Structured Inhibitory Activity Dynamics During Learning

Moises Arriaga[1] and Edward B. Han[1]*

[1]Department of Neuroscience, Washington University School of Medicine, St. Louis, MO, USA

*-Corresponding author: Edward B. Han (ehan23@wustl.edu)

Acknowledgements: This work was supported by the McDonnell Centers for Systems Neuroscience and Cellular and Molecular Neurobiology, and through the Cognitive, Computational, Systems Neuroscience Pathway at Washington University in St. Louis.

Abstract
Hippocampal network activity is tightly regulated by local inhibitory interneurons. Suppression of inhibition has been proposed to accelerate learning by enhancing network activity and plasticity; however, the activity dynamics of hippocampal interneurons during learning remain poorly understood. Furthermore, it is unknown if individual interneurons are stochastically suppressed across different learning episodes, mirroring the random remapping of place cells, or if instead they exhibit consistent patterns of activity suppression. These critical properties define how inhibition shapes and controls learning at a network level. To uncover the functional circuit dynamics of inhibition during novelty-induced learning, we recorded calcium activity from hippocampal CA1 interneurons using two-photon imaging as mice learned a virtual reality (VR) goal-directed spatial navigation task in new visual contexts. Here we focused on dendrite-targeting somatostatin-expressing interneurons (SOM-ints), which powerfully control burst firing and synaptic plasticity in excitatory neurons. We found robust activity suppression in SOM-ints upon exposure to novel virtual environments; activity then recovered over repeated exposures to the novel environment as the animal learned goal locations. At a population level, we found a continuum of activity suppression, from interneurons strongly suppressed to moderately activated during learning. Surprisingly, each interneuron exhibited a stable level of activity modulation: when animals were switched into a second novel environment, the magnitude of activity suppression was strongly correlated across remapping sessions. This work reveals dynamic inhibition suppression triggered by novel environments and the gradual return of inhibition with learning. Furthermore, unlike the stochastic remapping of place cells, inhibitory networks display a stable activity structure across learning episodes. This functional inhibitory circuit architecture suggests that individual interneurons play specialized and stereotyped roles during learning, perhaps by differentially regulating pyramidal subnetworks specialized for plasticity and stability.

Introduction

Excitation is balanced by inhibition in neuronal networks (Andersen et al., 1963). In cortical circuits, feedforward inhibition is rapidly and robustly recruited by excitatory inputs, while pyramidal neuron firing elicits feedback inhibition to further dampen excitability. Furthermore, inhibition strongly controls synaptic plasticity, a putative cellular mechanism of learning (Bliss and Lømo, 1973; Whitlock et al., 2006). Intact inhibition limits potentiation to relatively low levels while pharmacologically blocking inhibition facilitates both the induction and magnitude of potentiation (Artola and Singer, 1987; Bear et al., 1992; Steward, 1990; Wigström and Gustafsson, 1983). Thus, inhibition suppression is a potential mechanism for enhancing learning by favoring synaptic plasticity in excitatory neurons.

Notably, inhibition can be strongly modulated *in vivo* in freely moving rodents. CA1 fast-spiking putative interneurons are suppressed when rats explore a novel spatial environment, although these studies did not explicitly address task learning, as animals were foraging for random rewards (Frank et al., 2004; Nitz and McNaughton, 2004; Wilson and McNaughton, 1993). Learning new food locations in a familiar environment dynamically modulated fast-spiking interneuron activity and altered their associations with pyramidal cell ensembles (Dupret et al., 2013). Hippocampal CA1 parvalbumin-expressing interneurons (PV-ints) have decreased network connectivity during initial learning of the Morris Water Maze that then increases with task performance and this modulation is necessary for learning (Donato et al., 2013; Ruediger et al., 2011). Furthermore numerous studies have shown suppression of SOM- and/or PV-ints is necessary for certain types of cortical and amygdalar learning, often triggered by disinhibitory inputs from interneurons that preferentially target other interneurons for inhibition, including vasoactive intestinal peptide-expressing interneurons (VIP-ints) (Gentet et al., 2012; Lee et al., 2013; Letzkus et al., 2011; Wolff et al., 2014). This work demonstrates the dynamic nature of inhibition during learning and identifies the importance of inhibition suppression in certain types of learning.

Our understanding of *in vivo* inhibitory activity in the hippocampus is primarily driven by recordings of soma-targeting fast-spiking interneurons (likely PV-ints) since their distinctive firing characteristics make them identifiable in extracellular electrophysiology recordings. SOM-ints, while far less understood, are particularly interesting because they selectively innervate the dendrites of pyramidal neurons and can directly control dendritic excitability. Dendritic spikes, typically characterized by calcium entry through Ca^{2+} channels or NMDA receptors, generate burst firing of neurons and can mediate long-term plasticity, place field formation, and learning (Bittner et al., 2015, 2017; Cichon and Gan, 2015; Larkum et al., 1999). Formation of place fields is associated with dendritic spikes that occur during transient periods of SOM-int activity suppression in novel environments (Sheffield et al., 2017). However, there is less evidence for SOM-int suppression during learning in the hippocampal formation. SOM-int activation rather than suppression is required for fear learning, both in CA1 or the dentate gyrus (Lovett-Barron et al., 2014; Stefanelli et al., 2016). The activity dynamics of SOM-ints during more complex learning paradigms remains unknown.

While interneuron activity can be dynamically modulated during learning, the stability of these activity dynamics across learning episodes has not been studied. This is partially a technical issue since extracellular electrode recordings are typically stable for a few hours and have limited ability to identify interneurons, making longitudinal recording of single interneuron activity difficult. Conceptually, interneuronal activity is thought to be governed by pyramidal neurons, driven in a feedforward and/or feedback manner. Thus, if activated pyramidal neuron ensembles are stochastic, stochastic ensembles of interneurons should result. An alternate possibility is that individual interneurons play consistent and reproducible roles during learning, reflecting an underlying structure in how inhibition regulates the network. While inhibitory neurons are composed of multiple cell-types playing distinct network roles (Klausberger and

Somogyi, 2008; Pelkey et al., 2017), little is known about the functional specialization of interneurons within a defined cell-type (Arriaga and Han, 2017).

Here we examined the structure of SOM-int inhibition by calcium imaging over multiple days and learning episodes to determine both the learning-related changes and the consistency of individual interneurons. We found SOM-int activity was strongly suppressed in novel environments, with activity gradually returning as animals re-learned goal locations. In contrast, for animals in a no-learning environment (static visual surround with no task), inhibition remained persistently suppressed over days, suggesting that recovery of inhibition is tied to learning rather than habituation. Surprisingly, suppressed inhibition triggered by context changes showed a defined population structure during both learning in the novel environment and the no task condition. Each interneuron exhibited consistent activity suppression, with high correlation of suppression across multiple novel contexts as well as the no task condition. These data reveal interneuron activity suppression during remapping and learning in the hippocampus and reveal functional inhibitory structure that may route the encoding of information in the pyramidal network.

Results

Virtual reality behavior and learning

To study the activity dynamics of SOM-ints during learning, we used two-photon calcium imaging to stably record from neurons over weeks. Our goal was to study the activity dynamics of the same cells in a goal-directed, spatial learning task. Here, establishing an appropriate learning task in VR was critical. We required a task that was complex enough to engage hippocampal learning, yet simple enough to learn over a few days. Our task is a modification on the virtual track running task used by numerous groups to study place cell activity. Water scheduled mice run to alternating ends of the virtual track to receive water rewards, using their movement on a floating spherical treadmill (Styrofoam ball) to control their movement in VR. Mice need to physically rotate the ball at the ends of the track to turn around in VR and run forward on the ball to reach the opposite end. Initial training in this task typically takes from 2 – 5 weeks with 30 – 40 min sessions per day.

To adapt this task for studying learning, we ran animals for 7 min in the well-trained familiar environment (Fam1-1), immediately switched mice to a new visual virtual environment for 14 min (New), and then back for another 7 min session in the original familiar environment (Fam1-2) (schematic, Fig. 1A). The task was identical in familiar and New epochs but the visual textures of the walls and tracks change, as well as positions and textures of distal landmarks surrounding the track. We repeated this protocol over five days, with the same New environment each time.

We note that this task in VR is more difficult than its real world analogue, where animals run to alternating ends of a narrow track for food rewards which can be learned in a few trials and is not affected by hippocampal lesions (Kim and Frank, 2009). At the same time the complexity of the task, requiring goal-directed spatial navigation, makes it more suitable for studying learning than simple VR tasks where animals run in one direction.

This paradigm has numerous advantages for studying learning. First animals are already well-trained in the familiar environment, which entails a mixture of cognitive and motor learning. This prior expertise allows mice to learn at a much faster rate in a novel context (2 -5 days) than the initial learning of the task (2 – 5 weeks). Second, the immediate switch from a familiar to a novel context marks a highly salient and defined start to the learning period. Finally, a return to the familiar environment (Fam1-2) allows us to verify that impaired behavioral performance in New is not due to time-dependent effects such as satiation or fatigue.

To better understand the behavioral effects of the familiar to novel environment switch and investigate links to learning in New, we characterized behavioral performance using a large cohort of mice, a subset of which were used for SOM-int imaging (the remainder were used for other experiments and not discussed here). We quantified task performance as number of rewards per minute (rew/min) in Fam1-1, New, and Fam1-2 epochs. Upon initial exposure to the New virtual world, animal behavior was dramatically altered. Performance in New was significantly worse on Day 1, compared to the average performance of the flanking Fam epochs (Figure 1C, performance in new normalized to familiar on day 1=0.587±0.0728, $p<0.001$). This impairment was largest on Day 1 and gradually decreased over the next four days of exposure to the same “new” world (Spearman $\rho=0.406$, $p<0.001$).

If these performance increases indicate learning in the New environment, then animals should also exhibit signs of increased familiarity with goal locations and expectations. In familiar environments, animals typically slow down before entering the reward zone (marked by a period of deceleration several second prior to the reward), in anticipation of receiving reward and consuming water. In new environments, animals initially did not decelerate as they approached the reward zone, but over 5 days of repeated exposure to New, they decelerated more and more, approaching Fam levels (Figure 1D, Day 1, $p<0.01$, Day 2 and 3, $p<0.05$). Similarly, animals in familiar environments generally lick within 500ms of a reward (“correct” licks); in New,

animals initially lick outside of reward zones, with their behavior improving over repeated training (Figure 1E, Day 1, $p<0.01$). These metrics support the interpretation that progressive improvements in rew/min in the New environment are due to animal learning.

We characterized several other metrics to assess learning (Figure 1 – Figure Supplement 1). In familiar environments, animals lick in anticipation of reward even if they have mistakenly returned to the same reward zone twice in a row ("failed reward zone entries" which are unrewarded). In contrast, animals running in the New environment rarely licked when entering an unrewarded end zone (Figure 1—figure supplement 1A, Day 1, $p<0.05$). Over the course of subsequent exposures to the New environment, failed end zone entries were accompanied by more bouts of licking, indicating that the mice begin to associate the end zone locations with rewards. This difference in licking behavior cannot be explained by a difference in the amount animals licked in the New world, as there was no significant difference in the rate at which mice licked on the first day of remapping or for most of the other days (Figure 1—figure supplement 1B, Day 1).

Mice also learned to turn around closer to the correct end zone over the course of exposure to the New environment. In the familiar environment the location of these turns on a failed trial was much closer to the correct end zone than the failed trials in the New environment (Figure 1—figure supplement 1C Day 1 distance from end zone in Fam.=15.0±1.08 cm, location in New=29.6± 1.56 cm, $p<0.001$). Over the course of the remapping experiment, these early turns moved closer to the end zone, suggesting increased awareness of reward sites over time (Spearman's $\rho=-0.500$, $p<0.001$).

Do non-learning related mechanisms, such as changes in locomotion, contribute to task performance? While animals briefly stopped moving immediately after being switched into New, there was no significant difference in total stopped time between New and Fam on Day 1 (Figure 1—supplement figure 1D, percent time stopped in Fam=8.78±2.08%, percent time stopped in new=5.05±1.31%, $p=0.0593$). Similarly, on average, there was no decrease in locomotion speed between Fam. and New (measured as total ball speed), on the Day 1 of remapping (Figure 1—figure supplement 1E, Day 1, normalized speed in new=0.951±0.079, $p=0.719$). Speed in the New world decreased on later days of remapping (Day 3 and Day 4, $p<0.05$); however, task performance had recovered at this point and changes in mean speed were not associated with performance changes. Thus, non-specific changes in locomotion are unlikely to account for increasing behavioral performance over time.

Decreased behavioral performance in New could also result from fatigue, reward satiation, or decreased motivation. However, upon returning to Fam1-2 after New, performance increased to levels similar to Fam1-1 (Figure 1—figure supplement 1F, Day 1, Fam1-1 reward rate=3.16±0.392 rew/min, Fam1-2 reward rate=3.23±0.289 rew/min, $p>0.99$). Taken together, these data indicate that, over the five day remapping protocol, animals learn the reward location in the new track and suggest that increasing behavioral performance is, in part, dependent on this learning.

Characterization of neuronal calcium activity in novel virtual environments

To investigate *in vivo* interneuronal activity dynamics during exposure to novel environments, we used two-photon imaging of neuronal calcium activity during a spatial navigation task in visual virtual reality (VR) (Arriaga and Han, 2017). We used an electric tunable lens to image a 3-D volume of mouse hippocampal CA1 by capturing sequential imaging frames along the *z*-axis moving from ventral to dorsal. Images captured neuronal somata from *stratum pyramidale* and *oriens*, over four to six planes at a frame rate of 5.2– 7.8 Hz per plane. Cre-dependent AAV1-Syn-Flex-GCaMP6f was injected into SOM-cre$^+$ transgenic mice to drive a genetically encoded calcium sensor specifically in SOM$^+$ hippocampal interneurons. Calcium activity can be taken as a proxy for neuronal activity as multiple studies using simultaneous *in vivo* imaging and cell-attached patch electrophysiology on the same

neurons have found strong correlation between spiking and calcium signals (Chen et al., 2013; Dana et al., 2016). We, and others, have measured the activity dynamics and coding properties of hippocampal neurons in visual virtual environments and found significant similarity between VR and real world behavior in several aspects of function such as place coding, direction-specificity, place cell remapping in novel environments, and interneuron activity correlation and anti-correlation with locomotion (Arriaga and Han, 2017; Harvey et al., 2009; Sheffield et al., 2017).

SOM-int activity suppression during learning in new environment

To investigate the functional activity dynamics of SOM-ints during learning, we recorded calcium activity from the same cells as animals performed the VR track running task in Fam1-1, New, and Fam1-2 over five days. SOM-int neuronal activity, measured as ΔF/F, was strongly suppressed upon transition into New (ΔF/F of 6 sample cells from one imaging plane of one mouse, Figure 2A). Individual neurons were differentially suppressed with some relatively unaffected. On returning to Familiar in Fam1-2, calcium activity rapidly recovered, as did behavioral performance. Similar results can be seen in all cells from this animal (Figure 2B). We quantified neuronal activity for all cells in this example animal as mean ΔF/F and compared across Fam1-1, New, and Fam1-2. Activity suppression is calculated as the percent difference of mean ΔF/F for each cell between New and the average of Fam1-1 and Fam1-2 (Fam), using the formula $Percent\ Change = \frac{\Delta F/F_{New} - \Delta F/F_{Fam}}{\Delta F/F_{Fam}}\ x\ 100$. The histogram of percent differences for each cell in the sample mouse shows a distribution of cells that are suppressed in New (Figure 2C). The calcium activity from another sample mouse over five days exposure to New follows a similar pattern of activity dynamics (Figure 2—figure supplement 1). Across all animals, suppression histograms of SOM activity over the five day protocol show a large initial suppression in New that diminishes over days of exposure (Figure 2D).

To quantify this suppression over time, we calculated a percent difference for each mouse by averaging all cells from each mouse and then calculated a grand mean for all mice on each day (Figure 2E). Indeed, SOM-int activity exhibited significant suppression in New that gradually decreased over days (mean suppression on Day 1=-30.4±3.46%, $p<0.001$). This decrease in suppression paralleled the increase in behavioral performance in New, as normalized to the average performance in Fam1-1 and -2 (Figure 2F, mean normalized performance on Day 1=0.557±0.114, $p<0.05$). (Note that behavioral metrics in Fig.2 are from 8 SOM-cre imaged mice, a partially overlapping set of the larger behavioral cohort of 9 mice in Fig. 1. Mice in Fig. 1 all had lick data, while not all mice in this group did.) These data show, on average, strong suppression of SOM-int upon exposure to a New environment, with recovery of activity over repeated exposures. At the same time, behavioral performance was initially impaired and then increased over time (Figure 2G-H). These data are consistent with inhibition suppression acting as a permissive gate for learning; as behavior improves, increased inhibition may serve to stabilize learning in the network.

A possible confound to this interpretation would be decreased locomotion in New. Many SOM-ints have activity that is positively correlated with locomotion, although there is a distinct population that is anti-correlated (Arriaga and Han, 2017). Thus, it is possible that decreased interneuron activity is due to decreased locomotion in New that gradually recovers over repeated exposure. While average locomotion was the same in New and Fam (Figure 1—figure supplement 1D), it remains possible that more nuanced changes in movement could result in decreased SOM-int activity. To more thoroughly explore this possibility, we made a general linear model (GLM) to predict each cell's fluorescence based on movement. Using total ball speed (including both forward and yaw components of ball rotation) to represent locomotion and the timing of reward deliveries from Fam1-1, we trained the model to predict each cell's calcium activity from behavioral data (Figure 3A). In 6 sample cells from a single mouse, modeled ΔF/F

was very similar to actual ΔF/F in Fam1-1, as measured by Root Mean Square (RMS) error. In New, the fit of modeled ΔF/F was much worse than the actual ΔF/F suggesting that changes in locomotion are insufficient to explain the decreased ΔF/F in New. Finally, in Fam1-2, modeled ΔF/F was well fit to actual ΔF/F, although not as well as in Fam1-1 (Figure 3—figure supplement 1). To quantify these changes, we compared the RMS error of the modeled ΔF/F between New and the average of Fam1-1 and -2 and found that it was significantly different on all five days suggesting that changes in locomotion were unlikely to be responsible for decreased ΔF/F in New (Figure 3B, $p<0.001$). Cell activity in Fam1-2, while recovered from the suppressed state seen in the New world, was often lower than activity seen in Fam1-1. This lower fluorescent level in Fam1-2 is apparent in the higher degree of model RMS error seen in this environment. Measuring the correlation between modeled activity and fluorescent activity provides a measure of how well the behavioral data used to train the model explains cell activity despite the observed lower levels of fluorescence which may be the result of residual suppression from exposure to the New world, or imaging related photo-bleaching. Modeled activity was well correlated in familiar environments and poorly correlated in the New world (Figure 3C, Day 1, $p<0.001$), while activity in the New world becomes more strongly correlated with model activity over the course of repeated days of exposure. These results suggest that exposure to a novel environment can at least partially decouple the positive correlation between locomotion and SOM-int activity.

Another potential interpretation of our data is that decreased SOM-int activity is driven by surprise at the context switch, with habituation to this surprise gradually restoring inhibitory activity. To test this possibility, we dissociated surprise at the context switch from learning in the new environment by replacing the New environment with a no-task, no-reward epoch and a static visual scene (black screen). Under these conditions, if surprise drives inhibition suppression and recovery is due to habituation, we would see the same suppression and recovery over time as previously shown. On the other hand, if learning is necessary for the recovery of inhibitory activity and learning is prevented, we should see sustained inhibitory suppression.

Taking advantage of the long-term recording stability of two-photon calcium imaging, we recorded the same cells from a subset of the mice ($N=4$, $n=69$) used in the remapping experiment, allowing us to directly compare the kinetics of inhibition recovery when learning was present or absent. The no task environment evoked strong suppression of SOM-int activity, similar to the suppression seen when mice are switched into New (Figure 4 A-D, Mean percent change on Day 1=-48.1±4.18%, $p<0.01$). However, over five days of exposure to the same no task environment, inhibition remained strongly suppressed (Figure 4E Figure 4—figure supplement 1, Spearman's $\rho=0.0188$ $p=0.939$). In marked contrast, the same cells from the same mice showed strong recovery of inhibition over five days of exposure to New (Figure 4E, Spearman's $\rho=0.589$, $p<0.01$). Thus, recovery of inhibition is unlikely to be the result of habituation to surprise and is more likely tied to learning in the new context.

Consistent inhibitory structure during learning

On average, SOM+ neurons are suppressed in a new environment (Fig. 2E), but the degree of activity suppression is heterogeneous across neurons, ranging from strong inhibition to moderate activation in individual cells (Figure 2D). This variability could represent noise, dependent on stochastic ensembles of activated pyramidal neurons in different environments, or could indicate that different SOM-ints have different, but persistent, network roles. These two models have very different implications for interneuronal roles during learning. The first suggests a pooled inhibition model where each interneuron is "equi-potential" and has no particular specialized role in the network. The second signifies that network inhibition has a

consistent structure, with the intriguing possibility that different interneurons play distinct functional roles.

We tested whether the structure of inhibition suppression was stochastic or consistent by putting a subset of the animals previously described through a second remapping protocol where they are exposed to another distinct and novel visual virtual environment, labeled “New 2,” with the original novel environment now labeled “New 1.” By recording the same cells across the two remapping protocols, we could correlate the magnitude of each cell’s activity suppression in New 1 vs. New 2. If SOM-ints are stochastically recruited by network activity, there should be no correlation in activity suppression across New 1 vs. New 2. However, we found strong correlation between activity suppression in New 1 vs. New 2 in individual SOM-ints. This correlation was strong on Day 1, and strikingly, this correlation was significant across all days of the remapping protocol (Figure 5A, Day 1, Pearson’s Correlation=0.697, $p<0.001$; Day 2, $r=0.668$, $p<0.001$; Day 3, $r=0.764$ $p<0.001$; Day 5, $r=0.583$ $p<0.001$). We have previously verified place cell global remapping across different virtual environments in this task (Arriaga and Han, 2017), strongly suggesting that this consistent functional inhibitory network structure occurs despite differing ensembles of activated pyramidal neurons.

A trivial explanation for these results could be that this correlation results from general similarities across the two behavioral epochs, such as having equivalent tasks or shared layout of the virtual worlds. To probe the structure of inhibition suppression in a drastically different context, we used the no task epoch described previously (Figure 4). Here the visual scene is different and static, and there is no behavioral task. Even here, when comparing activity suppression in New 1 vs. no task epochs, we found significant correlation for each cell on day 1, indicating that the functional inhibitory network structure for these two very different behavioral epochs is very similar. We also measured correlation in activity suppression throughout the 5 day protocol. While the correlation was significant on day 1 (Figure 5B, $r=0.464$ $p<0.001$), it was not for days 2 – 5 (Figure 5B, Day 2, $r=0.0892$, $p>0.999$; Day 3 $r=0.0259$, $p>0.999$; Day 4 $r=0.0159$, $p>0.999$; Day 5, $r=-0.0774$, $p>0.999$). This was not surprising because inhibitory activity in New 1 recovers as animals learn the task, while suppression remains strong with no learning in no task epochs.

The rapid time-dependent loss of correlation between New 1 and no task epochs reinforces how striking the correlation structure is for New 1 vs. New 2, indicating that not only is initial suppression correlated but there is also consistency in the temporal dynamics of this structure in time across five days (Figure 5C). Similarly, the mean difference in percent change for each cell between New 1 and New 2 remains stable across the remapping paradigm (Spearman’s $\rho=0.074$, $p=0.173$), whereas the difference between New 1 and the no task epoch steadily increases over the five day course of exposure to each environment (Figure 5D, Spearman’s $\rho=0.385$, $p<0.001$).

These results show that each cell exhibits a consistent degree of activity suppression during learning. To understand more about how this structure arises, we looked for other factors that were associated with each cell’s magnitude of activity suppression. First, we examined the recovery of inhibition in cells as a function of initial level of activity suppression on day 1 of remapping into New 1. Cells stratified by the magnitude of their initial suppression continue to be stratified by activity suppression across the five day remapping protocol, with the most strongly suppressed on day 1 remaining the most suppressed on day 5, while the least suppressed remain the least suppressed (Figure 5—figure supplement 1A). This finding indicates that each cell’s initial activity suppression is predictive of future activity throughout the protocol, suggesting that neurons may not be drawn from the same population of functionally homogeneous interneurons. Similarly, cells which were most suppressed in the No Task epoch remained the most suppressed through subsequent days of exposure to this epoch (Figure 5—figure supplement 1B).

Does the degree of activity suppression differentially map onto distinct SOM-int cell types? SOM-ints are primarily composed of two functionally and anatomically distinct types, OLM and bistratified interneurons (Royer et al., 2012). The somata of OLM neurons lie in *stratum oriens* (SO), while bistratified interneurons are mostly in *stratum pyramidale* (SP). We found no difference in activity suppression between cells with somata in SO vs. SP (Figure 5—figure supplement 1C, $p=0.380$). Baseline fluorescence (which may be used, with limitations, as an indicator of basal firing rate) also exhibited no relationship to activity suppression (Figure 5—figure supplement 1D Spearman's $\rho=-0.164$ $p=0.171$).

We previously identified two distinct populations of SOM-ints, one whose activity was positively correlated with locomotion and another whose activity was anti-correlated (Arriaga and Han, 2017). These two populations, as measured by phase angle of the hilbert transform of the cell's correlation between stop-triggered mean activity and running speed, were also present in these data. However, there was no difference in activity suppression between the two (Figure 5—figure supplement 1F, $p>0.99$). We also found no relationship between soma area and activity suppression (Figure 5—figure supplement 1E, $p>0.99$). Thus, the degree of activity suppression was not readily explained by previously identified cell classes nor by cellular properties.

Finally, we examined whether suppression could be predicted by each cell's degree of coupling to locomotion. SOM-int activity is modulated by locomotion with most cells having a positive correlation, but the degree of coupling ranges from strongly positive correlated to weakly negatively correlated (Pearson Correlation from -0.300 to 0.750). . To examine whether locomotion coupling to activity was correlated to each cell's magnitude of activity suppression, we quantified locomotion-dependence of activity with the GLM trained by locomotion in Fam1-1 (prior and independent of the New epoch). The RMS error of the model fit to the actual $\Delta F/F$ for each cell (Figure 3) is indicative of how well the GLM predicts each cell's activity based on locomotion, with low error indicating high predictability and high error indicating lower predictability. Surprisingly, we found strong anti-correlation between percent change in activity in New (Day 1) vs. RMS error (Figure 5 –figure supplement 1G, Spearman's $\rho=-0.299$, $p<0.001$). These results indicate that cells whose activity is only weakly locomotion-dependent are more suppressed in a novel environment.

Discussion

In this work we addressed critical questions in neuronal network function: how is inhibition dynamically regulated during learning, and is there a persistent structure in functional inhibitory activity dynamics? We found that the activity of SOM-ints was initially strongly suppressed upon exposure to a novel virtual environment. Interneuron activity gradually recovered as animals learned the goal-directed spatial navigation task in the second, initially novel, environment. Furthermore, there was a persistent inhibitory network structure in the transition from familiar to novel environments. Each interneuron exhibited a characteristic amount of activity suppression over multiple remappings as well as in a drastically different no-task environment where there was no relevant learning.

Our findings are consistent with a model in which a new environment is associated with decreased network inhibition which then gradually recovers over the course of learning to stabilize the network, possibly via synaptic plasticity (Cohen et al., 2017; Wilson and McNaughton, 1993). We tie inhibitory suppression to learning in two ways. First, using learning of a goal-directed spatial navigation task, we were able to produce prolonged suppression of SOM-ints. This prolonged SOM-int suppression overlapped with learning of the new environment, in particular reward location, as evidenced on behavioral measures including task performance, reward-associated locomotion, and licking patterns. Second, we tested the role of learning in recovery of inhibitory activity by providing an epoch with no task, and therefore no learning. We found that inhibition remained suppressed when animals are not engaging in spatial learning.

Two other noteworthy observations arise from this no task experiment. First, inhibition was strongly suppressed even though the animal was not learning, suggesting that inhibition suppression is characteristic of a network primed for learning, rather than an indication of novelty. Second, locomotion speed is significantly greater in the no task epoch on day 1 than in Fam. Based on the general positive correlation between movement speed and interneuron activity (Arriaga and Han, 2017) we would expect inhibitory activity to be greater in the no task epoch, yet it is far lower. This result, in combination with the decoupling of locomotion and activity during learning in the new environment, highlights the importance of behavioral state in controlling the output mode of neurons.

We have shown long-lived SOM-int suppression that is distinct from transient suppression of SOM-ints during place cell global remapping (Sheffield et al., 2017). How much inhibition suppression is associated with novelty-induced place cell global remapping and how much is due to learning *per se*? In both real world and VR experiments, switching animals to new environments triggers a few minutes of inhibition suppression and formation of new place cell maps on the same timescale (Barbieri et al., 2005; Hainmueller and Bartos, 2018; Leutgeb et al., 2004; Muller and Kubie, 1987; Sheffield and Dombeck, 2015). In contrast, in a learning task with no global remapping where freely moving rats learned new goal locations in a familiar environment, fast-spiking putative interneurons both increased and decreased activity as performance increased. This reflected dissolution of ensembles encoding old reward maps and formation of new cell assemblies, comprising pyramidal and interneurons, representing new reward locations. Taken together, it is likely that the initial SOM-int activity suppression in our experiments is triggered by the switch into a new context and associated global remapping while slowly increasing inhibition thereafter is associated with additional map refinement related to task-learning. While it is likely that learning shares significant common network mechanisms with place cell remapping, further experiments using more detailed simultaneous recordings of pyramidal and interneuronal activity will refine this picture, while more subtle environmental manipulations will help dissociate global remapping effects from learning.

Here, we described the activity of interneurons on a time-averaged basis, in line with traditional methods for studying neuronal tuning. Our data are consistent with a simple model where levels of inhibition may control learning in the pyramidal network. Neuronal activity is

dynamically modulated on faster timescales than the first order analysis presented here. Future work will focus on higher resolution investigations of interneuron activity, particularly in relationship to pyramidal neuron ensembles that are likely to encode learning.

Another critical contribution of this study was the discovery of stable inhibitory activity dynamics during learning, enabled by our stable, long-term recording of identified interneurons. We found inhibitory population structure, with some SOM-ints strongly suppressed during learning while others were unaffected. Surprisingly we found that these activity dynamics were stable on a cell-by-cell basis, with each interneuron having a consistent level of activity suppression, both across multiple new environments and in a no task epoch. These findings reveal an underlying inhibitory circuit structure that is observable when the animal is primed for learning.

Our finding that activity dynamics of interneurons is at least partially independent of activated pyramidal neurons argues against the commonly accepted model where interneurons are simple passive followers of pyramidal activity. Each interneuron's stability across learning episodes, where we expect stochastic ensembles of activated pyramidal neurons in the hippocampus, suggests that their activity dynamics can be set at least partially independently of excitatory drivers. This inhibitory framework instead suggests that inhibition shapes the encoding of information in the network. It is possible that pyramidal neurons downstream of strongly suppressed interneurons are more likely to express plasticity during learning, as a direct result of increased activity due to release of inhibition. Conversely pyramidal neurons downstream of less suppressed interneurons may see relatively normal levels of inhibition during learning, making their activity more stable. Thus, this inhibitory structure may serve to balance plasticity and stability in the network by defining functionally specialized ensembles of pyramidal neurons.

One prediction of this model is preferential or targeted connectivity in the outputs of SOM-ints onto pyramidal neurons. While such a scenario is at odds with the prevailing view of "pooled" or "blanket" inhibition where interneurons make promiscuous and non-selective synapses, significant evidence exists for preferential connectivity both in the hippocampus and cortex. In the hippocampus, PV-expressing basket cells preferentially inhibit deep pyramidal neurons projecting to the amygdala while also being more likely to receive excitation from superficial pyramidal neurons or deep pyramidal neurons projecting to the prefrontal cortex (Lee et al., 2014). In the medial entorhinal cortex, cholecystokinin-expressing basket cells selectively target pyramidal neurons that project extra-hippocampally. Furthermore, in the hippocampus, interneurons participate in cell assemblies with pyramidal neurons and can share coding properties such as place fields (Ego-Stengel and Wilson, 2007; Kubie et al., 1990; Marshall et al., 2002). Similarly, functional subnetworks of interneurons and pyramidal neurons have been identified in the cortex. Finally, this work identifying specialization of interneuron function is complemented by evidence for functionally distinct subsets of CA1 pyramidal neurons (Danielson et al., 2016; Graves et al., 2012; Mizuseki et al., 2011).

We identified structured activity dynamics in the functional responses of interneurons as animals learned a task in novel virtual environments. What mechanisms might generate the activity dynamics and structure within the interneuron population? Neuromodulatory transmitters targeting G-protein coupled receptors are likely to play a significant role. Novelty produces strong changes in neuromodulation, with sharp increases in acetylcholine (Acquas et al., 1996), norepinephrine (Sara et al., 1995), and dopamine (Kempadoo et al., 2016; Takeuchi et al., 2016), among others. Furthermore, many of these neuromodulatory systems are necessary for the formation of new memories (Atherton et al., 2015; Bick and Eskandar, 2016). Differing levels of inhibitory activity suppression could be set by expression levels of neuromodulatory receptors in each cell. For example, interneurons show markedly divergent responses to acetylcholine depending on their composition and expression of receptors (McQuiston and Madison, 1999).

Another possible mechanism for suppressing inhibition is disinhibitory connections from other interneurons targeting SOM-ints. In the hippocampus, this disinhibitory input can come from local VIP, PV, and SOM interneurons (Francavilla et al., 2015). Indeed, in our experiments some SOM-ints were activated, although it remains unclear if these interneurons provide disinhibitory input. VIP interneurons are strongly associated with disinhibition and previous work showed that these neurons are necessary for hippocampal-dependent learning. Furthermore many forms of cortical learning depend on VIP interneuron function (Pi et al., 2013).

Finally, it is possible that decreased SOM-int activity is inherited from reduced upstream excitatory input. In this case feed-forward inhibition is driven by EC and CA2/3 input while feed-back excitation is driven by local CA1 neurons. However, during learning or novelty, CA3 and EC pyramidal neurons don't change firing rates, while CA1 pyramidal neurons increase activity (Csicsvari et al., 2007) making it unlikely that inhibitory suppression is purely a function of reduced excitatory drive.

Materials and Methods:

Animals. All experiments were approved by the Washington University Animal Care and Use Committee. Heterozygotes (+/-) from two cre-driver mice lines on a C57Bl/6J genetic background were used to label parvalbumin-expressing and somatostatin-expressing inhibitory interneurons: SST$^{tm2.1(cre)Zjh}$/J (SOM-cre) and Pvalb$^{tm1(cre)Arbr}$/J (PV-cre; Jackson Labs). All imaging data were from SOM-ints while behavioral data in Figure 1 were from PV- and SOM-ints.
Viral Injections and hippocampal window implantation. Surgical procedures, VR track running behavior, and two-photon imaging have been described previously (Arriaga and Han 2017). Briefly, mice were injected with adeno-associated virus (AAV) at 2-4 months of age. Mice were anesthetized with isoflurane, and a small (.5mm) craniotomy was opened above the left cortex. Virus (AAV1.Syn.Flex.GCaMP6f.WPRE.SV40, Penn Vector Core, University of Pennsylvania, 1.71 x 10^{13} genome copies, diluted 1:1–1:4 with PBS, ~50nL total volume) was pressure injected through a beveled micro-pipette targeting CA1 (-1.8 ML, -2.0 AP, -1.3 DV).
After virus injection, mice were water-scheduled for 1-3 weeks and an imaging cannula (2.8 mm diameter) was implanted above the hippocampus by aspirating the overlying cortex. Mice recovered for at least two weeks after surgery before beginning training.
VR track running behavior. The virtual reality display used a custom-built semi-cylindrical projection screen (1 ft radius) and two rear projectors (Optima 750ST). Projection screen was ~12 inches in front of the mouse and occupied 180º of horizontal, 16º below the horizon and 35º above. The mouse was head-fixed on a spherical Styrofoam treadmill supported on a cushion of air from a 3D printed base which allowed free ball rotation with mouse locomotion. Treadmill movement was tracked using a G400 Logitech mouse configured in LabView (National Instruments). The VR environment was rendered using ViRMEn (Virtual Reality Matlab Engine; Aronov and Tank, 2014). Mice were trained to run to alternating ends of a linear VR track (180 cm) for 2-5 weeks until they consistently achieved target performance (>2 rewards/min for one week). After training, mice were imaged during exposure to a new visual virtual world. Remapping experiments consisted of 7 min behavior in the familiar track (Fam1-1), an instantaneous switch to a novel track of the same length with different visual textures and landmarks (New) for 14 min, and then a return to the familiar (Fam1-2) for 7 min. This remapping protocol was repeated for 5 successive days with the same, decreasingly novel, New world. In a subset of animals, a second remapping task was also performed. This task was identical to the first with the exception of a different New environment (New 2). Additionally, this same subset of animals was imaged in a No Task session. This session consisted of 7 min of navigation in the familiar track, 14 min of exposure to a dark screen with no rewards, and a return to the initial familiar environment for 7 min.
Two-Photon Imaging. Calcium imaging was performed on a Neurolabware laser-scanning two-photon microscope, with the addition of an electric tunable lens (ETL; Optotune, EL-10-30-NIR-LD) and f=-100 mm offset lens to rapidly change axial focal length. We imaged 4-6 axial planes spanning up to 250 µm in the z-axis at a total frame rate of 31Hz, resulting in a per plane sampling rate of 5.2Hz for a 6 plane recording and 7.8Hz for a 4 plane recording. Field of view in *x-y* was 500 x 500µm.
Data Analysis. Data were analyzed using custom programs written in Matlab (RRID:SCR_001622). Images were motion-corrected using cross-correlation registration and rigid translation of individual frames. Slow fluctuations in fluorescence were removed from calculations of $\Delta F/F_0$ by calculating F_0 using the eighth percentile of fluorescence intensity from a sliding window 300 s around each time point. ROIs were selected using a semi-automated

process. Possible ROIs were identified as contiguous regions with SD>1.5 and an area >90 μm^2. Overlapping ROIs were manually separated, ROIs were redrawn by hand to separate adjacent cells into distinct ROIs. Unresponsive puncta, or those with low signal-to-noise ratios (initially identified as having a skewness of ΔF/F in the first familiar environment less than 0.3) were dropped from further analysis. When the same cell was recorded in multiple planes, the brightest ROI was used. Neuropil contamination was removed by subtracting a perisomatic fluorescence signal from an annulus between 5 and 20 µm from each ROI, excluding any other possible ROIs ($F_{Corrected\text{-}ROI}=F_{ROI} - .8 * F_{Neuropil}$).
The percent change in the New environment was calculated on each day for each cell as the ratio between the mean fluorescence in the 14 min New world exposure and the mean of the fluorescence from the two 7 min familiar worlds exposures, normalized by the sum of these means $Percent\ Change = 1 - \frac{\mu_{New}}{\mu_{(Fam1-1+Fam1-2)}}$.
Behavior Analysis. Ball movement data, sampled at 1kHz, was downsampled to match the imaging frame rate. All normalized behavioral metrics were normalized by taking the ratio of the metric in the New world to the mean of the metric in the Fam1-1 and 1-2. Task performance was calculated as the rewards received per minute. Speed was calculated as the Euclidean sum of forward and yaw components of ball velocity. Deceleration was calculated as the first derivative of the forward component of the ball speed during a three second window prior to reward. Location of trial failure was identified as the distance from the correct destination end zone where the animal turns around before reaching the end zone.
Lick behavior was detected using a 2-transistor lick detection circuit (Slotnick 2009). Individual licks were not well isolated, so lick responses were binned into periods of repeated licking termed bouts. A bout of licking was defined as a period of repeated lick responses with than 200 ms between repeated lick signals. The lick rate was calculated as the number of these licking bouts per minute, normalized to the mean rate of licking bouts in the familiar worlds. The fraction of correct licks was calculated as the fraction of licking bouts which overlap with or occur within 500 milliseconds of reward delivery. The fraction of licks in unrewarded end zones was calculated as the fraction of incorrect, unrewarded, entries into the track end zone which elicited a bout of licking.
General Linear Model of Activity. A general linear model was used to estimate fluorescence as a function of the behavioral parameters which are correlated with cell activity. The model predicts fluorescence as the linear combination of weighted, time-lagged behavior components. Modeling of interneuron fluorescence was done using the *glmfit* function in Matlab with a normal distribution and an identity link function. The oscillatory nature of interneuron fluorescent activity series, without the large transients typical in pyramidal cells, were better fit using a normal distribution than the Poisson distribution commonly used in generalized linear models of pyramidal cell activity. Models were trained using fluorescence data from Fam1-1 and fit using the forward and rotation components of running speed and the timing of rewards. Root mean square (RMS) error in Fam1-1 was calculated using 10-fold cross validation, successive models were trained on 9/10 of the data set and tested on 1/10 of the data, the average RMS error across these ten sessions was used as the error from Fam1-1. Error in other sessions was calculated by calculating the residuals from applying the model fits from Fam1-1 to speed and reward data from subsequent epochs. Correlation was measured as the Pearson correlation between modeled traces and ΔF/F in each context.
Experimental Design and Statistical Analysis. Behavior data is reported from 9 mice (5 male, 4 female). We recorded 162 somatostatin-cre positive cells (mean=20.3+/-4.9) from eight mice (6

male, 2 female). For the second remapping and no task paradigms we recorded from 4 of these mice (3 male, 1 female), tracking 69 cells across all 3 contexts.

Significance of normalized data metrics were calculated using 1 sample t-tests of mean values with Bonferroni-Holm corrections. Significance of differences between familiar and new epochs were calculated using paired sample t-tests of mean values with Bonferroni-Holm corrections. Significance of RMS error values were calculated on each cell in each epoch with paired-sample t-tests with Bonferroni-Holm corrections. Correlations of percent change across paradigms were calculated using Pearson Correlation, corrected with Bonferroni-Holm corrections.

Statistical analyses were performed in Matlab.

References

Acquas, E., Wilson, C., and Fibiger, H.C. (1996). Conditioned and Unconditioned Stimuli Increase Frontal Cortical and Hippocampal Acetylcholine Release: Effects of Novelty, Habituation, and Fear. The Journal of Neuroscience *16*, 3089–3096.

Andersen, P., Eccles, J.C., and LøYning, Y. (1963). Hippocampus of the Brain: Recurrent Inhibition in the Hippocampus with Identification of the Inhibitory Cell and its Synapses. Nature *198*, 540.

Arriaga, M., and Han, E.B. (2017). Dedicated hippocampal inhibitory networks for locomotion and immobility. J. Neurosci. 1076–17.

Artola, A., and Singer, W. (1987). Long-term potentiation and NMDA receptors in rat visual cortex. Nature *330*, 649.

Atherton, L.A., Dupret, D., and Mellor, J.R. (2015). Memory trace replay: the shaping of memory consolidation by neuromodulation. Trends in Neurosciences *38*, 560–570.

Barbieri, R., Wilson, M.A., Frank, L.M., and Brown, E.N. (2005). An analysis of hippocampal spatio-temporal representations using a Bayesian algorithm for neural spike train decoding. IEEE Transactions on Neural Systems and Rehabilitation Engineering *13*, 131–136.

Bear, M.F., Press, W.A., and Connors, B.W. (1992). Long-term potentiation in slices of kitten visual cortex and the effects of NMDA receptor blockade. Journal of Neurophysiology *67*, 841–851.

Bick, S.K.B., and Eskandar, E.N. (2016). Neuromodulation for restoring memory. Neurosurgical Focus *40*, E5.

Bittner, K.C., Grienberger, C., Vaidya, S.P., Milstein, A.D., Macklin, J.J., Suh, J., Tonegawa, S., and Magee, J.C. (2015). Conjunctive input processing drives feature selectivity in hippocampal CA1 neurons. Nat Neurosci *18*, 1133–1142.

Bittner, K.C., Milstein, A.D., Grienberger, C., Romani, S., and Magee, J.C. (2017). Behavioral time scale synaptic plasticity underlies CA1 place fields. Science *357*, 1033–1036.

Bliss, T.V.P., and Lømo, T. (1973). Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. The Journal of Physiology *232*, 331–356.

Chen, T.-W., Wardill, T.J., Sun, Y., Pulver, S.R., Renninger, S.L., Baohan, A., Schreiter, E.R., Kerr, R.A., Orger, M.B., Jayaraman, V., et al. (2013). Ultrasensitive fluorescent proteins for imaging neuronal activity. Nature *499*, 295–300.

Cichon, J., and Gan, W.-B. (2015). Branch-specific dendritic Ca2+ spikes cause persistent synaptic plasticity. Nature *520*, 180–185.

Cohen, J.D., Bolstad, M., and Lee, A.K. (2017). Experience-dependent shaping of hippocampal CA1 intracellular activity in novel and familiar environments. ELife *6*, e23040.

Csicsvari, J., O'Neill, J., Allen, K., and Senior, T. (2007). Place-selective firing contributes to the reverse-order reactivation of CA1 pyramidal cells during sharp waves in open-field exploration. European Journal of Neuroscience *26*, 704–716.

Dana, H., Mohar, B., Sun, Y., Narayan, S., Gordus, A., Hasseman, J.P., Tsegaye, G., Holt, G.T., Hu, A., Walpita, D., et al. (2016). Sensitive red protein calcium indicators for imaging neural activity. ELife *5*, e12727.

Danielson, N.B., Zaremba, J.D., Kaifosh, P., Bowler, J., Ladow, M., and Losonczy, A. (2016). Sublayer-Specific Coding Dynamics during Spatial Navigation and Learning in Hippocampal Area CA1. Neuron *0*.

Donato, F., Rompani, S.B., and Caroni, P. (2013). Parvalbumin-expressing basket-cell network plasticity induced by experience regulates adult learning. Nature *504*, 272–276.

Dupret, D., O'Neill, J., and Csicsvari, J. (2013). Dynamic Reconfiguration of Hippocampal Interneuron Circuits during Spatial Learning. Neuron *78*, 166–180.

Ego-Stengel, V., and Wilson, M.A. (2007). Spatial selectivity and theta phase precession in CA1 interneurons. Hippocampus *17*, 161–174.

Francavilla, R., Luo, X., Magnin, E., Tyan, L., and Topolnik, L. (2015). Coordination of dendritic inhibition through local disinhibitory circuits. Frontiers in Synaptic Neuroscience *7*.

Frank, L.M., Stanley, G.B., and Brown, E.N. (2004). Hippocampal Plasticity across Multiple Days of Exposure to Novel Environments 10.1523/JNEUROSCI.1958-04.2004. J. Neurosci. *24*, 7681–7689.

Gentet, L.J., Kremer, Y., Taniguchi, H., Huang, Z.J., Staiger, J.F., and Petersen, C.C.H. (2012). Unique functional properties of somatostatin-expressing GABAergic neurons in mouse barrel cortex. Nature Neuroscience *15*, 607.

Graves, A.R., Moore, S.J., Bloss, E.B., Mensh, B.D., Kath, W.L., and Spruston, N. (2012). Hippocampal pyramidal neurons comprise two distinct cell types that are countermodulated by metabotropic receptors. Neuron *76*, 776–789.

Hainmueller, T., and Bartos, M. (2018). Parallel emergence of stable and dynamic memory engrams in the hippocampus. Nature *558*, 292.

Harvey, C.D., Collman, F., Dombeck, D.A., and Tank, D.W. (2009). Intracellular dynamics of hippocampal place cells during virtual navigation. Nature *461*, 941--946.

Kempadoo, K.A., Mosharov, E.V., Choi, S.J., Sulzer, D., and Kandel, E.R. (2016). Dopamine release from the locus coeruleus to the dorsal hippocampus promotes spatial learning and memory. Proc. Natl. Acad. Sci. U.S.A. *113*, 14835–14840.

Kim, S.M., and Frank, L.M. (2009). Hippocampal lesions impair rapid learning of a continuous spatial alternation task. PLoS ONE *4*, e5494.

Klausberger, T., and Somogyi, P. (2008). Neuronal Diversity and Temporal Dynamics: The Unity of Hippocampal Circuit Operations. Science *321*, 53–57.

Kubie, J.L., Muller, R.U., and Bostock, E. (1990). Spatial firing properties of hippocampal theta cells. J. Neurosci. *10*, 1110–1123.

Larkum, M.E., Kaiser, K.M., and Sakmann, B. (1999). Calcium electrogenesis in distal apical dendrites of layer 5 pyramidal cells at a critical frequency of back-propagating action potentials. Proc. Natl. Acad. Sci. U.S.A. *96*, 14600–14604.

Lee, S., Kruglikov, I., Huang, Z.J., Fishell, G., and Rudy, B. (2013). A disinhibitory circuit mediates motor integration in the somatosensory cortex. Nat Neurosci *16*, 1662–1670.

Lee, S.-H., Marchionni, I., Bezaire, M., Varga, C., Danielson, N., Lovett-Barron, M., Losonczy, A., and Soltesz, I. (2014). Parvalbumin-Positive Basket Cells Differentiate among Hippocampal Pyramidal Cells. Neuron *82*, 1129–1144.

Letzkus, J.J., Wolff, S.B.E., Meyer, E.M.M., Tovote, P., Courtin, J., Herry, C., and Lüthi, A. (2011). A disinhibitory microcircuit for associative fear learning in the auditory cortex. Nature *480*, 331–335.

Leutgeb, S., Leutgeb, J.K., Treves, A., Moser, M.-B., and Moser, E.I. (2004). Distinct Ensemble Codes in Hippocampal Areas CA3 and CA1. Science *305*, 1295–1298.

Lovett-Barron, M., Kaifosh, P., Kheirbek, M.A., Danielson, N., Zaremba, J.D., Reardon, T.R., Turi, G.F., Hen, R., Zemelman, B.V., and Losonczy, A. (2014). Dendritic Inhibition in the Hippocampus Supports Fear Learning. Science *343*, 857–863.

Marshall, L., Henze, D.A., Hirase, H., Leinekugel, X., Dragoi, G., and Buzsáki, G. (2002). Hippocampal pyramidal cell-interneuron spike transmission is frequency dependent and responsible for place modulation of interneuron discharge. J. Neurosci. *22*, RC197.

McQuiston, A.R., and Madison, D.V. (1999). Muscarinic receptor activity induces an afterdepolarization in a subpopulation of hippocampal CA1 interneurons. J. Neurosci. *19*, 5703–5710.

Mizuseki, K., Diba, K., Pastalkova, E., and Buzsáki, G. (2011). Hippocampal CA1 pyramidal cells form functionally distinct sublayers. Nat Neurosci *14*, 1174–1181.

Muller, R.U., and Kubie, J.L. (1987). The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. J. Neurosci. *7*, 1951–1968.

Nitz, D., and McNaughton, B. (2004). Differential Modulation of CA1 and Dentate Gyrus Interneurons During Exploration of Novel Environments. Journal of Neurophysiology *91*, 863–872.

Pelkey, K.A., Chittajallu, R., Craig, M.T., Tricoire, L., Wester, J.C., and McBain, C.J. (2017). Hippocampal GABAergic Inhibitory Interneurons. Physiological Reviews *97*, 1619–1747.

Pi, H.-J., Hangya, B., Kvitsiani, D., Sanders, J.I., Huang, Z.J., and Kepecs, A. (2013). Cortical interneurons that specialize in disinhibitory control. Nature *503*, 521–524.

Royer, S., Zemelman, B.V., Losonczy, A., Kim, J., Chance, F., Magee, J.C., and Buzsáki, G. (2012). Control of timing, rate and bursts of hippocampal place cells by dendritic and somatic inhibition. Nat Neurosci *15*, 769–775.

Ruediger, S., Vittori, C., Bednarek, E., Genoud, C., Strata, P., Sacchetti, B., and Caroni, P. (2011). Learning-related feedforward inhibitory connectivity growth required for memory precision. Nature *473*, 514–518.

Sara, S.J., Dyon-Laurent, C., and Hervé, A. (1995). Novelty seeking behavior in the rat is dependent upon the integrity of the noradrenergic system. Cognitive Brain Research *2*, 181–187.

Sheffield, M.E.J., and Dombeck, D.A. (2015). Calcium transient prevalence across the dendritic arbour predicts place field properties. Nature *517*, 200–204.

Sheffield, M.E.J., Adoff, M.D., and Dombeck, D.A. (2017). Increased Prevalence of Calcium Transients across the Dendritic Arbor during Place Field Formation. Neuron *96*, 490-504.e5.

Stefanelli, T., Bertollini, C., Lüscher, C., Muller, D., and Mendez, P. (2016). Hippocampal Somatostatin Interneurons Control the Size of Neuronal Memory Ensembles. Neuron *89*, 1074–1085.

Steward, D.S. (1990). Interdisciplinary graduate programs which require academic and professional perspectives. Behavioral Sciences & the Law *8*, 291–300.

Takeuchi, T., Duszkiewicz, A.J., Sonneborn, A., Spooner, P.A., Yamasaki, M., Watanabe, M., Smith, C.C., Fernández, G., Deisseroth, K., Greene, R.W., et al. (2016). Locus coeruleus and dopaminergic consolidation of everyday memory. Nature *537*, 357–362.

Whitlock, J.R., Heynen, A.J., Shuler, M.G., and Bear, M.F. (2006). Learning induces long-term potentiation in the hippocampus. Science *313*, 1093–1097.

Wigström, H., and Gustafsson, B. (1983). Large long-lasting potentiation in the dentate gyrus in vitro during blockade of inhibition. Brain Research *275*, 153–158.

Wilson, M.A., and McNaughton, B.L. (1993). Dynamics of the hippocampal ensemble code for space. Science *261*, 1055–1058.

Wolff, S.B.E., Gründemann, J., Tovote, P., Krabbe, S., Jacobson, G.A., Müller, C., Herry, C., Ehrlich, I., Friedrich, R.W., Letzkus, J.J., et al. (2014). Amygdala interneuron subtypes control fear learning through disinhibition. Nature *509*, 453–458.

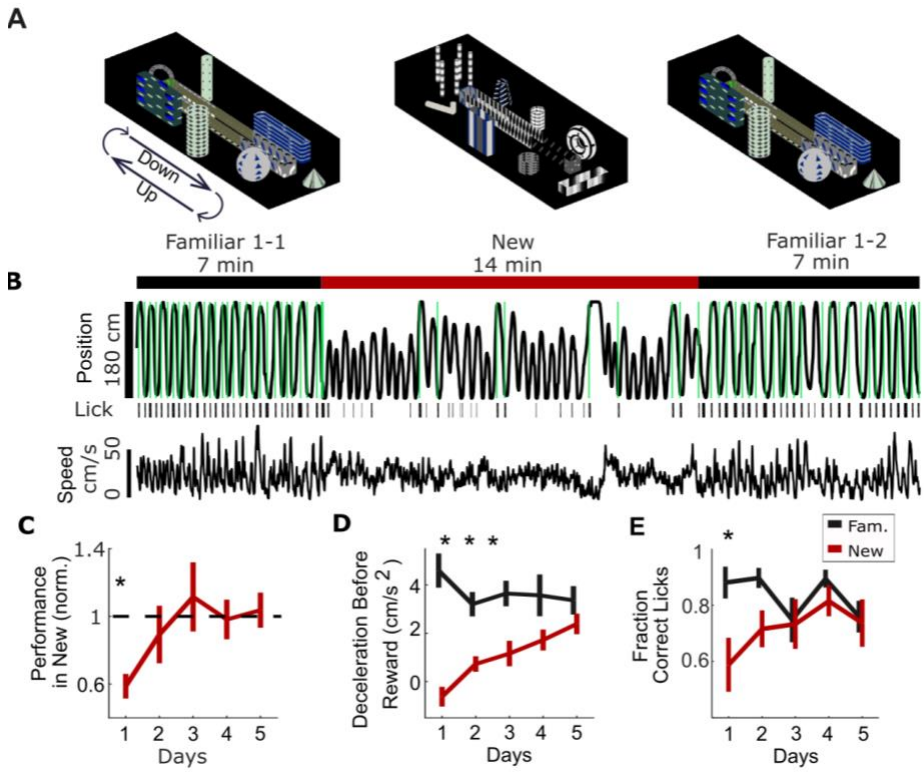

Figure 1. Learning in New Virtual Reality Environments. A, Overview of virtual reality task. Mice run to alternating ends of the VR track to receive water rewards. Animals spend 7 min in a familiar world (Familiar 1-1), which is instantaneously replaced with a new world (New), for 14 min before returning to the same familiar world (Familiar 1-2). **B,** Sample mouse behavior on first exposure to New showing high performance in flanking Familiar epochs and poor performance in New. Mouse position in VR track and timing of water rewards (green transients). Timing of bouts of licking water spout (dark bars are periods of licking). Ball speed during VR task. **C,** Mice were initially impaired in behavioral performance in the new environment but improved over days. Mean performance, measured in rewards per minute, in New world over 5 days of remapping. Normalized to mean reward rate in both familiar environments that day. **D,** Mice slow down in anticipation of reward in the familiar environment, measured as deceleration in the 3s window before reward. Over 5 days they begin to correctly anticipate rewards in the New world. **E,** Fewer licks in the New world are tied to reward delivery than in the familiar environment on day 1, this difference decrease over the 5 days of remapping. Fraction of bouts of licking which overlap with the delivery of a reward across five days of remapping. (N=9)

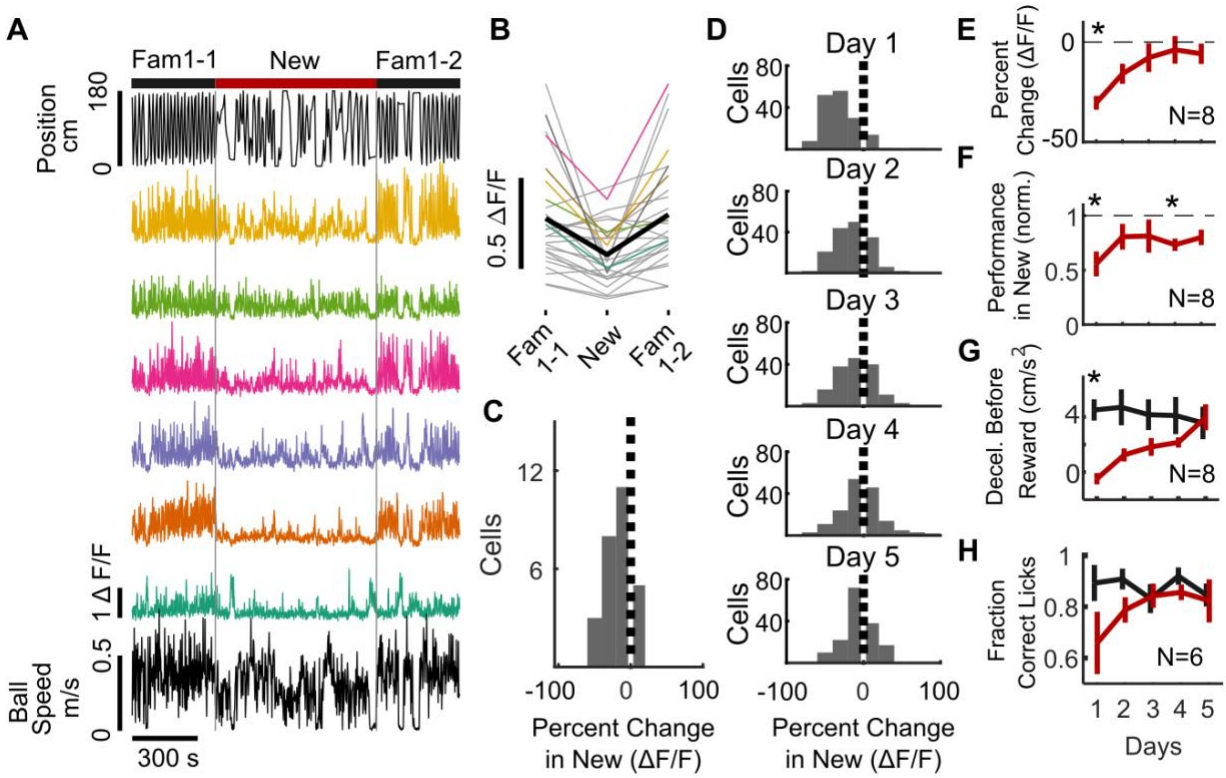

Figure 2. SOM interneuron activity suppression in new environments. A, SOM-ints are suppressed on first day exposure to New environment. Top, position in VR track, middle, ΔF/F of sample cells, bottom, ball speed. **B,** Mean ΔF/F of all cells from example mouse on day 1 of remapping showing general activity suppression (mean of all cells in black). **C,** Continuum of activity changes in New relative to Familiar, shown as histogram of percent change for cells from example mouse on day 1 of remapping. (n=28). **D,** Magnitude of suppression gradually decreases over days of exposure to New world. (N=8 mice, n=162 cells) **E,** SOM-int activity is initially suppressed but recovers over days of exposure to New. **F,** Performance in New world increases over several days. **G,** Mice increasingly slow down prior to a reward suggesting re-learning of goal location in New. (N=8) **H,** Fewer licks in the New world are tied to reward delivery than in the familiar environment on day 1, this difference decrease over the 5 days of remapping. Fraction of bouts of licking which overlap with the delivery of a reward across five days of remapping. (N=6)

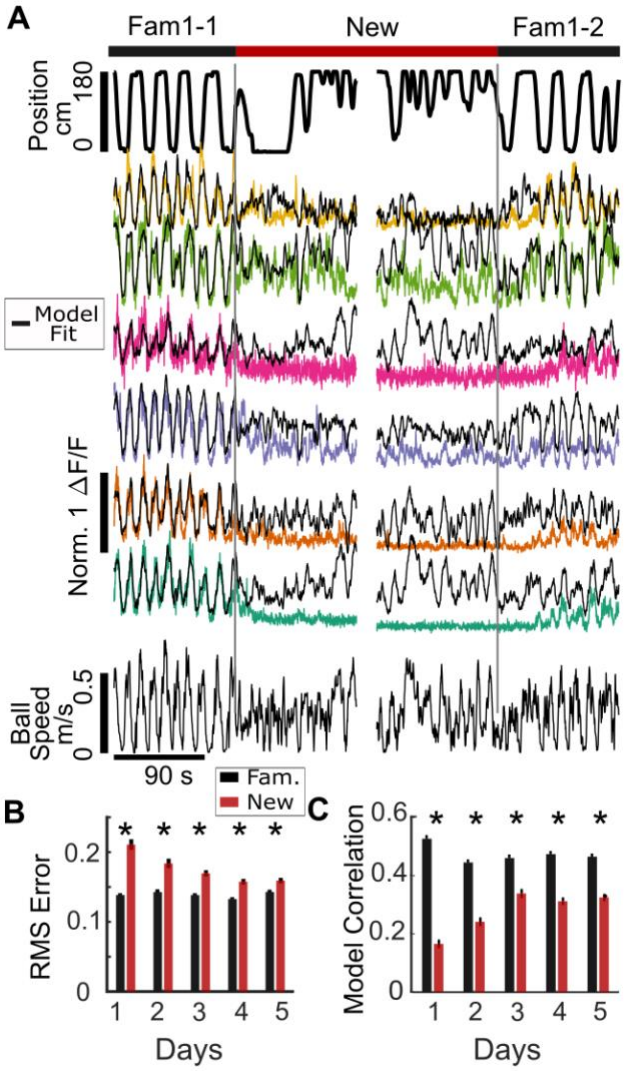

Figure 3. Decreased inhibition is not due to altered locomotion. **A,** A general linear model (GLM) was trained as a function of locomotion and rewards in the familiar environment for SOM+ cells. Modeled ΔF/F (black) is larger than actual ΔF/F (colored traces) in the novel environment, indicating that suppression of activity is greater than predicted from locomotion. **B,** Average Root Mean Square Error between modeled fluorescence and actual cell fluorescence in Fam and New. **C**, Average correlation between modeled fluorescence and cell fluorescence across environments and days. (N=8, n=162).

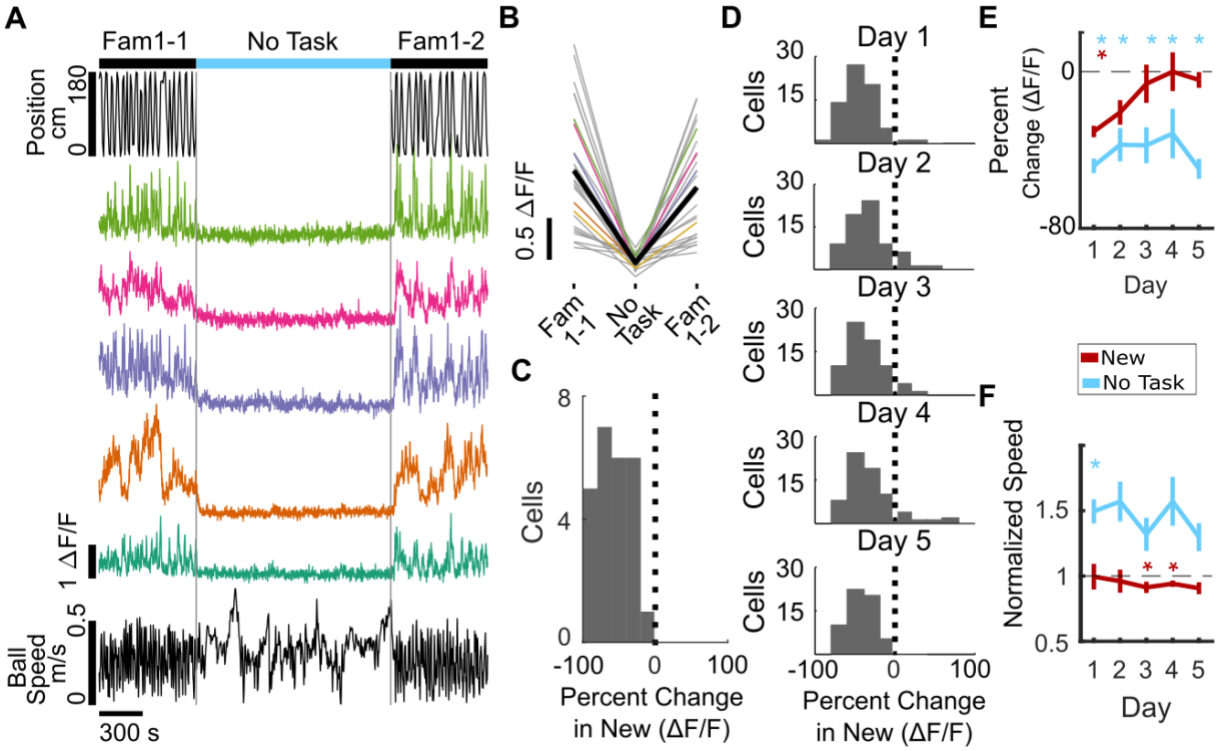

Figure 4. SOM-int Activity Suppression Does Not Habituate in No Task Environment. A, Cells from sample mouse are strongly suppressed in No Task epoch (static black screen). Top, position in VR track, middle, ΔF/F of sample cells, bottom, ball speed. **B,** Mean ΔF/F of all cells from example mouse on Day 1 of No Task exposure showing activity suppression (mean of all cells in black). **C,** Histogram of cell activity index from example mouse on day 1 of remapping showing strong suppression. (n=18). **D,** Interneurons remain suppressed over several days of No Task exposure. Histogram of percent change of cells. **E,** In No Task exposure, SOM-ints remain suppressed in contrast to recovery during exposure to New. **F,** Average speed in No Task environment increases relative to Familiar, in contrast to same or decreased speed in New. (N=4, n=69)

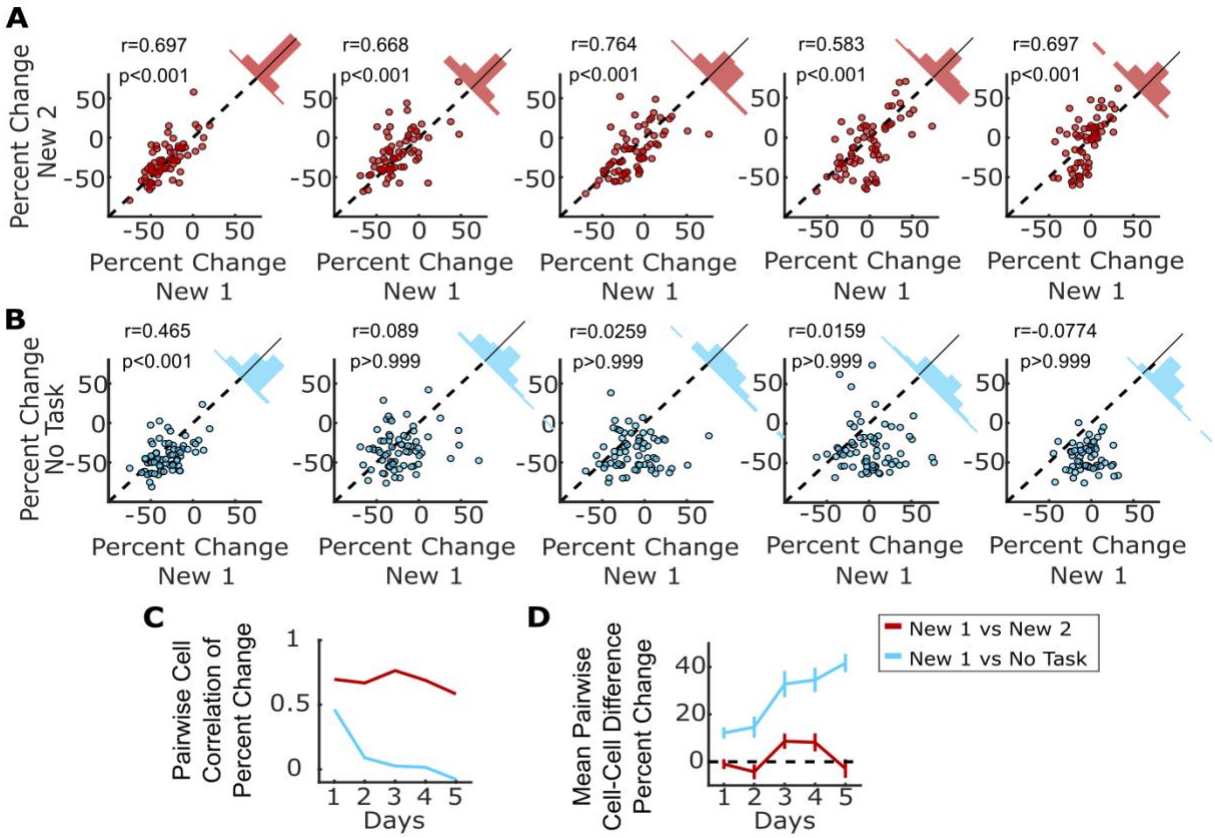

Figure 5. Consistent SOM-int Activity Responses Across Different New Environments and In No Task Epoch. A, Individual SOM-ints show correlation in activity modulation in two distinct New environments. **B,** A similar correlation of activity modulation is seen between New 1 and No Task exposures for Day 1. On subsequent days, correlation disappears because SOM-int activity begins to return in New 1 while remaining suppressed in No Task. **C,** Correlation between percent change of cells between two remapping sessions or remapping session 1 and no task exposure session. **D,** Mean difference in percent change in activity in cells between remapping and no task exposure settings (N=4, n=69).

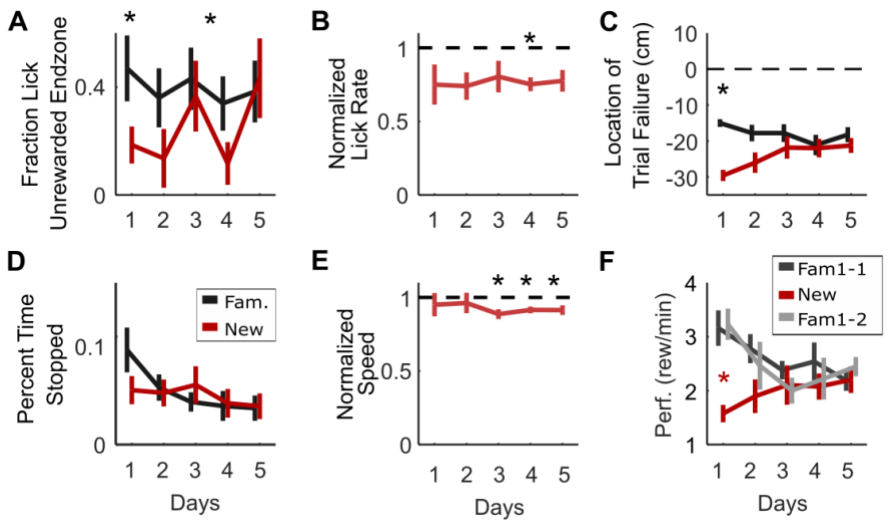

Figure 1-figure supplement 1. Behavior Metrics in New World. A, Mice are significantly more likely to lick in expectation of a reward upon incorrectly entering an endzone in the familiar environment. This difference decreases over the course of exposure to the New world. Fraction of incorrect, unrewarded, entries into a track endzone which coincide with a bout of licking across five days of remapping. **B,** Mice lick at similar rates in both the familiar environment and the New world. Frequency of bouts of licking in new world across five days of remapping, normalized to mean rate of licking in familiar worlds. **C,** Failed trials in the New world are farther from correct end zones than in the familiar environment. This difference decreases over 5 days of remapping. Mean distance from endzone of missed trial, identified as the location of a premature turn in the track. **D,** Mice spend similar amount of time stopped in both environments over 5 days. **E,** Average running speed is similar in both the familiar environment and the New world. Running speed in new world across five days of remapping, normalized to mean speed in familiar worlds. **F,** Mice were initially impaired in behavioral performance in the new environment but improved over days. Mean performance, measured in rewards per minute, in New world over 5 days of remapping. (N=9)

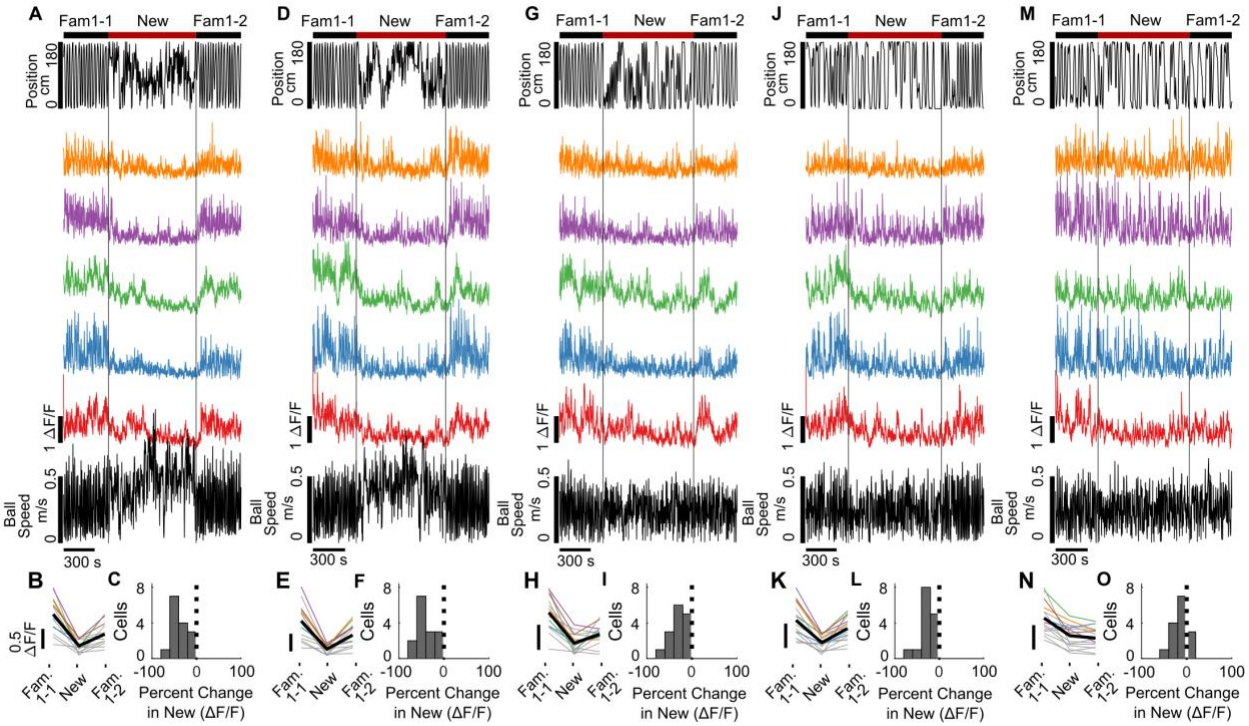

Figure 2-figure supplement 1. SOM interneuron novelty evoked suppression over five days of remapping

A, Top, position in VR track, middle, ΔF/F of sample cells, bottom, ball speed. **B,** Interneurons are markedly suppressed upon exposure to New world. Mean ΔF/F of all cells from example mouse on day 1 of remapping. **C,** Histogram of cell activity index from example mouse on day 1 of remapping (n=25).

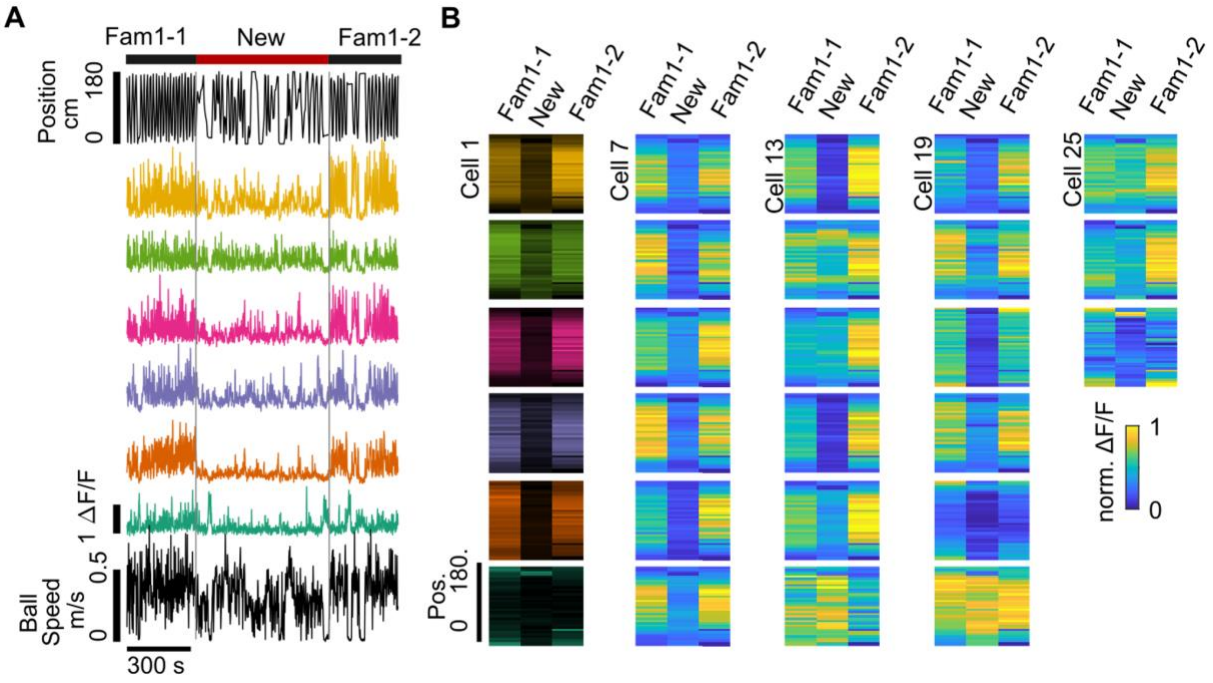

Figure 2-figure supplement 2. Cell firing fields A, Top, position in VR track, middle, ΔF/F of sample cells, bottom, ball speed. **B,** Spatial specificity changes in New World. Heatmaps of location specificity in each world on day 1 of remapping for the 28 cells in this example mouse.

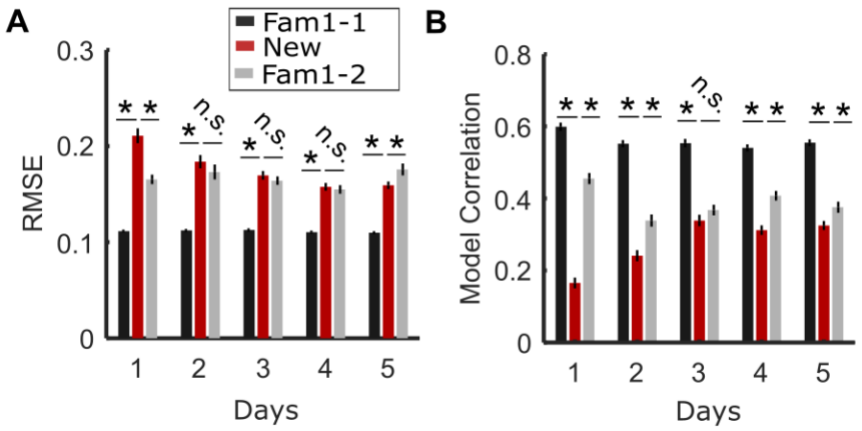

Figure 3-figure supplement 1. GLM performance in different environments

A, Average root mean square error in virtual environments. **B**, Average correlation between modeled fluorescence and cell fluorescence across environments and days. (N=8, n=162).

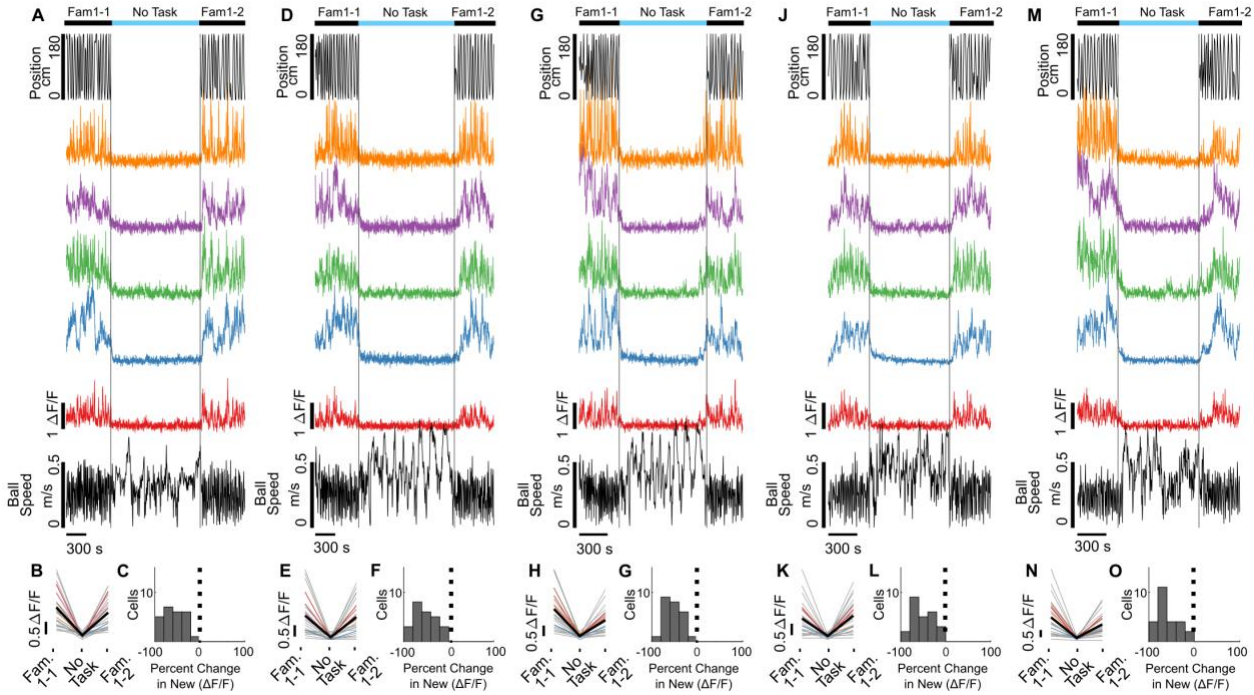

Figure 4-figure supplement 1. SOM interneuron novelty evoked suppression over five days of no task

A, Top, position in VR track, middle, ΔF/F of sample cells, bottom, ball speed. **B,** Mean ΔF/F of all cells from example mouse on day 1 of no task environment. **C,** Histogram of cell activity index from example mouse on day 1 of remapping. (n=24).

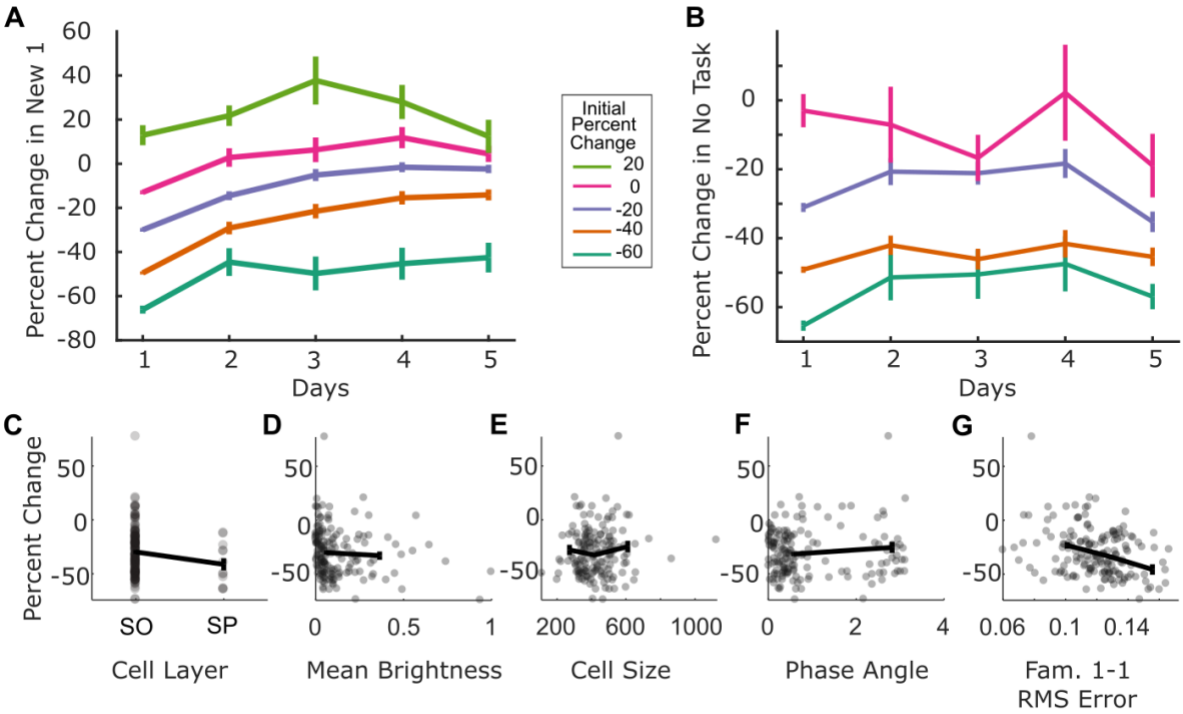

Figure 5-figure supplement 1: Organization of inhibition. A, Percent change in activity in new world, stratified by degree of suppression on day 1. **B,** Percent change in activity in no task session, stratified by degree of suppression on day 1. **C,** Percent change of activity in on day 1 of new world 1 organized by anatomic location of cell, stratum oriens (SO) or stratum pyramidale (SP). **D,** Percent change of activity in on day 1 of new world 1 organized by mean brightness. **E,** Percent change of activity in on day 1 of new world 1 organized by cell size. **F,** Percent change of activity in on day 1 of new world 1 organized by locomotion related activity in familiar 1-1. **G,** Percent change of activity in on day 1 of new world 1 organized by RMS error of LM regressed fluorescent activity in familiar 1-1. (N=8, n=162).