Goal-directed hippocampal theta sweeps during memory-guided navigation

Wenbo Tang[1], Xiyu Mei[1], Ryan E. Harvey[1], Estrella Carbajal-Leon[1], Talia Netzer[1], Hongyu Chang[1], Azahara Oliva[1], Antonio Fernandez-Ruiz[1*]

Affiliations:
[1]Department of Neurobiology and Behavior, Cornell University, Ithaca, NY, USA.

*Correspondence to: afr77@cornell.edu

Abstract

During navigation animals continuously sample their surrounding space and plan trajectories to distant goals. The brain mechanisms underlying these behaviors and how they coordinate to support memory-guided navigation in open environments is not understood. Using large scale recordings from behaving rats, we found two distinct types of place cell sequences within theta cycles that encoded trajectories sweeping beyond the animal's current location: stereotypic left-right alternating sweeps and learning-dependent goal-directed sweeps. Goal-directed sweeps predicted upcoming trajectories to remembered goal locations, were coordinated with prefrontal cortex activity and preferentially replayed during sharp-wave ripples. We describe a circuit mechanism for goal-directed theta sweeps involving an egocentric goal-direction signal in a subpopulation of CA1 cells together with reduced feedback inhibition. These results indicate a flexible mechanism to support different behavioral demands during navigation.

Introduction

Animals explore their environment constantly sampling their surroundings. They also navigate in a goal-directed manner to remembered locations such as food sources or shelters. Different neural mechanisms have been proposed to support these behaviors, with the hippocampal formation playing a central role (*1–4*). Hippocampal *place cells* fire at specific locations and generate a cognitive map of the environment as a population (*1*). In addition, place cell's firing is coordinated by the hippocampal theta rhythm, generating compressed temporal sequences that start at locations behind the animal and extend ahead of it ("theta sequences" or "sweeps";(*5–7*). These theta sequences have been profusely studied in rodents navigating linear mazes and found to sweep ahead of the path followed by the animal and to alternate between possible maze arms or goal locations at choice points, leading to the proposal that they represent possible behavioral alternatives to support goal-directed navigation (*8–13*). However, extending these studies to more naturalistic 2D navigation tasks has been challenging. A recent study analyzed theta sequences in rats free foraging in open arenas and reported that, instead of sweeping forward along the travelled path, they alternated stereotypically in left and right lateral sweeps

ahead of the animal (*14*). In contrast to previous claims (limited to experiments in linear mazes), the study proposed that theta sweeps are an experience-independent mechanism to sample the surrounding environment in an unbiased manner, regardless of goals and cognitive demands (*14*). However, to determine whether theta sequences are a mechanism for planning goal-directed trajectories, unbiased spatial sampling or both it is necessary to analyze them during both, free foraging as well as goal-directed navigation in open arenas.

If theta sequences support trajectory planning in open environments, they would need to be coupled with a vectorial signal that indicates the direction of the goal relative to the animal's position (in egocentric coordinates), even when the goal is at a distant location. They should point to the goal independently of heading direction and correlate with subsequent chosen trajectories. Furthermore, since goal-directed navigation does not only rely on the hippocampus but also on its interactions with multiple brain areas, including the prefrontal cortex (PFC; (*15*–*19*), other brain areas involved in goal-directed navigation and planning would be expected to coordinate with the hippocampus during goal-directed theta sweeps.

Results

Lateral alternating and goal-directed hippocampal theta sweeps during foraging

We performed large-scale neural ensemble recordings using ultra-high density silicon probes in the dorsal hippocampus and PFC of rats. Animals (n = 7) were trained to perform two different tasks in the same "cheeseboard maze" (*13*, *20*–*22*): random foraging (*rf*) and goal-directed navigation (*gdn*) to three hidden rewarded locations learned every session (Fig. 1A). Simultaneous recordings of hundreds of neurons allowed us to analyze hippocampal theta sequences in 2D during individual oscillation cycles (fig. S1). We used a Bayesian decoding approach to analyze spatial representations in theta sequences. During random foraging, rats explored the maze to search for randomly baited water rewards. During animal navigation we observed that theta sequences swept left and right ahead of the rat, alternating direction in successive theta cycles (Fig. 1B). These theta sweeps had an average angle of ~30 deg respect to the animal heading direction (Fig. 1, C and D) and an average length of 22.0 ± 2.5 cm (mean ± SD; *rf* = 23.2 ± 2.5 cm, *gdn* = 21.8 ± 2.5 cm, p = 0.32, rank sum test). Lateral alternating sweeps dominated theta sequence dynamics in all animals during random foraging sessions (Fig. 1, D and E and fig. S1), in agreement with a previous report (*14*, *23*).

In the goal-directed navigation task, rats had to discover three new hidden water reward locations in each session. They run 20-30 successive trials starting from a home box adjacent to the maze, finding the three rewards and returning to the box. All animals showed consistent learning performance as evidenced by the decreased latency and path length to retrieve all rewards across trials (fig. S1). In this task, we also observed lateral theta sweeps alternating between left and right directions during navigation, although to a smaller extent than during random foraging (Fig. 1, C and E, and fig. S1). In addition, we observed a new phenomenon that was absent during random foraging. During or preceding runs towards rewarded locations, some theta sequences swept forward in the direction of the goal (Fig. 1C). These goal-directed theta

sweeps were intermingled with left and right lateral sweeps (Fig. 1E and fig. S2). Across animals and sessions, we found that 16.3 ± 4.4% (mean ± SD) of theta sweeps in this task were goal-directed.

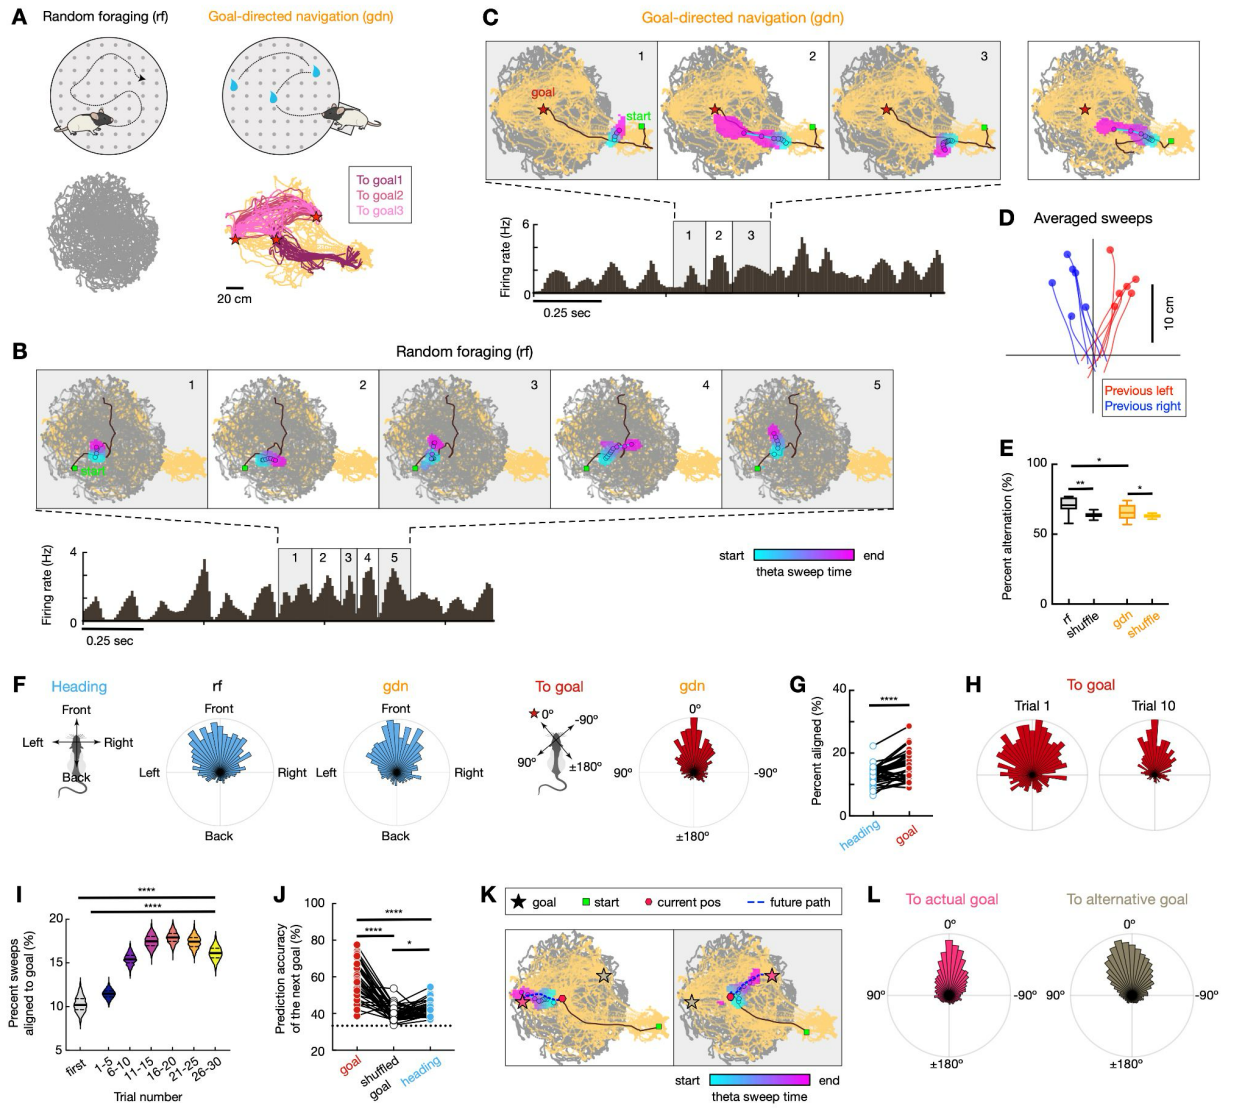

Fig. 1. Theta sweeps alternate laterally during random foraging and become goal-directed during memory-guided navigation

(A) Cheeseboard maze, showing random foraging (rf) and goal-directed navigation (*gdn*) tasks. Example animal trajectories are shown below, colored paths indicate goal-oriented trajectories. **(B)** Left–right alternating theta sweeps in CA1 during *rf*. *Top*: Sweeps decoded from CA1 population activity over five successive theta cycles (gray and yellow traces: positions visited during *rf* and *gdn*, respectively; black lines: current trajectory; green square: trajectory start; colored blobs: decoded positions, colored by time within sweep). *Bottom*: CA1 averaged firing rate showing 8–10 Hz theta modulation. **(C)** Goal-directed theta sweeps during *gdn* (data presented as in **B**). First and third sweeps are lateral; the second sweep points to the upcoming

goal (star). An additional goal-directed sweep example that does not align to the heading is shown on the right. **(D)** Sweeps rotated to head-centered coordinates and averaged across theta cycles in which the preceding sweep went left (red) or right (blue). Data from 6 example sessions. **(E)** Triplet sweep patterns (L–R–L or R–L–R) occurred more frequently than shuffled controls (*rf*: n = 12 sessions, 6 rats; *gdn*: n = 63 sessions, 5 rats; $^{**}p$ = 0.0078, $^{**}p$ = 0.032, Wilcoxon paired test; $^{*}p$(*rf* vs. *gdn*) = 0.029, rank-sum test). **(F)** Distributions of theta sweep directions relative to heading (left) and immediate next goal (right) **(G)** A greater percentage of theta sweeps aligned (≤ 10º) to the goal direction than to the heading direction ($^{****}p$ = 1.33e-5, Wilcoxon paired test). **(H)** Goal-centered sweep direction distributions during trial 1 vs. trial 10. **(I)** Percent of goal-directed sweeps across trials ($^{****}p$ < 1e-16, one-way ANOVA with *post hoc* test; bootstrapped n = 100). **(J)** Theta sweep directions better predicted the next goal chosen than heading direction or shuffled goals ($^{****}p$ = 3.29e-11, $^{*}p$ = 0.028, one-way repeated-measures ANOVA). **(K)** Examples of theta sweeps predicting the upcoming goal (red hexagon: current position; blue dashed line: future trajectory; pink star: actual goal; gray star: alternative goal). **(L)** When heading was biased toward a not chosen goal, theta sweep alignment to the chosen goal was stronger than to the alternative goal ($^{****}p$ = 4.0e-5, bootstrap test). N = 36 sessions (from 5 rats) in panels **F–L**.

Given that animals' heading direction tended to be correlated with goal direction in this task, we analyzed the alignment of theta sweeps to heading and goal directions. The percentage of theta sweeps aligned to the goal direction was significantly higher than that to the heading direction (Fig. 1F). The alignment of theta sweeps to the goal direction was not better than the alignment to heading direction during the first trial (mean ± SE = 10.3 ± 0.92 % aligned to goal, 10.4 ± 0.71% aligned to heading, p = 0.46, bootstrap test), but improved as a function of learning (Fig. 1H), while the fraction of lateral alternating sweeps remained constant (median ± SD = 66.7 ± 6.7 % for trial 1, and 66.7 ± 3.0 % for trials 26-30; n.s., p = 0.71, one-way ANOVA). Furthermore, the fraction of goal-directed sweeps gradually increased across learning trials (Fig. 1I).

As rats could run to any of three rewarded locations, we asked whether theta sweeps could predict which goal location would be visited next. Indeed, theta sweeps showed a high prediction accuracy of the next visited goal, significantly better than using goal-label shuffled data or the prediction of heading direction (Fig. 1J). We specifically analyzed theta sweeps preceding animal choice of trajectory to a specific goal and when the heading direction was biased to the goal that will not be chosen. In these cases, goal-directed sweeps were still biased to the chosen goal (Fig. 1, K and L), thus anticipating upcoming spatial trajectories independently of current heading.

Overall, our results show two different types of theta sweeps, left-right alternating lateral sweeps and goal-directed sweeps and show how their occurrence is modulated by behavioral demands and predictive of behavioral decisions during memory-guided navigation.

Egocentric goal-direction tuning in CA1 cells and theta sweeps

A mechanism involving inputs from head-direction cells together with spike-frequency adaptation has been proposed to explain left-right alternating theta sweeps (*23*, *24*). However, this mechanism cannot explain the goal-directed forward sweeps we found. While place cells and interneurons have been shown to prominently encode rewarded locations, this mechanism only works when an animal is close to the goal (*20*, *25*–*28*) and thus cannot account for goal-directed theta sweeps from locations distant to the goal (Fig. 1, C and J). We hypothesized that an egocentric goal-direction signal could bias theta sweeps towards the goal from any location, and independently of heading direction (Fig. 2, A and B).

To investigate the mechanisms of goal-directed theta sweeps, we first examined the tuning of individual CA1 cells to the goal direction. Rats sampled all goal-direction angles during the goal-directed navigation task (Fig. 2C). Given that most hippocampal cells are tuned to spatial location and heading direction, we used a linear-nonlinear Poisson (LN) model (*29*, *30*) to identify cells whose spike prediction was significantly improved by incorporating goal direction, beyond spatial location and heading direction alone. We found a subset of cells (10.7 ± 3.8%, mean ± SD) that displayed significant angular tuning to the goal direction (Fig. 2C-D). At the population level, goal-direction tuning spanned all angles but was more prominent when goal aligned with heading direction (Fig. 2C). Some goal-direction cells also showed clear place fields, and others exhibited angular tuning to the goal without clear place tuning (Fig. 2E). Although this type of response has not been previously described in rodents, it resembles similar responses observed in flying bats (*31*).

We next investigated whether the firing of goal-direction cells was coupled with goal-directed theta sweeps. We found that goal-directed cells fired more during goal-directed compared to lateral sweeps (Fig. 2, F-G). Moreover, goal-direction cells fired more on the late part of the theta cycle during goal-directed sweep, compared to non-goal-tunned cells (Fig. 2H). The late part of the cycle is when theta sequences sweep ahead towards the goal, while the earlier part of the cycle encodes locations closer to the animal (Fig. 2H). Thus, these observations support a potential role on goal-direction cells in biasing theta sweeps towards goals. Furthermore, the goal-direction tuning of CA1 cells was not significantly different from shuffles in the initial trials but increased rapidly afterwards (Fig. 2I), suggesting that goal-oriented theta sweeps require learning.

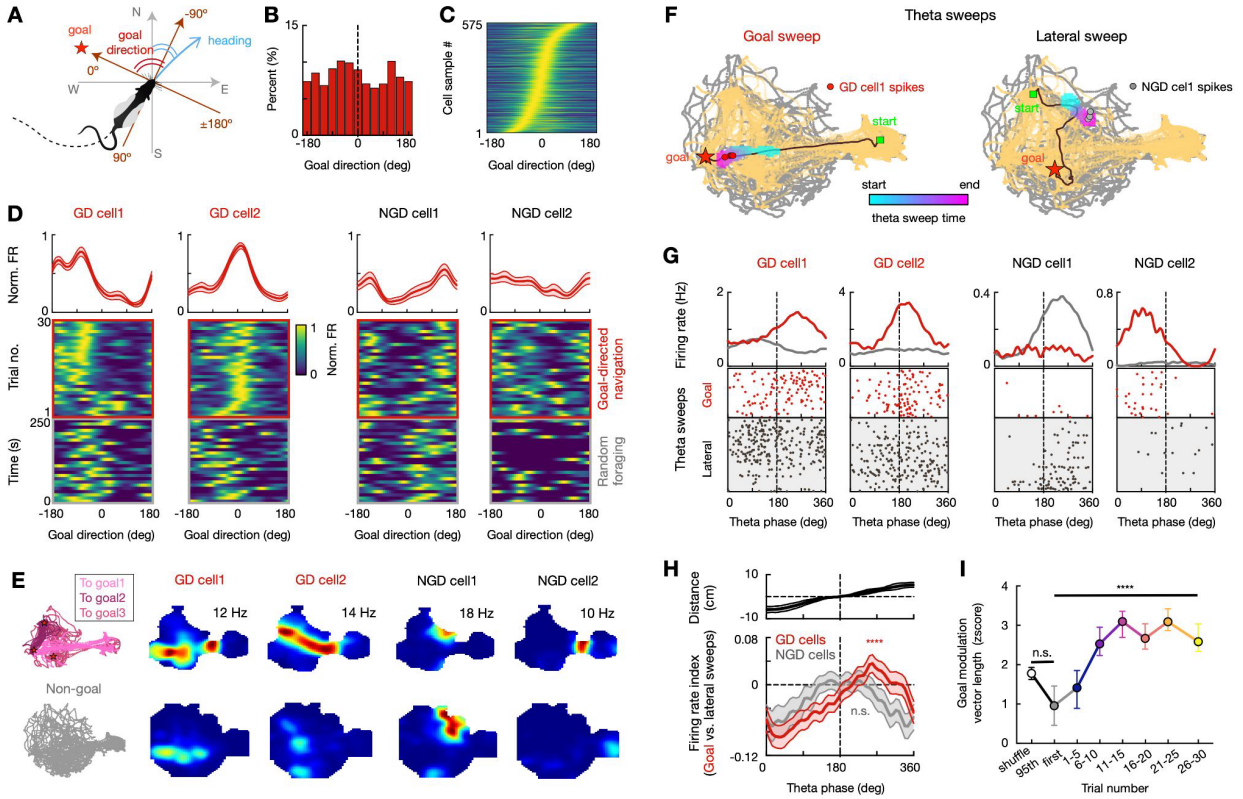

Fig. 2. An egocentric goal signal develops over learning during goal-directed theta sweeps.

(A) Goal-direction angle, defined as the azimuthal angle between the heading direction (blue arrowhead) and the animal-to-goal direction (red arrowhead; (*31*)). **(B)** Distribution of time spent by the animal in different goal-direction angles . **(C)** Normalized goal-direction tuning for all goal-tuned cells sorted by preferred direction (a given cell could contribute up to three cell samples with 3 different goals). **(D)** Four example CA1 cells (*columns*; GD cell: goal-direction cell; NGD cell: non-goal-direction cell). *Top*: Goal-direction tuning curves during *gdn* (mean ± SEM). *Middle*: Goal-direction angles across trials. *Bottom*: Goal-direction angles during rf (using goal locations from *gdn* task). **(E)** Spatial firing rate maps of the cells in **D**. GD cell 1: tuned to both goal direction and place. GD cell 2: tuned to goal direction, weak place tuning. NGD cell 1: classical place cell. NGD cell 2: remaps between gdn and rf without goal-direction tuning. **(F)** Example goal sweep (*left*) and lateral sweep (*right*) showing spike timing (circles) of GD and NGD cells within a theta sweep. **(G)** Theta-phase-aligned rasters and average firing rates of example GD and NGD cells during goal (red) and lateral (gray) sweeps. **(H)** Firing rate index (goal vs. lateral sweeps) as a function of theta phase (mean ± SEM). *Top*: Distance between decoded and actual position averaged over the theta cycles. *Bottom*, GD cells show elevated late thetaphase firing (**p = 0.0023, n = 191 cells from 5 rats), but NGD cells did not (n.s., p = 0.46, n = 126 cells from 5 rats). **(I)** Goal-direction modulation increased with learning, quantified as Rayleigh vector length z-scored relative to the shuffle distribution (median ± 95% CI; n.s., p > 0.99; ****p = 1.78e-12; Friedman test with Dunn's *post hoc*).

A circuit mechanism for goal-directed theta sweeps

We next sought to investigate the circuit mechanisms for goal-directed theta sweeps. In previous computational modeling work it has been proposed that spike-frequency adaptation can facilitate the generation of alternating lateral theta sweeps by preferentially inhibiting place cells along the travelled path (*23*, *24*). We thus hypothesized that a reduced spike-frequency adaptation together with a goal-direction signal could allow forward, goal-directed sweeps to emerge. To test this hypothesis, we incorporated variable spike-frequency adaptation and an egocentric goal-direction signal in a computational model that can generate alternating lateral theta sweeps (Fig. 3A; (*23*)) . With this modification, the model generated intermingled alternating lateral and goal-directed sweeps (Fig. 3B). In particular, the switch from lateral to goal-directed sweeps occurred in theta cycles with reduced spike-frequency adaptation and enhanced goal-direction activity during the late half of theta cycles (Fig. 3, C and D). These results suggest that a local egocentric goal-direction signal together with reduced spike-frequency adaptation could support the generation of goal-directed theta sweeps.

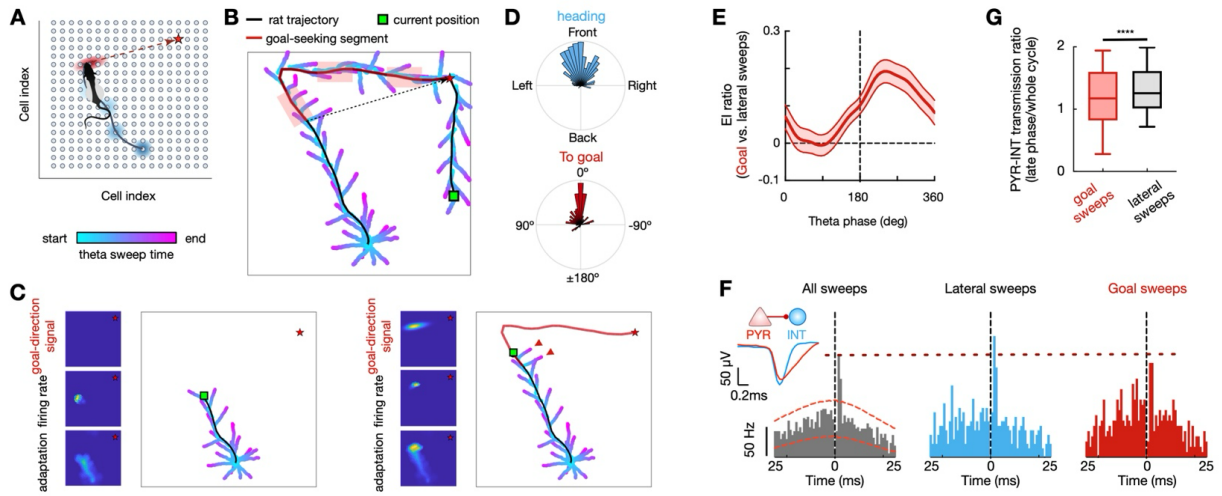

Fig. 3. Emergence of goal-directed theta sweeps mediated by feedback inhibition and egocentric goal signal.

(A–D) Computational model of goal-directed theta sweeps. **(A)** Schematic of the model, a single-bump, two-dimensional attractor network in which bump movement (blue circle) is driven by spike-frequency adaptation (*23*). Additionally, a subset of neurons exhibits goal-direction tuning (red ellipse), active during late theta phases. **(B)** Simulated theta sweeps, colored by time within each sweep. Goal-direction cells were active during the red segments. Black line: simulated rat trajectory. **(C)** Example snapshots showing goal signal strength, population firing rate, and spike-frequency adaptation strength before (*left*) and after (*right*) goal signal activation. Red arrowheads highlight goal-directed sweeps. **(E)** Firing rate ratio index of CA1 pyramidal cells (PYRs) vs. interneurons (INTs) during goal vs. lateral sweeps, across theta phases (mean ± SEM). Note the increased excitation-inhibition ratio during late theta phases (****p = 7.10e-6, rank-sum test, n = 36 sessions from 5 rats). **(F)** Cross-correlograms of a representative pre-

/postsynaptic PYR–INT pair during the late phase of all theta sweeps, lateral sweeps, and goal-directed sweeps. Dashed lines for reference Inset: averaged spike waveforms. **(G)** Decreased PYR–INT spike transmission probability during late theta phases of goal sweeps compared to lateral sweeps (**p = 0.0016, n = 242 pre-/postsynaptic PYR–INT pairs from 5 rats; Wilcoxon paired test).

We next turned to examine those mechanisms in our hippocampal data. A candidate to regulate spike-frequency adaptation in place cells is feedback inhibition from local interneurons (*32*, *33*). To determine the role of local inhibition during goal-directed and lateral theta sweeps, we analyzed it as a function of theta phase. Theta sequences sweep from positions behind the animal in the early half of the cycle to positions ahead of it in later theta phases (Fig. 2H,(*5*–*7*). Consistent with the model prediction, we found an enhanced excitation / inhibition balance (i.e., reduced inhibition) in the late phase of the theta cycle for goal-directed compared to lateral theta sweeps (Fig. 3E). To specifically examine the contribution of feedback inhibition, we analyzed functional monosynaptic pyramidal - interneuron (PYR-INT) cell pairs (*34*, *35*) PYR-INT spike-transmission probability in the late theta cycle was lower during goal-directed compared to lateral theta sweeps (Fig. 3, F and G), suggesting a reduced feedback inhibition. The observation of both, an increased firing of goal-directed cells (Fig. 2H) and reduced feedback inhibition (Fig. 3, E-G) during goal-directed compared to lateral sweeps, supported our model predictions for a mechanism for the generation of goal-directed theta sweeps.

Goal-directed theta sweeps are preferentially replayed and coordinated with the prefrontal cortex

During pauses in locomotion the hippocampus generates temporally compressed place-cell sequences that "replay" their ordered firing during theta oscillations in recent behavior (*36*–*38*). Hippocampal replay has been proposed to support learning and memory formation (*39*–*41*), and its generation is linked to plasticity during theta sequences (*13*, *42*). If goal-directed theta sweeps play a role in memory-guided navigation and are shaped by learning, in contrast to experience-independent lateral alternating sweeps (*14*), they might be preferentially replayed. To investigate the relationship between theta sweeps and replay, we detected sequential activity patterns during immobility periods in the maze that coincided by sharp-wave ripple (SWR) oscillations, and compared the spatial trajectories represented during these replay events with those represented during theta sweeps. Consistent with our hypothesis, replay sequences were preferentially biased towards goal locations (Fig. 4, A -B, and fig. S3), mirroring the structure of goal-directed theta sweeps (Fig. 4C). Moreover, replay sequences were significantly more similar to goal-directed theta sweeps than to lateral sweeps (Fig. 4D), suggesting that replay preferentially reinstates goal-directed theta sweeps.

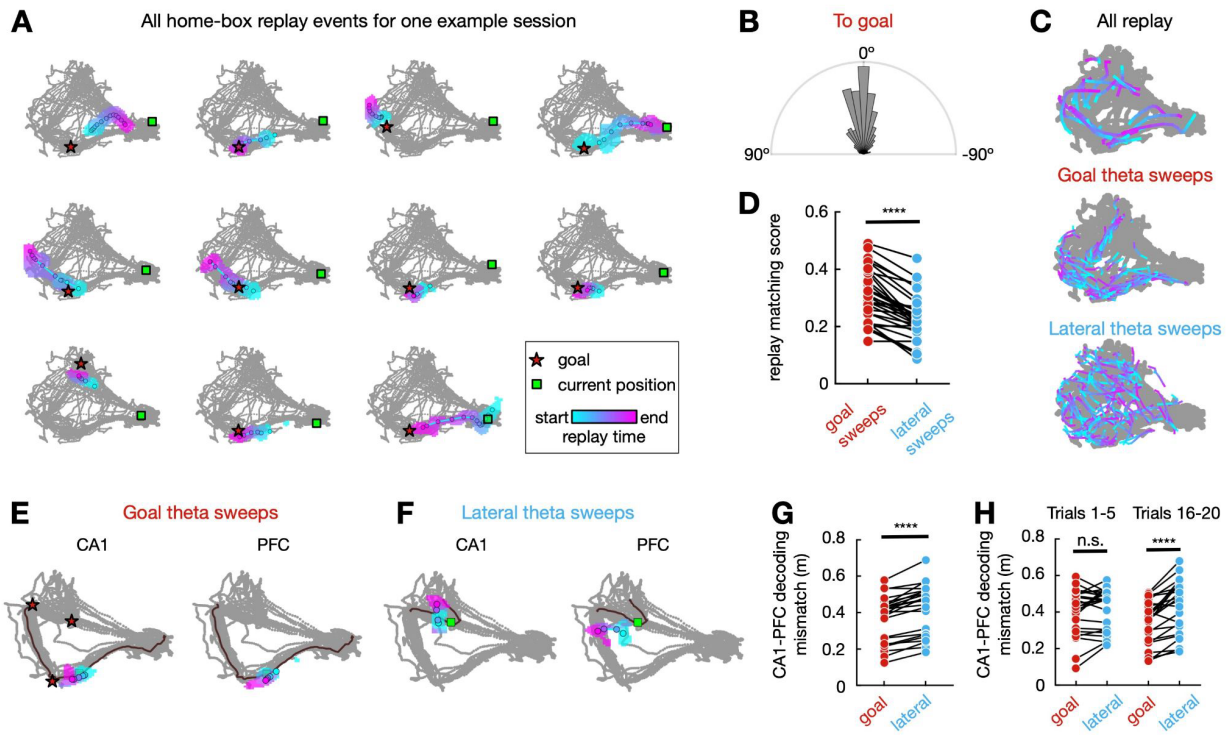

Fig. 4. Goal-directed theta sweeps are preferentially replayed and coordinated with the prefrontal cortex.

(A-D) Goal-directed theta sweeps were selectively replayed during awake SWRs. **(A)** All replay events (n = 11) occurring in the home box during a representative *gdn* session. Data are presented as in **Fig. 1B**, colored by time within SWRs. Note that most replay events were directed toward goal locations (red stars: closer goal location to replay decoded trajectory; filled circles: filled circles: positions with maximal decoded probability; green square: current animal position, gray lines: animal trajectories in the whole session). **(B)** Polar distribution of replay directions relative to goals during *gdn*. Plots span half a cycle, as forward and reverse replay directions were collapsed using absolute angles. **(C)** All awake replay events, goal-directed theta sweeps, and lateral sweeps from the example session shown in **A** (each color-coded line for one event). **(D)** Replay sequences more closely resembled goal-directed theta sweeps than lateral sweeps (****p = 4.66e-9, n = 31 sessions from 5 rats, Wilcoxon paired test). **(E-G)** CA1–PFC coordination was enhanced during goal-directed theta sweeps. **(E-F)** Example goal-directed **(E)** and lateral **(F)** CA1 (*left*) and PFC (*right*) theta sweeps showing decoded positions in. **(G)** Greater mismatch between CA1 and PFC decoded positions during goal-directed sweeps compared to lateral sweeps (****p = 2.38e-7, n = 24 sessions in 4 rats, Wilcoxon paired test). **(H)** CA1-PFC coordination during goal theta sweeps developed over learning (n.s., p = 0.30 for trials 1-5, and ****p = 1.05e-5 for trials 16-20, Wilcoxon paired test).

If goal-directed theta sweeps support planning of upcoming spatial trajectories, we would expect them to be coordinated with other brain regions also involved in planning and decision making, such as the PFC. PFC ensemble activity in theta cycles encoded compressed spatial trajectories

sweeping ahead of the animal in the goal-directed task (Fig. 4E; Tang 2021). These PFC theta sweeps were more coordinated with the hippocampus, i.e. represented more similar spatial trajectories, during goal-directed CA1 sweeps than during lateral sweeps (Fig. 4E and fig. S3). Furthermore, the coordination of PFC sweeps with hippocampal goal-directed sweeps increased with learning (Fig. 4H).

Theta sweeps exploit internal latent maps of the environment

In line with the traditional view that the hippocampus predominantly encodes physical space (*1*, *2*), theta sequences could be a mechanism for rapid sampling of the surrounding environment (*14*). An alternative possibility is that the hippocampus encodes different maps of the same environment that correspond to different behavioral rules or contexts, and theta sequences are a mechanism to explore these internal representations and guide flexible behavior. According to the first hypothesis, if an animal learns two different sets of rewarded locations in the same maze, its hippocampus may preserve the same spatial map in both conditions, albeit with some degree of place cell changes around rewarded locations (Fig. 5A) (*20*, *27*, *28*, *43*). Alternatively, the hippocampus may generate two different maps representing each reward configuration (or "latent" maps, since environmental features do not change; Fig. 5A). In the second case, theta sequences would be embedded within these latent representations and dynamically switch as behavioral context changes, rather than traversing physical space.

To distinguish between these two possibilities, we trained rats to perform the same goal-directed navigation task as before but with two different reward configurations in consecutive sessions on the same day (Fig. 5B). Place cell tuning remained largely stable across the two sessions (fig. S4). Some place fields near old and new rewarded locations changed their firing rates or shifted towards newly rewarded locations on the second session (fig. S4), as previously described (*20*, *21*). Overall, the hippocampal spatial representation was largely preserved across task configurations (fig. S4).

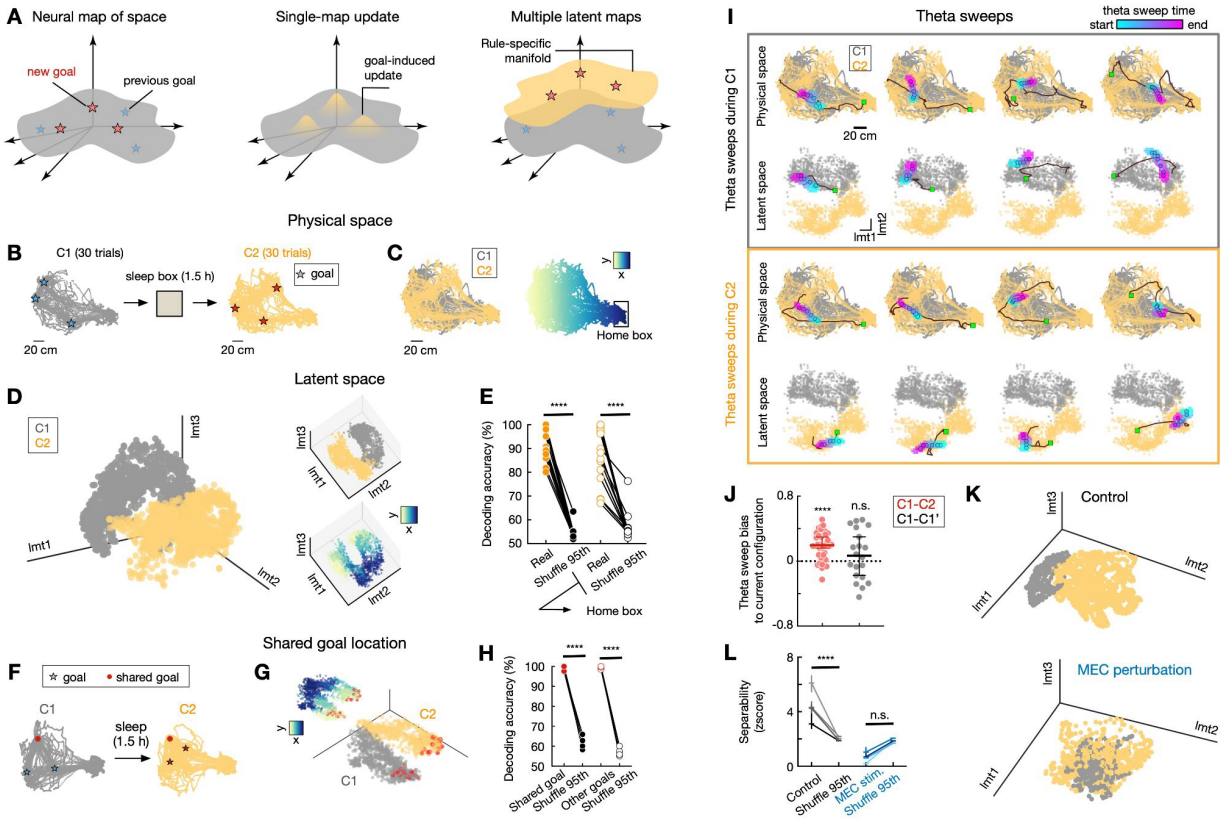

Fig. 5. Theta sequences exploit distinct rule-specific latent maps in the hippocampus.

(A) When learning a new set of hidden goals, the hippocampus may either update an existing map (*middle*, "single-map update") or form distinct, coexisting latent maps for each condition that are selectively engaged based on context (*right*, "multiple latent maps"). **(B)** Experimental design: animals learned configuration 1 (C1), followed by 1.5-hour rest, then configuration 2 (C2). **(C)** Animal trajectories, color-coded by configuration (*left*; C1, gray; C2, orange) and by position (*right*). **(D)** Three-dimensional embedding of neural activity using latent manifold tuning (LMT; see Methods), colored by condition and position (*inset*). Two distinct manifolds correspond to the two conditions. **(E)** Condition decoded from neural activity using *k*-nearest neighbor (kNN) and 10-fold cross-validation. Decoding accuracy exceeded that of shuffled rule labels (****p = 1e-4), even when decoding from home box activity alone (****p = 8.78e-9; n = 19 R1–R2 session pairs from 5 rats; one-way repeated-measures ANOVA). **(F–H)** Shared goals were represented in distinct latent states across conditions. **(F)** Task design with a shared goal between C1 and C2. **(G)** Manifold embedding showing separate latent states for visits to the shared goal (red circles). *Inset*: same embedding color-coded by position. **(H)** Condition decoding from neural activity at the shared goal location and other goals (****p = 6.00e-5 for shared goal, 6.54e-6 for other goals; n = 5 session pairs from 3 rats; one-way repeated-measures ANOVA). **(I)** Example theta sweeps decoded using standard spatial rate maps (top rows) and LMT-based latent tuning curves (bottom rows). Sweeps traversed similar spatial paths but exploited rule-specific latent maps. **(J)** Theta sweep decoding was strongly biased toward the current rule (****p = 4.63e-7; n = 38 sessions from 5 rats) and was not explained by temporal drift (C1' re-exposure: n.s., p = 0.33; n = 20

sessions from 4 rats; rank-sum test vs. 0). **(K–L)** Optogenetic disruption of MEC spike timing, which impaired theta sequences (*13*), also impaired latent map formation. **(K)** Manifold embeddings from an example control and MEC perturbation sessions. **(L)** Separability of condition-specific latent maps, measured by Dunn's index, was reduced during MEC perturbation (all $p > 0.05$, compared to rule-label shuffles, Wilcoxon paired test; $n = 4$ sessions from 3 rats per condition).

To distinguish between the "single-map updating" and "multiple latent maps" hypotheses (Fig. 5A), we analyzed the representational structure of ensemble neural activity, or neural manifold, across task configurations. To address this, we applied Latent Manifold Tuning (LMT; (*44*), an unsupervised method that simultaneously estimates neural manifolds and neuronal tuning to latent variables. If the hippocampus encodes multiple latent maps of the same physical maze, then distinct reward configurations should be separable based on population activity, even at identical physical locations throughout the maze, not just near the goals (Fig. 5A). Consistent with this prediction, we found that, despite only modest changes in single-cell place fields across sessions, population-level dynamics revealed distinct latent maps for each reward configuration within the same maze (Fig. 5D). Importantly, the same physical locations were represented as different latent states depending on the task context, including in regions far from the goals, such as the home box (Fig. 5D). As a result, the current reward configuration could be accurately decoded from neural activity (Fig. 5E). Furthermore, even when the two reward configurations shared a common goal location (Fig. 5F), neural activity at the shared goal diverged across the two latent maps (Fig. 5G-H), indicating a context-dependent coding of the same goal.

Place cell responses in the same environment gradually change over time, or drift, even in the absence of external changes (*29*, *45*, *46*). To test whether temporal drift can explain the generation of distinct maps of the same maze, we re-exposed rats to the second reward configuration in a new session after 1.5 hours of rest. The distance of the neural subspace in the second task session, with the previous one (with a different reward configuration) was larger than compared with the re-exposure session (with the same reward configuration) (fig. S5). This indicates that temporal drift alone cannot explain the generation of distinct maps for the different reward configurations.

Next, we examined the dynamics of theta sweeps across configurations. According to the latent-map hypothesis, we would expect theta sequences to sweep only within the relevant latent map for each configuration. To test this, we decoded theta sweeps in the latent space comprising the subspaces of both task configurations. Theta sweeps from one of the configurations were restricted to the corresponding neural subspace for that configuration (Fig. 5Ig,i). Even theta sweeps along similar spatial trajectories in both sessions, were completely segregated in neural space (Fig. 5I). In a subset of sessions, animals encounter the same reward configuration across two sessions. For those sessions, theta sweeps did not show a bias to the subspace of the current session (Fig. 5J), suggesting that the separation of theta sweeps for different task configurations could not be simply explained by temporal drift. Overall, these results suggest that theta sweeps are a mechanism to explore internal representations of space.

Finally, we analyzed previously published data in which we disrupted CA1 theta sequences by optogenetically perturbing the temporal structure of medial entorhinal activity while rats performed the same goal-directed cheeseboard maze task (*13*). Disrupting theta sequences during learning trials prevented the generation of two distinct maps for different reward configurations (Fig. 5, K and L). This suggests that synaptic plasticity during theta sequences may be a driving mechanism for the generation of latent spatial representations in the hippocampus.

Discussion

We found that the hippocampus generates two distinct types of temporally compressed place cell sequences within theta cycles during navigation. During random foraging in an open maze, theta sweeps alternated left and right of the animal's location in a highly stereotyped pattern. During goal-directed navigation to previously learned hidden goals on the same maze, some theta sequences swept towards one of the goal locations. These goal-directed sweeps were intermingled with lateral sweeps, and independent of heading direction. In contrast to the rigid nature of lateral theta sweeps, goal-directed sweeps were modulated by learning and predicted upcoming spatial trajectories to chosen goals. We propose that both types of patterns support complementary functions during navigation. Lateral sweeps might be a mechanism for unbiased sampling of the surrounding space, including unvisited locations, and contribute to the generation and updating of internal maps (*14*). On the other hand, goal-directed sweeps may facilitate memory-guided navigation and trajectory planning. In support of this interpretation, goal-directed hippocampal sweeps were more coordinated with trajectory representations in the PFC, another region implicated in behavioral planning and goal-directed navigation (*15*–*19*), than lateral theta sweeps. These results provide a unified framework to explain previous findings, including exclusive lateral sweeps during open field and linear maze tasks (*14*), sweeps to different arms at choice points in linear mazes (*8*, *9*, *11*) and forward sweeps to reward ports along linear tracks (*10*, *12*).

We observed that a subset of hippocampal cells developed goal-direction tuning selective over learning. These goal-direction cells provided a vectorial egocentric signal pointing to goals, even from distant locations in the maze. They thus provide a complementary mechanism for goal coding to the local goal overrepresentation by place cells commonly observed in rodents (*20*, *26*–*28*), and are reminiscent of similar goal vector signals found in flying bats (*31*). Goal-direction cells were increasingly active during goal-directed theta sweeps. Computational simulations indicated that a combination of variable spike frequency adaption (*23*, *24*) and an egocentric goal-direction signal in a continuous attractor network model can generate both lateral and goal-directed theta sweeps. In support of the model prediction, we found a decreased feedback inhibition (a mechanism for spike frequency adaptation- (*32*) to hippocampal place cells during goal-directed theta sweeps.

Hippocampal place cells replay experienced spatial trajectories during pauses in locomotion and sleep, coordinated by SWRs (*13*, *36*–*39*, *47*). These network events have been causally linked to learning, memory recall and consolidation (*48*–*51*). Previous work suggested that the generation of replay sequences depends on synaptic plasticity mediated by the fast activation of place cells

within theta cycles (*13*, *42*). In our goal-directed navigation task, we found that replay events during immobility periods in the maze were strongly biased to goal locations and more correlated with goal-directed than lateral theta sweeps. This observation suggests that goal-directed theta sweeps could be a mechanism promoting the prioritization of behavioral relevant spatial trajectory during replay and subsequent memory consolidation (*52*–*56*).

When rats learned two different sets of rewarded locations on the same maze, the hippocampus generated two distinct representations, or maps, of the maze. Despite small changes in place cell tuning, mostly localized around rewarded locations, at the population level, low dimensional representations of the same maze for both reward configurations were completely segregated. This result is in line with recent evidence suggesting that the hippocampal cognitive map does not only reflect physical space and associated sensory cues but also internal variables such as behavioral relevance, expectations or learned rules (*29*, *57*, *58*). Furthermore, theta sweeps along similar spatial trajectories during sessions with different reward configuration were restricted to the corresponding neural subspace. Therefore, theta sweeps could provide a mechanism for mental exploration of internal representation of the environment in support of planning, simulation and prediction of behavioral outcomes.

Acknowledgments
The authors thank members of the Oliva and Fernandez-Ruiz labs for providing useful feedback on the manuscript. This work was supported by NIH grant R01MH130367 and Whitehall Foundation (AO), NIH grant R01MH136355, DP2MH136496, Sloan Fellowship, Whitehall Research Grant, Klingenstein-Simons Fellowship and Pershing Square Foundation's MIND Prize (AFR), and Klarman Fellowship (WT).

Author contributions
W.T, X.M., R.E.H, E.L-C and T.N performed the experiments. W.T, X.M., R.E.H, H.C. and A.F.R. analyzed the data. W.T., A.F.R. and A.O. wrote the manuscript with input from all authors. A.F.R. and A.O. supervised the work.

Declaration of interests
The authors declare no competing interests.

References

1. J. O'Keefe, L. Nadel, *The Hippocampus as a Cognitive Map* (Clarendon Press ; Oxford University Press, Oxford : New York, 1978).

2. E. I. Moser, E. Kropff, M.-B. Moser, Place Cells, Grid Cells, and the Brain's Spatial Representation System. *Annual Review of Neuroscience* **31**, 69–89 (2008).

3. R. G. Morris, P. Garrud, J. N. Rawlins, J. O'Keefe, Place navigation impaired in rats with hippocampal lesions. *Nature* **297**, 681–683 (1982).

4. B. L. McNAUGHTON, C. A. Barnes, J. L. Gerrard, K. Gothard, M. W. Jung, J. J. Knierim, H. Kudrimoti, Y. Qin, W. E. Skaggs, M. Suster, K. L. Weaver, Deciphering The Hippocampal Polyglot: the Hippocampus as a Path Integration System. *J Exp Biol* **199**, 173–185 (1996).

5. W. E. Skaggs, B. L. McNaughton, M. A. Wilson, C. A. Barnes, Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus* **6**, 149–172 (1996).

6. G. Dragoi, G. Buzsáki, Temporal Encoding of Place Sequences by Hippocampal Cell Assemblies. *Neuron* **50**, 145–157 (2006).

7. D. J. Foster, M. A. Wilson, Hippocampal theta sequences. *Hippocampus* **17**, 1093–1099 (2007).

8. A. Johnson, A. D. Redish, Neural Ensembles in CA3 Transiently Encode Paths Forward of the Animal at a Decision Point. *J. Neurosci.* **27**, 12176–12189 (2007).

9. A. S. Gupta, M. A. A. van der Meer, D. S. Touretzky, A. D. Redish, Segmentation of spatial experience by hippocampal theta sequences. *Nat Neurosci* **15**, 1032–1039 (2012).

10. A. M. Wikenheiser, A. D. Redish, Hippocampal theta sequences reflect current goals. *Nat Neurosci* **18**, 289–294 (2015).

11. K. Kay, J. E. Chung, M. Sosa, J. S. Schor, M. P. Karlsson, M. C. Larkin, D. F. Liu, L. M. Frank, Constant Sub-second Cycling between Representations of Possible Futures in the Hippocampus. *Cell* **180**, 552-567.e25 (2020).

12. C. Zheng, E. Hwaun, C. A. Loza, L. L. Colgin, Hippocampal place cell sequences differ during correct and error trials in a spatial memory task. *Nat Commun* **12**, 3373 (2021).

13. C. Liu, R. Todorova, W. Tang, A. Oliva, A. Fernandez-Ruiz, Associative and predictive hippocampal codes support memory-guided behaviors. *Science* **382**, eadi8237 (2023).

14. A. Z. Vollan, R. J. Gardner, M.-B. Moser, E. I. Moser, Left–right-alternating theta sweeps in entorhinal–hippocampal maps of space. *Nature* **639**, 995–1005 (2025).

15. H. T. Ito, S.-J. Zhang, M. P. Witter, E. I. Moser, M.-B. Moser, A prefrontal–thalamo–hippocampal circuit for goal-directed spatial navigation. *Nature* **522**, 50–55 (2015).

16. J. D. Shin, W. Tang, S. P. Jadhav, Dynamics of Awake Hippocampal-Prefrontal Replay for Spatial Learning and Memory-Guided Decision Making. *Neuron* **104**, 1110-1125.e7 (2019).

17. R. E. Harvey, H. L. Robinson, C. Liu, A. Oliva, A. Fernandez-Ruiz, Hippocampo-cortical circuits for selective memory encoding, routing, and replay. *Neuron* **111**, 2076-2090.e9 (2023).

18. V. Hok, E. Save, P. P. Lenck-Santini, B. Poucet, Coding for spatial goals in the prelimbic/infralimbic area of the rat frontal cortex. *Proceedings of the National Academy of Sciences* **102**, 4602–4607 (2005).

19. E. Z. Patai, H. J. Spiers, The Versatile Wayfinder: Prefrontal Contributions to Spatial Navigation. *Trends Cogn Sci* **25**, 520–533 (2021).

20. D. Dupret, J. O'Neill, B. Pleydell-Bouverie, J. Csicsvari, The reorganization and reactivation of hippocampal maps predict spatial memory performance. *Nat Neurosci* **13**, 995–1002 (2010).

21. C. N. Boccara, M. Nardin, F. Stella, J. O'Neill, J. Csicsvari, The entorhinal cognitive map is attracted to goals. *Science* **363**, 1443–1447 (2019).

22. A. Fernández-Ruiz, A. Oliva, M. Soula, F. Rocha-Almeida, G. A. Nagy, G. Martin-Vazquez, G. Buzsáki, Gamma rhythm communication between entorhinal cortex and dentate gyrus neuronal assemblies. *Science* **372**, eabf3119 (2021).

23. J. Widloski, D. Theurel, D. J. Foster, Spontaneous alternation of place-cell sequences in the open field through spike frequency adaptation. *Cell Reports* **44**, 115475 (2025).

24. Z. Ji, T. Chu, S. Wu, N. Burgess, A systems model of alternating theta sweeps via firing rate adaptation. *Current Biology* **35**, 709-722.e5 (2025).

25. N. Jeong, X. Zheng, A. L. Paulson, S. M. Prince, V. P. Nguyen, S. R. Thomas, C. E. Gilpin, M. C. Goodson, A. C. Singer, Goal-specific hippocampal inhibition gates learning. *Nature*, 1–9 (2025).

26. C. Grienberger, J. C. Magee, Entorhinal cortex directs learning-related changes in CA1 representations. *Nature* **611**, 554–562 (2022).

27. M. Sosa, M. H. Plitt, L. M. Giocomo, A flexible hippocampal population code for experience relative to reward. *Nat Neurosci* **28**, 1497–1509 (2025).

28. J. L. Gauthier, D. W. Tank, A Dedicated Population for Reward Coding in the Hippocampus. *Neuron* **99**, 179-193.e7 (2018).

29. W. Tang, H. Chang, C. Liu, S. Perez-Hernandez, W. Y. Zheng, J. Park, A. Oliva, A. Fernandez-Ruiz, A hippocampal population code for rapid generalization. bioRxiv [Preprint] (2025). https://doi.org/10.1101/2025.03.15.643456.

30. K. Hardcastle, N. Maheswaranathan, S. Ganguli, L. M. Giocomo, A Multiplexed, Heterogeneous, and Adaptive Code for Navigation in Medial Entorhinal Cortex. *Neuron* **94**, 375-387.e7 (2017).

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

32. M. J. Higley, D. Contreras, Balanced Excitation and Inhibition Determine Spike Timing during Frequency Adaptation. *J. Neurosci.* **26**, 448–457 (2006).

33. A. Fernández-Ruiz, A. Oliva, G. A. Nagy, A. P. Maurer, A. Berényi, G. Buzsáki, Entorhinal-CA3 Dual-Input Control of Spike Timing in the Hippocampus by Theta-Gamma Coupling. *Neuron* **93**, 1213-1226.e5 (2017).

34. D. F. English, S. McKenzie, T. Evans, K. Kim, E. Yoon, G. Buzsáki, Pyramidal Cell-Interneuron Circuit Architecture and Dynamics in Hippocampal Networks. *Neuron* **96**, 505-520.e7 (2017).

35. F. Sharif, B. Tayebi, G. Buzsáki, S. Royer, A. Fernandez-Ruiz, Subcircuits of Deep and Superficial CA1 Place Cells Support Efficient Spatial Coding across Heterogeneous Environments. *Neuron* **109**, 363-376.e6 (2021).

36. D. J. Foster, M. A. Wilson, Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature* **440**, 680–683 (2006).

37. K. Diba, G. Buzsáki, Forward and reverse hippocampal place-cell sequences during ripples. *Nat Neurosci* **10**, 1241–1242 (2007).

38. M. P. Karlsson, L. M. Frank, Awake replay of remote experiences in the hippocampus. *Nat Neurosci* **12**, 913–918 (2009).

39. G. Buzsáki, Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and planning. *Hippocampus* **25**, 1073–1188 (2015).

40. D. J. Foster, Replay Comes of Age. *Annu. Rev. Neurosci.* **40**, 581–602 (2017).

41. H. R. Joo, L. M. Frank, The hippocampal sharp wave–ripple in memory retrieval for immediate use and consolidation. *Nat Rev Neurosci* **19**, 744–757 (2018).

42. C. Drieu, R. Todorova, M. Zugaro, Nested sequences of hippocampal assemblies during behavior support subsequent sleep replay. *Science* **362**, 675–679 (2018).

43. H. S. C. Chiossi, M. Nardin, G. Tkačik, J. Csicsvari, Learning reshapes the hippocampal representation hierarchy. *Proceedings of the National Academy of Sciences* **122**, e2417025122 (2025).

44. A. Wu, N. A. Roy, S. Keeley, J. W. Pillow, “Gaussian process based nonlinear latent structure discovery in multivariate spike train data” in *Advances in Neural Information Processing Systems* (Curran Associates, Inc., 2017;

45. Y. Ziv, L. D. Burns, E. D. Cocker, E. O. Hamel, K. K. Ghosh, L. J. Kitch, A. E. Gamal, M. J. Schnitzer, Long-term dynamics of CA1 hippocampal place codes. *Nat Neurosci* **16**, 264–266 (2013).

46. T. Hainmueller, A. Cazala, L.-W. Huang, M. Bartos, Subfield-specific interneuron circuits govern the hippocampal response to novelty in male mice. *Nat Commun* **15**, 714 (2024).

47. A. K. Lee, M. A. Wilson, Memory of Sequential Experience in the Hippocampus during Slow Wave Sleep. *Neuron* **36**, 1183–1194 (2002).

48. G. Girardeau, K. Benchenane, S. I. Wiener, G. Buzsáki, M. B. Zugaro, Selective suppression of hippocampal ripples impairs spatial memory. *Nat Neurosci* **12**, 1222–1223 (2009).

49. S. P. Jadhav, C. Kemere, P. W. German, L. M. Frank, Awake hippocampal sharp-wave ripples support spatial memory. *Science* **336**, 1454–1458 (2012).

50. A. Fernández-Ruiz, A. Oliva, E. Fermino de Oliveira, F. Rocha-Almeida, D. Tingley, G. Buzsáki, Long-duration hippocampal sharp wave ripples improve memory. *Science* **364**, 1082–1086 (2019).

51. H. Chang, W. Tang, A. M. Wulf, T. Nyasulu, M. E. Wolf, A. Fernandez-Ruiz, A. Oliva, Sleep microstructure organizes memory replay. *Nature* **637**, 1161–1169 (2025).

52. B. E. Pfeiffer, D. J. Foster, Hippocampal place-cell sequences depict future paths to remembered goals. *Nature* **497**, 74–79 (2013).

53. H. F. Ólafsdóttir, F. Carpenter, C. Barry, Task Demands Predict a Dynamic Switch in the Content of Awake Hippocampal Replay. *Neuron* **96**, 925-935.e6 (2017).

54. A. K. Gillespie, D. A. Astudillo Maya, E. L. Denovellis, D. F. Liu, D. B. Kastner, M. E. Coulter, D. K. Roumis, U. T. Eden, L. M. Frank, Hippocampal replay reflects specific past experiences rather than a plan for subsequent choice. *Neuron* **109**, 3149-3163.e6 (2021).

55. M. G. Mattar, N. D. Daw, Prioritized memory access explains planning and hippocampal replay. *Nat Neurosci* **21**, 1609–1617 (2018).

56. M. Huelin Gorriz, M. Takigawa, D. Bendor, The role of experience in prioritizing hippocampal replay. *Nat Commun* **14**, 8157 (2023).

57. A. A. Fenton, Remapping revisited: how the hippocampus represents different spaces. *Nat. Rev. Neurosci.* **25**, 428–448 (2024).

58. I. Zutshi, A. Apostolelli, W. Yang, Z. S. Zheng, T. Dohi, E. Balzani, A. H. Williams, C. Savin, G. Buzsáki, Hippocampal neuronal activity is aligned with action plans. *Nature*, 1–9 (2025).

59. A. Oliva, A. Fernández-Ruiz, G. Buzsáki, A. Berényi, Role of Hippocampal CA2 Region in Triggering Sharp-Wave Ripples. *Neuron* **91**, 1342–1355 (2016).

60. A. Mathis, P. Mamidanna, K. M. Cury, T. Abe, V. N. Murthy, M. W. Mathis, M. Bethge, DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nat Neurosci* **21**, 1281–1289 (2018).

61. J. Dimidschstein, Q. Chen, R. Tremblay, S. L. Rogers, G.-A. Saldi, L. Guo, Q. Xu, R. Liu, C. Lu, J. Chu, J. S. Grimley, A.-R. Krostag, A. Kaykas, M. C. Avery, M. S. Rashid, M. Baek, A. L. Jacob, G. B. Smith, D. E. Wilson, G. Kosche, I. Kruglikov, T. Rusielewicz, V. C. Kotak, T. M. Mowery, S. A. Anderson, E. M. Callaway, J. S. Dasen, D. Fitzpatrick, V. Fossati, M. A. Long, S. Noggle, J. H. Reynolds, D. H. Sanes, B. Rudy, G. Feng, G. Fishell, A viral strategy for targeting and manipulating interneurons across vertebrate species. *Nat Neurosci* **19**, 1743–1749 (2016).

62. P. C. Petersen, J. H. Siegle, N. A. Steinmetz, S. Mahallati, G. Buzsáki, CellExplorer: A framework for visualizing and characterizing single neurons. *Neuron* **109**, 3594-3608.e2 (2021).

63. E. L. Mackevicius, A. H. Bahle, A. H. Williams, S. Gu, N. I. Denisenko, M. S. Goldman, M. S. Fee, Unsupervised discovery of temporal sequences in high-dimensional datasets, with applications to neuroscience. *eLife* **8**, e38471 (2019).

Materials and Methods

Animals

All experiments conformed to guidelines established by the National Institutes of Health and have been approved by Cornell University's Institutional Animal Care and Use Committee. Rats (*n* = 7; adult male Long-Evans, 300-500 g, 3-6 months old) were kept in the vivarium on a 12-hour light/ dark cycle. Temperature and humidity in the room were kept at 68-72 F and 40-60%, respectively. They were housed 2 per cage before surgery and individually after it.

Surgical procedures

Silicon probes (NeuroNexus, Cambridge Neurotech, or Diagnostic Biosignals) were mounted on custom-made 3D-printed micro-drives to allow precise adjustment of the vertical position of sites after implantation. The probes were inserted above the target region. Craniotomies were sealed with sterile wax. Two stainless steel screws were placed bilaterally over the cerebellum to serve as ground and reference electrodes. Several additional screws were driven into the skull and covered with dental cement to strengthen the implant. Finally, a copper mesh mounted on a 3D-printed resin base was attached to the skull with dental cement and connected to the ground screw to act as a Faraday cage, attenuating the contamination of the recordings by environmental electric noise and protecting the headgear. After post-surgery recovery, probes were moved gradually in 50 to 150 µm steps per day until the desired position was reached. Hippocampal and cortical layers were identified physiologically by unit activity and characteristic LFP patterns (*17*, *22*, *59*). A variety of different silicon probes were implanted targeting dorsal hippocampal CA1 (-4.0-4.5 mm antero-posterior from Bregma, AP; ± 2.6 mm from midline, ML; 5 rats, right hemisphere; 2 rats, bilateral) and the right PFC (-3 mm AP and 0.7 mm ML). For optogenetic experiments (*13*), rats were also implanted with custom-made optic fiber arrays (three 200-µm core multimode fiber each, ~500 µm apart, connected to a single 2.5 mm steel ferrule; Doric Lenses) in both hemispheres over the medial entorhinal cortex (–7.7 / –8.4 / –9.1 AP; ± 4.6 ML and 4.7/ 4.3/3.2 mm from the surface of the brain, for each fiber respectively; see also *optogenetic manipulations*).

Behavioral tasks

After surgery, animals were handled daily and accommodated to the experimenter, recording room, and cables for one week before the start of the experiments. Prior to the start of the behavioral experiment, the animals were water restricted. Recordings were conducted using the Intan RHD2000 interface board or Intan Recording Controller, sampled at 20 kHz. Amplification and digitization were done on the head stage. In all the behavioral sessions, animals' position was tracked with an overhead camera (Basler) at 30 Hz (Pylon 5.0 software), and later extracted using DeepLabCut (DLC, v2.3.0) (*60*).

The cheeseboard maze was a circular platform (120-cm diameter), where the animals learned to find three goal wells that contained water rewards. A trial was completed once the animal had retrieved all rewards and returned to the start box to collect an additional food pellet reward. The locations of the goal wells changed daily but were fixed within a session. This strategy required the animals to update their memory for the new goal locations in each session but in an

otherwise familiar environment. Note that there was always a delay of approximately 30 seconds between trials. If the rat could not find all three rewards in 2 minutes, the trial was ended, and the animal was placed back in the start box. Each experimental day began with a pre-probe session consisting of five trials using the same reward configuration as the previous day, to assess whether the animal remembered the previous day's reward locations. This was followed by a 1.5-2 hour sleep session. Subsequently, the animal underwent a learning session to acquire a new set of three reward locations (rule 1, R1), conducted either with or without light stimulation. After a post-learning 1.5-2 hour sleep session, the animal underwent training involving either the acquisition of a novel reward configuration (rule 2, R2) or the recall of a previously learned configuration (re-exposure to rule 1, R1'). Each training session consisted of 20 to 30 trials. Learning performance was evaluated as the latency and the distance traveled from the start box to collect the three rewards. For the optogenetic experiments (Fig. 5, K-L, see also *optogenetic manipulations*), stimulation (Stim ON) and control (Stim OFF) sessions were alternated in a counterbalanced design across animals. In a subset of sessions, animals foraged for randomly baited water rewards on the same cheeseboard, either for 20 minutes following a 1.5-2 hour sleep session, or for 5-10 minutes immediately after the 20–30 trial learning session.

Optogenetic manipulations

For optogenetic experiments, rats were injected with custom-prepared AAV5-mDlx-hChR2(H134R)-mCherry from AddGene (plasmids were a gift from Dr. Gord Fishell; (*61*). Three injections per hemisphere were performed along the dorso-ventral MEC axes as follows: 1) -7.7 AP, ±4.6 ML, 4.7 mm depth, 200 nL; 2) -8.4 AP, ±4.6 ML, 4.3 mm depth, 400 nL; 3) -9.1 AP, ±4.6 ML, 3.2 depth, 700 nL. After injection, craniotomies were sealed and animals recovered in the vivarium for three weeks. Following this period, a second surgical procedure for implanting optic fibers and electrodes was performed, as described above. Optic fiber arrays were implanted in the same craniotomies performed previously for virus injection. For optogenetic stimulation, fiber array ferrules were connected with mating sleeves to 450-nm blue light-emitting laser diodes coupled to 2.5 mm steel ferrules (PL-450, Osram).

Optogenetic perturbations were performed by delivering blue light, modulated with a positive 53 Hz current sinusoid using an isolated current driver (Thorlabs). Light intensity was calibrated for each animal during home cage recordings by analyzing the suppression of LFP gamma power in the stratum lacunosum-moleculare during stimulation. A minimum power of 3 mW and a maximum of 6 mW was used. In the goal-directed navigation task on the cheeseboard, light stimulation was applied during all running periods in the maze during learning trials.

Recording system and data preprocessing

An Intan RHD2000 interface board or Intan Recording Controller was used for electrophysiological recordings. The sampling rate was set at 20 kHz. Both amplification and digitization were done in the head stage (Intan Technologies). Data were visualized online during recording using the Intan software and Neuroscope (Neurosuite). LFP signals were down-sampled at 1250 Hz for subsequent analysis.

Tissue processing and immunohistochemistry

Following the termination of the experiments, animals were deeply anesthetized and perfused transcardially first with 0.9% saline solution followed by 4% paraformaldehyde solution. The brains were sectioned into 70-µm thick slices (Leica Vibratome). The sections were washed and mounted on glass slides with a fluorescence medium (Fluoroshield with DAPI - F6057, Sigma, USA). A confocal microscope (Zeiss LSM 800) was used to obtain high-quality photos.

Quantification and statistical analyses

Spike sorting and single unit classification

Spike sorting was performed semi-automatically using KiloSort (https://github.com/cortex-lab/KiloSort), followed by manual curation using the software Phy (https://github.com/kwikteam/phy) and custom designed plugins (https://github.com/petersenpeter/phy-plugins) to obtain well-isolated single units. Cluster quality was assessed by manual inspection of waveforms and auto-correlograms, and by isolation distance metrics. Multi-units, noise clusters, or poorly isolated units were discarded from further analysis. Well isolated units were classified into putative cell types using the MATLAB package, Cell Explorer (*62*) (https://github.com/petersenpeter/CellExplorer). Spiking features such as auto-correlogram (ACG), spike waveform, and putative monosynaptic connections derived from short-term cross-correlograms (CCGs), were used to characterize and classify well-isolated units (*17*, *50*). Three cell types were assigned: putative pyramidal cells, narrow waveform interneurons, and wide waveform interneurons. The two key metrics used for this separation were burst index and trough-to-peak latency. Burst index was determined by calculating the average number of spikes in the 3-5 ms bins of the spike ACG divided by the average number of spikes in the 200-300 ms bins. To calculate the trough-to-peak latency, the average waveforms were taken from the recording site with the maximum amplitude for the averaged waveforms of a given unit. Only putative pyramidal cells were used for further analysis, unless otherwise specified.

Spatial fields and linearization

Spatial fields were calculated only during running periods (> 5 cm/s) at positions with sufficient occupancy (> 20 ms). 2D occupancy-normalized rate maps were constructed using spike counts and occupancies with 50 x 50 spatial bins, smoothed with a 2D Gaussian kernel (SD = 2).

Population vector correlation

To measure remapping in population activity (fig. S1), a population vector (PV) was constructed as the firing rate vector of all spatial-tuned cells in a certain spatial bin. The PV correlation was then defined as the Pearson correlation (r) between the PVs across all bins in the mazes between two sessions.

Theta phase estimation

To extract the theta phase, one channel around the CA1 *stratum oriens* was chosen and band-pass filtered in the range of 5-15 Hz. The theta phase was then computed using the Hilbert transform of the filtered LFP. Since the theta phase can vary depending on the extract position

of the channel within the hippocampus, theta phase was then re-aligned such that the distance between the decoded position (see *Bayesian decoding*) and the animal position was maximal at zero phase (Fig. 2H; (*14*, *23*).

SWR and HSE detection

For SWR detection, the wide-band signal from a CA1 pyramidal layer channel was filtered (difference of Gaussians; zero-lag, linear-phase finite impulse response (FIR) filter), and instantaneous power was calculated by clipping at 4 s.d., rectified and low-pass filtered (*13*, *50*, *51*, *59*). The low-pass filter cut-off was at a frequency corresponding to π cycles of the mean band-pass (for 80–250 Hz band pass, the low-pass was 55 Hz). Subsequently, the power of the non-clipped signal was computed, and all events exceeding 4 s.d. from the mean were detected. The events were then expanded until the non-clipped power fell below 1 s.d. Sharp waves were detected separately using LFP from a CA1 stratum radiatum channel, filtered with band-pass filter boundaries (5–40 Hz). LFP events of a minimum duration of 20 ms and a maximum of 400 ms exceeding 2.5 s.d. of the background signal were included as SWRs.

To refine the detection of SWR start and end for replay analyses (see *decoding theta and replay sweeps with Bayesian reconstruction*), the high synchrony events (HSEs) during immobility period (less than 5 cm/s) were first detected as previously reported(*13*, *51*). To detect HSEs, the combined activity of all pyramidal spikes was binned into 1-ms bins and smoothed with a Gaussian kernel (s.d. = 15 ms) to produce a population firing rate curve over time. HSEs were initially detected as contiguous periods when the population firing rate stayed above 3 s.d. of the mean, and further refined as times around the initially detected events during which the firing rate exceeded the mean. Only HSEs with a duration of 50 ms or more were included for further replay analysis.

Classification of goal-direction cells

To detect the goal-directional cells, we first compute the goal direction (γ; (*31*), defined as the angle between the azimuth heading direction of the rat (φ) and the direction of the vector (α) connecting the head position (x, y) with the goal position (x_G, y_G) as follow:

$$\alpha = angle\big((x_G - x) + (y_G - y) * i\big),$$
$$\gamma = modulo\big((\alpha - \varphi) + \pi, 2\pi\big) - \pi.$$

The goal-direction tuning curves of each cell were then computed by counting the number of spikes in each azimuthal bin (51 bins) and dividing it by the total time spent in that bin during all running periods in the maze during learning trials and smoothed with a Gaussian kernel (SD = 1.5) in a circular manner. The goal modulation was then measured as the Rayleigh vector length based on the goal-direction tuning curves, and the significance of goal-direction tuning was assessed by circularly time-shifting all spikes by a random interval, wrapping the end of the session to the beginning (*31*).

To further rule out potential confounds from place (p) and egocentric heading (φ) signals in the identification of goal-direction cells we used the LN model (*30*) to estimate the spike rate of a neuron at time bin t (r_t) as an exponential function of the sum of all variables (p_t, φ_t, α_t) projected onto a corresponding set of parameters (β_p, β_φ, β_α) as:

$$\boldsymbol{r} = exp(\sum_i X_i^N \beta_i)/\Delta t,$$

where $\boldsymbol{r}$ denotes a vector of firing rates of a neuron across N time bins, i indexes the variables ($i \in [p, \varphi, \alpha]$), and X_i^N is a state matrix for the i-th variable. Each column of X_i^N is a state vector x_i at each of the N time bins, whose elements are 0, except for the element (= 1) corresponding to the animal's current state of that variable.

To learn the parameters (β_p, β_φ, β_α), we maximized the Poisson log-likelihood of the observed spike train given the model-estimated spike number $r\Delta t$. The parameters were optimized using *fminunc* function from MATLAB, and the log-likelihoods were calculated from the held-out data via 10-fold cross-validation. To select the best LN model that characterizes the single-neuron response, we compared both the log-likelihoods and variances explained of the reduced (p, φ) and the full model (p, φ, α). We considered a cell as goal-directional if the full model had a significantly higher log-likelihood and explained variance than the reduced model. In all cases, significance was quantified through a one-sided rank sum test with $p < 0.05$ (*30*).

Decoding theta and replay sweeps with Bayesian reconstruction

We used a previously established method for detecting left-right lateral theta sweeps with Bayesian decoder to detect theta sweeps in our dataset (*14*). In brief, the spike times of CA1 pyramidal cells were binned into 10-ms bins, and the position was decoded in each time bin from the matrix of firing rate and tuning curves from N neurons, with an assumption of Poisson firing and a uniform prior:

$$P(\boldsymbol{x}|y) \propto exp\left(\sum_{i=1}^{N} y_i log\left(f_i(\boldsymbol{x})\right) - dt \sum_{i=1}^{N} f_i(\boldsymbol{x})\right),$$

where $P(\boldsymbol{x}|\boldsymbol{y})$ is the conditional probability for the rat's 2D location $\boldsymbol{x}$, given the observed spike count $\boldsymbol{y}$ and tuning curve $f(\boldsymbol{x})$. The decoded position was taken as the position bin that maximized $P(\boldsymbol{x}|\boldsymbol{y})$.

We then extracted individual sweeps from the decoded position trajectory by a simple sequence detection algorithm, as previously reported (*14*). Within each theta cycle, candidate sweeps were identified as the longest sequence of consecutive valid 10 ms time bins in which the decoded position shifted by less than 20 cm and the direction changed by less than 90° between adjacent bins. These candidate sweeps were then truncated to maximize the net Euclidean distance from start to end. A sweep vector $\boldsymbol{s}$ was defined as the vector connecting the low-pass filtered decoded position at the start of the theta cycle to the most distal point of the candidate sweep. To assess whether the decoded trajectory followed a coherent and straight path, we computed the goodness-of-fit (r^2) between the sweep vector axis and the series of [x, y] coordinates comprising the sweep, as previously described (*14*). Candidate sweeps were retained for further analysis if they consisted of at least four samples and exhibited a goodness-of-fit ($\boldsymbol{r^2}$) greater than 0.5. Sweep direction and length were defined as the angle and magnitude, respectively, of the corresponding sweep vector. Sweep prevalence was calculated as the proportion of theta cycles, during which the animal's running speed exceeded 15 cm/s, that contained sweeps meeting these criteria. Replay sweeps were extracted using the same method as for theta sweeps, but during periods of awake immobility (running speed < 5 cm/s) and HSEs coinciding with SWRs (see *SWR and HSE detection*).

The LMT model

Using the spatial firing rate maps, we can only decode position in one environment based on the tuning curves referenced to the physical space, making it impossible to detect context-specific representations and theta sweeps. To simultaneously extract the neural representations for different behavioral contexts in the same maze and characterize the neuronal tunings to latent task variables, we therefore adapted the latent manifold tuning (LMT) method (*14*, *44*). In this framework, all the simultaneously recorded CA1 pyramidal cells that pass the cell inclusion criteria (see spike sorting and single unit classification) were included in LMT analyses. For each behavioral session, spike counts were taken in 500-ms bins and only during running periods (> 5 cm/s). Bins with spikes from less than 5 cells were discarded. The LMT model assumes that neurons fire according to a Poisson process, and the observed Poisson spikes are governed by temporally smoothed latent variables and a set of nonlinear tuning curves, which are smoothly tuned to the latent variables. This model can simultaneously extract the latent dynamics and tuning curves in an unsupervised manner. The nonlinear tuning curves thus provide the mapping between the low-dimensional latent structure (i.e., neural manifold) and high-dimensional spiking responses. The dimensionality of the latent structure (i.e., number of latent variables) estimated by calculating the predictive log likelihood via 10-fold cross-validation (fig. S5).

To detect theta sweeps in the latent maps (Fig. 5I), we applied the same Bayesian decoding algorithm as in *Decoding theta and replay sweeps with Bayesian reconstruction*, but we replaced the spatial tuning curve $f(\boldsymbol{x})$ with the latent tuning curve $g(\boldsymbol{l})$, $\boldsymbol{l}$ is a set of latent variables estimated from the LMT model.

Monosynaptic connection analyses

Cross-correlograms (CCGs) between pairs of PYR and INT neurons were constructed (Fig. 3, F and G), and only connections from PYR to INT were considered. For visualization, CCGs were rate normalized (i.e., counts/number of reference spikes/Δt) in units of spikes per sec (Fig. 3F). Further, we identified putative monosynaptic connections of PYR-INT cell pairs by short-lag (2-3 ms) peaks in their CCGs as previously described (*34*, *35*). The peak in the CCG needed to exceed that from the slowly co-modulated baseline, and the peak in the causal direction (positive lags) needed to be significantly larger than the largest peak in the anti-causal direction (negative lags) for further analysis (*34*). To assess the effective strength of PYR-INT synaptic coupling, we calculated spike transmission probability of the PYR-INT cell pairs that exhibited monosynaptic connectivity. To calculate spike transmission probability, the lower frequency baseline λ_{slow} was estimated as previously described (*34*, *35*), and the spike transmission probability after N presynaptic spikes using raw CCG normalized by the number of reference spikes was defined as:

$$spike\ transmission\ probability = \frac{\sum_{\tau=0.8ms}^{2.8ms}\left(CCG(\tau)_{observed} - \lambda_{slow}(\tau)\right)}{N}$$

The spike transmission probabilities were further calculated and compared using the CCGs constructed from spikes occurring during goal-directed and lateral theta sweeps (Fig. 3G).

Simulation of 2D theta sweeps

We adapted a single-bump continuous attractor network (CAN) model, developed in a previous study (*23*) for modeling left-right alternating theta sweeps (Fig. 3, A-D). In brief (for a more detailed description, see (*23*), the instantaneous firing rate for a given cell in the network was $f(\boldsymbol{I}(t))$, where f is a rectified function, and $\boldsymbol{I}(t)$ was the input to the cell. Cells fired spikes according to a Poisson process based on $f(\boldsymbol{I}(t))$. The inputs to the cells were a combination of internal ($\boldsymbol{I}^{int}(t)$) and externally driven place input ($\boldsymbol{I}^{P}(t)$), each modulated at theta frequency.
The internal input was given by four components as $\boldsymbol{I}^{int}(t) = \boldsymbol{I}^{rec}(t) + \boldsymbol{I}^{adapt}(t) + \boldsymbol{I}^{goal}(t) + \boldsymbol{I}^{0}$, where $\boldsymbol{I}^{rec}(t)$ is the recurrent input derived from other cells, $\boldsymbol{I}^{adapt}(t)$ is the adaptive inhibitory input, $\boldsymbol{I}^{goal}(t)$ is the goal-direction signal, and $\boldsymbol{I}^{0}$ is the small positive constant bias common to all cells. The place input $\boldsymbol{I}^{P}(t)$ into the cell was given by a 2D Gaussian centered on the rat's position $x_{rat}(t)$. Both internal and place inputs were modulated at theta frequency. The $\boldsymbol{I}^{int}(t)$ was active during the late half of theta cycles, and the $\boldsymbol{I}^{P}(t)$ was active during the center 20% of theta cycles, which allows the place inputs to occur briefly around the time of the release of the internally driven dynamics.

Statistical Analyses

Data analysis was performed using custom routines in Python, MATLAB (MathWorks) and GraphPad Prism 10 (GraphPad Software). No specific analysis was used to estimate minimal population sample or group size, but the number of animals, sessions and recorded cells were larger or similar to those employed in previous related work (*10*, *14*, *17*, *20–22*). Unless otherwise noted, the non-parametric Wilcoxon rank-sum or Wilcoxon signed-rank test was used for unpaired and paired data comparisons, respectively, and ANOVA and Friedman test were used for multiple comparisons. All statistical tests were two-tailed with $p < 0.05$ as the cutoff for statistical significance, which is indicated by asterisks (*p<0.05, **p<0.01, ***p<0.001, and ****p<0.0001). Error bars show SEMs, and boxplots show median (central mark), 75th (box), and 90th (whiskers) percentile, unless indicated otherwise. For boxplots, points below the 10th percentile and above the 90th percentile are defined as outliers; outliers are not displayed in the plots but included in statistical analysis.