

Goal-specific hippocampal inhibition gates learning

<https://doi.org/10.1038/s41586-025-08868-5>

Received: 29 January 2024

Accepted: 25 February 2025

Published online: 9 April 2025

Open access

 Check for updates

Nuri Jeong^{1,2,4}, Xiao Zheng^{1,3,4}, Abigail L. Paulson¹, Stephanie M. Prince^{1,2}, Victor P. Nguyen¹, Sherina R. Thomas¹, Caroline E. Gilpin¹, Matthew C. Goodson¹ & Annabelle C. Singer¹✉

Goal-directed navigation in a new environment requires quickly identifying and exploiting important locations. Identifying new goal locations depends on neural computations that rapidly represent locations and connect location information to key outcomes such as food¹. However, the mechanisms to trigger these computations at behaviourally relevant locations are not well understood. Here we show that parvalbumin (PV)-positive interneurons in the mouse hippocampal CA3 have a causal role in identifying and exploiting new food locations such that decreases in inhibitory activity around goals enable reactivation to bind goal locations to food outcomes. PV interneurons in the CA3 substantially reduce firing on approach to and at goal locations while food-deprived mice learn to find food. These inhibitory decreases anticipate reward locations as the mice learn and are more prominent on correct trials. Sparse optogenetic stimulation to prevent goal-related decreases in PV interneuron firing impaired learning of goal locations. Disrupting goal-related decreases in PV interneuron activity impaired the reactivation of new goal locations after receipt of food, a process that associates previous locations to food outcomes such that the mice know where to seek food later. These results reveal that goal-selective and goal-predictive decreases in inhibitory activity enable learning, representations and outcome associations of crucial locations.

An animal's survival depends on accurately identifying and navigating to new food locations. Rapid spatial learning depends on the CA3 hippocampal region in which recurrently connected pyramidal cells form rapid associations, including associations between locations and other information, such as outcomes^{1–4}. In new environments, these pyramidal cells quickly form location-specific firing fields that tile the space to represent many positions with more pyramidal cells firing at goal locations^{5–9}. To associate location with crucial outcomes such as food, pyramidal firing patterns representing paths to and from food locations are reactivated when an animal pauses to eat^{9–12}. This reactivation, which is essential for spatial learning and memory, drives cells to fire together on short timescales and promotes plasticity between recurrently connected pyramidal cells^{13,14}. Thus, pyramidal cells over-represent and reactivate crucial locations, yet the mechanisms by which these changes occur selectively at important locations are unclear. Signals indicating that a location is crucial to the animal could act directly on excitatory cells or act through inhibitory activity that powerfully affects excitatory cells (Fig. 1a). We hypothesized that hippocampal interneurons have a key role in identifying and remembering goal locations through selective decreases in inhibitory activity around food locations. Decreases in inhibition after a reward would be well suited for enhancing representations and reactivation of reward locations, and decreases in inhibition that predict a reward would facilitate learning about locations that lead to reward.

Hippocampal inhibitory interneurons both respond to and control coordinated excitatory cell firing^{15–17}. Previous studies have shown that inhibitory cells have location-selective decreases in firing that tile the environment like reverse place fields, and these decreases can regulate place-field formation and refinement^{18–23}. We therefore hypothesized that preceding and at reward locations, coordinated decreases in inhibitory activity across multiple cells enable the identification and learning of goal locations. Previous research has shown changes in interneuron firing in the CA1 after animals received a reward, but these changes included a mix of both increases and decreases in interneuron firing^{21,24}. It is not known whether coordinated spatially specific decreases in inhibitory activity occur preferentially at important locations, such as food sources, nor whether decreases in inhibitory activity precede reward locations. Reward-selective coordinated decreases in interneuron firing preceding reward locations could facilitate learning not only locations in which reward was found but also locations that lead to reward, a crucial process to finding reward. The CA3 is an essential nexus of rapid spatial learning in a new environment and is required for the reactivation of previous patterns of activity^{2–4,12,25,26}. Accordingly, we determined whether CA3 interneurons decrease firing preferentially at reward locations during learning and whether such decreases enable learning and learning-associated plasticity events, like reactivation of goal locations after reward.

¹Coulter Department of Biomedical Engineering, Emory University and Georgia Institute of Technology, Atlanta, GA, USA. ²Neuroscience Graduate Program, Graduate Division of Biological and Biomedical Sciences, Emory University, Atlanta, GA, USA. ³Bioinformatics Graduate Program, Georgia Institute of Technology, Atlanta, GA, USA. ⁴These authors contributed equally: Nuri Jeong, Xiao Zheng. ✉e-mail: asinger@gatech.edu

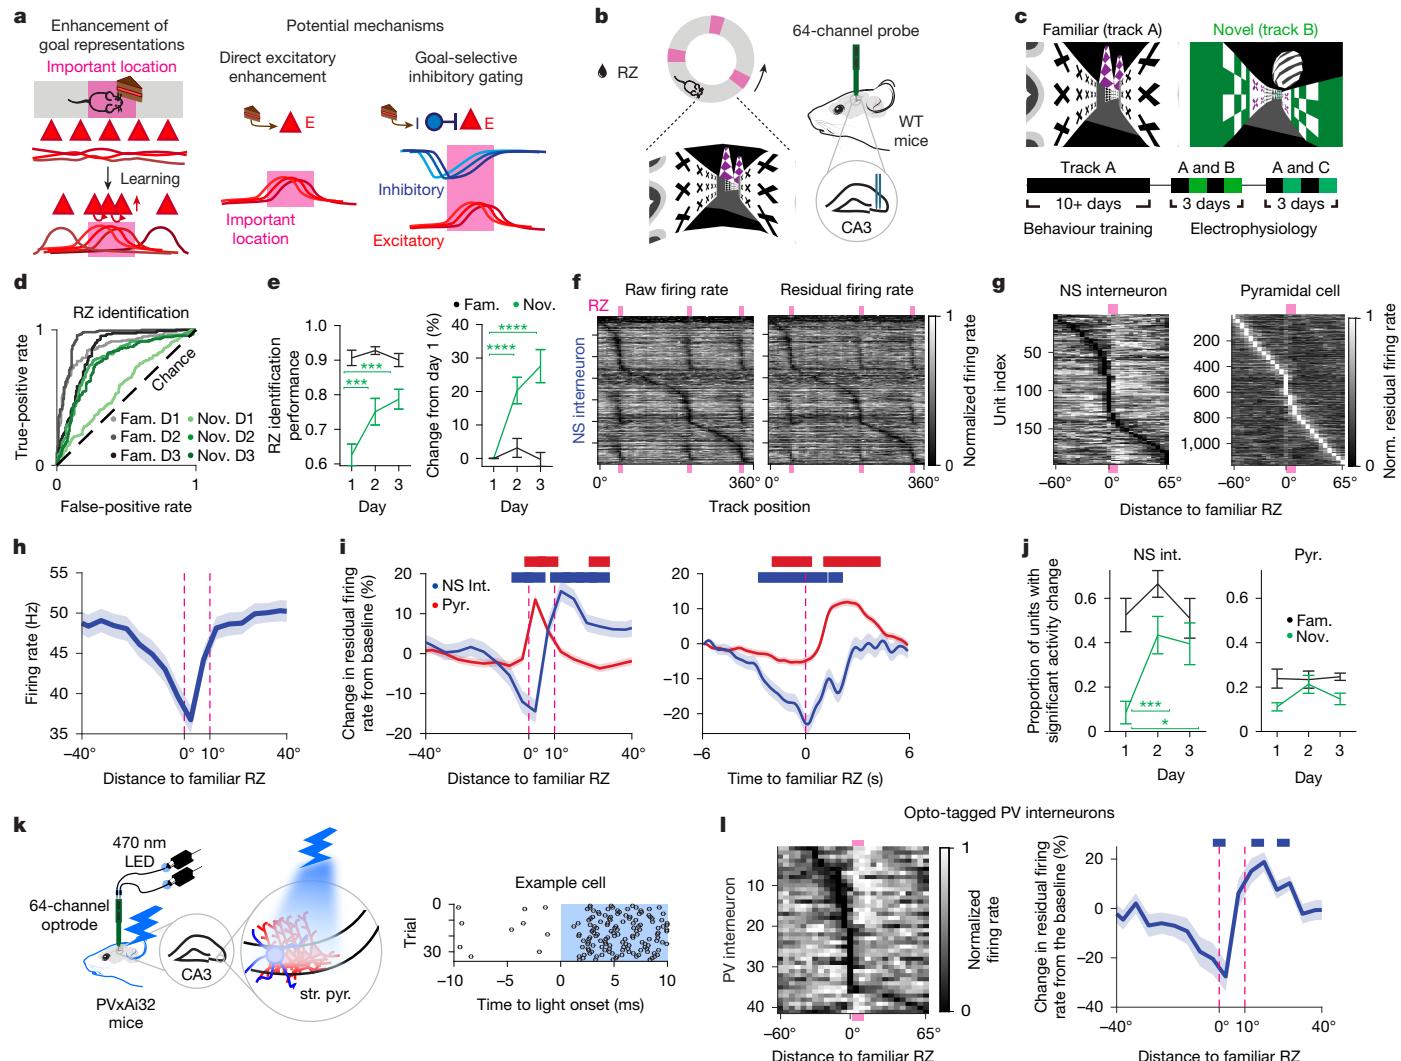


Fig. 1 | Interneuron firing decreases as mice approach learned goals.

a, Potential mechanisms of goal learning. **b**, The task with electrophysiology recording. **c**, The timeline (track C is not shown). **d**, Example receiver operating characteristic (ROC) curves indicating a decrease in mouse movement speed in the zone prior to the RZ of familiar (fam.; black) and novel (nov.; green) tracks. **e**, Performance based on the raw AUC (left, day 1 versus 2, $P = 0.001$; day 1 versus 3, $P = 0.001$) and the percentage change in AUC (right, day 1 versus 2, $P < 0.0001$; day 1 versus 3, $P < 0.0001$; linear mixed-effects model (LMM) with Tukey correction). $n = 14$ (familiar) and 13 (novel) sessions from 7 mice. **f**, Raw (left) and residual (right) NS interneuron normalized firing rates according to the position in the familiar track with RZs (pink). **g**, Normalized firing rates as in **f** according to the distance to the RZ for each cell type. **h**, The raw firing rates averaged across NS interneurons. **i**, Normalized residual firing over distance (left) or time (right) to the RZ for NS interneurons (NS Int., blue, $n = 196$ cells) and pyramidal cells (Pyr., red, $n = 1,157$ cells). Significant effect of position

(NS interneurons, $P < 2.2 \times 10^{-16}$; pyramidal cells, $P < 2.2 \times 10^{-16}$) and time (NS interneurons, $P < 2.2 \times 10^{-16}$; pyramidal cells, $P < 2.2 \times 10^{-16}$; LMM). **j**, The proportion of units with significantly decreased (NS interneurons, left, $n = 12$, 14 and 14 sessions, days 1–3) or increased (pyramidal cells, right, $n = 13$, 14 and 14 sessions, days 1–3) firing (novel day 1 versus 2, $P = 0.003$; day 1 versus 3, $P = 0.02$; LMM with Tukey correction). **k**, Example spikes of an opto-tagged unit. str. pyr., stratum pyramidale. **l**, The residual firing (left) and the cell-averaged change in firing (right) for PV cells ($n = 41$ cells). Significant effect of position ($P < 2.2 \times 10^{-16}$, LMM). For **e–h** and **j–l**, data are mean \pm s.e.m. For **i**, the coloured bars indicate significantly different bins (one-sided *t*-tests with Bonferroni correction). The mouse illustration was adapted from scidraw.io. * $P < 0.05$, ** $P < 0.005$, *** $P < 0.001$. The brain (F. Cauda) in the mouse head (E. Tyler and L. Kravitz) diagrams in **b**, **k** were adapted under a CC BY 4.0 licence; see the Methods for further details.

Rapid learning of new goal locations

To investigate hippocampal activity during new spatial learning, we developed a virtual reality (VR) behavioural paradigm in which we could quantify progressive learning of goal locations in a controlled environment. In an annular track with 36 equally sized zones with unique wall patterns, mice had to lick in three reward zones (RZs) to receive a reward (Fig. 1b and Extended Data Fig. 1a). We used timeouts to discourage mice from licking indiscriminately and placed three RZs unevenly around the track to prevent mice from using a distance-based strategy. After 7–14 days of training in a familiar track (track A), mice reduced their movement speed by an average of 29% and increased their lick rate by

an average of 21% when they approached the zone immediately preceding the RZ, the anticipatory zone (AZ), indicating that they had learned the RZ locations (Extended Data Fig. 1a). Successful task performance required visual cues; mice did not show a significant change in movement speed or licking around well-learned goal areas if visual cues were removed (Extended Data Fig. 1b).

On the first day of recording, mice alternated between sessions navigating on the familiar track and a novel track (track B or C) where they learned three new RZs through trial and error over 3 days (Fig. 1c and Extended Data Fig. 1a). Consistent with knowing the reward locations, the mouse movement speed in the familiar AZs was lower than in non-reward zones (NRZs)—control zones 30° after each RZ within

the same environment (Fig. 1d and Extended Data Fig. 1c). On the first day of novel-track exposure, mice showed no significant difference in movement speed between the novel AZ and NRZ, as is expected when the mouse does not know the reward locations (Fig. 1d and Extended Data Fig. 1d). By day 2 or 3, mice slowed before arriving at the RZs, indicating that they had learned to differentiate reward areas (Fig. 1d and Extended Data Fig. 1d). The performance increased by approximately 20% on day 2 and 30% on day 3 compared with day 1 in the novel environment on average (Fig. 1e and Extended Data Fig. 1e).

Inhibitory decreases at goal locations

We hypothesized that CA3 interneurons have a role in selectively gating excitatory activity around goals. In such a case, reward information would lead to coordinated decreases in inhibitory activity across many interneurons as animals approach and enter important locations. Decreases in inhibitory firing would precede increases in excitatory activity over seconds and promote representations and reactivation of important locations over days (Fig. 1a). To test this hypothesis, we first examined whether CA3 interneurons show a spatially selective firing pattern around learned goal locations. We recorded 3,489 classified single units in the CA3 during active spatial navigation in head-fixed mice and identified putative pyramidal cells and a fast-spiking group of interneurons with narrow spike widths (NS interneurons; Supplementary Fig. 1 and Supplementary Tables 1 and 2). Using opto-tagging, we confirmed PV interneurons overlap with the classification of NS interneurons recorded (Supplementary Fig. 2).

We found a marked goal-specific decrease in the firing activity of the majority of NS interneurons in the CA3 around three different RZs along the familiar track (Fig. 1f–h). Raw firing rates of NS interneurons decreased by 14.01 Hz on average (Fig. 1h and Extended Data Fig. 2a,e). To account for changes due to speed, we regressed out the estimated contributions of movement speed and lick rate to the raw firing rates (Methods and Extended Data Fig. 2). Regressing out the effects of speed, we observed an approximately 15–20% decrease in firing on average during the approach to familiar RZs, starting several seconds before the RZ (Fig. 1i). The observed patterns in the raw firing data were preserved in residuals, indicating that these results are not explained by position-dependent changes in movement speed or lick rate (Fig. 1f,h,i). At the population level, the decrease in inhibition preceded both RZ entry and increases in excitatory activity (Fig. 1i). We found a significant firing reduction in 40–70% of all recorded NS interneurons in the familiar environment (Fig. 1j). As expected, individual pyramidal cells tended to fire at one or a few locations that collectively tiled environment, with approximately 15–25% of units with increased activity around the RZ (Fig. 1j). Thus, a majority of NS interneurons showed goal-specific reductions in firing. The reduction in NS interneuron activity around goals developed rapidly over days in a novel environment as animals learned where to find new goal locations (Fig. 1j).

Multiple control analyses and experiments indicate that NS interneurons reduce firing rates around goals beyond what is expected by behavioural changes. First, the amount of variance explained by movement speed and lick rate was less than 25% for more than 99% of units (Extended Data Fig. 2b–d,f,g). Thus, speed modulation of neuronal firing is present, but our analysis examines changes in firing that are not explained by speed. Second, we found significant firing rate differences in individual interneurons between the AZ and RZ and baseline no-VR periods even when animals moved at similar speeds in 7 out of 8 comparisons made in both environments (Extended Data Fig. 3a). Third, we found no significant difference in the firing rates of the same neurons at the familiar AZ when animals moved at different speeds (Extended Data Fig. 3b). Thus, the firing rate of NS interneurons depends on the task context and on the animal's proximity to the RZ (Extended Data Fig. 3a,b). Finally, we introduced a new trial type in the familiar environment in a subset of animals ($n = 4$). To dissociate position-related

visual cues in VR from the animal's movement on the treadmill, on RZ entry, the visual cues on the screen froze for 3–10 s. This screen freeze occurred even if the animal was moving on the treadmill, when the visual cues would normally progress forward through the track (Extended Data Fig. 3c). After RZ entry, when the movement speed and lick rates were similar but the visual cues dissimilar, the firing rates of NS interneurons were lower on average when reward-associated visual cues stayed fixed during screen-freeze trials than during normal trials. These results show that inhibitory firing is reduced for goal-related cues and not purely for speed changes (Extended Data Fig. 3c). Furthermore, these goal-selective inhibitory decreases were not solely due to local sharp wave ripple (SWR) activity, as significant decreases in NS interneuron activity around the RZ occurred during trials with and without SWRs (Extended Data Fig. 4c,d).

We found significant decreases in activity around goals in multiple putative subtypes of NS interneurons, including basket cells, axo-axonic cells (AAC), cholecystokinin (CCK) cells and bistratified cells, identified on the basis of these subtypes' different patterns of firing relative to theta phase and SWRs^{27–30} (Extended Data Fig. 4). As all of these subtypes express PV, including a subset of CCK interneurons³¹, we further investigated PV interneurons optogenetically (Supplementary Fig. 2). Opto-tagged PV interneurons had goal-selective decreases in firing on approach to and in RZs, with firing decreasing by 30% on average (Fig. 1k,l). Previous studies have found direct inhibitory connections from PV interneurons to pyramidal cells in the CA3, but PV cells also synapse onto other interneurons³². We therefore used optogenetic stimulation of PV cells to determine the causal effect of PV firing on pyramidal cell activity. We found that optogenetically stimulating PV cells resulted in decreased pyramidal firing, showing that CA3 PV cells inhibit the activity of CA3 pyramidal cells (Extended Data Fig. 5).

Our data are consistent with goal-selective inhibitory gating in which inhibitory activity is selectively decreased around reward locations and occurs before increases in firing of pyramidal cells during the approach to and traversal of the RZ. Furthermore, the magnitude of NS firing reduction approaching and just after entry to the RZ was higher on correct trials than on incorrect trials (Extended Data Fig. 6a,b). These results show that a goal-specific reduction in inhibitory activity precedes and coincides with successful identification of a goal location.

New goal learning requires PV decreases

We identified that the goal-associated reduction in inhibitory activity is necessary for learning new goal locations. We used optogenetic stimulation to disrupt the normal inhibitory reduction around goals during active spatial navigation. We focused our stimulation on PV interneurons because we observed a robust pre-RZ activity reduction in the majority of NS interneurons and specifically in opto-tagged PV interneurons (Fig. 1i,k,l). Furthermore, PV interneurons are strong inhibitors of pyramidal cells at the soma^{33,34}. For PV-specific targeting *in vivo*, we crossed the PV-Cre mouse line with Ai32 mice to achieve Cre-dependent expression of the blue-light-sensitive opsin channelrhodopsin-2 (ChR2) in PV interneurons. Our light delivery approach stimulated PV cells in a small volume of CA3 (ref. 35) (Methods and Supplementary Fig. 2a).

We compared goal stimulation, delivered in the AZ and RZ, to sham stimulation, delivered in a control NRZ of the same size, with each delivered in separate novel environments (Fig. 2a). As different genotypes of mice behave differently, and to control for non-specific effects of expressing opsins and stimulating PV cells, we compared PV stimulation around goals to sham stimulation away from goals within PVxAi32 mice in two different environments. In one novel environment, animals underwent goal stimulation in which PV activity was stimulated specifically in the AZ (one zone before the RZ) and RZ. In the other novel environment, the same animals underwent sham stimulation in which PV activity was stimulated in two consecutive control zones, away from the RZ.

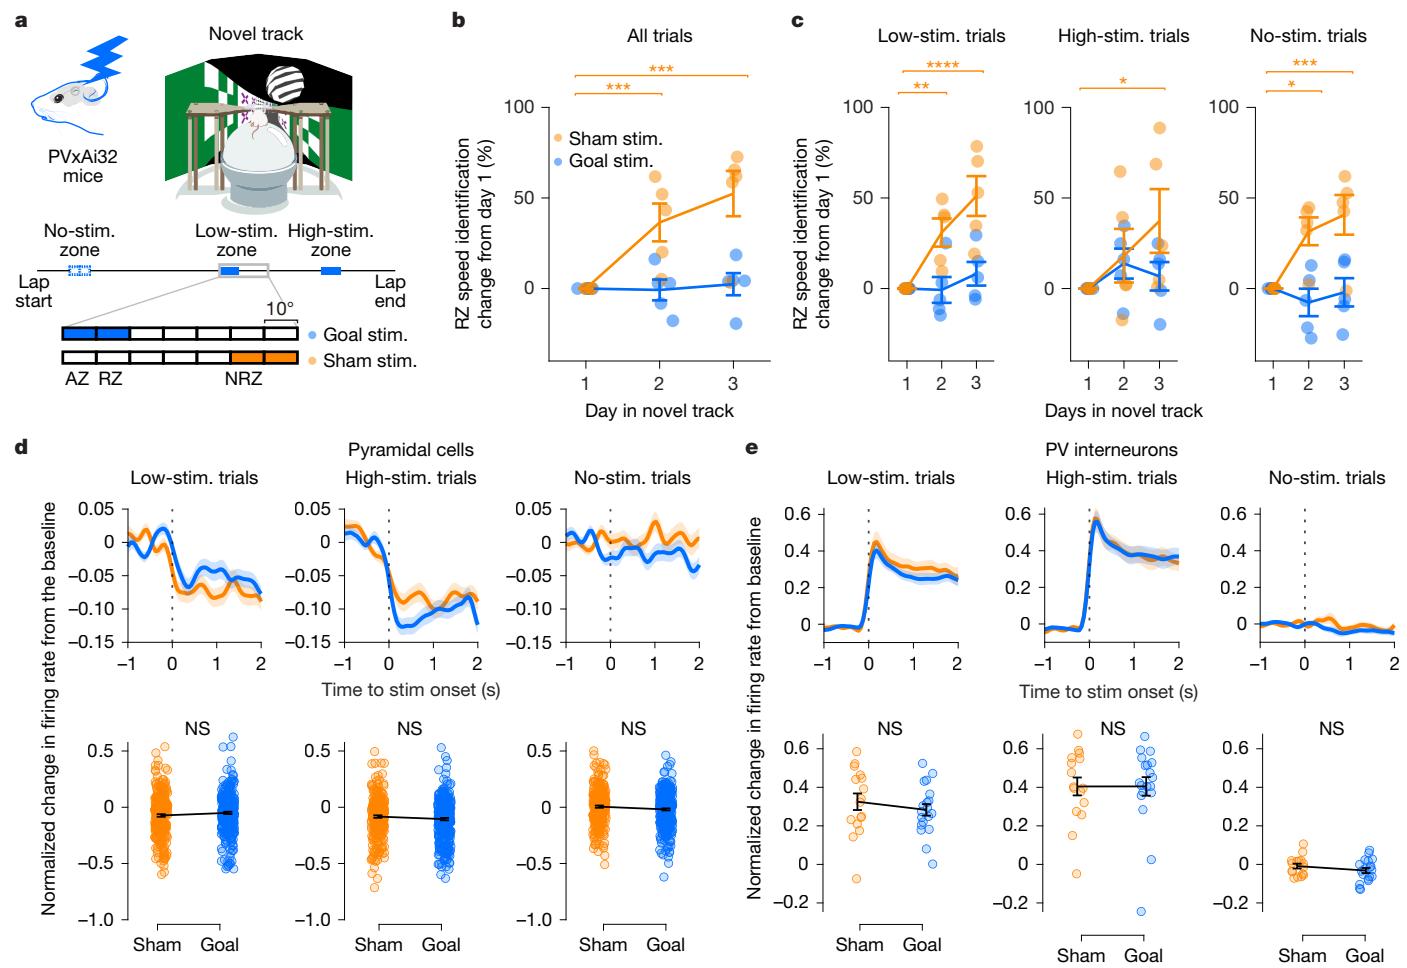


Fig. 2 | Goal-associated inhibitory reduction is required for new goal learning. **a**, Optogenetic PV stimulation protocol with stimulation locations and intensities for goal and sham stimulation (stim.). The mouse illustration was adapted from scidraw.io. **b**, Disrupting normal goal-associated inhibitory firing reduction (blue) impaired the learning of new goal locations. Performance based on speed differentiation over 3 days of learning for all trial types (goal stimulation, blue, $n = 5$ mice; sham stimulation, orange, $n = 5$ mice; day 1 versus day 2, $P = 0.0045$; day 1 versus day 3, $P = 0.0001$; LMM followed by Tukey correction). **c**, The RZ identification speed as in **b**, but for trials with low (left), high (middle) or no (right) stimulation. There was a significant performance increase from day 1 with sham, but not goal, stimulation at all intensities (low stimulation, day 1 versus day 2, $P = 0.006$; day 1 versus day 3, $P < 0.0001$; high

stimulation, day 1 versus day 3, $P = 0.03$; no stimulation, day 1 versus day 2, $P = 0.01$, day 1 versus day 3, $P = 0.002$; LMM with Tukey correction). **d**, The normalized change in pyramidal cell firing from the baseline over time with goal (blue) or sham (orange) stimulation across stimulation intensities (top). The dashed line indicates the stimulation onset. Bottom row, there were no significant differences in the mean firing rates in the 2-s period after the start of goal ($n = 307$ cells) or sham ($n = 270$ cells) stimulation across intensities. **e**, The change in firing rate as in **d**, but for PV interneurons. $n = 19$ (goal stimulation) and 16 (sham stimulation) cells. For **b–e**, data are mean \pm s.e.m. NS, not significant. ** $P < 0.01$. The brain (F. Claudi) in the mouse head (E. Tyler and L. Kravitz) diagrams in **a** were adapted under a CC BY 4.0 licence.

This sham stimulation covered the same length of track as the goal stimulation but was away from the reward locations. This design was used to determine whether the selective decreases in NS interneuron activity around RZs are important for learning and to control for the potential off-target effects of this artificial optogenetic PV stimulation.

Notably, disrupting the reduction in goal-associated firing in a small subset of CA3 PV cells by goal stimulation substantially impaired the learning of new goal locations. Mice that underwent PV goal stimulation did not show a significant speed difference between the AZ and NRZ over 3 days but did have intact learning and improvement over days when they received sham stimulation in a different novel environment (Fig. 2b,c, Extended Data Fig. 1f and Supplementary Figs. 3 and 4). We found significant effects of learning day, stimulation condition and the interaction between day and stimulation condition, confirming performance deficits with the goal-related PV stimulation.

Importantly, goal and sham stimulation conditions did not differ in how they affected firing of PV and pyramidal cells. The average reduction

in normalized firing from the baseline did not differ significantly between the goal and sham stimulation conditions with low or high stimulation for pyramidal cells nor for PV cells (Fig. 2d,e). At light offset, there were no significant differences between goal and sham stimulation (PV cells, $P = 0.83$; pyramidal cells, $P = 0.22$; two-sided Wilcoxon rank-sum test). Goal stimulation increased PV firing relative to the baseline, therefore preventing decreases in PV activity at the goal location; yet it did not result in the complete shut-down of excitatory activity like a temporary lesion (Extended Data Figs. 5f and 6c,d). We used relatively low light intensities of a small subset of PV cells (Methods). Indeed, high stimulation resulted in small decreases in pyramidal firing, from 10 to 15%, with even smaller decreases in response to low stimulation (Extended Data Fig. 6d). Furthermore, SWRs still occurred during PV stimulation.

Impaired learning could not be explained by PV stimulation directly affecting speed or licking behaviour. First, we took advantage of the fact that there were three RZs in each environment and varied our stimulation intensity across RZs. Notably, we found that goal stimulation at

other RZs disrupted inhibitory decreases in the AZ and RZ, where no stimulation was applied (Extended Data Fig. 6c,d). Furthermore, mice did not learn the RZ that was not stimulated (Fig. 2c (no-stimulation trials only)). Second, we found no significant differences in the licking and movement speed between no stimulation and low stimulation in both the familiar and novel environments (Supplementary Fig. 5a,b). Thus, the direct effects of stimulation on speed or licking behaviour cannot explain the deficits in new goal learning even at a low stimulation intensity. Third, there was no significant difference in the overall rate of licking outside the goal or stimulation locations between goal and sham conditions, ruling out the possibility that PV stimulation altered licking behaviour (goal stimulation, 1.3 ± 0.14 licks per s versus sham stimulation, 1.5 ± 0.09 licks per s; $P = 0.47$, two-sided Wilcoxon rank-sum test). Thus, we conclude that disrupting interneuron firing reduction specifically at goal locations impairs new goal learning that is not explained by differences in licking, running speed or task engagement.

At the end of the session, when mice no longer received stimulation at any goal location, mice continued to perform poorly, suggesting that there was no latent learning that emerged even when PV firing reduction was no longer disrupted. Goal-selective NS interneuron responses did not immediately appear after the stimulation ended, although we did observe a trend of a decrease (Supplementary Fig. 6). These findings suggest that stimulation prevents the formation of goal-selective NS interneuron responses and that changes in firing of these neurons over trials are required to form goal-selective responses.

Disrupting goal-associated inhibitory reduction with PV stimulation at well-learned goals in the familiar environment did not compromise behavioural performance at any stimulation intensity (Supplementary Fig. 5c). This finding is consistent with the observations that our goal stimulation did not completely silence pyramidal cell activity or impair the animal's running or licking ability. Taken together, these results show that goal-associated reduction in CA3 PV inhibition is required for learning new goal information but not for retrieving previously stored goal information.

Reduction develops during learning

As decreased inhibitory activity at goals is essential for learning, we hypothesized that goal-selective decreases in inhibition develop early during learning when place cell activity around goals undergoes refinement and stabilization. To capture RZ-related changes in firing, we selected units with a significant increase or decrease in activity around goals and determined when the population showed significant coordinated changes in activity around the RZ. Consistent with animals not having learned the new RZs yet, NS interneurons with a significant change around the RZ did not show a significant coordinated activity decrease around the RZ on day 1 (Fig. 3a,b and Extended Data Fig. 7a). By day 2, these NS interneurons exhibited a clear and persistent reduction in activity that appeared even before arriving at the RZ, and this pattern continued on day 3 (Fig. 3a,b and Extended Data Fig. 7a).

We found that the development of goal-associated inhibitory reduction coincides with the development of goal-representing cells, or pyramidal cells that have at least one receptive field with significant spatial modulation around goals. Previous studies have shown that such goal representation predicts successful goal-directed behaviour and improves with learning^{5,7,24,36,37}. During learning, field stability, measured through rate-map correlations across trials, and spatial information of goal-representing cells increased over days as animals performed better in the task (Fig. 3c,d). Field stability and spatial information of goal-representing cells were significantly higher in the familiar environment than the novel environment (Fig. 3c,d). The proportion of goal-representing cells did not change over days (Fig. 3e). As these responses develop over days, inhibitory activity decreases before excitatory increases on the timescale of seconds on days 2 and 3 (Extended Data Fig. 7).

As we found that putative PV cells have a causal role in goal learning, we examined how goal responses develop over trials in putative PV cells (Methods). On the first day in the first novel environment, putative PV interneuron responses developed over trials from initially no decrease around the reward locations in the first 25 trials to a significant decrease on approach to and in the reward location in the last 25 trials (Fig. 3f,g and Extended Data Fig. 7b-d). This development of inhibitory decreases coincided with improvements in behaviour (Fig. 3h). These findings reveal that decreases in interneuron activity develop from no decrease initially to anticipating the reward locations as animals learn on the first day of exposure to a new environment.

Together, these results show that goal-selective decreases in inhibitory activity develop rapidly during learning, when excitatory cells are undergoing refinement and stabilization of goal representations. Over seconds, NS interneuron activity decreases before pyramidal activity increases. Over days, pyramidal place fields and NS interneuron responses around reward locations develop together.

Stable place codes require PV decreases

We determined that inhibitory decreases around goals are required for stabilization and refinement of excitatory goal representations (Extended Data Fig. 8a-c). Although the stabilization of new goal fields, or field stability, increased over days with sham stimulation, it did not with goal stimulation (Extended Data Fig. 8d). Furthermore, spatial information of goal-representing cells increased over days with sham stimulation but not with goal stimulation (Extended Data Fig. 8e). The proportion of goal-representing cells over days did not change significantly (Extended Data Fig. 8f). Furthermore, goal-representing place cells had lower peak firing rates with goal stimulation than with the sham stimulation, and a higher vector strength as a function of theta, potentially from narrower preferred theta phases (Supplementary Fig. 7). Theta was altered in power but still clearly present during goal stimulation (Supplementary Fig. 7). Thus, disrupting inhibitory decreases around goals impaired the stability, refinement and amplitude of pyramidal goal representations.

As pyramidal cells within the CA3 form sequences, goal representations could influence spatial coding away from goals. We found that goal stimulation disrupted the stabilization of place cells in positions away from goals, outside the stimulated portion of the track (Extended Data Fig. 8g). The amount of spatial information of place fields away from goals also did not increase over days with goal stimulation but did with sham stimulation (Extended Data Fig. 8h). The proportion of non-goal-representing cells did not change significantly (Extended Data Fig. 8i). These results support a goal-selective gating role of inhibition that recruits excitatory neurons to sharpen and stabilize new spatial representations during learning.

Goal reactivation requires PV decreases

When animals pause to consume a reward, pyramidal firing patterns representing paths to and from the reward location are reactivated and this reactivation enhances plasticity and binds spatial locations to key outcomes such as food^{9,10,12}. Reactivation occurs during SWRs, which are required for rapid learning¹³, and these bursts of population activity affect synapses related to new spatial learning³⁸. NS interneurons decreased firing around the AZ and RZ over seconds, then rapidly and transiently increased firing during SWRs on average, consistent with previous work^{39,40} (Extended Data Fig. 9a,b and Supplementary Fig. 8). We hypothesized that this decrease in inhibition is required for intact SWR activity at goal locations in the novel environment. To control for animal position and position-related variability in behaviour, we included only SWRs that occurred while the animal was in the AZ or RZ. Disrupting goal-associated decreases in PV firing resulted in a lower rate of SWRs occurring around goal locations than in the sham stimulation

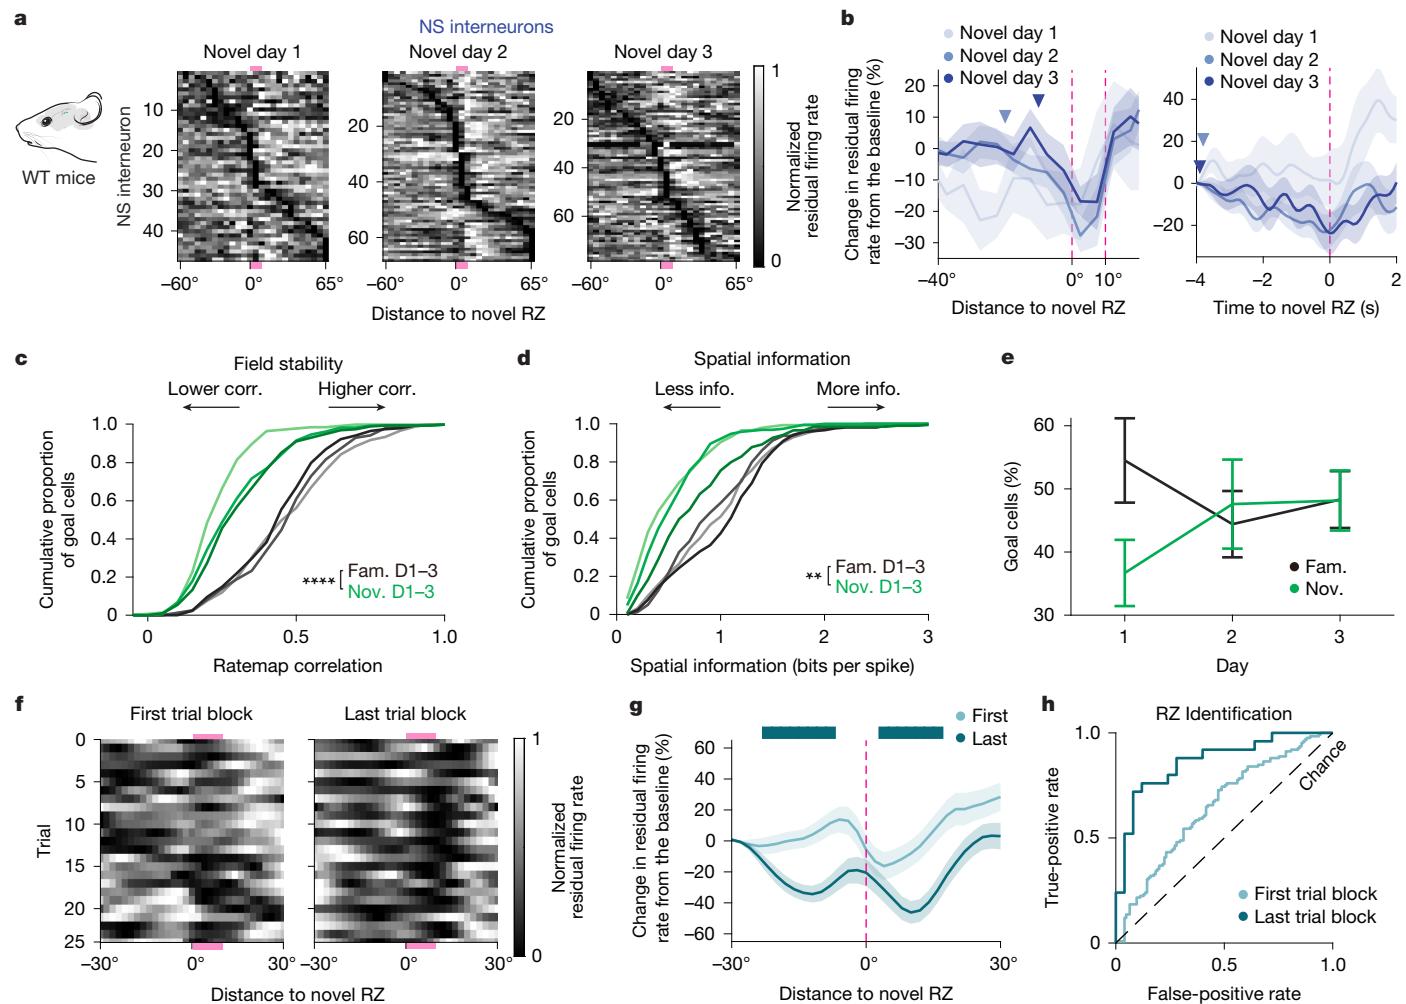


Fig. 3 | Goal-associated decrease in inhibitory activity develops over learning coinciding with refinement and stabilization of excitatory goal representations. **a**, The normalized residual firing rates of NS interneurons according to the distance to the novel RZ (pink) per day. The mouse illustration was adapted from scidraw.io. **b**, The percentage change in firing as a function of the distance (left) or time (right) to the novel RZ (pink), averaged across cells with a significant increase or decrease at goals (Methods; $n = 9, 20$ and 20 cells for days 1–3). The triangles indicate the start of the persistent firing decrease. **c,d**, The field stability (**c**) and spatial information (**d**) of goal cells in familiar (shades of black, $n = 185, 159$ and 209 cells for days 1–3) and novel (shades of green, $n = 146, 191$ and 204 cells for days 1–3) environments. There is a significant interaction of novel day and environment (field stability, $P = 5.4 \times 10^{-6}$; spatial information, $P = 0.009$; LMM). **e**, The proportion of pyramidal cells that are

goal cells over days (familiar, black, $n = 13, 15$ and 14 sessions for days 1–3; novel, green, $n = 13, 14$ and 14 sessions for days 1–3; LMM). **f**, The normalized residual firing rates averaged over putative PV cells in the first (left) and last (right) blocks of 25 trials during the initial novel exposure ($n = 10$ and 3 cells for the first and last blocks). **g**, The change in residual firing from the baseline, trial-averaged for the first (light turquoise) and last (dark turquoise) trial blocks in **f**. There is a significant effect of position (first block, $P < 2.2 \times 10^{-16}$; last block, $P < 2.2 \times 10^{-16}$; LMM). The bars above indicate bins that significantly differ from the baseline (one-sided t -tests with Bonferroni correction; $n = 25$ trials). **h**, Speed-based ROC curves for the first and last trial blocks ($n = 25$ trials per block, 7 mice). For **b,e,g**, data are mean \pm s.e.m. The brain (F. Claudi) in the mouse head (E. Tyler and L. Kravitz) diagrams in **a** were adapted under a CC BY 4.0 licence.

group (Fig. 4a,b and Extended Data Fig. 9e). The coactivation probability among simultaneously recorded pairs of goal-representing pyramidal cells during SWRs was lower in the goal stimulation group than in the sham stimulation group, as was SWR power and duration (Fig. 4c–e and Extended Data Fig. 9f–h). Neither SWR power nor duration in the familiar environment differed significantly with or without goal stimulation in the same PVxAi32 mice (Extended Data Fig. 9c,d). These results show that inhibitory reduction specifically at new goal locations is required for intact SWR activity during learning in novel environments.

Notably, the goal-related inhibitory reduction is essential for preferential reactivation of goal locations during SWRs. We used sequenceless decoding of SWR content to identify the most likely location represented during each SWR. We found significant differences in the information bias of SWRs between goal and sham stimulation.

SWRs in the sham stimulation sessions were much more likely to represent locations around goals than locations far from goals, similar to what we observed with no stimulation (Fig. 4f,g). On the other hand, when PV firing reduction was disrupted in the same animals, there was no significant difference in the proportion of SWRs representing locations near or far from goals, indicating that goal-related locations were not over-represented during SWRs (Fig. 4f,g). These findings show that goal-associated inhibitory reduction is required for preferential reactivation of goal-related information. Consistent with this idea, we observed that SWRs in wild-type (WT) mice were much more likely to code for locations near goals than far from goals in both familiar and novel environments (Fig. 4h). Overall, our results show that goal-associated reduction in PV interneuron firing gates SWRs and goal-informative SWR reactivation that represents locations most pertinent to task performance.

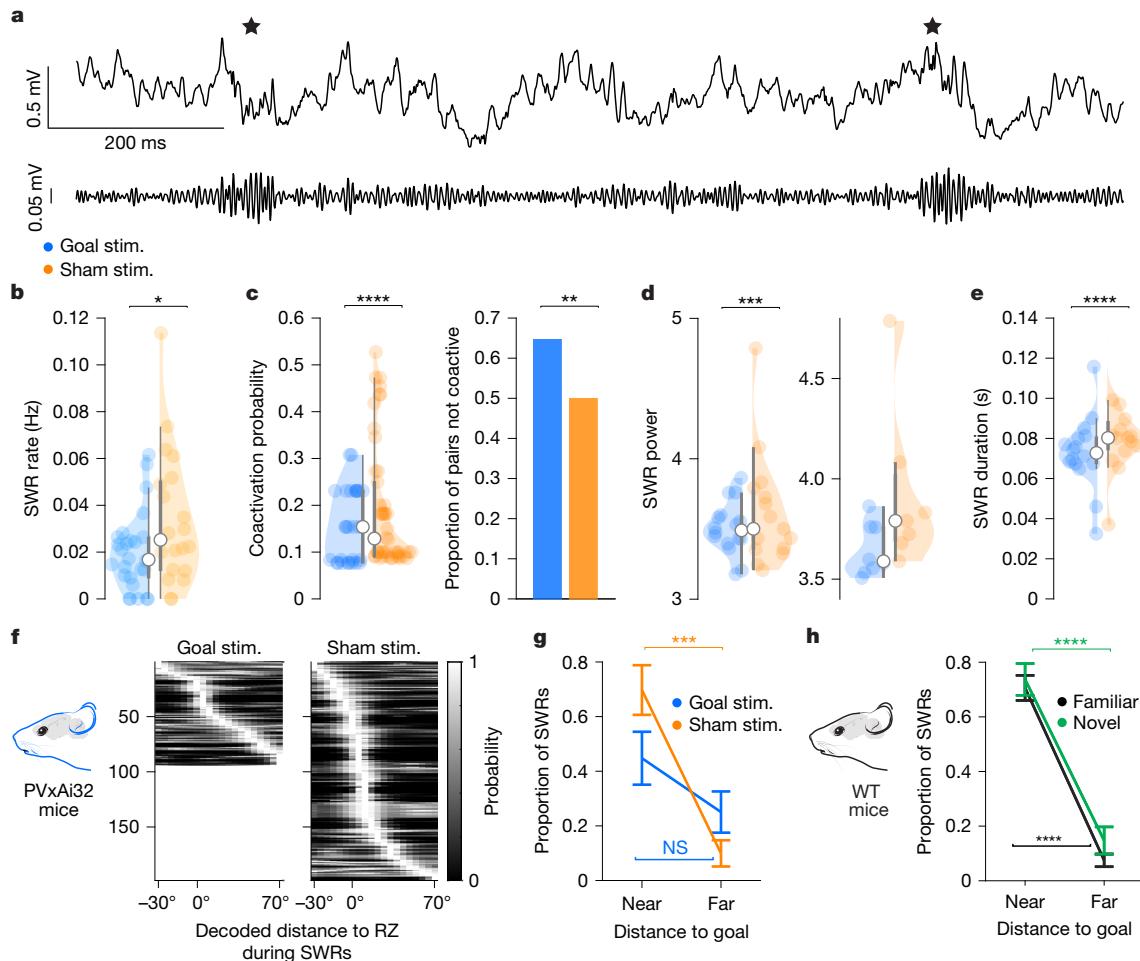


Fig. 4 | Inhibitory reduction is required for goal-related sharp-wave ripple reactivation. **a**, Raw (top) and band-pass-filtered (bottom) local field potential (LFP) traces with SWRs (stars). **b**, The SWR rate in the goal (blue) and sham (orange) stimulation groups in a novel environment (goal, $n = 26$ sessions; sham, $n = 20$ sessions; $P = 0.04$, LMM). **c**, The coactivation probability of goal cell pairs during SWRs with coactivity greater than 0.05 (left) (goal, $n = 27$ pairs; sham, $n = 46$ pairs; $P = 8.1 \times 10^{-12}$, LMM). Right, the fraction of cell pairs with a coactivation probability equal to or less than 0.05 ($P = 0.007$, χ^2 test; cell pairs were subsampled to be equal across stimulation conditions). **d**, SWR power (goal, $n = 71$ SWRs from 16 sessions; sham, $n = 78$ SWRs from 15 sessions; $P = 0.002$, LMM). Right, sessions with high SWR power (top 50% of all sessions) were selected for visualization purposes only. **e**, As in **d**, but for SWR duration ($P = 2.5 \times 10^{-8}$, LMM). **f**, The spatial probability of decoded ripple content for the goal ($n = 94$ SWRs) and sham ($n = 198$ SWRs) stimulation groups. **g**, The proportions of SWRs carrying near-goal versus far-goal content (goal, $n = 11$ sessions, $P = 0.27$; sham, $n = 12$ sessions, $P = 0.003$; two-sided Wilcoxon signed-ranked test). **h**, The proportions of SWRs as in **g**, but for WT mice (familiar, $n = 36$ sessions; $P = 1.06 \times 10^{-6}$; novel, $n = 26$ sessions, $P = 0.0004$; two-sided Wilcoxon signed-ranked tests). For **b–e**, the distribution, median (white circle), quartiles (thicker lines) and $1.5 \times$ the interquartile range (whiskers) are shown. The mouse illustration was adapted from scidraw.io. For **g, h**, data are mean \pm s.e.m. The brain (F. Claudi) in the mouse head (E. Tyler and L. Kravitz) diagrams in **f, h** were adapted under a CC BY 4.0 licence.

LMM). **f**, The spatial probability of decoded ripple content for the goal ($n = 94$ SWRs) and sham ($n = 198$ SWRs) stimulation groups. **g**, The proportions of SWRs carrying near-goal versus far-goal content (goal, $n = 11$ sessions, $P = 0.27$; sham, $n = 12$ sessions, $P = 0.003$; two-sided Wilcoxon signed-ranked test). **h**, The proportions of SWRs as in **g**, but for WT mice (familiar, $n = 36$ sessions; $P = 1.06 \times 10^{-6}$; novel, $n = 26$ sessions, $P = 0.0004$; two-sided Wilcoxon signed-ranked tests). For **b–e**, the distribution, median (white circle), quartiles (thicker lines) and $1.5 \times$ the interquartile range (whiskers) are shown. The mouse illustration was adapted from scidraw.io. For **g, h**, data are mean \pm s.e.m. The brain (F. Claudi) in the mouse head (E. Tyler and L. Kravitz) diagrams in **f, h** were adapted under a CC BY 4.0 licence.

Inhibitory decreases not due to salience

Goal-selective decreases in interneuron activity also occurred in a decision-making task in a familiar environment in the CA1 hippocampal subregion. Furthermore, we identified that NS interneuron activity decreases do not occur in response to other non-rewarding but salient features of the environment, like cues that instruct navigation behaviour. Mice were trained to navigate a Y-maze using visual cues displayed on the wall in the central arm of the track⁴¹ (Fig. 5a,b). On most trials, a cue presented at the start of the central arm indicated which arm of the track (left or right) was the rewarded location. On a subset of trials, a second visual cue appeared when mice reached a specific location after a short delay period. During the second cue, called the update cue, the visual patterns appeared on the opposite wall from the original cue indicating that the reward location switched from the initial arm, and the animals must change their initial planned goal arm maintained in memory to the opposite choice. We found significant decreases in CA1 NS interneuron activity leading to, and in, the reward locations in this task (Fig. 5c and

Supplementary Table 3). In this task, the update cue is highly salient to the animal as it instructs them to change their planned trajectory. NS interneuron activity did not decrease significantly around the update cue; indeed, there was a trend of an increase (Fig. 5d). These findings establish that goal-selective decreases in interneuron activity occur in a decision-making task in a familiar environment in the CA1 hippocampal subregion. Importantly, these results show that decreases in NS interneuron activity in the CA1 are selective for goal locations and not for other salient features of the environment.

Discussion

Together, our data show that reward-predictive coordinated decreases in CA3 inhibition gate new learning of crucial information in service of goal-directed navigation (Extended Data Fig. 10). By selectively promoting plasticity, this goal-associated disinhibition facilitates learning of reward locations and paths to reward locations. Our findings provide a direct inhibitory link to previously reported hippocampal signatures

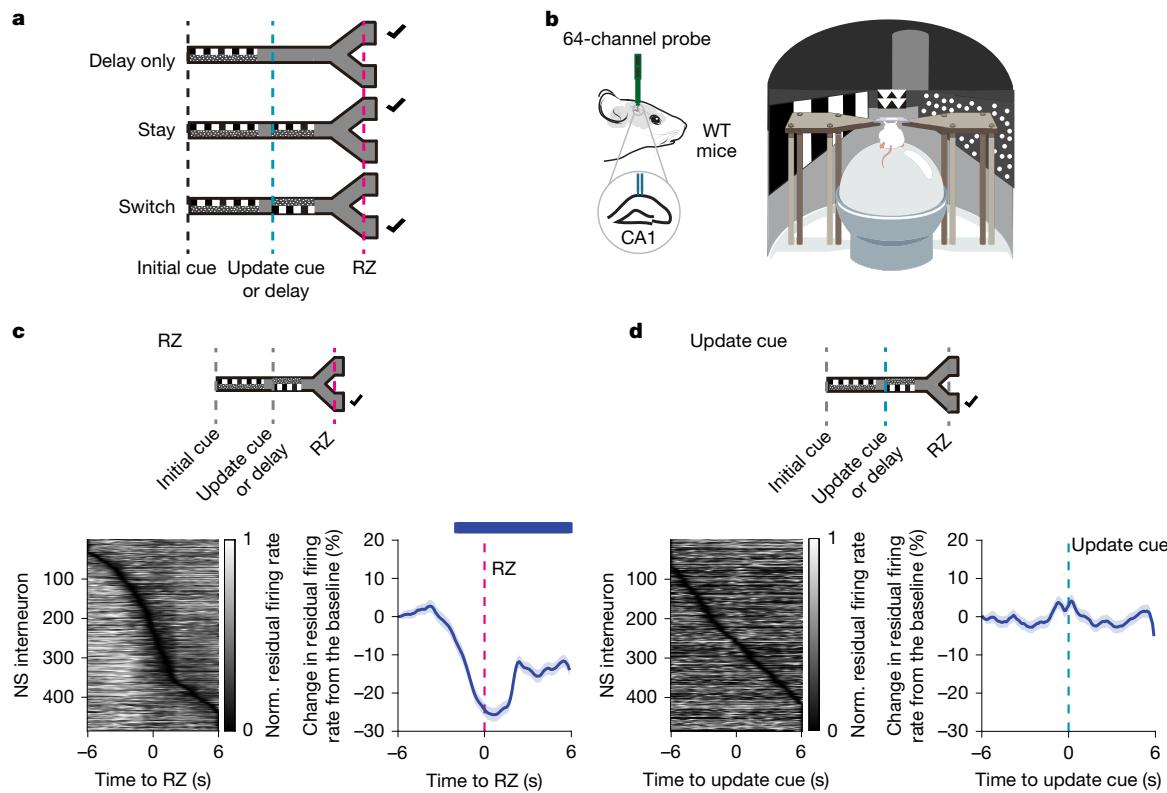


Fig. 5 | Goal-associated decreases in interneuron activity occur in the CA1 during a choice task. **a**, WT mice ($n = 7$) chose between left and right arms based on visual cues displayed on the centre walls of a virtual Y-maze. On most trials (top, delay only), the initial cue (the black dashed line marks start of initial cue presentation) indicated the rewarded arm (checkmark). On a subset of trials, an additional (update) cue was presented (teal dashed line), which indicated whether the rewarded arm stayed the same as the initial cue (stay, middle), or switched to the opposite arm (switch, bottom). The mice received a reward after entering the correct RZ at the end of the rewarded arm (pink dashed line). **b**, Electrophysiology analysis during Y-maze navigation. The mouse illustration was adapted from scidraw.io. **c**, Top, schematic indicating the rewarded areas. Left, the normalized residual firing rate of NS interneurons. Right, the

cell-averaged percentage change in normalized residual firing as a function of the time to the RZ (pink) for NS interneurons. $n = 485$ cells. There was a significant effect of time to RZ ($P < 2.2 \times 10^{-16}$, LMM). The blue bar indicates bins that significantly differ from zero or the baseline (two-sided t -test followed by Bonferroni correction). **d**, The normalized firing rate and the percentage change in firing rate as in **c**, but for the time around the update cue ($n = 485$ cells). There was a significant effect of time to the update cue ($P = 5.19 \times 10^{-13}$, LMM); however, no individual time bins were significantly different. For **c,d**, data are mean \pm s.e.m. The update task and trial type^{a1} (**a,c,d**), VR schematic^{a1} (**b**), mouse head (E. Tyler and L. Kravitz; **b**) and the brain in mouse head (F. Claudi; **b**) diagrams were adapted under a CC BY 4.0 licence.

of learning about goal locations: enhanced goal representation by pyramidal cells and SWR reactivation. Goal-associated inhibitory decreases were required for new goal learning and for stable representations and reactivation of goal locations, which predict stable memory formation^{6,37,42}. Previous studies show that inhibition regulates place-field formation or refinement and that changes in interneuron firing precede SWRs^{22–24,43–45}. However, these studies do not demonstrate coordinated inhibitory decreases preceding or within reward locations. Disrupting goal-associated decreases in inhibition around goals impaired SWR reactivation of goals, information that must be learned for successful task performance. Furthermore, we found that NS interneuron activity did not decrease around a salient wall cue that directs the animal to change its previously planned path. These results indicate that decreases in NS interneuron activity are specific to goal locations rather than generally salient features. Previous research has shown mixed increases and decreases in interneuron firing after reward. A study²⁴ found that most CA1 vasoactive intestinal peptide cells, which are disinhibitory to CA1, decreased activity after reward delivery, which would inhibit pyramidal cells. In analyses that controlled for the effects of speed, vasoactive intestinal peptide cells showed a wide range of pre-reward responses with some increasing and some decreasing activity. In another study examining interneuron firing around reward locations²¹, it was found that bistratified cells and somatostatin-positive cells ramped up their activity during approach to a familiar reward

location, a different pattern from our observations. In analyses that controlled for animal speed, they found interneuron activity decreased after animals entered the RZ, but not before. Ultimately, our findings are different because we show coordinated decreases in interneuron activity that precedes the RZ and we see these patterns in the CA3, upstream of the CA1. As the decreases in interneuron activity in our study predict rather than respond to reward, they are well positioned to enhance representations of reward-predictive locations.

On the timescale of seconds, decreases in NS interneuron activity clearly occur before pyramidal activity increases. On the timescale of days, the timescale over which these responses are refined, the pyramidal and NS interneuron responses develop concurrently. These different timescales are linked by previous work showing that artificial place fields can be induced by depolarizing a cell in a particular location^{46,47}. Thus, decreases in NS interneuron firing and resulting disinhibition of pyramidal cells preceding and in the RZ may enhance place-field development and stabilization in that location. We found interneuron firing activity occurred over long, behaviourally relevant timescales (seconds), similar to recently discovered behavioural time-scale synaptic plasticity (BTSP)^{47–50}. While BTSP appears to be present without goal enrichment, reduced PV activity may have a role in BTSP. We show that decreases in interneuron activity shift to anticipating the reward location as animals learn reward locations in the new environment. In this way, goal-selective disinhibition facilitates learning and

representations of both reward locations and locations that lead to reward, which is crucial for navigating to goals.

Online content

Any methods, additional references, Nature Portfolio reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at <https://doi.org/10.1038/s41586-025-08868-5>.

1. Kesner, R. P. Behavioral functions of the CA3 subregion of the hippocampus. *Learn. Mem.* **14**, 771–781 (2007).
2. Lee, I. & Kesner, R. P. Differential contribution of NMDA receptors in hippocampal subregions to spatial working memory. *Nat. Neurosci.* **5**, 162–168 (2002).
3. Rolls, E. T. The storage and recall of memories in the hippocampo-cortical system. *Cell Tissue Res.* **373**, 577–604 (2018).
4. Nakazawa, K. et al. Hippocampal CA3 NMDA receptors are crucial for memory acquisition of one-time experience. *Neuron* **38**, 305–315 (2003).
5. Hollup, S. A., Molden, S., Donnett, J. G., Moser, M. B. & Moser, E. I. Accumulation of hippocampal place fields at the goal location in an annular watermaze task. *J. Neurosci.* **21**, 1635–1644 (2001).
6. Danielson, N. B. et al. Sublayer-specific coding dynamics during spatial navigation and learning in hippocampal area CA1. *Neuron* **91**, 652–665 (2016).
7. Gauthier, J. L. & Tank, D. W. A dedicated population for reward coding in the hippocampus. *Neuron* **99**, 179–193 (2018).
8. Grienberger, C. & Magee, J. C. Entorhinal cortex directs learning-related changes in CA1 representations. *Nature* **611**, 554–562 (2022).
9. Pfeiffer, B. E. & Foster, D. J. Hippocampal place-cell sequences depict future paths to remembered goals. *Nature* **497**, 74–79 (2013).
10. Xu, H., Baracskay, P., O'Neill, J. & Csicsvari, J. Assembly responses of hippocampal CA1 place cells predict learned behavior in goal-directed spatial tasks on the radial eight-arm maze. *Neuron* **101**, 119–132 (2019).
11. Foster, D. J. & Wilson, M. A. Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature* **440**, 680–683 (2006).
12. Singer, A. C. & Frank, L. M. Rewarded outcomes enhance reactivation of experience in the hippocampus. *Neuron* **64**, 910–921 (2009).
13. Jadhav, S. P., Kemere, C., German, P. W. & Frank, L. M. Awake hippocampal sharp-wave ripples support spatial memory. *Science* **336**, 1454–1458 (2012).
14. Carr, M. F., Jadhav, S. P. & Frank, L. M. Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval. *Nat. Neurosci.* **14**, 147–153 (2011).
15. Royer, S. et al. Control of timing, rate and bursts of hippocampal place cells by dendritic and somatic inhibition. *Nat. Neurosci.* **15**, 769–775 (2012).
16. Stark, E. et al. Pyramidal cell-interneuron interactions underlie hippocampal ripple oscillations. *Neuron* **83**, 467–480 (2014).
17. English, D. F. et al. Pyramidal cell-interneuron circuit architecture and dynamics in hippocampal networks. *Neuron* **96**, 505–520 (2017).
18. Ego-Stengel, V. & Wilson, M. A. Spatial selectivity and theta phase precession in CA1 interneurons. *Hippocampus* **17**, 161–174 (2007).
19. Arriaga, M. & Han, E. B. Structured inhibitory activity dynamics in new virtual environments. *eLife* **8**, e47611 (2019).
20. Nitz, D. & McNaughton, B. Differential modulation of CA1 and dentate gyrus interneurons during exploration of novel environments. *J. Neurophysiol.* **91**, 863–872 (2004).
21. Geiller, T. et al. Large-scale 3D two-photon imaging of molecularly identified CA1 interneuron dynamics in behaving mice. *Neuron* <https://doi.org/10.1016/j.neuron.2020.09.013> (2020).
22. Dudok, B. et al. Recruitment and inhibitory action of hippocampal axo-axonic cells during behavior. *Neuron* **109**, 3838–3850 (2021).
23. Geiller, T. et al. Local circuit amplification of spatial selectivity in the hippocampus. *Nature* <https://doi.org/10.1038/s41586-021-04169-9> (2021).
24. Turi, G. F. et al. Vasoactive intestinal polypeptide-expressing interneurons in the hippocampus support goal-oriented spatial learning. *Neuron* **101**, 1150–1165 (2019).
25. Rolls, E. T. & Kesner, R. P. A computational theory of hippocampal function, and empirical tests of the theory. *Prog. Neurobiol.* **79**, 1–48 (2006).
26. Luo, A. H., Tahsili-Fahadan, P., Wise, R. A., Lupica, C. R. & Aston-Jones, G. Linking context with reward: a functional circuit from hippocampal CA3 to ventral tegmental area. *Science* **333**, 353–357 (2011).
27. Klausberger, T. & Somogyi, P. Neuronal diversity and temporal dynamics: the unity of hippocampal circuit operations. *Science* **321**, 53–57 (2008).
28. Klausberger, T. et al. Brain-state- and cell-type-specific firing of hippocampal interneurons in vivo. *Nature* <https://doi.org/10.1038/nature01374> (2003).
29. Klausberger, T. et al. Spike timing of dendrite-targeting bistratified cells during hippocampal network oscillations in vivo. *Nat. Neurosci.* **7**, 41–47 (2004).
30. Klausberger, T. et al. Complementary roles of cholecystokinin-and parvalbumin-expressing GABAergic neurons in hippocampal network oscillations. *J. Neurosci.* **25**, 9782–9793 (2005).
31. Grieco, S. F. et al. Anatomical and molecular characterization of parvalbumin-cholecystokinin co-expressing inhibitory interneurons: implications for neuropsychiatric conditions. *Mol. Psychiatry* **28**, 5293–5308 (2023).
32. Kohus, Z. et al. Properties and dynamics of inhibitory synaptic communication within the CA3 microcircuits of pyramidal cells and interneurons expressing parvalbumin or cholecystokinin. *J. Physiol.* **594**, 3745–3774 (2016).
33. Pelkey, K. A. et al. Hippocampal GABAergic inhibitory interneurons. *Physiol. Rev.* **97**, 1619–1747 (2017).
34. Lee, S.-H. et al. Parvalbumin-positive basket cells differentiate among hippocampal pyramidal cells. *Neuron* **82**, 1129–1144 (2014).
35. Madisen, L. et al. A toolbox of Cre-dependent optogenetic transgenic mice for light-induced activation and silencing. *Nat. Neurosci.* **15**, 793 (2012).
36. Dupret, D., O'Neill, J., Pleydell-Bouyer, B. & Csicsvari, J. The reorganization and reactivation of hippocampal maps predict spatial memory performance. *Nat. Neurosci.* **13**, 995–1002 (2010).
37. Zaremba, J. D. et al. Impaired hippocampal place cell dynamics in a mouse model of the 22q11.2 deletion. *Nat. Neurosci.* **20**, 1612–1623 (2017).
38. Norimoto, H. et al. Hippocampal ripples down-regulate synapses. *Science* **359**, 1524–1527 (2018).
39. O'Neill, J., Senior, T. & Csicsvari, J. Place-selective firing of CA1 pyramidal cells during sharp wave/ripple network patterns in exploratory behavior. *Neuron* **49**, 143–155 (2006).
40. Klausberger, T. GABAergic interneurons targeting dendrites of pyramidal cells in the CA1 area of the hippocampus. *Eur. J. Neurosci.* **30**, 947–957 (2009).
41. Prince, S. M., Yassine, T. A., Katragadda, N., Roberts, T. C. & Singer, A. C. New information triggers prospective codes to adapt for flexible navigation. Preprint at *bioRxiv* <https://doi.org/10.1101/2023.10.31.564814> (2023).
42. Roux, L., Hu, B., Eichler, R., Stark, E. & Buzsáki, G. Sharp wave ripples during learning stabilize the hippocampal spatial map. *Nat. Neurosci.* **20**, 845–853 (2017).
43. Vancura, B., Geiller, T., Grosmark, A., Zhao, V. & Losonczy, A. Inhibitory control of sharp-wave ripple duration during learning in hippocampal recurrent networks. *Nat. Neurosci.* **26**, 788 (2023).
44. Doron, A. et al. Hippocampal astrocytes encode reward location. *Nature* **609**, 772–778 (2022).
45. Szabo, G. G. et al. Extended interneuronal network of the dentate gyrus. *Cell Rep.* **20**, 1262 (2017).
46. Rolott, S. V. et al. Local feedback inhibition tightly controls rapid formation of hippocampal place fields. *Neuron* **110**, 783–794 (2022).
47. Bittner, K. C., Milstein, A. D., Grienberger, C., Romani, S. & Magee, J. C. Behavioral time scale synaptic plasticity underlies CA1 place fields. *Science* **357**, 1033–1036 (2017).
48. Bittner, K. C. et al. Conjunctive input processing drives feature selectivity in hippocampal CA1 neurons. *Nat. Neurosci.* **18**, 1133–1142 (2015).
49. Milstein, A. D. et al. Bidirectional synaptic plasticity rapidly modifies hippocampal representations. *eLife* **10**, e73046 (2021).
50. Li, Y., Briguglio, J. J., Romani, S. & Magee, J. C. Mechanisms of memory-supporting neuronal dynamics in hippocampal area CA3. *Cell* **187**, 6804–6819 (2024).

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2025

Methods

Animals

All procedures involving animals were performed in accordance with the guidelines provided by the Institutional Animal Care and Use Committee at the Georgia Institute of Technology. Our study used C57BL/6J adult male WT mice ($n = 9$ mice) and PVxAi32 mice ($n = 11$ mice) at 10–12 weeks of age (3.5–5 months at the time of the recordings). For optogenetics experiments, we crossed male PV-Cre knockin (The Jackson Laboratory, 017320) mice with female homozygous Ai32 mice (The Jackson Laboratory, 024109) to generate PVxAi32 mice ($n = 11$ mice) that express ChR2 specifically in PV-positive interneurons. The mice were housed in a reverse dark–light-cycle room (07:00 light off, 19:00 light on) with ad libitum access to food and water. Animal housing rooms are equipped with a ventilation system that provides 12 air changes per hour, temperature range of 64–79 °F and 30–70% relative humidity. We performed all behavioural training during the dark cycle. Sample sizes were determined based on sample sizes used in previous studies^{23,24,42}. We aimed for sample sizes of 5–7 mice per group with a large number of single units per animal. No power analysis was performed. Owing to the nature of experiment monitoring during electrophysiology recording, the experimenter could not be blinded to the condition. Experimenters were blinded to stimulation and novelty conditions for analyses. We used a within-subject design, in which individual animals were exposed to both conditions (novel and familiar or goal and sham stimulation), thus randomization to groups was not relevant.

Graphical illustrations of a mouse in VR were created by B. Mariner for the Singer laboratory. The drawing of mouse head profile (E. Tyler and L. Kravitz, mouse drinking, Zenodo, <https://doi.org/10.5281/zenodo.3925985>) and brain (F. Claudi, mouse brain sagittal, Zenodo, <https://doi.org/10.5281/zenodo.3925911>) were adapted from scidraw.io.

Surgery

Mice were handled for at least 3 days before stereotaxic surgery. For head-plate implantations, mice were deeply anaesthetized with isoflurane, and head-plates were affixed to the skull with dental cement (Parkell C&B Metabond). For acute extracellular recordings, craniotomies of 600–900 µm in diameter were made on either hemisphere of well-trained mice 1 day before the first recording. We used the following coordinates to target CA3 stratum pyramidale (from bregma): −1.8 mm anteroposterior (AP), 2.4 mm mediolateral (ML) and about 2.5 mm dorsoventral (DV).

VR spatial-learning task

Mice began habituation to head fixation on a treadmill at least 7 days after head-plate implantation. We trained mice to run on either a linear treadmill ($n = 4$ WT mice; PhenoSys SpeedBelt) or a spherical Styrofoam treadmill floating on air ($n = 5$ WT mice, 11 PVxAi32 mice) to determine whether changes in interneuron activity were consistent across systems. To increase motivation for running, animals were food-deprived gradually to 85–90% of their original body weight. Mice were head-fixed and trained daily to run unidirectionally on the virtual track for progressively longer periods. Mice received drops of sweetened condensed milk as a reward, delivered through a plastic needle. Licks were detected using either a piezoelectric sensor attached to the reward needle or a custom-built photointerruptor-based system. All virtual tracks were designed and interfaced with animals using the open-source software Virtual Reality MATLAB Engine (ViRME) open-source software (in MATLAB v.2015b) as previously described⁵¹. Proximal and distal cues were projected onto a cylindrical screen, creating a rich and immersive environment for virtual navigation. Voluntary movement of a mouse on the linear or spherical treadmill automatically advanced movement in the virtual track environment.

We used three annular tracks (tracks A, B and C; Fig. 1c) of the same size but with distinct visual cues, reward locations and distances

between the reward locations. All of the mice were initially trained on track A as the familiar environment and later during recordings were introduced to tracks B and C (not shown) as the novel condition. Speed gain was set such that 1° advancement through the virtual environment was equivalent to the movement of about 1.5 cm on the belt for a total distance of approximately 540 cm around each track. Animals completed a full lap (360°) in 66 s on average (range of 12–346 s). Gain on the spherical treadmill was adjusted to ensure similar virtual experience as with the linear treadmill. Each of the tracks had 36 non-overlapping, equally sized wall cues (10° per cue), and three of them were associated with reward (RZs, each 10°). The RZs were irregularly placed to prevent animals from solely using the inter-reward distances to perform the task across multiple environments.

For the first 3–5 days of training in the virtual environment, mice received rewards automatically when they arrived at the three RZs on the track. After this first phase of training, the mice showed anticipatory licking immediately before the reward delivery. In the subsequent training and recording sessions, the mice were required to lick in the RZs to trigger a reward delivery. To prevent generalized licking irrespective of position on training days, licking more than 25–50 times cumulatively outside the anticipatory and RZs triggered a 4-s time-out period. During a time-out period, the animal was teleported to a grey box and received no reward. After 4 s, the mouse resumed running from the location it had left off before the time-out. We introduced mice to novel tracks on days of recording only after a mouse was deemed to have learned and reached behavioural criteria indicative of good performance in the familiar environment. Mice were required to have 85% correct or higher performance in the familiar environment for at least two consecutive sessions before recording commenced. Behavioural performance criteria were set based on an animal's propensity to slow down and lick more in the AZs.

Behavioural data analysis

Raw behavioural data were divided based on an animal's position (in degrees) into either 360° laps around the whole environment or RZ-centred trials that spanned the area 60° before and after each 10° RZ for analyses. Mean speed ($\text{in } \text{° s}^{-1}$) was calculated by dividing the total distance travelled in each 2° position bin by the total time spent in that bin. The lick probability was calculated by taking the ratio of the number of licks per position bin over the total number of licks within the lap or trial. The lick rate (in licks per s) was calculated by dividing the number of total licks per position by the total time spent in that bin over all position bins. Speed and lick behaviour for each lap or trial was smoothed with a Gaussian-weighted moving-average filter (s.d. = 2 bins). Once learned, mice tended to show stereotypical behaviour in all three RZs. Thus, speed and lick behavioural data around all three RZs were concatenated and averaged for each session. We used ROC curves to quantify behavioural performance based on how well mice differentiated between the pre-RZ (AZ) and the non-reward control zone that appeared 30° after the end of each RZ. For lick-latency-based performance, we quantified the time between when an animal entered the RZ and the first lick compared with the control lick latency at the NRZ. We quantified learning over days using area under the receiver-operating characteristic curves in which larger AUC values indicate better behavioural performance. The performance metric using movement speed data tended to be better at position differentiation than licking-based metrics, probably because mice were not required to lick in the AZ to receive a reward and well-trained animals licked very sparsely (Extended Data Fig. 1e).

Electrophysiology

All extracellular electrophysiology recordings were performed using a poly 5 two-shank 64-channel silicone probe or a 64-channel optoelectrode of the same channel geometry (NeuroNexus). Neural data were acquired using either two 32-channel Upright Headstages with

Trodes software (v.2-2-3, SpikeGadgets) or two RHD 32-Channel Recording Headstages and Intan RHD2000 Evaluation system (version 1.5.4 with MATLAB file read_Intan_RHD2000_file v.2.0, Intan) at a sampling rate of 30 kHz with a ground pellet used as reference. The location was determined by stereotaxic coordinates, depth and electrophysiological signatures. Electrophysiological features indicative of the CA3 were closely monitored: high-amplitude theta activity during running, prominent sharp-wave ripples during stillness and high-amplitude (100 + μ V) action potentials appearing on many channels.

Each recording day consisted of at least two sessions of 15 to 30 min in the familiar (track A) and novel (track B or C) environments (always starting with the familiar environment). Behavioural sessions in the VR environment were separated by a shorter (5–10 min) baseline recording period in the dark. The movement direction was consistent (clockwise) across environments. We performed recordings from the same craniotomy (typically from the right hemisphere first) for three consecutive days, using the same novel track during this period. After the first set of novel learning days, a second craniotomy was made on the contralateral hemisphere on the fourth or fifth day, and mice learned a brand new track for three consecutive days starting on the following day (details on the number of cells recorded per mouse are provided in Supplementary Tables 1 and 2).

Optogenetic stimulation during behaviour

We recorded and stimulated neural activity simultaneously using a customized two-shank 64-channel optoelectrode with two optical fibres terminating 100–200 μ m above the top recording site on each shank. Ferrule patch cables (\varnothing 105 μ m core, 0.22 NA SMA905 to \varnothing 1.25 mm; Thorlabs, M63L01) were connected to optical fibres on one end through ceramic split mating sleeves (Thorlabs, ADAL1-5) and 470-nm fibre-coupled LEDs (Thorlabs, M470F3) with T-Cube LED driver (Thorlabs, LEDD1B) with a 15 V power supply unit (Thorlabs, KPS101) on the other end. This approach delivered truncated cones of light in the dorsal CA3 that were approximately 2 mm³. On the basis of light spread and attenuation measured in brain tissue in previous work, we estimate a volume of 0.005–0.02 mm³ received light at intensities needed to drive channelrhodopsin (greater than 1 mW mm⁻²)^{52,53}. Previous studies estimate between 500 and 2,500 PV cells per mm³ in the CA3 (refs. 54,55). We confirmed the optical targeting of 2–3 PV units on average in a single recording. Consistent with light-induced activation, we observed an increase in PV firing activity within around 3 ms of blue LED light (470 nm) onset (Supplementary Fig. 2c). We controlled the onset, offset and intensity of blue light with a National Instruments data acquisition system and custom MATLAB code. Both fibres were illuminated at the same time and intensity. Position-specific stimulation was triggered in real time by animal position in the virtual environment. In the goal location-specific condition (goal stimulation), light turned on when an animal entered the AZ, the zone immediately before the RZ and stayed on for up to 10 s or until the animal left the RZ, whichever came first. In the sham stimulation condition, the light turned on when the animal entered the NRZ and stayed on for up to 10 s or until the animal left the zone immediately after the NRZ. Sham stimulation controlled for non-specific effects of optogenetic stimulation and for the effects of PV activity on novelty in general as both types of stimulation were in novel environments. A subset of animals was stimulated at the anticipatory or RZ only, or for a fixed duration of 3 s regardless of position. For each of the three stimulation zones, we stimulated at the light intensity of either 0, 5 or 13 mW mm⁻², referred to as no-stim, low-stim or high-stim, respectively, from the fibre tip. The stimulation intensity remained the same for each goal location in the same novel environment over 3 days. We randomized the order of stimulation intensities across three stimulated areas in each environment across animals. Each stimulation session consisted of 20 min of stimulation trials followed by 5 min of stimulation-free trials to assess the potentially lasting effects of stimulation at the end of a session.

To examine the effects of goal stimulation in the familiar environment, all PVxAi32 mice received a shorter (about 10 min) session with goal location-specific stimulation trials in the familiar environment (track A) as their final session of the last day of recording.

Behaviour and recording in Y-maze with update cue

To test flexible decision-making in rodents in response to new information, we designed a VR update task that requires animals to perform a memory-guided decision-making task. On most trials, the task presents animals with an initial visual cue that indicates the correct goal location in a Y-maze environment, and the animals must then run down the track to the correct arm to receive a reward. On a subset of trials, a second visual cue appears that indicates that the reward location has either changed (switch trials) or stayed the same (stay trials). On these trials, the mice must choose to keep their original goal destination or switch to the other goal destination. The mice were rewarded at the end of the track if they selected the correct arm, followed by a VR screen freeze for 3 s and then an intertrial interval period of 6 s with a grey screen. On incorrect trials, no reward was delivered and the intertrial interval was 12 s, a longer interval as a form of punishment. Over the course of a session, animals successfully performed the behaviour across all trial types.

Behaviour training and recording methods are described in detail previously⁴¹. Additional eight-week-old C57Bl/6 WT mice ($n = 7$) were implanted and recovered as described above. Over a period of about 8 weeks (55.43 \pm 7.38 days of training, mean \pm s.e.m.), food-restricted mice underwent several phases of training for about 1 h per day 5–7 days per week to ultimately learn the task. In brief, animals first learned to run on a linear track, then trained in a short Y-maze, followed by a long Y-maze. In the Y-maze, the visual cues on the walls indicated which goal arm was rewarded. At first, the cues were visible for the entire track. A delay was then introduced between the cue turning off and the choice point when animals had to select a goal arm. This delay got progressively longer. After animals performed well with a long delay, the update cue was introduced. After animals demonstrated mastery of the task with the update cue, recordings were made in CA1 using a 64-channel, dual-shank NeuroNexus probe (targeting -1.8 to -2.0 mm AP, 1.5–1.8 mm ML and about 1.4 mm DV). Recordings were made over 6–12 session per animal (details on the number of cells recorded and trial types performed per animal are provided in Supplementary Table 3). Data were analysed as described above. The RZ onset was defined as when the animal received reward, and the update cue onset was defined as when the update cue was presented.

Histological verification of the probe location

On the last day of recording, the neural probe was dipped in fluorescent dye dil (0.9 mg ml⁻¹) before recording. After recording, the mice were deeply anaesthetized with isoflurane and perfused with 4% paraformaldehyde. The brains were extracted and drop-fixed in 4% paraformaldehyde for 24 h then rinsed in 1× phosphate-buffered saline (PBS) and either sectioned by a vibratome or cryoprotected in 20% sucrose solution overnight before being frozen for sectioning on a cryostat. Fixed brain tissues were cut coronally on a vibratome or a cryostat. Tissue sections were stained for nuclei with DAPI, mounted (Vectashield Antifade Mounting Media) and confocal imaged at $\times 10$ using the LSM 700 laser-scanning confocal microscope (Zeiss).

VR screen freeze manipulation at RZs

A small group of animals ($n = 4$) was exposed to a behavioural manipulation in VR on the last recording day to control for position-dependent changes in speed and licking activity. In this session, we introduced an automatic screen freeze at the three RZs in the familiar environment for a fixed duration of 3 to 10 s regardless of the animal's movement on the ball. The goal of this manipulation was to dissociate movement speed from position-specific visual cues. After VR manipulation, we examined

Article

the effects of distinct visual cues on firing rates of the same neurons within the same animals at similar running speeds and licking activity.

Preprocessing of LFPs and detection of sharp-wave ripples and theta periods

To obtain LFPs, recorded signals were downsampled to 2 kHz, band-pass filtered between 1 and 300 Hz and interpolated over outliers (noise) defined as 15 s.d. above the mean of the prefiltered signal. In all of our LFP analyses, we used the channel with the highest envelope amplitude for the ripple band signal as the proxy site of the stratum pyramidale⁵⁶. For analyses of theta, gamma and SWR periods, the LFP was band-pass-filtered based on frequency bands (4–12 Hz for theta, 1–4 Hz for delta, 12–30 for beta and 125–250 Hz for SWR) using a finite impulse response equiripple filter. SWR events were detected when the envelope amplitude of the filtered SWR trace (125–250 Hz) was greater than 3 s.d. above the mean for at least 20 ms (refs. 12, 57, 58). We excluded any events with a power ratio (power from 100 to 250 Hz/power from 250 to 400 Hz) less than four based on the typically observed frequency range of SWRs⁵⁹. We only included in our ripple analyses the periods with at least one multi-unit spike and excluded periods during which the movement speed was above 5° s⁻¹ during the 2-s time window, 1 s before and after the mid-point of each SWR event. Multi-unit spikes were extracted by band-pass filtering the raw recorded signal between 300 and 6,000 Hz and thresholding the filtered signal above 5 s.d. from the mean. Extracted periods of interest were then visually inspected to ensure accurate detection. The duration of SWRs was defined as the length of time that the SWR envelope was greater than the threshold of 3 s.d. above the mean. SWR power was quantified as s.d. above the mean ripple power for the entire recording session. The SWR rate was quantified by dividing the total number of SWR events by the total duration of stopped periods (movement speed fell below 2° s⁻¹ for at least 2 s) for individual recording sessions with at least ten ripples detected. Coactivation probability during SWRs was defined as the number of SWRs during which both neurons in a pair had spikes, divided by the total number of SWRs⁵⁸. To show the spread of data with individual datapoints, we used violin plots. Violin plots show the kernel density estimate of the distribution (lighter shaded area), individual data points (coloured circles) and box plot indicating the median (white circle), first and third quartiles (thicker lines) and 1.5× the interquartile range (whiskers).

Single-unit isolation

We identified and sorted putative single units using the automatic clustering software Kilosort2⁶⁰, followed by visualization and manual curation using Phy 2.0. Only well-isolated units with the signal-to-noise ratio of greater than 1 and <0.01% refractory period violations (inter-spike interval < 1 ms) were included in the study. We visually verified the firing rate stability of single units by ensuring the firing rates during periods in between VR sessions did not fall below the threshold set at 10% of the peak firing rate of the entire duration of the recording.

Cell-type classification

We classified single units as pyramidal cells and interneurons with narrow and wide waveforms based on the spike waveform's trough-to-peak latency, inter-spike interval statistics such as burst index¹⁵ and the autocorrelogram (ACG) fitted with a triple-exponential equation using CellExplorer (v.1.2) software⁶¹:

$$\text{ACG}_{\text{fit}} = \max \left(c \left(\exp \left(\frac{-(x - t_{\text{refrac}})}{\tau_{\text{rise}}} \right) - d \times \exp \left(\frac{-(x - t_{\text{refrac}})}{\tau_{\text{rise}}} \right) \right) + h \times \exp \left(\frac{-(x - t_{\text{refrac}})}{\tau_{\text{rise}}} \right) + \text{rate}_{\text{asymptote}}, 0 \right)$$

where c is the ACG τ decay amplitude, d is the ACG τ rise amplitude, h is the burst amplitude and t_{refrac} is the ACG refractory period (ms). The

CellExplorer software automatically identifies putative NS interneurons with trough-to-peak latency ≤ 0.425 ms, wide spike-width (WS) interneurons with trough-to-peak latency > 0.425 ms and ACG τ rise amplitude > 6 ms and the remaining units assigned as pyramidal cells. We visualized all pre-labelled units, and manually excluded bad units or relabelled a cell type based on the ACG and firing characteristics. This process led to some overlap in the distributions of identified cell types, consistent with examples in CellExplorer. We recorded a total of 3,489 well-isolated single units recorded in this study (Supplementary Fig. 1). Among these, we identified a total of 437 NS interneurons, 254 WS interneurons and 2,798 pyramidal cells. We confirmed that PV interneurons verified by opto-tagging had similar waveform properties and firing rates to NS interneurons that we recorded. To identify optically tagged PV-positive interneurons, we used the stimulus-associated spike latency test as previously described⁶² and performed manual verification by visualizing light-evoked firing activity within 1–3 ms of light onset. Putative PV cells in WT animals were identified as NS interneurons with mean firing rate greater than lowest quarter of opto-tagged PV cells in familiar environment (about 20.1 Hz) to capture the population of fast-spiking interneurons contributing to behaviourally relevant network oscillations like theta (intrinsic firing rate, 21 ± 5 Hz)⁶³, gamma (intrinsic firing rate, 32.70 ± 0.793 Hz)⁶⁴ and SWRs (intrinsic firing rate, 122 ± 32 Hz)⁶³.

Previous work has shown that different subtypes of NS interneurons in hippocampus have different firing patterns in relation to theta oscillations and SWRs^{27–30}. Thus, to investigate how reward-related firing patterns are related to these network-state related patterns and NS interneuron subtypes, we assessed their preferred spiking phase during theta oscillations and their firing rate around SWRs^{27–30}. We assigned cells to putative NS interneuron subtypes using a multi-step process first assessing preferred theta firing phase, then firing patterns around SWR. To assess firing as a function of theta phase, spikes during theta periods (see the 'Preprocessing of LFPs and detection of sharp-wave ripples and theta periods' section) for each cell were sorted into 18° theta phase bins, including all theta cycles. For each potential NS interneuron subtype, we performed a one-tailed t -test to compare the raw spike counts between that subtype's preferred phase (203–339° for PV-expressing basket cells (PVBCs), 130–240° for axo-axonic cells (AACs), 74–236° for cholecystokinin-expressing (CCK) cells and 347–171° for bistratified cells) and the non-preferred phases. Cells with a false-discovery rate (FDR)-corrected $P < 0.05$, indicating strong preferred theta phase within one or more of the tested subtypes, were sorted by their t -scores for each subtype. For each cell, if the top t -score was more than 0.5 above the next highest t -score, the cell was assigned the single subtype associated with the top t -score. If any of the top three t -scores were within 0.5 of each other, then a cell was assigned multiple subtypes associated with the highest t -scores. The theta-phase assignment was then curated by firing rate around SWRs. For each cell, we assessed the firing patterns as a function of binned time periods during and around SWRs, comparing the spike counts around the SWRs to shuffled non-theta, non-SWR baseline periods, as described previously²⁸. As the SWRs are asymmetrical, we normalized the time around the SWR as follows: the period between the start and middle of the ripple was divided into four bins, and the same bin size was used for the pre-ripple period. Similarly, the period between the middle and end of the ripple was divided into four bins, and the same bin size was used for the post-ripple period. For comparison, we randomly selected 100 non-theta, non-SWR periods baseline periods of the same length as the SWRs. A cell was classified as a PVBC or a bistratified cell if spiking was elevated during SWRs, specifically the spike counts in six bins around the middle of the ripple were higher than the mean + 2 s.d. of the baseline. If a cell exhibited a biphasic pattern, meaning that its spike counts in the four bins around the ripple start was higher than the mean + 1 s.d. of the baseline, and the spike counts in ten following bins was lower than the mean – 2 s.d. of the baseline, then this cell was classified as AAC. As a previous study showed that CCK cells have

little change in firing rate during the SWRs³⁰, we assigned CCK to the cell if the spike counts in four bins around the ripple start, and four bins around the ripple end, were both within the range of mean ± 2 s.d. of the baseline. For cells with multiple subtype assignments based on theta phase, if their theta-assigned subtypes overlapped with the SWR-assigned subtypes, then the overlapping subtype was the final subtype assignment. On the basis of these criteria, we identified 137 out of 203 NS interneurons with a single subtype assignment, among which there were 41 AACs, 21 PVBCs, 27 bistratified cells and 48 CCK cells.

Place-field analyses

The occupancy-normalized firing rate map of putative pyramidal cells was quantified from spike counts and time spent in 5° position bins each smoothed with a Gaussian kernel (s.d. = 2 bins), excluding periods of stillness (movement speed $<2^\circ \text{ s}^{-1}$). The rate map was constructed by dividing the smoothed spike count by the smoothed occupancy for each bin. The following criteria were used to identify place cells among putative pyramidal cells: (1) mean firing rate of less than 10 Hz to exclude potential interneurons; (2) peak firing rate of at least 1 Hz; and (3) spatial information content greater than the 95th percentile of the information content generated from shuffled (repeated 1,000 times) data. Spatial information, expressed in bits per spike, was computed using the formula⁶⁵:

$$I = \sum_{i=1}^N p_i \frac{\lambda_i}{\lambda} \log_2 \frac{\lambda_i}{\lambda}$$

where p_i is the probability of the animal occupying the position bin i , λ_i is the mean firing rate of the cell in bin i and λ is the mean firing rate across all position bins. Goal-representing cells (goal cells) were pyramidal cells that had a place field peak in the AZ or RZ. Non-goal representing cells (non-goal cells) were defined as pyramidal cells without a place field peak within 20° (2 zones) of the start of the RZ. The field stability of goal or non-goal cells was measured using the Pearson correlation coefficient of trial-by-trial rate map correlations.

Multiple linear regression

To control for the well-documented relationship between movement speed and interneuron activity, we used multiple linear regression. We also accounted for potential effects of position-specific lick rates on firing activity. For each unit, we fit a multiple linear regression model with position-binned trial-by-trial firing rates as the response and position-binned speed, lick rate, and the interaction between speed and lick rate as predictors. We then subtracted the expected effects of speed and lick behaviour from the observed firing rates. The remaining data (residuals) were used in further analyses as estimated neural activity that could not be attributed to position-related changes in speed or lick rate across trials. For visualizing population activity, we first normalized residual firing rates over the binned position by dividing each bin by the peak firing rate per unit and averaged rewarded trials per unit. Thus, the normalized firing was scaled from 0 to 1 with 1 being the peak residual firing of each cell. To illustrate the change in residual firing activity from the baseline, we subtracted the baseline firing rate defined as the mean normalized firing rate of the first two bins (in the position range of [-60, -50] degrees where zero refers to the start of the RZ). The resulting change in normalized firing is a fractional change (or percentage change) and was shown as the mean and s.e.m. averaged over cells or shown per cell. In heat maps, we then sorted units based on the binned position with the peak or trough activity, for pyramidal cells or interneurons, respectively. We also tested other nonlinear regression approaches and found that the linear regression performed similarly to or better than other methods (Supplementary Table 4).

Quantifying firing activity around goal locations

To quantify the proportion of units with significantly increased or decreased activity in the familiar environment, we first created an

occupancy normalized firing rate map (as described above) using a 2° bin size for each RZ-centred trial (from -60° to 70° with zero being the start of the RZ) per unit. We then used a generalized linear model to find the best linear fit for the rate map over the pre-RZ position bins (-60° to 0°) and obtain a P value and estimated coefficient (slope) per unit. Only units with significant adjusted P values after correcting for the FDR of 0.01 were identified as units with significant activity change. This approach was applied to both pyramidal cells and interneurons to find significant changes in firing around goals. On the basis of our observation of decreased activity on average in interneurons and increased activity on average in pyramidal cells around goals, we aimed to identify interneurons with significant decreases in activity and pyramidal cells with significantly increased activity around goals. We included interneurons with significant negative coefficients as interneurons with significantly reduced activity and we included pyramidal cells with significant positive coefficients as pyramidal cells with significantly increased activity.

The firing rate changes around the new RZs in the novel environment tended to be more variable than in the familiar environment. We therefore chose a different approach to identify units with a significant increase or decrease in firing in the novel environment to average over in Fig. 3b. We randomly shuffled both the position-binned trial-by-trial maps of firing rates, movement speed and lick rates used for multiple linear regression. This process was repeated 1,000 times for each unit to generate a shuffled distribution. We compared these shuffled distributions in each position bin against the observed session-averaged residual firing rate in the same position bin. We identified units with a significant decrease or increase at each bin if the observed residual was either less than the 10th percentile (for interneurons) or greater than 90th percentile (for pyramidal cells), respectively, of the shuffled distribution at each position bin. The population averages (Fig. 3b) are the average of units that had at least one bin determined to be significant 10° before and after the RZ. We indicated the first bin within -40° of the RZ with a persistent decrease in firing, meaning that firing decreased below the baseline and stayed below the baseline until the RZ (Fig. 3b (triangles)). The persistent decrease in firing on day 1 started at -45°.

We assessed how goal-selective inhibitory decreases develop within the first day of exposure to the first novel environment. On this first day, animals have already learned the parameters of the VR environment and task from the familiar environment, and now must learn where to lick for reward in a new environment for the first time. We analysed interneuron firing over blocks of 25 trials as animals learned the novel track. ROC curves were generated for the same trial blocks using pooled speeds in the AZ and NRZ across 7 WT animals. The trial-level speeds were min-max normalized per day and aggregated across animals for group-level visualization of the first and last trial blocks.

Ripple content decoding

To control for animal position and position-related variability in behaviour, we included only SWRs that occurred while the animal was in the AZ or RZ in this analysis. To determine the content of individual ripple events, we performed sequenceless decoding to decode each ripple as a single time window⁶⁶. Before the decoding position during ripples, we confirmed decoding of current position during running in the familiar environment as this neural activity would be expected to represent current position. For each ripple, we calculated the spatial probability distribution using a simple Bayesian decoder as previously described^{57,66}. In brief, the probability of particular positions given the spiking activity at each timepoint (expressed in $n_{\text{PositionBins}} \times n_{\text{Time}}$) was calculated using the formula:

$$P(X|\mathbf{N}_1^C) = \frac{\left(\prod_{i=1}^C P(N_i|X) \times P(X) \right)}{P(\mathbf{N}_1^C)}$$

where $P(X|\mathbf{N}_1^C)$ is the probability of position X (distances relative to the RZ using 5° bins) given \mathbf{N}_1^C , a vector of spike counts for all C cells

Article

recorded simultaneously within the specified time window. For all ripple events, we used a single 250-ms time window centred at the midpoint of each event regardless of its duration, and the decoder was applied to spiking within this window. $P(N_i|X)$, the probability that cell i fires N_i spikes at position X , was computed using the population firing rate map averaged across trials (a $n_{\text{Units}} \times n_{\text{Bins}}$ matrix using 5° position bins as a function of distance to RZ, from -60 to 70°) multiplied by the decoding time window. $P(X|N_i)$ was then normalized across position bins to sum to 1. From the single probability estimate, we identified the decoded position bin with the highest spatial probability as the most likely spatial information (expressed as the relative distance to RZ) carried by the population activity during each event.

Theta modulation of firing

Theta modulation was analysed during running periods when the movement speed was greater than 2° s⁻¹ in the AZ. For each cell, the theta phase of spikes occurring in the defined running periods was extracted. The preferred theta phase was calculated as the circular mean of the theta phases for the spikes. The vector strength was calculated for each cell by implementing the 'circ_r' function from the 'Circular Statistics Toolbox (Directional Statistics)' in MATLAB (v.R2023a).

Immunohistochemistry

We used an additional six PVxAi32 mice (3 females) to verify the expression of Chr2 localized in PV-expressing interneurons. Thirty-micrometre coronal sections were blocked in 5% goat serum in 1× PBS for 30 min followed by 30 min in 5% goat serum with 0.3% Triton X-100. After blocking, the sections were incubated in primary antibody rabbit anti-parvalbumin (Swant, PV27, 1:5,000) at 4 °C overnight, washed in 1× PBS three times for 5 min each, incubated with secondary antibody goat anti-rabbit IgG Alexa Fluor Plus 647 (Invitrogen, A55055, 1:2,000) at room temperature for 1 h and washed three times in 1× PBS. The sections were mounted and imaged at ×10 or ×20 with an LSM 700 laser-scanning confocal microscope (Zeiss). Images were analysed using Zen Blue (v.3.3) microscopy software.

Statistical analysis

For data with repeated samples from the same animal and day, we used LMM analysis in R (v.4.2.2) and lme4 package⁶⁷ (v.1.1.35.1) to evaluate significant differences while controlling for repeated measures from sessions or animals. This approach was used for most analyses including data from multiple cells recorded from the same animal on the same day. Statistically significant differences were first estimated with an ANOVA with Kenward–Roger's methods using the lmerTest package (v.3.1.3) with F -statistics reported⁶⁸. The emmeans package (v.1.8.9; <https://cran.r-project.org/web/packages/emmeans/index.html>) was used to adjust P values for multiple comparisons. For pairwise comparisons, we assessed significant differences using estimated marginal means and reported Tukey-adjusted P values and T -ratios (indicating $\beta/\text{s.e.}(\beta)$, where β is the regression coefficient). To determine statistically significant learning over days in WT mice, we included the areas under the behavioural ROC curves as the dependent variable; day, environment and day by environment interaction terms as fixed effects; and animal IDs as a random effect. Thus, the model specification for WT mice was as follows: AUC - day + environment + day × environment + (1|AnimalID). For PVxAi32 mice, our main comparison was between goal stimulation and sham stimulation conditions that occurred in the novel environment only; we therefore included AUC as the dependent variable; day, stimulation condition (not environment) and the interaction between the two as fixed effects; and animal ID as a random effect. Thus, for PVxAi32 mice, we used the following model specification: AUC ~ day + stimulation condition + day × stimulation condition + (1|AnimalID). Comparison of spatial firing (including spatial information, rate-map correlation, effects of position and temporal bins, theta modulation, peak firing rate) between groups

were compared using the following model specification: for WT mice, spatial information - day × environment + (1|AnimalID/CellID); for PVxAi32 mice with goal and sham stimulation comparisons, spatial information - day × stimulation_condition + (1|AnimalID/CellID). To show significant interaction effects, asterisks are indicated between novel days 1–3 and familiar days 1–3 or sham stimulus days 1–3 and goal stimulus days 1–3 in the figures.

For data that were sampled per session, we tested for statistically significant differences using a nonparametric two-sided Wilcoxon rank-sum test for non-uniformly distributed data. This approach applied to things like the proportion of ripples with near versus far goal content, which was computed per day and normalized within animal. Note that $P < 2.2 \times 10^{-16}$ indicates that the computed P value is below the minimum possible value reported by R. Similarly $P < 0.001$ indicates that the P value is the minimum possible value for tests that control for multiple comparisons.

For nonparametric paired comparisons, we used Wilcoxon signed-rank tests with the Bonferroni method to correct for multiple corrections. To identify position or time bins with firing rate changes significantly different from zero, we used one-sample permutation t -tests (5,000 times) and indicated these with colour coded horizontal bars above plots of cell-averaged firing as a function of position or time.

Data with error bars were reported as mean ± s.e.m. Details on the statistical analyses for each figure panel are reported in Supplementary Tables 5–22.

Reporting summary

Further information on research design is available in the Nature Portfolio Reporting Summary linked to this article.

Data availability

Processed data used in this study are available at FigShare⁶⁹ (<https://figshare.com/s/421cf9870016f963fcf5>). Data from this study are available for research purposes. Source data are provided with this paper.

Code availability

Custom code created for this paper is available at GitHub (https://github.com/singerlabgt/Inhibitory_Gating_Project.git).

51. Aronov, D. & Tank, D. W. Engagement of neural circuits underlying 2D spatial navigation in a rodent virtual reality system. *Neuron* **84**, 442–456 (2014).
52. Zhang, F., Wang, L. P., Boyden, E. S. & Deisseroth, K. Channelrhodopsin-2 and optical control of excitable cells. *Nat. Methods* **3**, 785–792 (2006).
53. Aravanis, A. M. et al. An optical neural interface: in vivo control of rodent motor cortex with integrated fiberoptic and optogenetic technology. *J. Neural Eng.* **4**, S143 (2007).
54. Morellini, F. et al. Improved reversal learning and working memory and enhanced reactivity to novelty in mice with enhanced GABAergic innervation in the dentate gyrus. *Cereb. Cortex* **20**, 2712–2727 (2010).
55. Bjerke, I. E. et al. Densities and numbers of calbindin and parvalbumin positive neurons across the rat and mouse brain. *iScience* **24**, 101906 (2020).
56. Gordon, J. A., Lacefield, C. O., Kentros, C. G. & Hen, R. State-dependent alterations in hippocampal oscillations in serotonin 1A receptor-deficient mice. *J. Neurosci.* **25**, 6509–6519 (2005).
57. Karlsson, M. P. & Frank, L. M. Awake replay of remote experiences in the hippocampus. *Nat. Neurosci.* **12**, 913–918 (2009).
58. Singer, A. C., Carr, M. F., Karlsson, M. P. & Frank, L. M. Hippocampal SWR activity predicts correct decisions during the initial learning of an alternation task. *Neuron* **77**, 1163–1173 (2013).
59. Ylinen, A. et al. Sharp wave-associated high-frequency oscillation (200Hz) in the intact hippocampus: network and intracellular mechanisms. *J. Neurosci.* **15**, 30–46 (1995).
60. Pachitariu, M., Steinmetz, N., Kadir, S., Carandini, M. & Harris, K. D. Kilosort: realtime spike-sorting for extracellular electrophysiology with hundreds of channels. Preprint at bioRxiv <https://doi.org/10.1101/061481> (2016).
61. Petersen, P. C., Siegle, J. H., Steinmetz, N. A., Mahallati, S. & Buzsáki, G. CellExplorer: a framework for visualizing and characterizing single neurons. *Neuron* **109**, 3594–3608 (2021).
62. Kvitsiani, D. et al. Distinct behavioural and network correlates of two interneuron types in prefrontal cortex. *Nature* **498**, 363–366 (2013).
63. Lapray, D. et al. Behavior-dependent specialization of identified hippocampal interneurons. *Nat. Neurosci.* <https://doi.org/10.1038/nn.3176> (2012).

64. Antonoudiou, P., Tan, Y. L., Kontou, G., Louise Upton, A. & Mann, E. O. Parvalbumin and somatostatin interneurons contribute to the generation of hippocampal gamma oscillations. *J. Neurosci.* **40**, 7668–7687 (2020).
65. Skaggs, W. E., McNaughton, B. L., Wilson, M. A. & Barnes, C. A. Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus* **6**, 149–172 (1996).
66. Carey, A. A., Tanaka, Y. & van der Meer, M. A. A. Reward revaluation biases hippocampal replay content away from the preferred outcome. *Nat. Neurosci.* **22**, 1450–1459 (2019).
67. Bates, D., Mächler, M., Bolker, B. M. & Walker, S. C. Fitting linear mixed-effects models using lme4. *J. Stat. Softw.* **67**, 1–48 (2015).
68. Kuznetsova, A., Brockhoff, P. B. & Christensen, R. H. B. lmerTest package: tests in linear mixed effects models. *J. Stat. Softw.* **82**, 1–26 (2017).
69. Jeong, N. et al. Data for 'Goal-specific hippocampal inhibition gates learning'. *Figshare* figshare.com/s/421cf9870016f963fcf5 (2025).

Acknowledgements We thank B. Mariner and the members of scidraw.io for graphical illustrations; and L. Walker, J. E. Markowitz, M. J. M. Rowan, L. Frank, B. Mensh and the members of the Singer laboratory for their comments on the manuscript. A.C.S. was supported by the Packard Foundation, the National Institutes of Health (NIH)-National Institute of Neurological Disorders and Stroke Grants R01 NS109226 and 2RF1 NS109226, the NIH National Institute of Aging Grant RF1 AG078736-01, McCamish Foundation, Friends and Alumni of Georgia Tech, and the Lane Family; N.J. by the Michael Kuhar Neuroscience Research Fellowship; A.L.P. by NIH National Research Service Award Grant 5F31AG066410, NIH Grant T32 NS007480-18,

the Fulton County Elder Health Science Fellowship, the Wright Family and the J. Norman and Rosalyn Wells Fellowship; and S.M.P. by NSF GRFP grant DGE-1444932, NIH grant T32NS096050-22 and the Emory University Women's Club Graduate Research Fellowship.

Author contributions N.J. and A.C.S. conceived the study and designed the experiments. N.J. performed all the behaviour, recording and stimulation experiments in the CA3 in the annular maze with help from X.Z. S.M.P. performed all of the behaviour and recording experiments in the CA1 for the decision-making task and developed the task. N.J. developed recording and behaviour methods with help from X.Z. and A.L.P.; N.J., X.Z. and A.L.P. wrote code and analysed the data. V.P.N., S.R.T., C.E.G. and X.Z. trained animals and analysed behaviour during training. S.R.T., M.C.G. and N.J. performed histology and histology analysis. A.C.S. provided funding, training and supervision. N.J., X.Z., A.L.P. and A.C.S. wrote the manuscript with input and approval from all of the authors.

Competing interests The authors declare no competing interests.

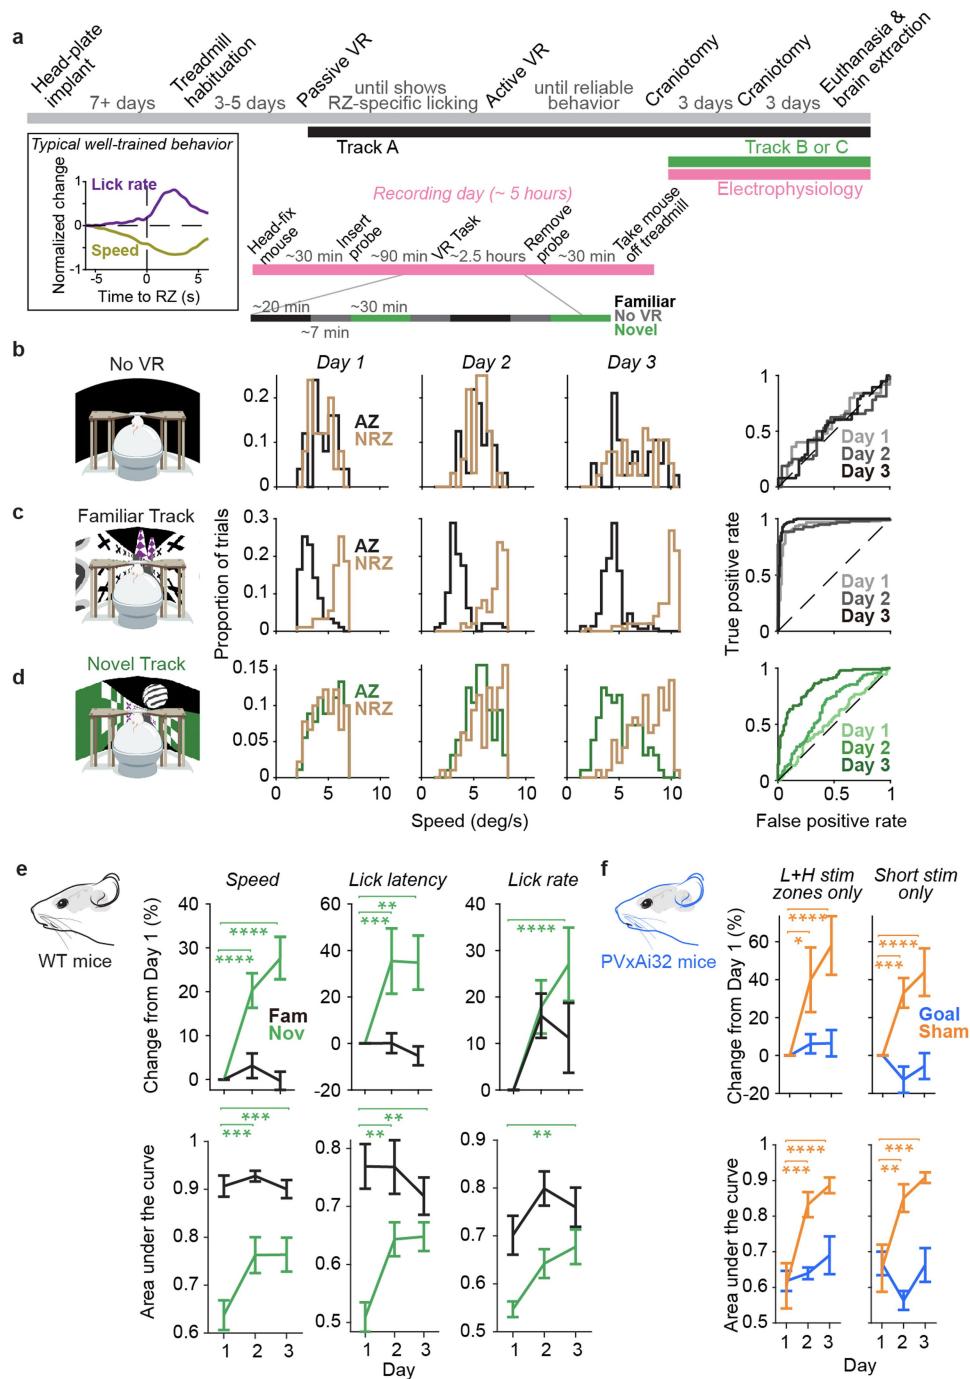
Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41586-025-08868-5>.

Correspondence and requests for materials should be addressed to Annabelle C. Singer.

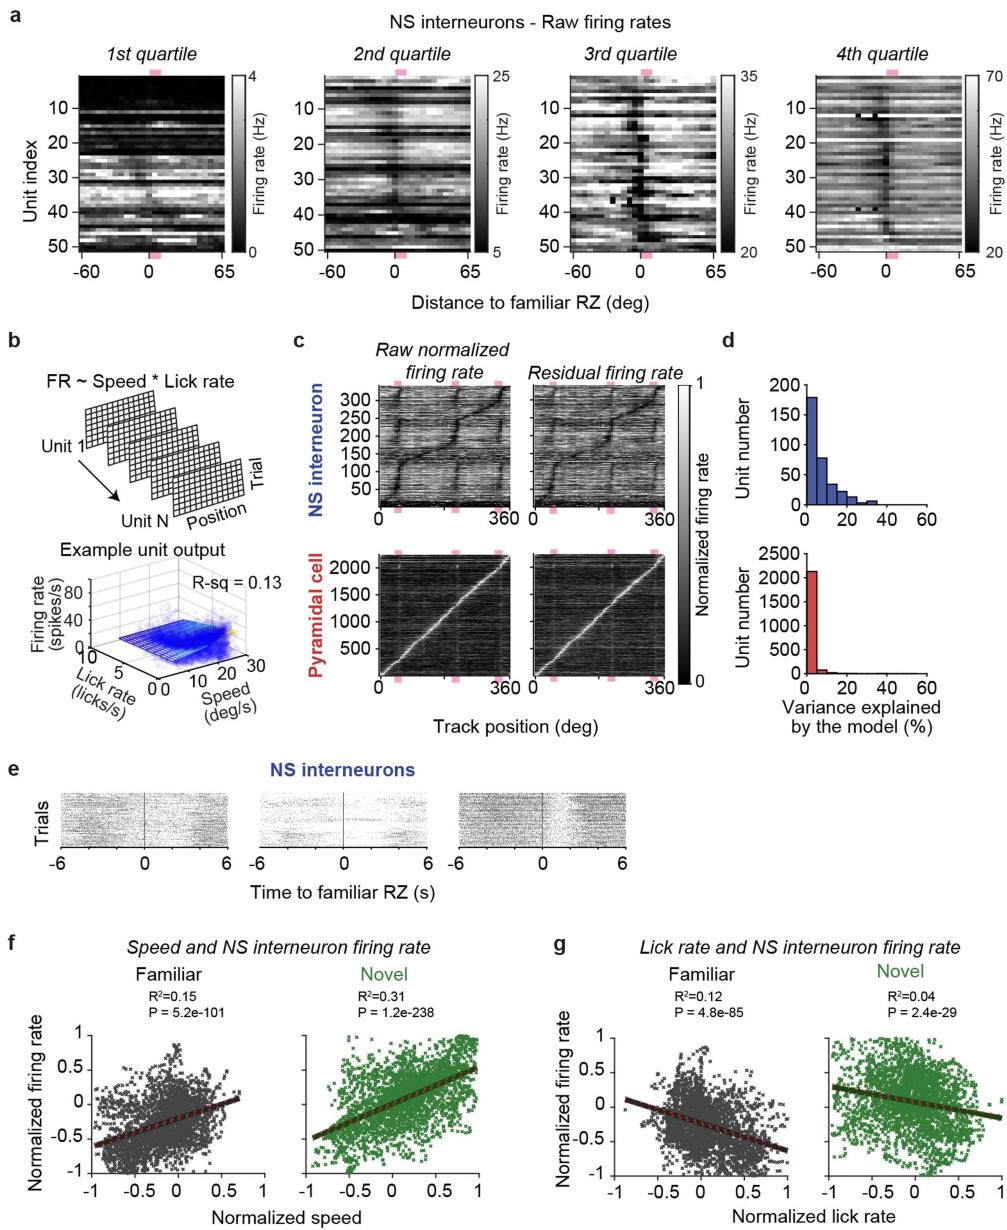
Peer review information *Nature* thanks James Carmichael and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. Peer review reports are available.

Reprints and permissions information is available at <http://www.nature.com/reprints>.



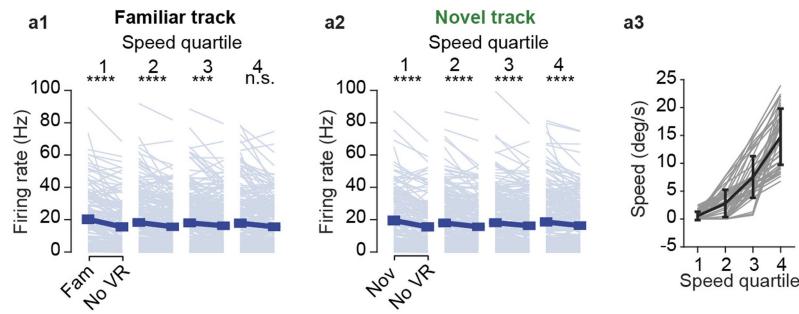
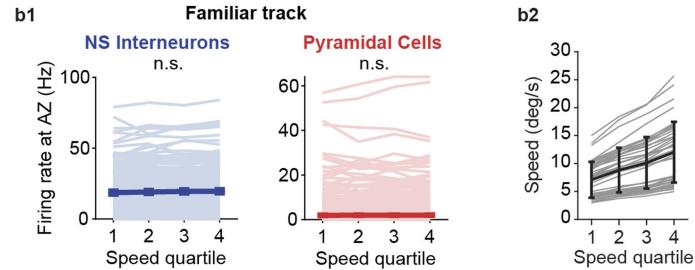
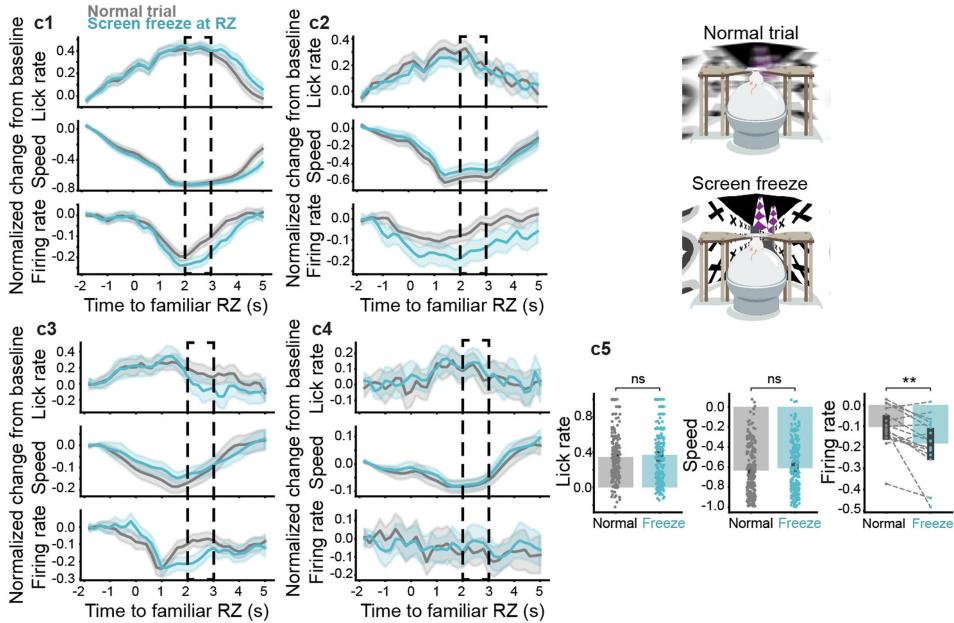
Extended Data Fig. 1 | Mice rapidly learn new goal locations in a novel virtual reality spatial learning task. **a**, Experimental timeline. Box, typical behaviour of a well-trained animal. Mean \pm SEM change in lick rate (purple) and movement speed (yellow) of an example animal, averaged over trials ($n = 80$) in a single session in the familiar environment. **b-d**, Example animal's speed distributions in AZ (black in familiar or green in novel) or NRZ (tan) and respective ROC curves over three days during no-VR baseline periods (**b**), familiar navigation (**c**), and novel navigation (**d**). **e**, Raw (top row) or percent

change (bottom row) in AUC from Day 1, based on speed (left), lick latency (middle), or lick rate (right). LMM followed by Tukey correction ($n = 14$ familiar sessions and 13 novel sessions from 7 mice). **f**, Same as **e** for PVxAi32 mice with sub-selected trials with low- or high-intensity stimulation zones only (left, “L + H stim zones only”), or trials where stimulation duration did not exceed 5 s (right, “Short stim only”). LMM followed by Tukey correction ($n = 15$ goal and 15 sham stimulation sessions from 5 mice each). Data in **e-f** represent mean \pm SEM. * $P < 0.05$; ** $P < 0.01$; *** $P < 0.005$; **** $P < 0.001$.



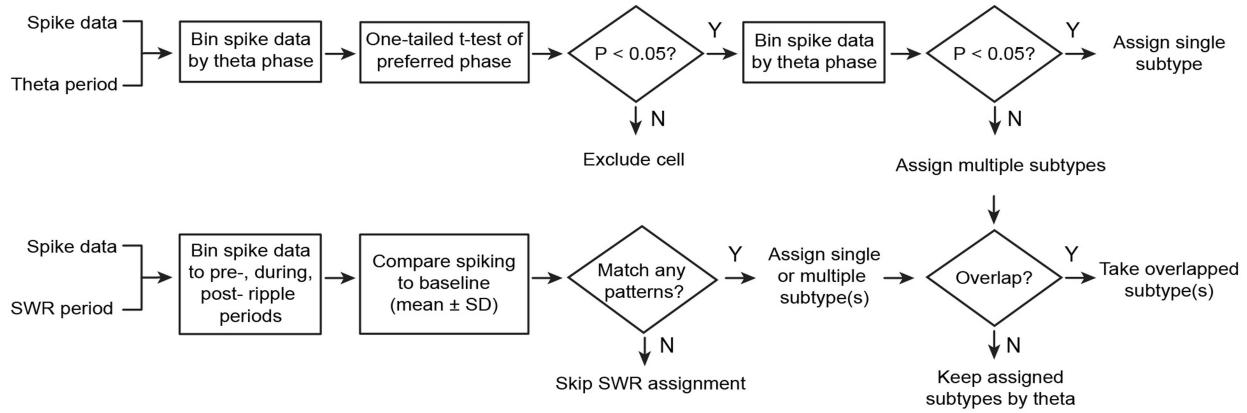
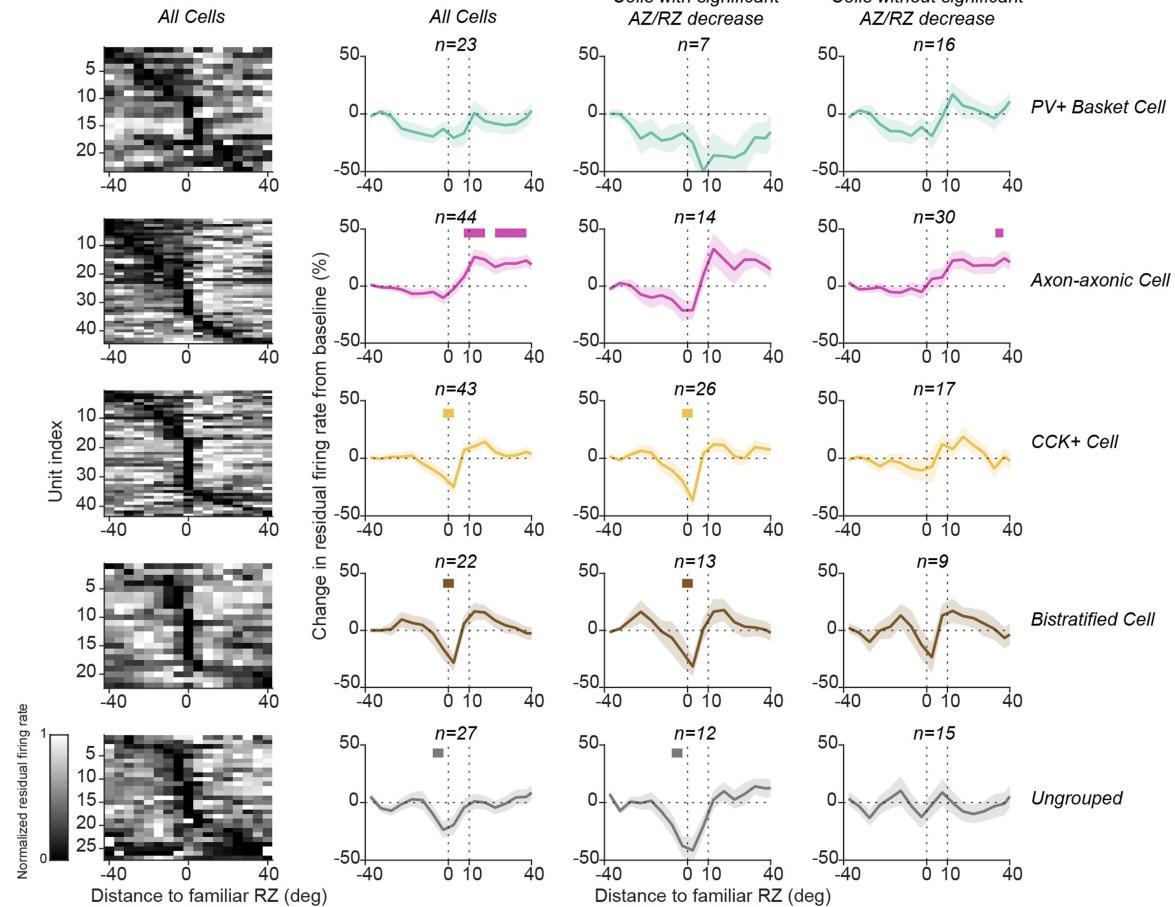
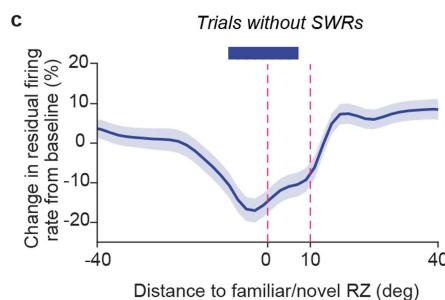
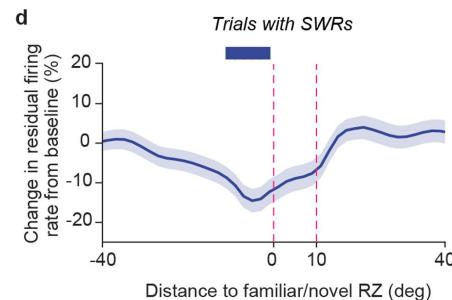
Extended Data Fig. 2 | Residual firing after estimating and removing contributions of speed and licking activity on neuronal firing rates. **a**, Raw firing rates for each NS interneuron (row) as a function of distance to RZ (pink rectangles) across movement speed quartiles. **b**, Model construction for performing multiple linear regression with an example unit's model (illustrated in a grid) plotted against the trial data (blue circles) with its R^2 -squared value on top right. **c**, Normalized raw (left) and residual (right) firing rates for all recorded NS interneurons (top, blue) and pyramidal cells (bottom, red). Each row is a single unit. Some pyramidal cells that fired around one RZ, also fired at other RZs. **d**, Percent variance explained by the multiple linear regression model fit for each unit for NS interneurons (top, blue) and pyramidal cells (bottom, red). For $>99\%$ of neurons, less than 25% of the variance in the neural firing rate data is explained by movement speed and lick rate. **e**, Raster plots of

three example NS interneurons as a function of time to the RZ. Each dot is a spike. **f**, Linear correlation between normalized NS interneuron firing rate and normalized speed in the familiar (black) and novel (green) environments. Each dot represents a spatial bin, with the y-value indicating the average firing rate of all NS interneurons recorded within a day across trials and x-value indicating the average speed across trials within a day. Familiar Days 1, 2, and 3 and Novel Days 1, 2, and 3 were pooled from WT animals ($n = 7$ mice). Speed predicted 15–31% of the variability in NS interneuron firing consistent with our hypothesis that other variables like proximity to reward also affect NS interneuron firing. **g**, Same as **f** for correlation with lick rate. The brain (F. Caudate) in mouse head (E. Tyler and L. Kravitz) diagrams in **e** were adapted under a CC BY 4.0 licence; see the Methods for further details.

a Comparing firing rates between no-task and during-task periods at similar speeds**b Comparing firing rates during reward approach across speed quartiles****c Comparing firing rates between goal vs non-goal cues at similar speeds**

Extended Data Fig. 3 | Inhibitory firing reduction cannot be explained by position-related changes in movement speed or lick rate. **a.** Firing rates of NS interneurons in speed-matched trials during no-VR versus navigation in the familiar (**a₁**) and novel (**a₂**) environments across speed quartiles (two-sided Wilcoxon sign-rank test with Bonferroni correction). Faint lines are individual neurons. Dark lines indicate mean ± SEM. **a₃**, Absolute speed values corresponding to the speed quartiles in **a₁** and **a₂**. Faint lines indicate each individual day. **b.** As in **a** for firing rates at AZ for NS interneurons (**b₁**, blue) and pyramidal cells (**b₁**, red), with absolute speed quartiles (**b₂**). LMM with Tukey correction. **c.** Normalized change from baseline, for trial-averaged lick rate (top), movement speed (middle), or NS inhibitory firing rate (bottom) as a function of time to RZ or freeze onset between normal (grey) and screen-freeze

trials (teal) across four animals (**c₁₋₄**). NS interneuron firing was not corrected for speed or licking behaviour. Dashed boxes indicate quantification periods used in **c₅** when speed was not significantly different. Only NS interneurons with stable firing rates and significant decreases at AZ and RZ were selected and averaged across trials in each animal ($n = 7, 2, 7$, and 1 NS interneurons in $224, 130, 30$, and 38 normal and screen freeze trials for animals $1-4$ respectively). Trials were subsampled to be equal across trial types in each animal). Significant difference in NS interneuron firing rates (freeze vs. normal, $P = 0.0018$, one-sided paired t-test with Bonferroni correction). Bar plots with error bars indicate mean ± SEM; two-sided paired t-test for velocity and lick rate; one-sided paired t-test for firing rate; the p-values were corrected for multiple testing using the Bonferroni method. n.s. (not significant); ** $P < 0.01$; *** $P < 0.005$, **** $P < 0.001$.

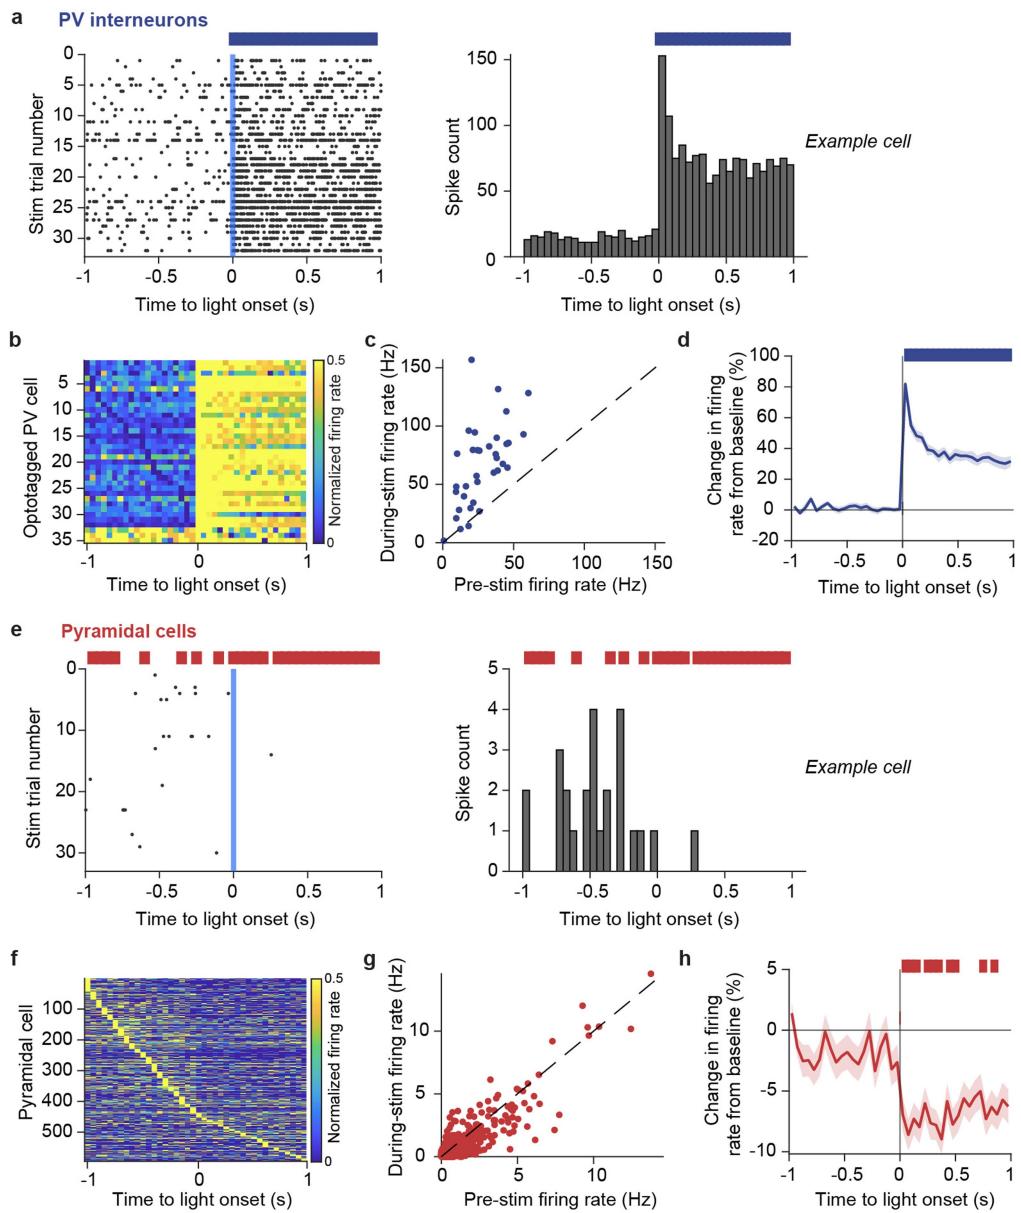
a**b****c****d****Extended Data Fig. 4** | See next page for caption.

Article

Extended Data Fig. 4 | Goal-associated inhibitory decreases generalize across subtypes of NS interneurons and are not due to SWR occurrence.

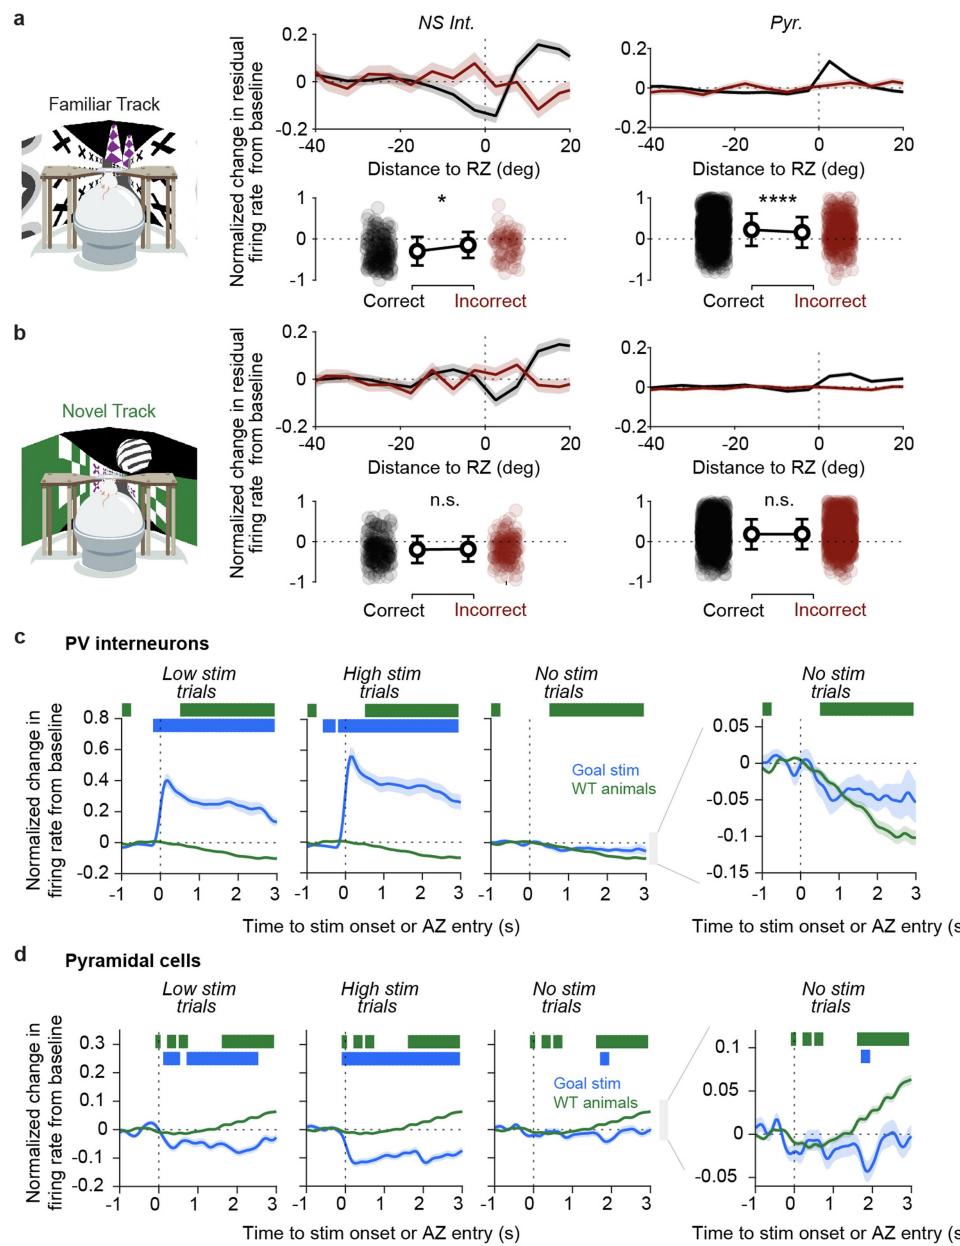
a, Flow diagram for assigning putative NS interneuron subtypes (see Methods). **b**, NS interneurons of WT animals ($n = 7$ mice) were putatively classified as PV+ basket cells (top row), axo-axonic cells (second row), CCK+ cells (middle row), bistratified cells (fourth row), or ungrouped (last row). We found decreases in firing in the AZ/RZ in all putative subtypes (far left and centre left) with some cells of each subtype showing significant decreases (centre right) and some showing no significant change (far right). Even among cells with no significant

change in firing at the AZ/RZ there were often similar trends of a decrease around the AZ/RZ on average (far right). Familiar Days 1, 2, and 3 were pooled. **c-d**, Population average percent change in normalized residual firing during trials without (**c**) or with (**d**) detected SWRs as a function of distance to RZ (pink dashed lines) for NS interneurons ($n = 203$ cells). Significant effect of position (with SWRs, $P < 2.2e-16$; without SWRs, $P < 2.2e-16$; LMM). Data in **b-d** represent mean \pm SEM with colour-coded lines above plots indicate bins that are significantly ($P < 0.05$) different from zero or baseline using data pooled from all days (two-sided t-test followed by Bonferroni correction).



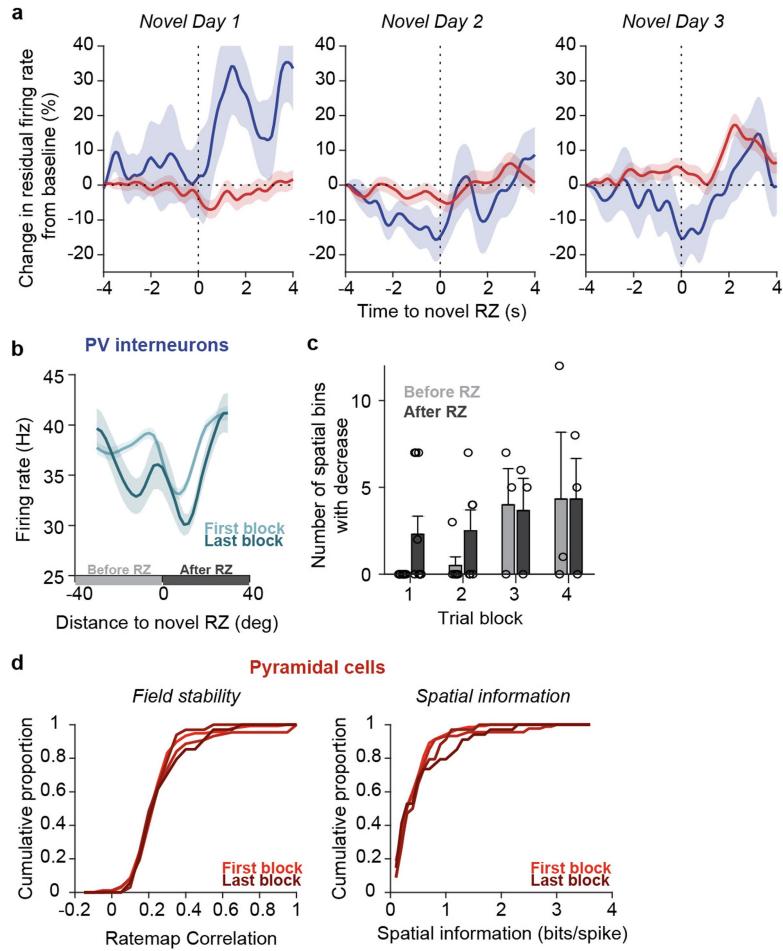
Extended Data Fig. 5 | PV interneuron firing following optogenetic stimulation inhibits pyramidal cell activity. **a**, Raster plot (left) and histogram (right) showing example PV interneuron firing as a function of time to light onset (light blue bar). In all panels, light remained on for over 1 s. **b**, Normalized raw firing rates of all optically tagged PV interneurons ($n = 35$ cells) as a function of time to light onset for all stimulation conditions. **c**, Firing rates of PV interneurons (blue dots) from **b**, before (“pre-stim”) and during stimulation (“during-stim”). **d**, Population average percent change in normalized raw firing

as a function of time to light onset (dashed line) for PV interneurons. Significant effect of time to light onset ($P < 2.2e-16$, LMM). **e**, Same as **a** for an example pyramidal cell. **f-h**, As in **b-d** for all recorded pyramidal cells ($n = 594$ pyramidal cells). Significant effect of time to light onset in **h** ($P < 2.2e-16$, LMM). In **a**, **d**, **e**, and **h** colour-coded lines above the raster plot and histogram indicate bins that are significantly ($P < 0.05$) different from zero or baseline using data pooled from all stimulation conditions (one-sided t-test followed by Bonferroni’s correction). Data in **d** and **h** represent mean \pm SEM.



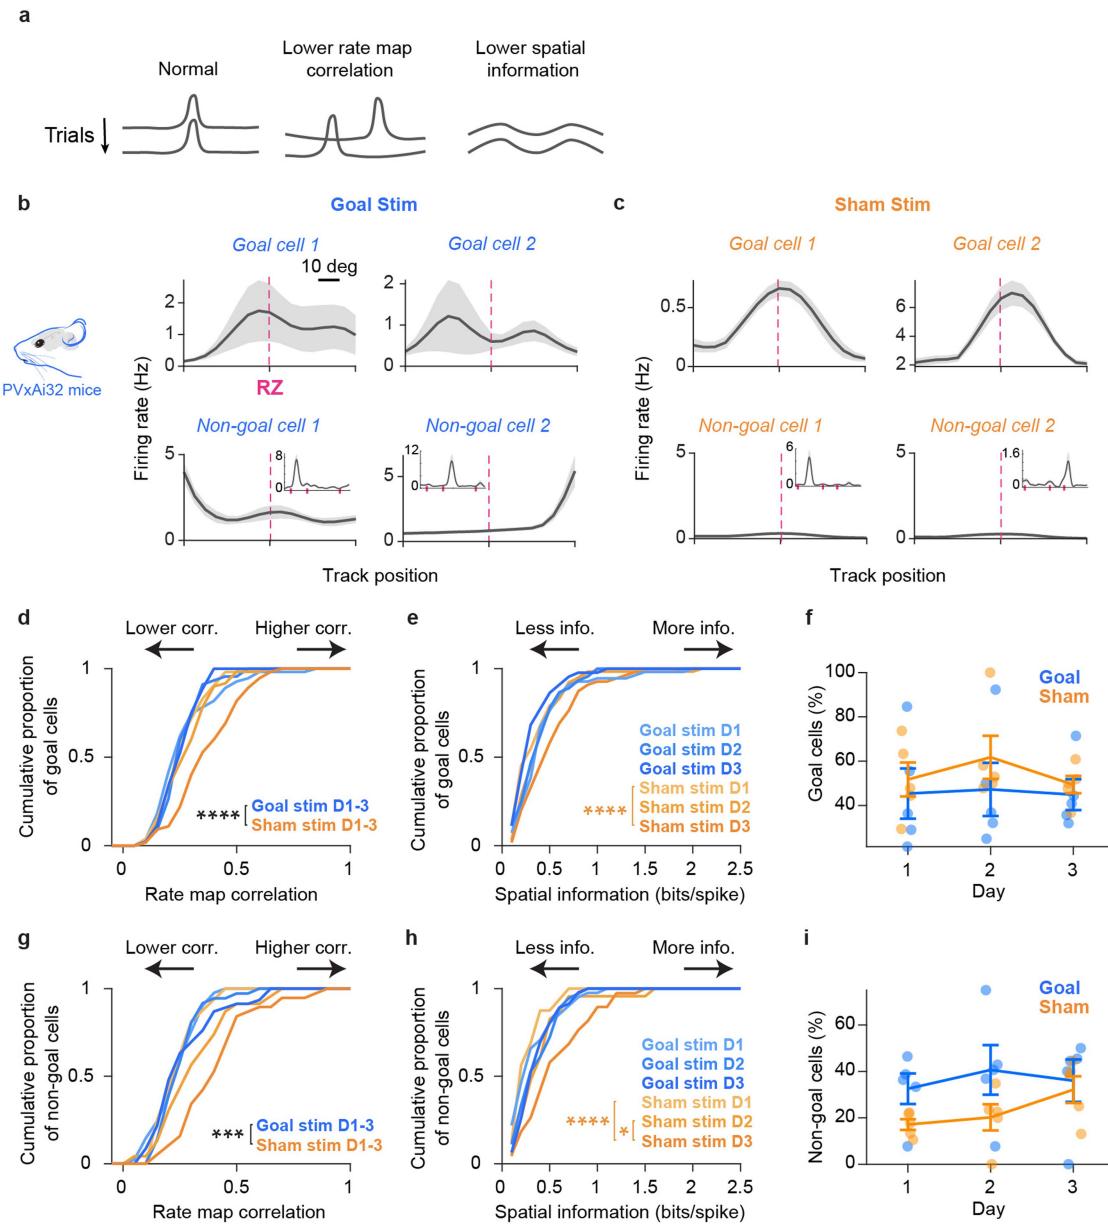
Extended Data Fig. 6 | Goal-associated inhibitory firing reductions on correct and incorrect trials and in different stimulation conditions. **a.**, *Top row*, Normalized change in residual firing rates of NS interneurons (“NS Int.” left) and pyramidal cells (“Pyr.” right) across distance to the familiar RZ during correct (black) or incorrect (burgundy) trials from 7 WT mice. *Bottom row*, Trough (for NS interneurons) or peak (for pyramidal cells) residual firing rate change around the RZ (from -10 to 5 degrees) during correct and incorrect trials. Each dot represents a single unit. Significant difference between correct and incorrect trials (NS interneurons $n = 196$ and 75 cells, $P = 0.013$; pyramidal cells $n = 1168$ and 375 cells, $P = 6.7 \times 10^{-5}$; LMM). **b.**, As in **a** for novel environments with no significant difference (NS interneurons $n = 193$ and 146 cells, pyramidal cells $n = 1144$ and 898 cells for correct and incorrect trials, respectively).

c-d., Normalized change in firing rates of opto-tagged PV cells (**c**) or pyramidal cells (**d**) around novel AZ entry (goal stimulation onset) per stimulation intensity (blue). Cells from unstimulated WT animals (green) are shown as reference. Colour-coded lines above indicate bins significantly different from zero. Significant effect of time around RZ entry with goal stimulation for both cell types ($n = 29$ PV interneurons, $P < 2.2 \times 10^{-16}$, $P < 2.2 \times 10^{-16}$, $P = 4.6 \times 10^{-11}$ for low, high, no stimulation; $n = 330$ pyramidal cells, $P < 2.2 \times 10^{-16}$, $P < 2.2 \times 10^{-16}$, $P = 3.7 \times 10^{-6}$ for low, high, no stimulation; LMM) and in WT animals ($n = 110$ putative PV interneurons, $P < 2.2 \times 10^{-16}$; $n = 1178$ pyramidal cells, $P < 2.2 \times 10^{-16}$; LMM). Data in **a-d** were pooled from all days per environment and represent mean \pm SEM. n.s. (not significant); * $P < 0.05$; *** $P < 0.001$.



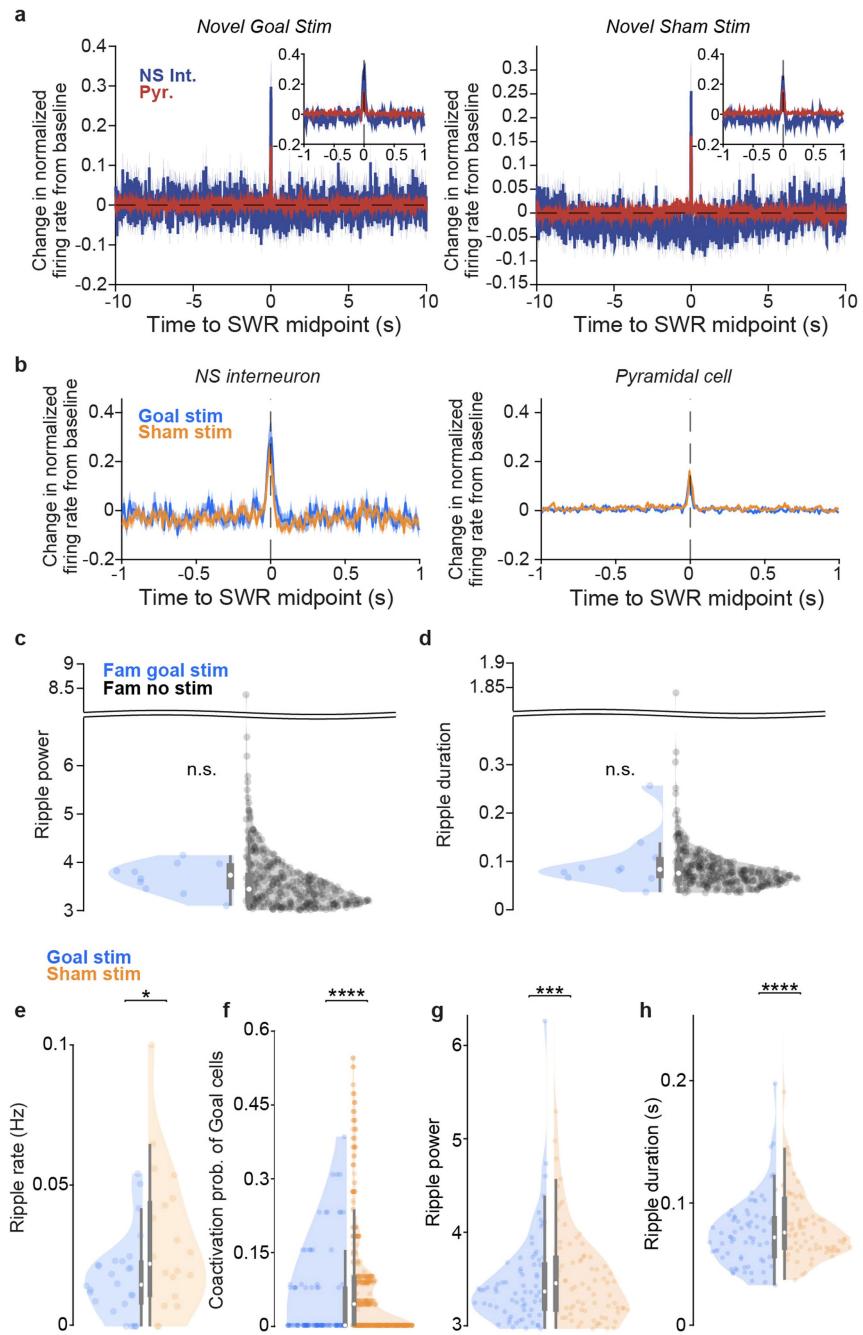
Extended Data Fig. 7 | Goal-associated inhibitory decreases emerge during the first day of navigation of a novel environment. **a**, Percent change of normalized residual firing rate from baseline for NS interneurons (blue) and pyramidal cells (red) that had a significant change in firing around the RZ as a function of time to RZ over three days of navigation of the first novel environment. (NS interneurons; Day 1, $n = 9$ cells, no significant effect of time; Day 2, $n = 20$ cells, $P < 2.2\text{e-}16$, main effect of time; Day 3, $n = 20$ cells, $P = 3.6\text{e-}08$, main effect of time; LMM. Pyramidal cells; Day 1, $n = 184$ cells, $P = 0.0005$, main effect of time; Day 2, $n = 175$ cells, $P = 2.5\text{e-}10$, main effect of time; Day 3,

$n = 199$ cells, $P < 2.2\text{e-}16$, main effect of time; LMM). **b**, Raw firing rates of putative PV interneurons as a function of distance to the novel RZ during the first (light turquoise) and last block of 25 trials (dark turquoise) on the first day of navigation in the first novel environment. **c**, Cell-averaged bin counts with a significant PV firing reduction ($n = 10, 6, 3, 3$ cells for trial blocks 1-4). **d**, Field stability (left) and spatial information (right) for all pyramidal cells ($n = 138$ cells) on the first day of navigation in the first novel environment were not significantly different from early (lightest red) to late (darkest red) blocks of 25 trials. Data in **a-c** represent mean \pm SEM.



Extended Data Fig. 8 | Disrupting goal-associated inhibitory decreases impairs stabilization and refinement of new spatial representations by pyramidal cells. **a**, Schematic of possible effects of goal stimulation on stabilization and refinement of place cells. **b**, Example firing rates of goal-representing cells (“Goal cell”, top) or non-goal cells (bottom) during sessions with goal stimulation (“Goal stim,” blue) in a novel environment. The firing rate over the full track is shown inset for non-goal cells, to visualize the location of the place field away from the RZ (pink). **c**, As in **b** for sham stimulation (“Sham stim”, orange). **d-e**, field stability (**d**) and spatial information (**e**) of goal-representing pyramidal cells (“goal cells”) across days with goal stimulation (“Goal stim D1-3”, shades of blue, $n = 55, 45, 44$ cells for Days 1-3) and sham stimulation (“Sham stim D1-3”, shades of orange, $n = 53, 60, 55$ cells for Days 1-3).

Significant interaction of day and stimulation condition for field stability ($P = 3.9e-06$) and spatial information ($P = 4.0e-06$; LMM with Tukey correction). **f**, percent of “goal cells” across days with goal stimulation and sham stimulation ($n = 5$ mice per day for each stimulation condition). **g-i**, As in **d-e** for pyramidal cells with significant spatial modulation outside goals (“non-goal cells,” $n = 41, 36, 46$ cells for Day 1-3 with goal stimulation; $n = 16, 23, 38$ cells for Day 1-3 with sham stimulation). **i**, As in **f** for number of “non-goal cells” ($n = 5$ mice per day for each stimulation condition). Significant interaction of day and stimulation condition for field stability ($P = 0.0002$) and spatial information ($P = 0.01$; LMM with Tukey correction). Data in **b-c,f,i** represent mean \pm SEM. * $P < 0.05$; ** $P < 0.01$; *** $P < 0.005$; **** $P < 0.001$.

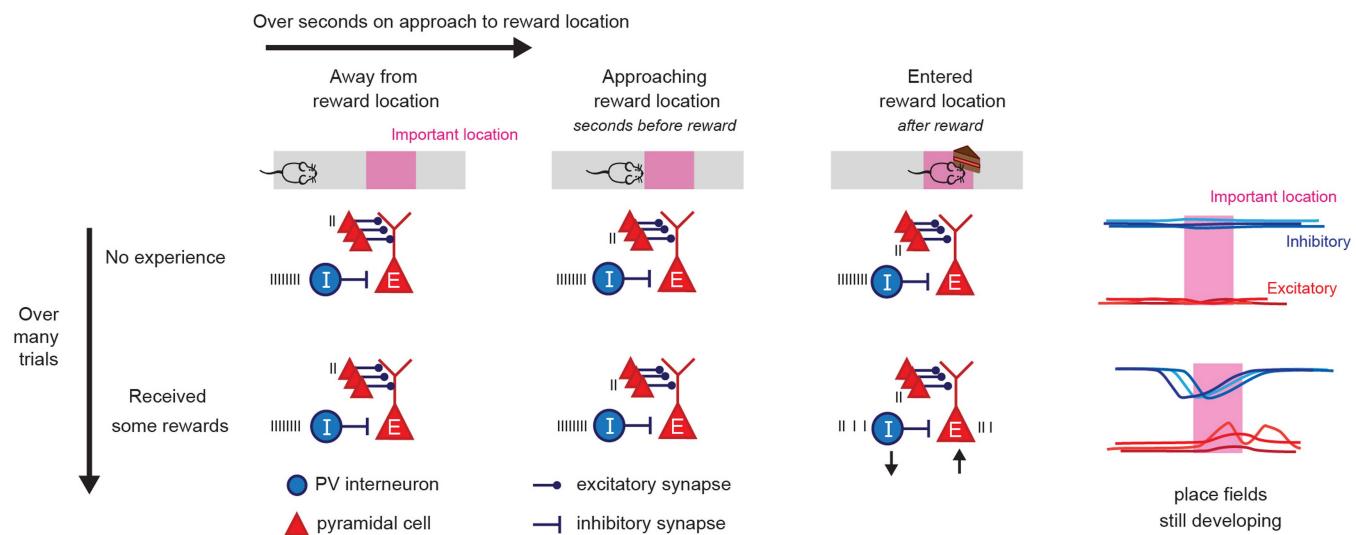


Extended Data Fig. 9 | Sharp-wave ripples during goal and sham stimulation in familiar and novel environments. **a**, Change in normalized firing rate of NS interneurons (“NS Int.” dark blue) or pyramidal cells (“Pyr.” red) around the SWR midpoint, averaged across SWR events in sessions with goal (“Novel Goal Stim,” left) or sham stimulation (“Novel Sham Stim,” right) in the novel environment. Inset, zoomed-in to show 1 sec. **b**, Normalized firing rate change between goal stimulation (“Goal stim,” blue) and sham stimulation (“Sham stim,” orange) for NS interneurons (left) and pyramidal cells (right) in the familiar environment. **c-d**, No significant difference in SWR power (**c**) or duration (**d**) between goal stimulation (blue, $n = 10$ SWRs) and no stimulation (black, $n = 390$ SWRs) in the familiar environment (LMM). Note that part of the y-axis is removed to show full distributions; no data points were removed. **e**, SWR rate during sessions with goal (blue, $n = 26$ sessions) and sham (orange, $n = 20$

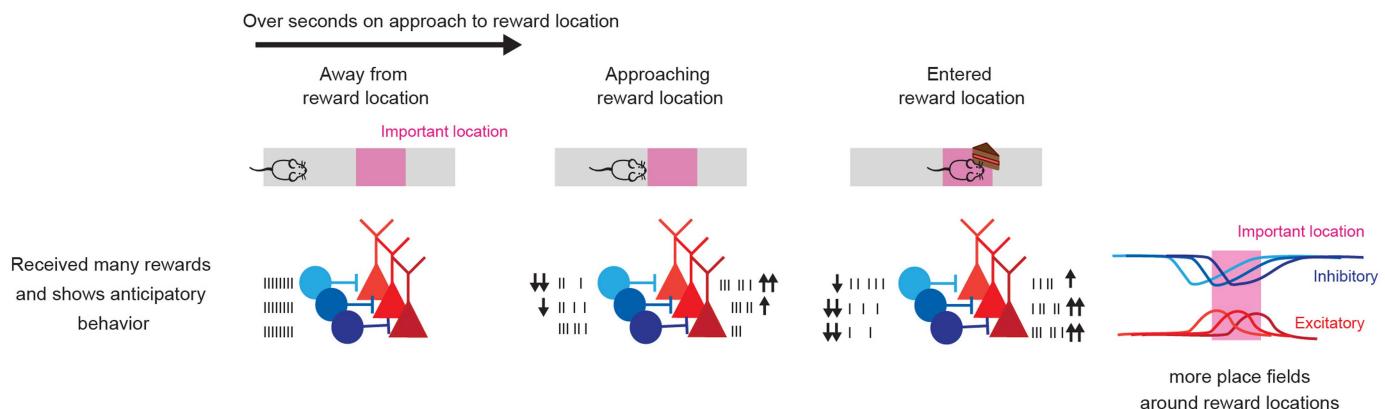
sessions) stimulation in a novel environment. Main effect of stimulation condition ($P = 0.04$, LMM). **f**, Coactivation probability of goal cell pairs during SWRs in goal ($n = 111$ pairs) and sham ($n = 516$ pairs) stimulation session ($P = 8.1e-12$, LMM). **g-h**, As in **e** for SWR power (**g**) and duration (**h**) for goal ($n = 71$ SWRs) and sham ($n = 78$ SWRs) stimulation (power: $P = 0.002$, LMM, not shown: 10.87 in sham stimulation; duration: $P = 2.5e-09$, LMM, not shown: 0.34, 0.50, 0.66, 1.00 in sham stimulation). Data in **a-b** represent mean \pm SEM. **c-h** show distribution, median (white circle), quartiles (thicker lines), and whiskers (thin lines). In **e-h**, all goal stimulation days and all sham stimulation days pooled. n.s. (not significant); * $P < 0.05$; ** $P < 0.005$; *** $P < 0.001$. The brain (F. Caudate) in mouse head (E. Tyler and L. Kravitz) diagrams in **b** were adapted under a CC BY 4.0 licence; see the Methods for further details.

Article

Coordinated Goal-Selective Inhibitory Firing Decreases Develop Rapidly During Learning



Inhibitory Firing Decreases Predict Reward & Gate Excitatory Responses At Reward Locations



Extended Data Fig. 10 | Model of Goal-Selective Disinhibition to Promote Learning of Reward Locations. This simplified schematic shows how PV interneurons (blue) and pyramidal (red) cells' firing changes around reward locations as animals learn reward locations (pink). Our data is consistent with a model in which coordinated decrease in inhibition around reward locations develop rapidly to promote excitatory representations and reactivation of reward locations. *Top*, When a naïve animal approaches and traverse a new reward location (track schematic above shows animals location in the track), NS interneurons inhibit pyramidal cells with no clear reliable response to the reward location ("No experience" row). After an animal has received a few

rewards, NS interneurons exhibit coordinated decreases in firing around the reward locations ("Received some rewards" row). These goal-selective decreases co-occur with and are required for the refinement and stabilization of excitatory place fields around goal locations. Over seconds NS interneuron activity decreases before increases in pyramidal activity. Over trials and days, NS interneurons and pyramidal cells response around reward locations develop together. *Bottom*, once animals have learned the reward locations, NS interneuron firing not only decreases in reward locations but also before reward locations. Such reward-predictive decreases in inhibition are well-suited to facilitate learning about locations that lead to reward.

Reporting Summary

Nature Portfolio wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Portfolio policies, see our [Editorial Policies](#) and the [Editorial Policy Checklist](#).

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a Confirmed

- The exact sample size (n) for each experimental group/condition, given as a discrete number and unit of measurement
- A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
- The statistical test(s) used AND whether they are one- or two-sided
Only common tests should be described solely by name; describe more complex techniques in the Methods section.
- A description of all covariates tested
- A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
- A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
- For null hypothesis testing, the test statistic (e.g. F , t , r) with confidence intervals, effect sizes, degrees of freedom and P value noted
Give P values as exact values whenever suitable.
- For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
- For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
- Estimates of effect sizes (e.g. Cohen's d , Pearson's r), indicating how they were calculated

Our web collection on [statistics for biologists](#) contains articles on many of the points above.

Software and code

Policy information about [availability of computer code](#)

Data collection	Data was collected using SpikeGadgets acquisition system with Trodes software version 2-2-3 or Intan RHD2000 Evaluation system version 1.5.4 with MATLAB file <code>read_Intan_RHD2000_file</code> version 2.0, ViRMEn open-source software in Matlab 2015b, and on a Zeiss confocal microscope with Zen Blue version 3.3 microscopy software,
Data analysis	Analyses were performed using custom pipelines in Matlab version R2023a and R version 4.2.2 <code>lme4</code> package with the following libraries: <code>lme4</code> version 1.1.35.1, <code>lmerTest</code> version 3.1.3, <code>emmeans</code> version 1.8.9, <code>Kilosort2</code> , <code>Phy</code> 2.0, <code>Cell Explorer</code> version 1.2

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors and reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Portfolio [guidelines for submitting code & software](#) for further information.

Data

Policy information about [availability of data](#)

All manuscripts must include a [data availability statement](#). This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A description of any restrictions on data availability
- For clinical datasets or third party data, please ensure that the statement adheres to our [policy](#)

Data availability: Processed data used in this study is available on FigShare: <https://figshare.com/s/421cf9870016f963fcf5>. Source data are provided with the paper. Data from this study are available for research purposes

Research involving human participants, their data, or biological material

Policy information about studies with [human participants or human data](#). See also policy information about [sex, gender \(identity/presentation\), and sexual orientation](#) and [race, ethnicity and racism](#).

Reporting on sex and gender	Not applicable
Reporting on race, ethnicity, or other socially relevant groupings	Not applicable
Population characteristics	Not applicable
Recruitment	Not applicable
Ethics oversight	Not applicable

Note that full information on the approval of the study protocol must also be provided in the manuscript.

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

Life sciences Behavioural & social sciences Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see nature.com/documents/nr-reporting-summary-flat.pdf

Life sciences study design

All studies must disclose on these points even when the disclosure is negative.

Sample size	Sample sizes were determined based on sample sizes used in prior studies that inspired our work including Geiller et al. <i>Nature</i> 2022, Roux et al. <i>Nature Neuroscience</i> 2017, and Turi et al. <i>Neuron</i> 2019. We aimed for sample sizes of 5-7 mice per group. No power analysis was performed. We recorded from a large number of single units from each animal.
Data exclusions	Animals were trained for electrophysiological recordings but only those that successfully learned the familiar track were included based on pre-established performance metrics. Two mice were excluded due to wrong targeting of the brain region for electrophysiology.
Replication	We repeated the experiments and analyses across multiple animals, serving as biological replicates. Experiments were replicated in 7 WT mice for characterization of interneuron activity around reward locations in CA3 during learning, 4 WT mice for screen freeze manipulation, 5 PVxAi32 mice for manipulation of interneuron activity in CA3 in novel environment, 5 PVxAi32 mice for manipulation of interneuron activity in CA3 in familiar environment, 7 WT mice for characterization of interneuron activity around reward locations in CA1 in a Y-maze task, 6 PVxAi32 mice for histological analysis of PV-ChR2 coexpression. All attempts were successful if the experiment succeeded. One goal stimulation condition in one animal was excluded due to probe being off-target
Randomization	We used a within subject design, in which individual animals were exposed to both conditions (eg Novel and Familiar or Goal and Sham stimulation). Thus randomization to each group was not relevant
Blinding	As described above, animals were exposed to both conditions eg Novel and Familiar or Goal and Sham stimulation. Due to the nature of experiment monitoring during electrophysiology recording, the experimenter could not be blind to condition. Experimenters were blinded to stimulation and novelty conditions for analyses.

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

n/a	<input type="checkbox"/> Involved in the study <input checked="" type="checkbox"/> Antibodies <input checked="" type="checkbox"/> Eukaryotic cell lines <input checked="" type="checkbox"/> Palaeontology and archaeology <input type="checkbox"/> Animals and other organisms <input checked="" type="checkbox"/> Clinical data <input checked="" type="checkbox"/> Dual use research of concern <input checked="" type="checkbox"/> Plants
-----	---

Methods

n/a	<input type="checkbox"/> Involved in the study <input checked="" type="checkbox"/> ChIP-seq <input checked="" type="checkbox"/> Flow cytometry <input checked="" type="checkbox"/> MRI-based neuroimaging
-----	--

Antibodies

Antibodies used

We used primary antibody rabbit anti-parvalbumin (Swant, PV27, 1:5,000) and secondary antibody goat anti-rabbit IgG Alexa Fluor Plus 647 (Invitrogen, A55055, 1:2,000)

Validation

The Swant PV27 antibody had previously been validated for optogenetic experiments (Kaneko et al., 2022). Scn1a.PV-Cre.tdT and WT.PV-Cre.tdT mice were injected in the subdural space with AAV9.hSyn.DIO.ChrimsonR.mRuby2.5T and AAV9.EF1a.DIO.hChR2(E123A).EYFP at P0 and processed for immunohistochemistry at P35-56. Sections were mounted on glass slides and confocal microscopy was performed to visualize cells immunopositive for PV and co-expressing hChR2-eYFP and ChrimsonR-mRuby2.

Reference: Kaneko K, Currin CB, Goff KM, et al. Developmentally regulated impairment of parvalbumin interneuron synaptic transmission in an experimental model of Dravet syndrome. *Cell Rep.* 2022;38(13):110580. doi:10.1016/j.celrep.2022.110580

Animals and other research organisms

Policy information about [studies involving animals](#); [ARRIVE guidelines](#) recommended for reporting animal research, and [Sex and Gender in Research](#)

Laboratory animals

Our study used C57BL/6J adult male wild-type (WT) mice (N = 7 mice) and PVxAi32 mice (N = 10 mice) at 10-12 weeks of age (3.5-5 months at the time of recordings). Animals were housed in a reverse dark-light cycle room (07:00 light off, 19:00 light on) with ad libitum access to food and water. Animal housing rooms are equipped with a ventilation system that provides 12 air changes per hour, temperature range of 64-79°F and 30-70% relative humidity.

Wild animals

No wild animals were used.

Reporting on sex

All animals used in this study were male.

Field-collected samples

The study did not involve samples collected from the field.

Ethics oversight

All animal work was approved by the Institutional Animal Care and Use Committee at the Georgia Institute of Technology.

Note that full information on the approval of the study protocol must also be provided in the manuscript.

Plants

Seed stocks

Report on the source of all seed stocks or other plant material used. If applicable, state the seed stock centre and catalogue number. If plant specimens were collected from the field, describe the collection location, date and sampling procedures.

Novel plant genotypes

Describe the methods by which all novel plant genotypes were produced. This includes those generated by transgenic approaches, gene editing, chemical/radiation-based mutagenesis and hybridization. For transgenic lines, describe the transformation method, the number of independent lines analyzed and the generation upon which experiments were performed. For gene-edited lines, describe the editor used, the endogenous sequence targeted for editing, the targeting guide RNA sequence (if applicable) and how the editor was applied.

Describe any authentication procedures for each seed stock used or novel genotype generated. Describe any experiments used to assess the effect of a mutation and, where applicable, how potential secondary effects (e.g. second site T-DNA insertions, mosaicism, off-target gene editing) were examined.