visual flow input. Mice did not receive reward on this track, but nevertheless continued to run. Recordings started after two days of exposure to the visual flow only track. Prior to the start of dark recordings, mice typically ran for 1 hour on the visual flow only track.

Electrophysiological recording

Extracellular recordings were made by acutely inserting silicon probes into a craniotomy above the target brain region while the animal was head-fixed in the VR setup. One insertion was performed per day per mouse, after which the craniotomy was covered with Kwik Cast and returned to on subsequent days for additional recordings. Typically, we performed 3 insertions on each side for MEC recordings, and 2 insertions per craniotomy for recordings targeting V1 and RSC (all in similar sites within the same craniotomy) for a total of 6 recordings per mouse over 6 days.

Most recordings were made using single Neuropixels probes (Jun et al., 2017; Lopez et al., 2016). For MEC recordings, 17 recordings from 7 mice were made with Phase 3A Option 4 probes (276 recording channels), and 62 recordings from 11 mice were made with Neuropixels 1.0 probes (384 recording channels, <https://www.neuropixels.org/>). For recordings targeting V1 (24 sessions, 9 mice) and RSC (18 sessions, 9 mice), all recordings were made with Neuropixels 1.0 probes. Raw voltage traces were filtered, amplified, multiplexed, and digitized on-probe and recorded using SpikeGLX (<https://billkarsh.github.io/SpikeGLX>). Voltage traces were filtered between 300 Hz and 10 kHz and sampled at 30 kHz with gain = 500 (AP band) or filtered between 0.5 Hz and 1 kHz and sampled at 2.5 kHz with gain = 250 (LFP band). All recordings were made using the contiguous set of recording channels closest to the tip of the probe (Bank 0). The probe's ground and reference pads were shorted together and soldered to a gold pin, which was then connected to a gold pin that had been soldered to a skull screw on the animal. For Phase 3A Option 4 probes, recordings were made in external reference mode, using the skull screw as external reference. For Neuropixels 1.0 probes, recordings were made in internal reference mode, with the probe tip set as reference, but the ground and reference pads were still connected to the skull screw in order to ground the animal and prevent electrostatic damage to the probe.

In some recordings (11 recordings from 3 mice, 597 units total), we used 64-channel silicon probes from Cambridge Neurotech (model H3, <https://www.cambridgeurotech.com/>). Data were recorded using a Neuralynx Digital Lynx SX data acquisition system.

Prior to insertion, probes were first dipped in one of three different colors of dye (DiI, DiD, and DiO). Different colors were used to allow up to three different penetrations in the same craniotomy to be distinguished in histology. The mouse was head-fixed on the VR rig and the KwikCast was removed from above the craniotomy. For MEC recordings the probe was mounted on the rig at a slight angle to maximize alignment between the probe and MEC (approx. 10 degrees). The probe was then lowered down to the level of the skull, aligned mediolaterally with the fiducial mark, and inserted as close to the transverse sinus as possible. The well was filled with saline (0.9% NaCl) and covered with a layer of silicone oil to prevent drying. The probe was inserted slowly (~10 μ m/sec) until there was no longer any activity at the tip or until the probe started to bend. The probe was retracted 100–200 μ m from its deepest penetration point and left to settle for 30 minutes before starting the recording. For recordings targeting V1 and RSC, we followed a similar strategy.

Recordings lasted between 40 and 150 minutes, during which the mouse ran one or more VR sessions under various gain, contrast, and cue removal conditions. Based on previous work, we chose to focus on gain values < 1, as these tended to show more strongly the influence of path integration on spatial firing in MEC (Campbell et al., 2018; Chen et al., 2019). The VR emitted a TTL pulse every frame from an Arduino UNO, and these pulses were recorded in SpikeGLX using an auxiliary National Instruments data acquisition card (NI PXIe-6341 with NI BNC-2110) to synchronize VR traces with neurophysiological data.

After each recording session, the probe was rinsed with deionized water, soaked in deionized water for at least 15 minutes, soaked in an enzymatic detergent (2% Tergazyme in deionized water) for at least 60 minutes, and rinsed and soaked again in deionized water for at least 15 minutes.

Dark and passive playback recordings

There were two separate experiments in which mouse ran in the dark. In the first experiment, mice ran in the dark with rewards delivered every 400 cm (baseline trials with monitors turned off; Figures S1G and S1H). These dark sessions occurred in between two VR

sessions with the monitors on. In the second experiment, mice ran freely in the dark for ~30 minutes and no rewards were delivered (Figure 6). These dark sessions occurred after mice ran VR sessions with gain or contrast manipulations (~200 trials).

In a separate set of experiments, mice passively sat on the running wheel while visual stimuli from a previous VR session were played back (Figures S1E and S1F). The running wheel was braked during passive playback. Mice used in these experiments were habituated to the braked condition during training. Passive playback sessions occurred in between two 50-trial VR sessions. The first 50-trial VR session was used to generate the playback stimuli.

Histology

After the last recording session, the mouse was euthanized with an injection of pentobarbital (0.2 mL i.p. of 100mg/ml beuthanasia) and perfused transcardially with 1X phosphate-buffered saline (PBS) followed by 4% paraformaldehyde (PFA). Brains were extracted and stored in 4% PFA for at least 24 hours after which they were transferred to 30% sucrose and stored for at least 48 hours. Finally, brains were frozen and cut into 65 μ m sections (sagittal sections for MEC recordings, coronal sections for V1 and RSC recordings) with a cryostat, stained with DAPI, and imaged with a widefield fluorescence microscope (Axio Imager 2, Zeiss). To localize MEC recordings we identified the sections where the probe entered MEC and where the tip of the probe was localized. From these two points, we determined the section of the probe within MEC and only used units recorded within this span for further analysis. To reconstruct the probe tracts for V1 and RSC recordings, we used software written in MATLAB (<https://www.biorxiv.org/content/10.1101/447995v1>) (Shamash et al., 2018). This allowed us to first register the coronal sections to a standard atlas and then reconstruct the probe tract within the standard atlas framework. This information was used to assign each cluster to a brain region. Note that for RSC, we only used units from A29 (denoted as RSPd and RSPv in the Allen Mouse Brain Atlas).

QUANTIFICATION AND STATISTICAL ANALYSIS

All analyses were performed in MATLAB or Python using custom scripts.

Spike sorting and synchronization

Spike sorting was performed offline using Kilosort2 (<https://github.com/MouseLand/Kilosort>) (Pachitariu et al., 2016). Data from the Phase 3A Option 4 Neuropixels probes and the Cambridge Neurotech H3 probes, which were gathered prior to the release of Kilosort2, were sorted with Kilosort1 (<https://github.com/cortex-lab/KiloSort>). Clusters were manually inspected and curated in Phy (<https://github.com/cortex-lab/phy>). Clusters were identified as 'good' if they met the criteria for contamination, signal to noise ratio and firing rate. To estimate contamination, we calculated the autocorrelation of spike times with 1 ms bins and calculated the ratio between the 0-1ms bin and the peak of the autocorrelation. This ratio had to be smaller than 0.2 for a cluster to be considered. To estimate noise, we extracted voltage traces around 1000 randomly selected single spikes for each cluster using built in functions in Phy. For each spike, we then extracted 1 ms of data (30 samples) prior to the detected onset of the spike waveform and calculated the standard deviation across all concatenated samples to estimate the noise. The spike amplitude was divided by this value to estimate signal to noise ratio. This ratio had to be larger than 3 for a cluster to be considered. Finally clusters containing less than 1000 spikes were discarded. For data sorted with Kilosort2, these criteria were applied units classified as 'good' by Kilosort2. Spike times and cluster identities were extracted from the output of KiloSort and Phy using code from the Spikes repository (<https://github.com/cortex-lab/spikes>). VR data (VR position, lick times, and contrast and gain values for each trial) were synchronized to spiking data by substituting the time of each VR frame with the time of the corresponding detected TTL pulse. Synchronization was checked by comparing the difference between subsequent VR frame times with the difference between subsequent TTL pulse times and confirming that these were highly correlated ($\rho > 0.95$). For data acquired with Neuropixel 1.0 probes, we additionally corrected any drift between the IMEC PXle acquisition module and the auxiliary acquisition module (NI PXle-6341) by aligning the sync waveforms generated by the IMEC PXle. This was necessary as the two modules tended to drift apart by about 5ms/hr. Because the frame rate of the VR was not constant but instead fluctuated slightly around 60 Hz, we used linear interpolation to resample VR data to a constant 50 Hz for ease of subsequent analysis. Finally, spike times, cluster identities, and synchronized VR data were saved together in a single file for analysis in MATLAB (one file per VR session).

Spatial firing rate maps

To compute spatial firing rate maps, we binned the position into 2cm bins, computed spike counts for each position bin, divided the counts by the occupancy of each spatial bin and smoothed the resulting spatial firing rate (Gaussian kernel, standard deviation = 4 cm). To calculate average firing rate as a function of position (Figure 1J), we first averaged the spatial firing rate maps across trials for a neuron and across spatially stable neurons for a recording session. We then finally averaged these per-session firing rates, excluding blocks with fewer than 5 stable neurons (see below), for each brain region.

Running speed and lick rate

To calculate the running speed of the mouse, we calculated the difference in VR position between consecutive VR frames. The resulting running speed trace was smoothed with a Gaussian kernel (sigma = 0.2 s). To calculate running speed as a function

of position, we averaged the smoothed running speed for each position bin. We measured licking behavior with an infrared light barrier. Thus, licking resulted in drops of photodiode voltage output. Photodiode voltage was monitored by an Arduino UNO. The VR computer queried the voltage on each frame and individual licks were defined as the voltage dropping below a predefined threshold.

Spatial stability and similarity matrix

To compute the spatial stability measure and the spatial similarity matrix for a neuron, we first computed trial-by-trial spatial firing rate maps as described above, subtracted the mean firing rate from each trial, and then computed the normalized cross correlation with a maximal lag of 30 cm for each pair of trials. The similarity for a pair of trials was defined as the maximum value of this cross correlation. To compute a spatial stability score for a given set of trials for each cell, we averaged all pairwise cross correlations. We defined a neuron as stable if the resulting average was ≥ 0.5 . To calculate a similarity matrix for a block of trials in a recording session, we averaged individual similarity matrices of each neuron that fulfilled the stability criteria for a set of trials (usually the 6 trials preceding a gain change). Note that in this way, neurons did not have to fulfill the stability criteria across the whole session. Blocks with fewer than 5 cells that fulfilled the stability criteria were excluded.

To quantify trial-by-trial similarity as a function of inter-trial distance (Figure S2H), we calculated pairwise stability as described above for all pairs of co-recorded units for a block of 16 baseline trials. To be included units were required to have average spatial stability ≥ 0.5 across all pairs of neighboring trials (15 pairs for 16-trial baseline block).

To calculate spatial similarity of baseline maps (Figure S2I), we first calculated the spatial firing rate maps for each block by averaging across 6 trials before the onset of the gain change. For each unit, we then calculated the spatial similarity as described above between these two blocks and averaged across all co-recorded units of a block. We only included units that were stable in both baseline blocks, i.e., the average spatial similarity was ≥ 0.5 in both blocks of baseline trials.

To calculate spatial similarity for blocks of 200 cm (Figure S7B), we first calculated the baseline firing rate maps for each block by averaging across 6 trials before onset of the gain change. For each unit, we then calculated the spatial similarity for each section of 200 cm by calculating the cross correlation between the mean subtracted spatial firing rate of the current section, and the corresponding baseline section. Note that, for gain change trials and baseline-post trials, the baseline map was calculated by averaging across all 6 baseline-pre trials. For the similarity to baseline for baseline-pre trials, the baseline spatial firing rate map was obtained by averaging across the remaining 5 baseline-pre trials.

Map shift

To calculate the shift between spatial firing rate maps, we computed spatial similarity matrix as described above. We estimated the map shift between pairs of trials as the location of the maximum cross correlation. As for the similarity matrices, we averaged these shift matrices across all stable neurons of a block to estimate the population trial-by-trial map shifts. To estimate the average shift between baseline trials and gain change trials, we averaged the population shift across all pairs of baseline trials and gain change trials (indicated by red rectangle in Figure 3C).

Position decoding

To decode position from the firing rates of simultaneously recorded neurons (Figures 1K and 7G; Figures S3F and S7E) we first chose a set of encoding (training) trials. Cells were included for analysis if they had a stability score ≥ 0.5 in encoding trials (see above for definition of stability score). For these neurons, we calculated the z-scored firing rate over time (time bin: 0.02 s, Gaussian smoothing sigma: 0.2 s) and used the z-scored firing rate to compute spatial tuning curves for encoding trials by averaging the population activity within each position bin (2 cm bins). This defined an average N-dimensional trial trajectory during encoding trials, where N is the number of neurons. For each time point during decoding (test) trials, we determined the closest point on the average encoding trajectory by minimizing the Euclidean distance. The decoded position then corresponded to the VR position for this point on the average trajectory. In both encoding (training) and decoding (test) data, we only considered times when the mouse was running (cutoff: 2 cm/s).

We used this approach to study the time course of MEC remapping in response to gain = 0.5 (Figures S7D–S7F). We inferred the location and time of remapping by finding the first instance when the absolute decoding error became larger than 50 cm, preceded by a period of at least 3 s with absolute error smaller than 50 cm.

We also used this approach to estimate the dependence of decoded position on running speed (Figure S3F). For blocks of 10 baseline trials with at least 5 stable cells (stability score ≥ 0.5), we decoded position in each trial by treating the other 9 trials as the encoding trials. We computed decoding error as the difference between decoded position and actual position, respecting the circularity of the track. For example, if decoded position were 399 cm and true position were 1 cm, the error would be -2 cm, not 398 cm. We removed times when this difference exceeded 50 cm. We concatenated decoder error, true position, and running speed traces across trials, and used multiple linear regression to estimate the influence of running speed and position bin (bin size = 20 cm) on decoder error. For each session, we averaged coefficients over all 10 trial blocks with at least 5 stable cells (Figure S3F). The sign of the coefficient was flipped in Figure S3F to match the signs of retrospective/prospective coding conventions in Figure 4.

To determine the decoding error as a function of number of included neurons in baseline trials (Figure 1K), we followed a similar decoding procedure. First, we calculated spatial firing rate maps for trials 5–20 for each neuron in each recording session. Note that these were always baseline trials. We excluded the first 4 trials because we noted that spatial firing rates tended to be more variable at

the beginning of a recording. For further analysis, we only included neurons with a stability score ≥ 0.5 across all 16 trials. To estimate decoding accuracy for N included neurons, we drew a random sample of spatial maps of N neurons. We then averaged the spatial firing rates across 15 training trials for each neuron (leaving out one trial for test). This gave us the training firing rate trajectory in N dimensions. To decode the position from the test trial, we found the closest point on the training trajectory for each position bin by minimizing the Euclidean distance, as above. In this way, we found the decoded position for each position bin of the test trial. We repeated this procedure for all 16 trials. This gave us an estimated position for each position bin in each of the 16 trials. We computed the decoding error by taking the average of the absolute difference between actual position and decoded position, across 16 trials, respecting the circularity of the track as above. To estimate the distribution of the absolute error, we also calculated the histogram of the absolute error. We repeated this approach for each set of N neurons by drawing random sets of N neurons 100 times.

Classification of gain change responses as “stable” or “remapping”

To classify gain change responses, we first vectorized the upper half of each trial-trial similarity matrix and stacked the vectorized trial blocks into a matrix, resulting in a $N \times M$ matrix, where N is the total number of trial blocks across all three brain regions, and M is the number of trial pairs. We then calculated the principal components of this matrix and used the first principal component to divide the blocks into 2 clusters by applying a threshold of 0.4, which corresponded to the location of the minimum between the two peaks of the score distribution (Figure 2F). To estimate the remapping behavior in baseline trials in MEC, we calculated spatial similarity matrices of blocks of 16 baseline trials, including only neurons that had a spatial stability score ≥ 0.5 across the first 6 trials. We then vectorized the upper half of each similarity matrix as described above and subtracted from each vector its mean to center the data, and projected onto the first principal component of the gain change block similarity matrix to get a projection score for each baseline-only block. We assigned each baseline-only block as either stable or remapping by thresholding this score (> 0.4 = stable, < 0.4 = remapping).

Estimating coordination of remapping

To calculate within-block correlation between the spatial similarity map of individual units and the remainder of the co-recorded population (Figure 2J), we calculated the correlation between the spatial similarity map of each unit and the population similarity map. The population similarity map was obtained by averaging across the remaining group of simultaneously recorded spatial units. This was repeated for each unit and then averaged across all co-recorded spatial units for a block to obtain an average within-block correlation of the individual spatial similarity maps to the population similarity map. To calculate across-block correlation, we correlated the spatial similarity map of each unit for a block, with the population maps of all other ‘remap’ blocks, and averaged across all co-recorded spatial units to obtain the average across-block correlation.

Estimating temporal delay factors

To visualize the effect of running speed on spike location (Figure 4C), we calculated the average spatial firing rate maps for each neuron as described above for blocks of 16 baseline trials. Gain change trials were not included in this average. We identified the VR position with maximum average firing rate for each cell (indicated by a red line in Figure 4C). Finally, we created raster plots for each neuron in a region surrounding the maximum firing rate location (-50 cm to $+40$ cm) sorting trials by the average running speed (average over 10 cm preceding location of maximum firing rate).

To estimate the temporal delay for a neuron, we employed a grid search over a range of values d (-300 ms to 300 ms, in 10 ms increments). We calculated time-shifted spatial firing rate maps for each trial as described above, but for 1 cm spatial bins and a smoothing kernel with $\sigma = 3.5$ cm, and shifting the time of each spike by d . To estimate the trial-by-trial alignment of firing rate maps, we computed trial-by-trial correlations of the time-shifted spatial firing rate maps for each pair of trials. To minimize boundary effects, we only used positions between 10 cm and 390 cm. We estimated the temporal delay d for a neuron for a block of trials as the value d that resulted in the highest average correlation across trials (e.g., Figure 4D). As for map shifts, we ignored d for a block of trials if the maximum occurred at either limit. To reduce overfitting and minimize the effects of remapping and drifts in spatial firing, we repeated this procedure for blocks of continuous baseline trials (blocks were non-overlapping). For each neuron, the final estimate of temporal delay was the average across multiple blocks ($8 - 18$ blocks, median: 11 , depending on trial structure of the recording session). Blocks of trials were ignored if the spatial stability of the neuron was smaller than 0.5 . Values were used for further calculations if at least 3 blocks fulfilled the range criterion (d in $(-300, 300)$) and stability criterion (MEC: 26% of neurons, V1: 38% of neurons, RSC: 45% of neurons). The same approach was used to estimate temporal delay at different contrast levels (Figure S4G).

Map shift after temporal delay correction

For each block of the ‘stable’ cluster, we shifted the spike times of each neuron by its estimated temporal delay, and calculated similarity matrices, shift matrices, and map shifts as described above. We then calculated the difference in map shift between uncorrected and corrected pairs of blocks (Figure 4F).

Spatial firing rates in low contrast

In a subset of mice, we reduced the contrast of the visual landmarks to test the influence of visibility of the landmarks on neural activity. We started with a block of 24 trials of high contrast. Two sets of 10 trials with baseline gain were interleaved with gain = 0.8. This was followed by a block of 24 trials with low contrast, with same gain structure as for the high contrast block. This allowed us to test the influence of visual contrast on neural firing both during baseline gain as well as gain = 0.8. We calculated the shift of spatial maps during gain = 0.8 as described above, including only trials with spatial similarity ≥ 0.5 to the baseline map, and only considering neurons with spatial stability ≥ 0.5 in the 6 trials preceding the gain change both during high contrast and low contrast (Figure 5). To visualize the effect of contrast on the location of peak firing rate, we calculated the average spatial firing rate maps for each neuron as described above for blocks of 6 trials preceding gain change in high contrast. We then determined the location of peak firing and averaged spatial firing rates during gain change trials around the location of peak firing rate (± 30 cm), normalizing to the maximum firing rate within each trial. We repeated this for gain change during low contrast. We then averaged the peak firing rate curves for high and low contrast for all neurons with spatial stability ≥ 0.5 in the 6 trials preceding the gain change both during high contrast and low contrast.

Spatial autocorrelation and distance tuning

To compute the spatial autocorrelation of firing rate traces during dark sessions (Figure 6), we first computed the firing rate as a function of distance, similar to the spatial firing rate but replacing position by distance run since start. We then computed the autocorrelation of this firing rate trace, including lags up to 800 cm. To identify the preferred distance of a neuron, we used MATLAB's *findpeaks* function to detect the maximum peak in the spatial autocorrelation. In addition, we computed a shuffled spatial autocorrelation by shifting spike times relative to elapsed distance by random offsets ($n = 300$ shuffles, offset drawn from a uniform random distribution with the interval $(20 \text{ s}, \text{max_t})$, where max_t was the duration of the recording). For each shuffle, we computed the maximum autocorrelation peak up to a maximum lag of 800 cm. We determined that a neuron was distance-tuned if the prominence of the peak was greater than 0.1, and the height of the autocorrelation at this preferred distance was greater than the 99th percentile of the shuffled distribution of autocorrelation peaks. This threshold was used to compare gain change responses, temporal delay factors, and dimensionality between distance and non-distance neurons (Figures 6J–6Q). In the histogram of preferred distances (Figure 6H), we used 2 cm bins and smoothed the resulting histogram with a Gaussian kernel with standard deviation 1 bin.

Estimating dimensionality of distance and non-distance cell activity using principal components analysis

We used dimensionality reduction with PCA to compare dimensionality of distance and non-distance cell activity during running in the dark (Figure 6Q). We first chose dark sessions with at least 30 of each cell type and randomly selected 30 distance and non-distance cells for each of these sessions. We computed smoothed firing rate traces for each neuron (time bin = 20 ms, standard deviation of Gaussian filter = 200 ms), z-scored each neuron's firing rate trace separately, and performed PCA on the resulting firing matrix (separately for distance and non-distance cells). We compared cumulative variance explained by the top N principal components between distance and non-distance cells. To control for differences in mean firing rate and anatomical location between distance and non-distance cells, we performed the same analysis but chose, for each distance neuron, the non-distance neuron with the closest mean firing rate or anatomical location (requiring a minimum distance of 25 μm between cells to avoid cross-contaminated clusters) (Figure S5C).

Receptive field mapping in V1

For recordings targeting V1, we mapped receptive fields of V1 neurons by playing movies based on the cosine of Gaussian noise for 15–20 minutes, presented to the eye contralateral to the recording hemisphere, using the same screen arrangement as during the VR task. To map spatio-temporal receptive fields, we calculated the spike triggered average for each neuron by first identifying for each spike the stimulus frames preceding the spike by 120ms to 190ms (frames 8 to 13), and then averaging across all spikes. This resulted in spike-triggered average stimulus maps for each of the 5 frames preceding a spike. We then averaged the maps and smoothed the resulting average with a Gaussian smoothing kernel with $\sigma = 0.5$. The average map was then Z-scored and all pixels with a $|z\text{-score}| < 2.5$ set to 0. If any non-zero pixels remained, we then fit a 2D Gaussian distribution to the remaining non-zero pixels, separately for pixels with positive and negative z-score. For neurons with a single positive or negative receptive field, its center was used as the estimate for the receptive field center. For neurons with multiple fields, we averaged the center of all subfields to generate an estimate for the receptive field center.

Dimensionality reduction of spatial firing rates using UMAP

We used UMAP 0.5.1 (McInnes et al., 2018) with Python 3.8 to visualize population state space trajectories (Figures 1L, S2K, and S7D). First, we calculated spatial firing rates as described above for all co-recorded neurons in a session, for 16 baseline trials (Figure 1L) or blocks consisting of 6 baseline-pre trials, 4 trials of gain = 0.5 or 0.8 and 6 baseline-post trials (Figures S2K and S7D). These spatial firing rates were concatenated, resulting in a $3200 \times N$ matrix, where N was the number of co-recorded neurons and 3200 was the number of spatial bins in 16 trials. We then used UMAP with cosine similarity as distance metric, and otherwise default parameters, to embed the neural activity into 2 dimensions.

Kalman filter model

The Kalman filter is a method for tracking variables of interest by balancing external inputs with updates from an internal model describing how they evolve over time. The Kalman filter assumes that the variables are normally distributed and the processes updating the internal estimate are linear. The assumption that variables are Gaussian is convenient because it means we only have to track their means and variances to fully describe them, and the assumption that the internal process is linear makes it easy to model. Importantly, when these assumptions are met, the Kalman filter is the optimal Bayesian solution to this problem.

We consider the evolution of a position estimate, x_A , which is updated by both internal path integration and external input from landmarks located at x_L according to a continuous Kalman filter. In the Kalman filter formulation, the internal estimate x_A and external inputs x_L both have associated uncertainties, which we call σ_A and σ_L . Going forward, it is mathematically convenient to work with the *difference* between x_A and x_L , which we call $\Delta x = x_A - x_L$. We now consider how Δx evolves in time according to a Kalman filter.

In the Kalman interpretation, the relative strengths of landmark and self-motion inputs are determined by the different levels of uncertainty attached to them. In the real world, these uncertainties correspond to the magnitudes of random errors in landmark perception and path integration. Here, for simplicity, we ignore these random errors and only consider the levels of uncertainty. When random noise is incorporated into the model, the distribution of position self-estimates will be a Gaussian distribution centered around the results derived here.

To build up to a continuous Kalman model, we first consider a series of closely spaced, weak landmark input events, separated by time Δt . At each landmark event, the landmark input tells the system it is at $\Delta x = 0$ with some uncertainty σ_L . Because the landmark events are frequent and weak, we have $\sigma_L \gg \sigma_A$. The position estimate Δx becomes a weighted average of the original position estimate (Δx) and the position estimate given by new landmark information ($\Delta x = 0$) where the relative weighting is given by the uncertainties associated with each signal:

$$\Delta x \rightarrow \Delta x \cdot \left(\frac{\frac{1}{\sigma_A^2}}{\frac{1}{\sigma_A^2} + \frac{1}{\sigma_L^2}} \right) = \Delta x \cdot \left(\frac{\sigma_L^2}{\sigma_A^2 + \sigma_L^2} \right) = \Delta x \cdot \left(1 - \frac{\sigma_A^2}{\sigma_A^2 + \sigma_L^2} \right) \approx \Delta x \cdot \left(1 - \frac{\sigma_A^2}{\sigma_L^2} \right)$$

And the uncertainty σ_A of the position estimate after the landmark event will decrease in light of the increased information as follows:

$$\sigma_A^2 \rightarrow \left(\frac{1}{\sigma_A^2} + \frac{1}{\sigma_L^2} \right)^{-1} = \frac{\sigma_A^2}{1 + \frac{\sigma_A^2}{\sigma_L^2}} = \sigma_A^2 \left(\frac{\sigma_L^2}{\sigma_A^2 + \sigma_L^2} \right) = \sigma_A^2 \left(1 - \frac{\sigma_A^2}{\sigma_A^2 + \sigma_L^2} \right) \approx \sigma_A^2 \left(1 - \frac{\sigma_A^2}{\sigma_L^2} \right)$$

In both cases, the last simplification was made using the fact that $\sigma_L \gg \sigma_A$.

In between the landmark events, the animal performs linear path integration:

$$\Delta x \rightarrow \Delta x + \Delta t \cdot (v_A - v_L)$$

where v_L is the velocity of the landmarks and v_A is the animal's internal velocity estimate. For simplicity, we are not modeling noise in the velocity estimate. When there is no gain change, $v_A - v_L = 0$. During a gain change v_A differs from v_L and so this term is non-zero (Figure S8C).

We model the position uncertainty to grow linearly with time according to a diffusivity constant α :

$$\sigma_A^2 \rightarrow \sigma_A^2 + \Delta t \cdot \alpha$$

Now we can take the limit $\Delta t \rightarrow 0$. We can define the "Kalman landmark strength," $w = \lim_{\Delta t \rightarrow 0} \frac{1}{\Delta t \cdot \sigma_L^2}$, which has units of $1 / (cm^2 \cdot time)$.

The remainder of the derivation assumes that this limit is finite and non-zero. In this limit, the difference equations become:

$$\Delta x \rightarrow \Delta x \cdot \left(1 - \frac{\sigma_A^2}{\sigma_L^2} \right) + \Delta t \cdot (v_A - v_L) = \Delta x + \Delta t \cdot ((v_A - v_L) - \Delta x \cdot (\sigma_A^2 w))$$

$$\sigma_A^2 \rightarrow \sigma_A^2 \left(1 - \frac{\sigma_A^2}{\sigma_L^2} \right) + \Delta t \cdot \alpha = \sigma_A^2 + \Delta t \cdot (\alpha - (\sigma_A^2)^2 w)$$

The second equation reaches an equilibrium at

$$(\sigma_A^2)_{EQ} = \sqrt{\frac{\alpha}{w}}$$

Which makes the first equation reach equilibrium at

$$(\Delta x)_{EQ} = \frac{v_A - v_L}{\sqrt{\alpha w}}$$

This steady state solution is exactly analogous to the linearization of the sub-critical regime of a previously published coupled oscillator attractor model (Campbell et al., 2018) around $D = 0$, where $\omega = \sqrt{\alpha w}$ (Figure S8).

While a Kalman filter captures the equilibrium behavior of the attractor model's sub-critical regime, it differs from the attractor model in two key ways. First, the attractor model as it stands has no representation of attractor uncertainty. One implication of this difference is that the two models make different predictions about the impact of landmarks on the future behavior of the system after the landmarks have been encountered. To see this, note that the Kalman model only reduces to the attractor model when we assume σ_A^2 has reached an equilibrium, and to reach equilibrium, the Kalman model requires uniform values of σ_L^2 . In reality, different landmarks can have different strengths. In this situation, the Kalman model and the attractor model make different predictions. If a Bayesian agent observes very strong landmarks, its certainty about its position self-estimate will go up, and it should weight subsequent landmarks less. Thus, if the animal encounters a particular strong/weak landmark during a gain change, σ_A^2 will decrease/increase, causing *future* landmarks to have less/more of an effect on the position estimate. The attractor model does not have any representation of uncertainty, so it does not predict this memory-like effect.

Second, because a Kalman filter is linear, it does not yield a super-critical regime. Instead of de-cohering after a sharp threshold, the position estimate will continue to shift indefinitely. Note that the equilibrium value of the position uncertainty $(\sigma_A^2)_{EQ} = \sqrt{\frac{\alpha}{w}}$ only depends on the diffusivity α and the Kalman landmark strength w . Thus, for the Kalman model, all values of gain will lead to pure shifts of the firing patterns, without any remapping, rescaling, or warping of any kind. Therefore, the Kalman model only applies to the "stable" (Cluster 1) responses (Figure 2G) and cannot capture the remapping phenomena we observed in the data (Figures 2H–2J and 7).