Cell-Type Specific Responses to Associative Learning in the Primary Motor Cortex

Candice Lee¹, Emerson Harkin¹, Xuming Yin¹, Richard Naud¹⁻⁴, Simon X. Chen^{1,3,4†}

¹Department of Cellular and Molecular Medicine, University of Ottawa, Ottawa K1H 8M5, Canada

²Department of Physics, STEM Complex, University of Ottawa, Ottawa, ON K1N 6N5, Canada

³Brain and Mind Research Institute, University of Ottawa, Ottawa K1H 8M5, Canada

⁴Center for Neural Dynamics, University of Ottawa, Ottawa K1H 8M5, Canada

†To whom correspondence should be addressed. Email: schen2@uottawa.ca

Abstract:

The primary motor cortex (M1) is known to be a critical site for movement initiation and motor learning. Surprisingly, it has also been shown to possess reward-related activity, presumably to facilitate reward-based learning of new movements. However, whether reward-related signals are represented among different cell types in M1, and whether their response properties change after cue-reward conditioning remains unclear. Here, we performed longitudinal *in vivo* two-photon Ca^{2+} imaging to monitor the activity of different neuronal cell types in M1 while mice engaged in a classical conditioning task. Our results demonstrate that most of the major neuronal cell types in M1 showed robust but differential responses to both the conditioned cue stimulus (CS) and reward, and their response properties undergo cell-type specific modifications after associative learning. PV-INs' responses became more reliable to the CS, while VIP-INs' responses became more reliable to reward. PNs only showed robust responses to novel reward, and they habituated to it after associative learning. Lastly, SOM-INs' responses emerged and became more reliable to both the CS and reward after conditioning. These observations suggest that cue- and reward-related signals are preferentially represented among different neuronal cell types in M1, and the distinct modifications they undergo during associative learning could be essential in triggering different aspects of local circuit reorganization in M1 during reward-based motor skill learning.

INTRODUCTION

The primary motor cortex (M1) is an essential site for movement execution and motor learning. Within M1, neurons encode movement goals and movement kinematics (Georgopoulos et al., 1992; Moran & Schwartz, 1999; Peters et al., 2014). Intriguingly, neurons in M1 have also been reported to show reward-related activity. *In vivo* recording studies performed in nonhuman primates found neurons in M1 that encode reward anticipation, reward delivery, and mismatches between the two (Marsh et al., 2015; Ramakrishnan et al., 2017; Ramkumar et al., 2016). In human subjects, reward has also been shown to modulate M1 activity, likely through an inhibitory circuit-dependent mechanism (Thabit et al., 2011). However, it remains unclear how reward-related responses are represented in M1, and if the representation changes with associative learning.

It was recently shown that in well-trained mice performing a skilled reaching task, a subset of layer 2/3 (L2/3) pyramidal neurons (PNs) in M1 specifically report successful, but not failed, reach-and-grasp movements. In contrast, a different subset of PNs report only failed reach-and-grasp movements (Levy et al., 2020). Since the ability to use past experience to learn action-outcome associations is critical to survival, encoding the outcome in M1 may be an important part of motor skill learning. It is widely accepted that associative learning using reinforcement can accelerate and enhance learning (Abe et al., 2011; Nikooyan & Ahmed, 2015). In the case of motor learning, studies have demonstrated that positive feedback (reward) facilitates motor memory retention and negative feedback (punishment) speeds up the learning process (Galea et al., 2015). One hypothesis is that during learning, reward signals in the brain, together with neuromodulators and synaptic plasticity, are involved in potentiating and optimizing the neural circuitry in M1 that underlies the rewarded movement. Implementing such a learning process would necessitate the interplay between different cell types within the local microcircuitry (Richards et al., 2019).

M1, like other cortical areas, is densely packed with PNs and diverse inhibitory interneuron (IN) types and is wired in a delicately balanced and intricate circuit. Different IN subtypes have been shown to have distinct gene expression profiles, electrophysiological properties, and connectivity motifs (Fishell & Rudy, 2011; Markram et al., 2004). Somatostatin-, parvalbumin-, and vasoactive intestinal peptide-expressing inhibitory neurons (SOM-INs, PV-INs and VIP-INs, respectively) are three major non-overlapping subtypes of GABAergic neurons that broadly form a common microcircuit motif in the cortex. Some studies have demonstrated that SOM-INs preferentially target distal dendrites of PNs to filter synaptic inputs, fast-spiking PV-INs preferentially target perisomatic regions of PNs enabling strong inhibition of spiking, and VIP-INs regulate local microcircuits by controlling other local INs (Pfeffer et al., 2013). Due to their diverse properties and strategic connectivity motifs, these INs exert fine control over local network activity and provide a potential mechanism for how the brain processes reward signals and ultimately uses this information to optimize neural activity related to learned motor skills.

Multiple studies using *in vivo* opto-recordings in the primary visual cortex have shown that visual orientation selectivity in PNs is modulated and sharpened by PV- and SOM-INs (Atallah et al., 2012; Lee et al., 2012; Wilson et al., 2012). In the primary auditory cortex, PV- and SOM-INs exert analogous control over PN frequency tuning (Seybold et al., 2015). Moreover, in the auditory cortex, prefrontal cortex, and basolateral amygdala, reinforcement signals such as reward and punishment have been shown to recruit VIP-INs, which in turn, inhibit SOM- and PV-INs (Krabbe et al., 2019; Pi et al., 2013). The subtype-specific roles of these INs have long been elusive, but a complex picture is emerging where INs are not only responsible for maintaining a delicate balance of excitation and inhibition, but are also actively involved in processing activity in the cortex (Lee et al., 2020; Wood et al., 2017).

Here, we employed chronic *in vivo* two-photon imaging, combined with a head-fixed classical conditioning task, to monitor the activity of the same population of PNs, PV-INs, SOM-INs, or VIP-INs before and after associative learning to investigate whether and how conditioned cue stimulus (CS) and unconditioned reward are represented among different neuronal cell types in M1. Our results demonstrate

that all four major cell types in M1 show distinct responses to CS and reward, and their response properties undergo cell-type specific modifications after associative learning. Notably, PV-INs and VIP-INs exhibited stimulus-specific modifications, in which PV-INs became more reliably responsive to the CS but not to the reward, whereas VIP-INs became more reliable to the reward but not to the CS. PNs initially showed robust responses to novel reward but became habituated to it after associative learning. Lastly, SOM-IN responses emerged with learning and responded more reliably to both the CS and reward. Taken together, these results show that cue- and reward-related signals are preferentially represented among major neuronal cell types in M1, and they undergo cell-type specific modifications during associative learning, indicating they may have distinct roles in integrating reinforcement signals to promote circuit reorganization in M1 during motor skill learning.

RESULTS

To understand how reward-associated signals are represented within the local microcircuitry in M1 before and after associative learning, we established a head-fixed auditory cued reward conditioning task, which allowed us to combine the task with *in vivo* two-photon Ca^{2+} imaging to examine the response properties of different neuronal cell type populations in awake and behaving mice (**Figure 1A**). In this task, water-restricted mice were exposed to a conditioned stimulus (CS; auditory tone, 1s duration), followed by a 1.5s delay and then the delivery of the unconditioned stimulus (US; water reward, ~10 μL). Mice were trained for ~30-35 trials/session (1 session/day for 7 days) with randomly varied inter-trial intervals (ITI) between 60 – 120s (**Figure 1B**). Since M1 is known to be involved in movement initiation and motor skill learning, we chose to use a simple classical conditioning task with just an auditory tone paired with reward and omitted any additional training where mice would be required to learn a new movement. The rationale for this is that many neuronal cell types, including PNs, PV- and SOM-INs, have been shown to undergo modifications when mice acquire new movements (Chen et al., 2015; Cichon & Gan, 2015; Donato et al., 2013; Xu et al., 2009). However, since licking is an innate movement that does not induce plastic changes in adult mice (Chen et al., 2015; Komiyama et al., 2010; Peters et al., 2014), we can reliably attribute changes in neuronal activity over the course of the task to associative learning, rather than motor learning.

Mice learned to associate the CS with the reward after 7 days, shown by an increase in anticipatory lick rate, a conditioned response, following the cue onset on day 7 compared to day 1 (**Figure 1C**). On a trial-by-trial basis, anticipatory lick rate did not change significantly within a single session on both day 1 and day 7, implying limited within-session improvements (**Figure 1D**). To ensure the increase in lick rate was specific to the CS, we compared the mean lick rate during the CS and reward period to the lick rate during self-initiated spontaneous lick bouts in the ITI (in the absence of the CS or reward). To be consistent with the 2.5s analysis window for CS responses, we analysed the first 2.5s of ITI lick bouts and 2.5s following reward delivery. On day 1, the mean lick rate during ITI lick bouts ($1.47 \pm 0.17/\text{s}$) and the CS ($1.37 \pm 0.17/\text{s}$) were similar, while the lick rate following reward delivery was significantly higher ($7.16 \pm 0.48/\text{s}$). In contrast, on day 7 following associative learning, the lick rate during the CS period ($3.42 \pm 0.37/\text{s}$) was significantly higher than during ITI lick bouts ($1.71 \pm 0.2/\text{s}$), demonstrating that the mice effectively learned the CS-reward association by day 7 (**Figure 1E, F**).

To investigate the activity of different neuronal cell types during this task, we used *in vivo* two-photon Ca^{2+} imaging of different cell type populations. To target PNs in M1, we injected an adeno-associated virus (AAV) carrying a Ca^{2+} indicator (GCaMP6f) driven by the CaMKII promoter (AAV1.CaMKIIa.GCaMP6f) into M1 of wildtype B6129S mice. After 3-5 weeks, we recorded the activity of hundreds of L2/3 PNs using two-photon microscopy in awake mice while they underwent the head-fixed conditioning task, and we tracked the same population of neurons on day 1 and day 7 (**Figure**

2A). We identified all the active neurons within a session, irrespective of the behavioural task (see Methods), and sorted neurons by the timing of their peak activity relative to the CS onset. It was apparent that there were subpopulations of neurons more responsive to CS, reward, or both (**Figure 2B, C**). We also repeated the experiments to examine if the major IN subtypes in M1 also respond to the CS and reward during the conditioning task. To do this, we injected AAV-Syn-Flex-GCaMP6f in PV-Cre, SOM-Cre or VIP-Cre transgenic mice to selectively express GCaMP6f in PV-INs, SOM-INs, or VIP-INs, respectively, and then performed *in vivo* two-photon Ca^{2+} imaging to monitor the response properties of the same population of INs on day 1 and day 7 after associative learning (**Figure 2 – figure supplement 1 and 2, Supplementary File 1**). We compared the mean percentage of active cells within the entire session to ensure all cell types had a similar proportion of active cells (irrespective of the behavioural task) on day 1 and day 7 (**Figure 2D**).

To examine task-related activity in each cell type, we first compared the mean percent of active cells during the CS and reward to a null distribution made by randomly sampling the session irrespective of the behavioural task, and then calculating the mean percentage of active neurons during the sampled period. By repeating this 1,000 times for each cell type on day 1 and day 7, we created a distribution of the percentage of active neurons that were present at baseline levels or by chance. Surprisingly, we found that only PV-IN and VIP-IN cell types had a percent of CS- and reward-responsive cells that were significantly greater than chance level on both day 1 and day 7 (PV-IN CS: day 1: $15.26 \pm 2.11\%$, day 7: $24.09 \pm 2.98\%$; PV-IN reward: day 1: $20.17 \pm 3.27\%$, day 7: $27.47 \pm 3.66\%$; VIP-IN CS: day 1: $11.29 \pm 3.23\%$, day 7: $16.59 \pm 2.01\%$, VIP-IN reward: day 1: $18.65 \pm 5.91\%$, day 7: $26.3 \pm 6.04\%$; **Figure 2F, G**). PN responses to the CS were not different from the null distribution on both day 1 and 7 (day 1: $18.42 \pm 1.05\%$, day 7: $16.41 \pm 1.33\%$; **Figure 2E**); in contrast, PN responses to the reward were significantly higher on day 1 but significantly lower than the null distribution on day 7 (day 1: $23.95 \pm 2.49\%$, day 7: $12.12 \pm 0.95\%$; **Figure 2E**). Lastly, SOM-INs showed significant responses to the CS and reward only on day 7 following associative learning, while on day 1, they demonstrated no response to the CS and a modest response to the reward (CS: day 1: $5.53 \pm 2.7\%$, day 7: $13.56 \pm 3.17\%$; Reward: day 1: $9 \pm 3.66\%$, day 7: $12.15 \pm 3.93\%$; **Figure 2H**). Based on these findings, we decided to only examine the sessions where the percent of active cells were significantly greater than the null distribution in our subsequent analysis, as a non-significant percent of active cells during the stimulus period cannot be readily distinguished from non-task related baseline noise.

We began our analysis on PV-INs and VIP-INs because they both showed significant responses to both the CS and reward on day 1 and day 7. To understand how their representations of reward and reward-associated cues changed over the course of learning, we first analyzed the tuning of individual cells to unbiasedly identify their response properties. By quantifying the tuning of each cell's average response during the CS and reward response periods (2.5s window) using the non-parametric Spearman correlation ρ (see Methods), we observed a wide range of tuning coefficients to the CS and reward, with a small proportion that was strongly positively or negatively tuned to the CS or reward stimulus (tuning coefficient near -1 or 1; **Figure 3A-D**), consistent with our earlier analyses demonstrating that neurons in M1 show activity associated with the CS or reward during the conditioning task. We next examined whether the tuning coefficient within each cell type changed after associative learning by calculating the change in tuning coefficients for each cell between day 1 and day 7. Again, to validate our findings, we compared these values to a null distribution of $\Delta\rho$ values obtained by randomly sampling the two sessions (see Method Details). The PV-IN population did not show any significant changes in either CS or reward tuning between day 1 and 7 ($\overline{\Delta\rho}_{\text{CS}} = -0.049 \pm 0.046$, $\overline{\Delta\rho}_{\text{reward}} = 0.014 \pm 0.054$; **Figure 3E, F**), indicating that neither CS- nor reward-related tuning became stronger after associative learning. In

contrast, VIP-INs' CS tuning did not change significantly between day 1 and 7 ($\overline{\Delta\rho}_{CS} = -0.065 \pm 0.048$, but VIP-INs' reward tuning significantly increased on day 7 ($\overline{\Delta\rho}_{\text{reward}} = 0.161 \pm 0.086$; **Figure 3G, H**), suggesting a strengthening of VIP-IN responsivity to reward following associative learning.

Although the tuning properties can reveal changes in task-related responsivity, this analysis is limited in identifying changes at the trial-by-trial level. When we assessed population activity following the CS onset (**Figure 4A**), it was apparent that a group of PV-INs and VIP-INs were responsive to CS on both day 1 and day 7 (**Figure 4B, 5B**). Hence, by identifying and tracking the same neurons from day 1 to day 7, we were able to ask if there was (1) an increase in the number of neurons being recruited as CS- or reward-responsive during associative learning or (2) a change in the trial-by-trial reliability of CS and reward responses. When we compared the mean percent of CS-responsive neurons on day 1 and day 7, we found that the average percent of CS-responsive PV-INs during a trial increased significantly by day 7 (day 1: $15.26 \pm 2.11\%$, day 7: $24.09 \pm 2.98\%$; **Figure 4C**), while the percent of CS-responsive VIP-INs did not change (day 1: $11.29 \pm 3.23\%$, day 7: $16.59 \pm 2.01\%$; **Figure 4D**), demonstrating that more PV-INs became responsive to the CS after associative learning. We then assessed the reliability of the responses, defined as the percent of trials within a session where a neuron was responsive to the CS. This measure quantifies how consistently a neuron responded to the CS within a session. We first plotted the cumulative distribution function of reliabilities among all PV-INs and VIP-INs. We observed that PV-INs, as a population, were significantly more reliable in their CS responses than VIP-INs on day 1 (**Figure 4E**). When we sorted neurons based on their day 1 reliability values and followed them to day 7, we observed that many of the PV-INs that initially had low reliability to CS became more responsive on day 7 (**Figure 4H**); therefore, we grouped neurons into 'High Reliability' if they were among the top 50th percentile, while neurons in the lower 50th percentile were deemed 'Low Reliability'. We found that PV-INs that began as highly reliable maintained their reliability to the CS (day 1: $29.8 \pm 1.51\%$, day 7: $33.87 \pm 4.72\%$), while PV-INs that began as low reliability became significantly more reliable ($8.47 \pm 0.46\%$, day 7: $18.99 \pm 3.76\%$; **Figure 4F**). In contrast, the reliability of both High and Low VIP-INs did not change (High Reliability: day 1: $26.55 \pm 2.62\%$, day 7: $25.93 \pm 3.81\%$; Low Reliability: day 1: $6.32 \pm 0.76\%$, day 7: $14.24 \pm 2.6\%$; **Figure 4G**). We then followed individual Low Reliability PV-INs and calculated the change in reliability to the CS ($\Delta\text{reliability}_{CS}$) from day 1 to day 7. As a control, we randomly sampled the day 7 session irrespective of the behavioural task and calculated a reliability value. We then subtracted that value from the actual day 1 CS reliability to generate a randomized change in reliability ($\Delta\text{reliability}_{\text{random}}$) for each Low Reliability neuron (see Methods). When we compared the two distributions, we found that among the Low Reliability PV-INs, $\Delta\text{reliability}_{CS}$ was significantly greater than the $\Delta\text{reliability}_{\text{random}}$ control (**Figure 4I**). Lastly, we examined if the onset of neuronal activity after the CS changed after associative learning, and if any neurons that were previously responsive to the CS became responsive to the reward only. We did not observe a change in the onset of neuronal activity following CS (**Figure 4 – figure supplement 1A, see Methods**). Furthermore, 97% of the CS-responsive PV-INs from day 1 still showed responsivity to CS on day 7 (**Figure 4 – figure supplement 1C-D**). Altogether, these results show that as a population, more PV-INs became responsive to the CS, and this is mainly due to Low Reliability PV-INs that became more reliable following associative learning.

To demonstrate that changes in PV-INs' representation of the CS resulted from associative learning, we conducted additional control experiments to examine their responses to tone when mice received no rewards or non-paired randomly timed rewards. In the first experiment, water-restricted PV-Cre mice were exposed to the same auditory tone used as the CS (1s duration) but all water rewards were omitted (randomly varied ITI between 60 – 120s; ~30 trials/session; 1 session/day for 7 days). We imaged

PV-IN activity on both day 1 and day 7 and assessed their population responses following the tone onset (**Figure 4 – figure supplement 2A**). We first compared the mean percent of active cells during tone to a null distribution made by randomly sampling the session irrespective of the behavioral task (as in Figure 2). Surprisingly, PV-INs did not show significant tone-responsive cells compared to the chance level on either day 1 or day 7 (**Figure 4 – figure supplement 2B**); hence, the average percent of tone-responsive PV-INs per trial also did not increase from day 1 to day 7 (**Figure 4 – figure supplement 2C**). Next, in a separate cohort of mice, we exposed water-restricted PV-Cre mice to tone, followed by a ‘non-paired’ water reward that was given at randomly varied time intervals (40 – 80s). Mice were also trained for ~30 trials/session (1 session/day for 7 days) with a randomly varied ITI between 15 – 25s (**Figure 4 – figure supplement 2D**). We found PV-INs were significantly responsive to the tone stimuli on day 1, similar to what we observed earlier in the CS-reward task (**Figure 2F**). Interestingly, by day 7, PV-INs no longer responded to the tone stimulus (**Figure 4 – figure supplement 2G**). We next examined if mice that received the tone stimulus with non-paired water reward learned to associate the two after 7 days. We found the animals did not learn the association, as their conditioned response (tone-evoked anticipatory licking) did not increase at day 7 (**Figure 4 – figure supplement 2E-F**). Unlike the mice that learned the association in the CS-reward task, we did not observe a change in the mean percent of tone-responsive PV-INs from day 1 to day 7 in the non-paired paradigm, and the reliability of Low Reliability PV-INs to tone also did not change (**Figure 4 – figure supplement 2H-I**). Together, these results suggest that PV-INs in M1 do not respond to auditory tone in general, but instead only respond to the tone when the animal actively associates it with reward. Moreover, the changes among PV-INs to the CS tone from day 1 to day 7 are specific to associative learning.

We next assessed reward responses among PV-INs and VIP-INs in the same manner but now looked for responses within 2.5s of the reward delivery time (**Figure 4B, 5A-B**). We tracked the same neurons from day 1 to day 7 and compared the mean percent of reward-responsive neurons. PV-INs and VIP-INs did not show a significant change in the percent of responsive cells per trial (PV-IN: day 1: $20.17 \pm 3.27\%$, day 7: $27.47 \pm 3.66\%$; VIP-IN: day 1: $18.65 \pm 5.91\%$, day 7: $26.3 \pm 6.04\%$; **Figure 5C, D**). When we examined the cumulative distribution of reliabilities for reward responses between the two cell types, VIP-INs as a population were significantly more reliable than PV-INs on day 1 (**Figure 5E**). By dividing the cells into High and Low Reliability groups, we found the High Reliability VIP-INs maintained their reliability (VIP-IN High Reliability: day 1: $38.79 \pm 2.71\%$, day 7: $35.88 \pm 4.32\%$), and the Low Reliability VIP-INs became significantly more reliable on day 7 (VIP-IN Low Reliability: day 1: $10.25 \pm 1.67\%$ day 7: $24.06 \pm 4.87\%$; **Figure 5G-H**). In contrast, both High and Low Reliability PV-INs maintained their reliability to reward (High Reliability: day 1: $35.59 \pm 2.81\%$, day 7: $41.94 \pm 5.7\%$; Low Reliability: day 1: $10.58 \pm 0.74\%$, day 7: $19.59 \pm 3.08\%$; **Figure 5F**). We also followed individual Low Reliability VIP-INs and calculated the change in reliability to reward ($\Delta reliability_{reward}$) from day 1 to day 7. As a control, we randomly sampled the day 7 session irrespective of the behavioural task and calculated a random $\Delta reliability_{random}$ for each neuron as described above. When we compared the two distributions, the $\Delta reliability_{reward}$ was significantly greater than the $\Delta reliability_{random}$, demonstrating that Low Reliability VIP-INs became more reliably responsive to reward (**Figure 5I**). We also examined if the onset of neuronal activity after reward consumption changed after associative learning, and if neurons that were previously responsive to reward became responsive to the CS only. We did not observe a change in the onset of neuronal activity following reward (**Figure 4 – figure supplement 1B**). Furthermore, 95% of the VIP-INs from day 1 still showed responsivity to reward on day 7 (**Figure 4 – figure supplement 1E-F**). Lastly, to demonstrate that changes in VIP-INs’ representation of reward resulted from associative learning, we examined both PV-IN and VIP-IN responses to reward in mice that were exposed to the non-paired behavioral paradigm (tone + randomly

timed water rewards; **Figure 4 – figure supplement 2D-F**). Consistent to what we observed in the CS-reward paradigm, both PV-INs and VIP-INs consistently showed higher mean percent of active cells during reward (2.5s from the first lick after reward delivery) compared to the chance level on both day 1 and day 7 (**Figure 4 – figure supplement 2J-M**). However, because these mice did not learn the association between the auditory tone and randomly delivered water reward, we did not see an increase in the reliability of the Low Reliability VIP-INs from day 1 to day 7 (**Figure 4 – figure supplement 2O**), or a change in the percent of responsive cells in either PV-INs or VIP-INs (**Figure 4 – figure supplement 2K-N**). Altogether, we found that during associative learning, while the proportion of reward-responsive VIP-INs during a given trial did not change, a subset of VIP-INs that were largely unresponsive to reward on day 1 became more reliably responsive on day 7.

Although PV-INs and VIP-INs were the only cell types that were significantly responsive to both CS and reward on both day 1 and day 7, PNs and SOM-INs also had significant responses to specific stimuli on certain days. While PNs did not show significant CS responses when compared to baseline, their reward responses on day 1 were significantly above the null distribution, and they became significantly lower than the null distribution on day 7 (**Figure 2E**). This result is in line with the change in tuning coefficient ($\Delta\rho_{reward}$), which showed a significant decrease in reward tuning between day 1 and 7 ($\overline{\Delta\rho}_{reward} = -0.141 \pm 0.067$; **Figure 6A**). Moreover, the cumulative distribution function of PN reliability also shifted significantly to lower reliabilities on day 7 compared to day 1 (**Figure 6B**). These results indicate that PNs initially responded to novel reward; however, they habituated to the reward following associative learning.

SOM-INs initially had no response to the CS on day 1, but their responses became significant on day 7 (**Figure 2H**). The change in CS tuning coefficient ($\Delta\rho_{tone}$) was not significant ($\Delta\rho_{tone} = .059 \pm 0.031$, **Figure 6C**), suggesting their responsivity did not change with learning. Interestingly, when we assessed the cumulative distribution of CS response reliability on day 1 and day 7, the cumulative distribution function shifted to significantly higher reliability values on day 7 (**Figure 6C, D**). Notably, by day 7, there was a visible reduction in the number of SOM-INs that had 0% reliability to CS on day 1, indicating they were completely unresponsive to the CS on day 1 but not on day 7. Finally, SOM-INs showed modest but significant responses to reward on day 1 and 7 (**Figure 2H**). When we assessed the reward tuning among the SOM-IN population, $\Delta\rho_{reward}$ did not show a significant change between day 1 and 7 ($\Delta\rho_{reward} = .017 \pm 0.040$, **Figure 6E**). However, SOM-IN reliabilities also shifted to higher values on day 7 (**Figure 6F**). Altogether, these results suggest that in naïve mice, SOM-INs were unresponsive to the CS and modestly responsive to the novel reward; however, following associative learning, SOM-INs became more reliably responsive to both the CS and reward.

Lastly, reward consumption requires innate tongue movements during licking, and since microstimulation of M1 in mice has been shown to evoke tongue and jaw movements (Komiyama et al., 2010), it is crucial to distinguish whether the observed CS and reward responses resulted from task-related stimuli or if the activity is simply associated with licking movements. We demonstrated earlier that head-fixed mice learned the CS-reward association by displaying the conditioned response (anticipatory licking) following the CS on day 7 (**Figure 1**). To address this potential confound, we identified all the self-initiated licking bouts during ITIs, when no reward was present (**Figure 7A-C; Figure 2 – figure supplement 1**). We first assessed all the significantly active cells in each cell type (identified in **Figure 2D**) during the first lick bout of each ITI on day 1 and day 7. We observed that in each cell type, the majority of the neurons were non-lick neurons on both day 1 and day 7 (**Figure 7 – figure supplement 1A**). We then tracked individual lick and non-lick neurons and examined if they shifted their responses after associative learning. We found that most of the neurons maintained the same responses, and non-lick neurons were still the majority in all cell types (**Figure 7 – figure supplement**

1B). Next, we examined whether the lick neurons also showed mixed responses to CS, reward, or CS + reward. Indeed, lick neurons exhibited mixed responses to CS, reward, or CS + reward (**Figure 7 – figure supplement 1C**). We then further divided them into three categories – ‘CS cells’, ‘reward cells’, and ‘CS + reward cells’ and compared the percentage of neurons in each category between day 1 and day 7; we did not observe a significant difference (**Figure 7 – figure supplement 1D**). Lastly, at the population level, we examined the response reliability index of all the active neurons during all ITI lick bouts and compared them to the response reliability index for the CS and reward. On both day 1 (when there was minimal anticipatory licking during the CS) and day 7 (when mice showed anticipatory licking), all cell types exhibited lower reliability index values for the ITI lick bouts compared to the CS and reward, indicating that the increase in task-related responses following water rewards was specific to the reward stimulus, and not licking movements (**Figure 7D-J**). Together, these results suggest that the cell-type specific modifications observed between day 1 and 7 were not caused by licking movements.

DISCUSSION

M1 is known to be involved in motor initiation, movement kinematics and motor learning. Recent studies have demonstrated reward-related activity in M1 using *in vivo* electrophysiological recordings in nonhuman primates (Marsh et al., 2015; Ramakrishnan et al., 2017; Ramkumar et al., 2016) and transcranial magnetic stimulation in human subjects (Thabit et al., 2011). However, whether CS- and reward-associated signals are represented among different neuronal cell types within the microcircuit in M1 is still unclear. Using chronic two-photon Ca^{2+} imaging, combined with transgenic mouse lines and viral strategies to target different neuronal cell types, we demonstrated that during a conditioning task, all major cell types in M1 responded to either the CS, the reward, or both. Most notably, each cell type underwent distinct modifications after association learning. By tracking the same population of neurons, we revealed that the CS-responding population increased among PV-INs and individual cells responded more reliably to the CS following associative learning. On the contrary, VIP-INs became more reliable in response to reward. When mice underwent control behavioural paradigms where tone was not paired with reward and no associative learning occurred, PV-INs and VIP-INs did not undergo these changes. Additionally, PNs had a drastically reduced response to reward, while SOM-INs became more reliable to both the CS and the reward. Our findings suggest that each cell type has a distinct role in processing information related to the cue-reward association in M1, and they may work together to provide the reinforcement signals in M1 that are important for motor skill learning.

Previous studies in trained rhesus monkeys performing a joystick center-out task have shown a widespread representation of reward anticipation and reward-related activity among cortical neurons in M1 (Ramakrishnan et al., 2017). Consistent with earlier work, we also observed reward-related activity in all four major cell types in M1, even in naïve mice on day 1 when they were first exposed to the CS and reward. It has been reported that in sensory cortices, repeated passive exposure to a sensory stimulus leads to a long-lasting reduction in PN responsivity, but when animals are engaged in learning, PNs maintain their responsivity to the repeated stimulus (Kato et al., 2015; Makino & Komiyama, 2015). However, we found in M1, when water-restricted mice were engaged in a conditioning task to learn the association between the CS and water reward, PNs still showed a drastic habituation to the reward stimulus. A recent study that imaged neuronal activity in expert mice performing a head-fixed pellet reaching task demonstrated that L2/3 PNs in M1 are involved in encoding movement outcome (success vs. failure) but not the appetitive outcome (reward vs. no reward). However, the authors did not image the mice at the naïve stage (Levy et al., 2020). Hence, one possibility is that L2/3 PNs in M1 encode reward signals during the naïve stage, but after associative learning, they habituate and become unresponsive to the reward stimulus. In addition, in the sensory cortices, the flexibility to either respond to or ignore sensory stimuli is based on the stimulus’ behavioral relevance and is gated by local SOM-INs (Kato et al., 2015;

Makino & Komiyama, 2015; Poort et al., 2021). In line with these findings, we found that SOM-INs became more reliably responsive to both the CS and the reward with associative learning. We also observed stimulus-specific increases in the reliability of PV-INs' response to the CS and VIP-INs' response to the reward after associative learning, and these changes to tone and reward were not observed in the absence of associative learning. When mice were exposed to tone alone, PV-INs did not show significant responses to tone on either day 1 or day 7, while mice exposed to a non-paired tone and reward task showed significant responses to tone on day 1, but not on day 7. This suggests that in M1, PV-INs only respond to behaviourally relevant cues such as those that predict reward. Finally, while VIP-INs remained responsive to reward in the non-paired paradigm, VIP-INs did not show changes in reliability in the absence of associative learning. Together, our results suggest that different IN subtypes may have distinct roles in processing CS- and reward-related information in M1 during motivated associative learning.

One hypothesis is that PV-INs are recruited by the CS to control the behavioral responses (anticipatory licking) during reward anticipation since PV-INs are known to regulate PN firing through both feedforward and feedback inhibition (Fishell & Rudy, 2011; Xu & Callaway, 2009; Xue et al., 2014). Similar observations have been reported in the striatum, in which optogenetic activation or suppression of PV-INs during a similar conditioning task impaired anticipatory licking, demonstrating the importance of PV-INs in the expression of conditioned responses (Lee et al., 2017). Likewise, PV-INs in the basolateral amygdala, are also recruited during the CS and subsequently inhibited during the US in an auditory fear conditioning task. Optogenetic activation of PV-INs during the CS increased conditioned freezing behaviour while PV-IN suppression reduced freezing, indicating bidirectional control of the conditioned response (Wolff et al., 2014). Our results demonstrate that in a naïve animal, a subset of PV-INs in M1 are responsive to the CS only when rewards are present, and more PV-INs are recruited by the CS if the animal learns that the CS predicts reward. This suggests that in M1, PV-IN responses to the CS are not purely sensory, but rather, they may play an important role in controlling the behavioral responses to the CS.

VIP-INs, on the other hand, were significantly less reliable in responding to the CS compared to PV-INs, and their responses to the CS remained low. However, VIP-INs' responses to the reward were more reliable than those of PV-INs, and they became more closely tuned and reliably responsive to the reward with learning. Due to the disinhibitory position of VIP-INs in the microcircuit, activation of VIP-INs can lead to widespread increases in local excitability and contribute to regulating cortical gain (Fu et al., 2014; Jackson et al., 2016; Pfeffer et al., 2013). Furthermore, a growing body of evidence suggests a general principle across brain regions, in which VIP-INs receive long-range inputs (Duan et al., 2020; Krabbe et al., 2019; Turi et al., 2019; Zhang et al., 2014; Gasselín et al., 2021), respond to reinforcement signals (Krabbe et al., 2019; Pi et al., 2013), and play an important role in goal-oriented learning (Krabbe et al., 2019; Turi et al., 2019). Taken together, our results suggest that during CS-reward conditioning, PV-INs in M1 encode the CS association, and may regulate local circuit activity related to reward anticipation, whereas VIP-INs act as a context-dependent switch following the reward delivery (Muñoz et al., 2017; Turi et al., 2019) to instruct and disinhibit local PNs to enable learning-induced plastic changes critical for the acquisition of new movements. An interesting point to note is that since PV-INs are only responsive to tone when it is paired with reward, and neither PV-INs nor VIP-INs undergo plastic changes in the absence of associative learning, M1 is unlikely to be a primary site for learning reward-predictions. We hypothesize that other brain regions are responsible for learning relevant CS-reward associations while filtering out behaviourally irrelevant stimuli, and these regions subsequently send long-range inputs to M1 to instruct motor responses to the CS and the reward. In summary, this study provides insight on how different IN subtypes in M1 integrate incoming inputs from various brain regions and

orchestrate local circuit plasticity. Future work will be important to identify the origin of these putative long-range inputs to different cell types in M1.

FIGURE LEGENDS:

Figure 1. Associative learning during a head-fixed classical conditioning task.

(A) Schematic of head-fixed classical conditioning task.

(B) Trial structure.

(C) Mean lick rate per second on day 1 and day 7. Binned over 0.5s intervals. Lick rate following CS onset up to reward delivery time is higher on day 7. Two-way ANOVA, *** $p < 1 \times 10^{-3}$, effect of time: $p < 1 \times 10^{-3}$, effect of day: $p < 1 \times 10^{-3}$.

(D) Mean anticipatory lick rate across trials within day 1 and day 7 sessions. Mean anticipatory lick rate was calculated from CS onset to end of delay period. Two-way ANOVA, effect of trial number: $p = 0.91$, effect of day: $p < 1 \times 10^{-3}$.

(E) Mean lick rate during the first 2.5s of ITI lick bouts, 2.5s following CS onset and 2.5s following reward delivery on day 1. Each point is the mean from an individual mouse. One-way ANOVA with Tukey-Kramer correction for multiple comparisons, ITI vs. CS: $p = 0.97$, ITI vs. reward: $p < 1 \times 10^{-3}$, CS vs. reward: $p < 1 \times 10^{-3}$.

(F) Mean lick rate during the first 2.5s of ITI lick bouts, 2.5s following CS onset and 2.5s following reward delivery on day 7. Each point is the mean from an individual mouse. One-way ANOVA with Tukey-Kramer correction for multiple comparisons. ITI vs. CS: $p = 1.08 \times 10^{-3}$, ITI vs. reward: $p < 1 \times 10^{-3}$.

n = 23 mice. Error bars show SEM.

Figure 2. Longitudinal *in vivo* Ca^{2+} imaging of neuronal responses in M1 during a classical conditioning task

(A) Experimental timeline (top). *In vivo* two-photon imaging of L2/3 PNs expressing GCaMP6f in M1 (bottom left). The same population can be tracked from day 1 to day 7 (bottom right). Yellow arrows indicate example tracked neurons across days.

(B) Z-scored fluorescence traces from 13 neurons (top), and the corresponding licking measured with the lick-o-meter (bottom) from the same mouse and same trial on day 1. Grey bar represents the timing of the CS. Dotted red line indicates the onset of water reward delivery.

(C) Z-scored activity of all the active neurons from an example mouse during one representative trial on day 1 and day 7, sorted by timing of maximum activity following the CS onset. Grey bar represents the timing of the CS. White line indicates the onset of water reward delivery.

(D) Mean percent of active neurons within a session, irrespective of the behavioural task, for PN, PV-IN, VIP-IN and SOM-INs on day 1 and day 7. All cell types showed a similar percentage of active neurons. One-way ANOVA, n.s., non-significant, day 1: $p = 0.38$, day 7: $p = 0.22$. Error bars show SEM.

(E-H) Mean percent of responsive neurons to CS (top) and reward (bottom) within 2.5s of CS/reward onset for each cell type. Violin plots show null distribution of percentage of responsive neurons made by randomly re-sampling mice and shuffling the session, 1,000 times (see Methods). The circle represents the mean percentage of tone- or reward- responsive neurons. Monte Carlo with Bonferroni correction, n.s., non-significant, * $p < 0.05$, *** $p < 1 \times 10^{-3}$. PN CS day 1: $p = 0.582$, CS day 7: $p = 0.423$, Reward day 1: $p < 1 \times 10^{-3}$, Reward day 7: $p = 0.015$, $n = 1029$ cells from 6 mice **(E)**, PV-IN CS day 1 $p < 1 \times 10^{-3}$, CS day 7: $p < 1 \times 10^{-3}$, Reward day 1: $p < 1 \times 10^{-3}$, Reward day 7: $p < 1 \times 10^{-3}$, $n = 316$ cells from 6 mice **(F)**, VIP-IN CS day 1: $p = 0.039$, CS day 7: $p < 1 \times 10^{-3}$, Reward day 1: $p < 1 \times 10^{-3}$, Reward day 7: $p < 1 \times 10^{-3}$, $n = 407$ cells from 4 mice **(G)**, SOM-IN CS day 1: $p = 0.47$, CS day 7: $p < 1 \times 10^{-3}$, Reward day 1: $p = 0.033$, Reward day 7: $p < 1 \times 10^{-3}$, $n = 189$ cells from 7 mice **(H)**.

Figure 3. Learning-associated changes in single-neuron tuning properties in M1

(A-B) Example fluorescence traces, color-coded based on non-parametric Spearman correlation with CS **(A)** or reward **(B)**. Each trace is a trial averaged response from different neurons on day 7.

(C-D) Trial-averaged fluorescence of all the neurons recorded on day 7 and sorted to the value of the Spearman correlation (-1 to 1) to CS **(C)** or reward **(D)**. Active neurons during the CS or reward responsive period showed higher tuning coefficient.

(E-H) Left, distribution of changes in Spearman correlation $\Delta\rho$ with CS **(E, G)** or reward **(F, H)** for PV-INs (top) and VIP-INs (bottom). Each curve represents a Gaussian kernel density estimate of the distribution of $\Delta\rho$ in a single mouse. Right, mean change in Spearman correlation $\overline{\Delta\rho}$ for PV-INs and VIP-INs. Null distributions (gray) were estimated by re-sampling each mouse and shuffling trials 1,000 times (see description of calculation of tuning coefficients in Methods). VIP-IN reward tuning significantly increased with associative learning. Monte-Carlo, *** $p < 1 \times 10^{-3}$, n.s., non-significant, PV-IN CS $p = 0.12$ **(E)**, PV-IN Reward $p = 0.61$ **(F)**, VIP-IN CS $p = 0.082$ **(G)**, VIP-IN Reward $p < 1 \times 10^{-3}$ **(H)**

PV-IN: $n = 316$ cells from 6 mice. VIP-IN: $n = 407$ cells from 4 mice.

Figure 4. PV-IN and VIP-IN CS-related responses before and after associative learning

(A) Trial structure. Gray shaded bar represents the response period analyzed for CS-responsive activity.

(B) Z-scored activity of all the active PV-INs from an example mouse during one representative trial on day 1 and day 7, sorted by timing of maximum activity following the CS onset. Grey bar represents the timing of the CS. White line indicates the onset of water reward delivery.

(C-D) Mean percent of cells that are responsive to the CS for PV-INs **(C)** and VIP-INs **(D)**. PV-INs showed an increase in the percent of CS-responsive neurons after reinforcement learning, while VIP-INs did not show any change. Paired t-test, * $p < 0.05$, n.s., non-significant, PV-IN: $p = 0.031$ **(C)**, VIP-IN: $p = 0.38$ **(D)**

(E) Cumulative probability plots showing the percent of trials that each neuron responded to the CS for PV-INs and VIP-INs on day 1. Neurons from each cell type were pooled across mice. PV-INs showed significantly greater reliability to the CS than VIP-INs. Kolmogorov-Smirnov test, $p < 1 \times 10^{-3}$.

(F-G) Mean reliability index of cells that are responsive to the CS for PV-INs **(F)** and VIP-INs **(G)**. Each cell type is divided into High or Low Reliability Group based on the 50th percentile from the cumulative probability plots in **(E)**. High Reliability PV-INs maintained their consistency, while the Low Reliability

group became more consistent in their responses to CS. **(G)** VIP-INs did not show a change in either group after learning. Paired t-test, PV-IN high reliability: $p = 0.38$, PV-IN low reliability: $p = 0.044$, VIP-IN high reliability: $p = 0.79$, VIP-IN low reliability: $p = 0.060$.

(H) CS-responses from all tracked PV-INs in one example mouse on day 1 and day 7. Left, cumulative distribution of CS response reliability among all tracked cells within the example mouse. Right, binary map of each cell's CS response (active or not) across all trials on day 1 and day 7. Cells were sorted by their day 1 reliability shown on the left and the order is maintained on day 7. Cells with low response reliability to CS on day 1 became more reliable on day 7.

(I) Cumulative probability plots of the change in reliability from day 1 to 7 ($\Delta reliability_{CS}$) among PV-INs with Low Reliability group to CS on day 1. Bold red, $\Delta reliability_{CS}$ of the population of PV-INs. Thin red lines show the $\Delta reliability_{CS}$ distribution within individual PV-Cre mice. As a control, day 7 session was randomly sampled and a random reliability was calculated. $\Delta reliability_{random}$ was calculated by subtracting day 1 CS reliability from the day 7 random reliability (Grey, $\Delta reliability_{random}$ from the same population of PV-INs). Kolmogorov-Smirnov test, *** $p < 1 \times 10^{-3}$ PV-IN: $n = 316$ cells from 6 mice. VIP-IN: $n = 407$ cells from 4 mice. Error bars show SEM.

Figure 5. PV-IN and VIP-IN reward-related responses before and after associative learning

(A) Trial structure. Gray shaded bar represents the response period analyzed for reward-responsive activity.

(B) Z-scored activity of all the active VIP-INs from an example mouse during one representative trial on day 1 and day 7, sorted by timing of maximum activity following the CS onset. Grey bar represents the timing of the CS. White line indicates the onset of water reward delivery.

(C-D) Mean percent of cells that are responsive to the reward for PV-INs **(C)** and VIP-INs **(D)**. Neither PV- or VIP-INs showed a significant change. Paired t-test, n.s., non-significant, PV-IN: $p = 0.16$ **(C)**, VIP-IN: $p = 0.16$ **(D)**

(E) Cumulative probability plots showing the percent of trials that each neuron responded to reward for PV-INs and VIP-INs on day 1. Neurons from each cell type were pooled across mice. VIP-INs showed a significantly greater response reliability to the reward than PV-INs. Kolmogorov-Smirnov test, * $p < 0.05$, $p = 5.2 \times 10^{-3}$.

(F-G) Mean reliability index of cells that are responsive to the reward for PV-INs **(F)** and VIP-INs **(G)**. Each cell type is divided into High or Low Reliability Group based on the 50th percentile from the cumulative probability plots in **(E)**. High and Low Reliability PV-INs maintained their consistency. **(G)** High Reliability VIP-INs did not show a change, while Low Reliability VIP-INs significantly increased in reliability following associative learning. Paired t-test, PV-IN high reliability: $p = 0.36$, PV-IN low reliability: $p = 0.058$, VIP-IN high reliability: $p = 0.090$, VIP-IN low reliability: $p = 0.045$.

(H) Reward-responses from all tracked VIP-INs in an example mouse on day 1 and day 7. Left, cumulative distribution of reward response reliability among all tracked cells within the example mouse. Right, binary map of each cell's reward response (active or not) across all trials on day 1 and day 7. Cells were sorted by their day 1 reliability shown on the left and the order is maintained on day 7. Cells with low response reliability to reward on day 1 became more reliable on day 7.

(I) Cumulative probability plots of the change in reliability from day 1 to 7 ($\Delta reliability_{reward}$) among VIP-INs with low reliability to reward on day 1. Bold blue, $\Delta reliability_{reward}$ of the population. Thin blue lines show the $\Delta reliability_{reward}$ distribution within individual VIP-Cre mice. As a control, day 7

session was randomly sampled and a random reliability was calculated. $\Delta reliability_{random}$ was calculated by subtracting day 1 reward reliability from the day 7 random reliability (Grey, $\Delta reliability_{random}$ from the same population of VIP-INs). Kolmogorov-Smirnov test, *** $p < 1 \times 10^{-3}$

PV-IN: n = 316 cells from 6 mice. VIP-IN: n = 407 cells from 4 mice. Error bars show SEM.

Figure 6. PN and SOM-IN reliability is altered after associative learning.

(A) Left, distribution of changes in PN Spearman correlation $\Delta\rho$ for reward. Each curve represents a Gaussian kernel density estimate of the distribution of $\Delta\rho$ in a single mouse. Right, mean change in Spearman correlation $\overline{\Delta\rho}$. Null distributions (gray) were estimated by re-sampling each mouse and shuffling trials 1,000 times. Reward tuning among PNs decreased after associative learning, Monte Carlo, *** $p < 1 \times 10^{-3}$.

(B) Cumulative probability plots showing the percent of trials that each neuron responded to reward for PNs on day 1 and day 7. Neurons were pooled across mice. Day 7 reliability was significantly lower than day 1. Kolmogorov-Smirnov test, *** $p < 1 \times 10^{-3}$.

(C) Left, distribution of changes in SOM-IN Spearman correlation $\Delta\rho$ with CS. Each curve represents a Gaussian kernel density estimate of the distribution of $\Delta\rho$ in a single mouse. Right, mean change in Spearman correlation $\overline{\Delta\rho}$. SOM-INs did not show a change in tone tuning. Monte Carlo, n.s., non-significant, $p = 0.128$.

(D) Cumulative probability plots showing the percent of trials that each neuron responded to the CS for SOM-INs on day 1 and day 7. Neurons were pooled across mice. Day 7 reliability was significantly greater than day 1. Kolmogorov-Smirnov test, $p < 1 \times 10^{-3}$.

(E) Left, distribution of changes in SOM-IN Spearman correlation $\Delta\rho$ with reward. Each curve represents a Gaussian kernel density estimate of the distribution of $\Delta\rho$ in a single mouse. Right, mean change in Spearman correlation $\overline{\Delta\rho}$. SOM-INs did not show a change in reward tuning. Monte Carlo, n.s., non-significant, $p = 0.598$.

(F) Cumulative probability plots showing the percent of trials that each neuron responded to the reward for SOM-INs on day 1 and day 7. Neurons were pooled across mice. Day 7 reliability was significantly greater than day 1. Kolmogorov-Smirnov test, ** $p = 0.012$.

PN: n = 1029 cells from 6 mice. SOM-INs: n = 189 cells from 7 mice.

Figure 7. Cell-type specific cue and reward activity are not due to licking movements.

(A) Schematic of the cued reward conditioning task (top) and the trial structure (bottom).

(B) Example licking behaviour during the ITIs from one mouse on day 1. Purple shading shows licks that were considered to be an individual lick bout. Gray shaded bar shows the CS timing and the red dotted line shows the reward timing.

(C) Z-scored activity of all the active PNs from an example mouse during one representative ITI lick bout. Left: maximum activity aligned to the lick bout onset. Right: maximum activity aligned to the CS onset. Grey bar represents the timing of the CS. White line indicates the time of water reward delivery.

(D-F) Mean reliability index for all cell types with significant CS-related responses during ITI lick bouts with no water reward present, compared to the mean reliability index during the CS and up to but not including the reward delivery time. As shown in figure 1, mice display anticipatory licking during the CS on day 7, but not day 1. PV-INs, VIP-INs and SOM-INs were more reliably responsive during the CS than during licking movement alone. Only sessions/cell-types with significant CS-related responses were analyzed. Paired t-test, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. PV-INs day 1: $p < 1 \times 10^{-3}$, PV-IN day 7: $p = 6.6 \times 10^{-3}$ **(D)**, VIP-IN day 1: $p = 0.0355$, VIP-IN day 7: $p < 1 \times 10^{-3}$ **(E)**, SOM-IN day 7: $p = 3.9 \times 10^{-3}$ **(F)**.

(G-I) Mean reliability index for all cell types with significant reward-related responses during ITI lick bouts with no water reward present, compared to the mean reliability index following reward timing. All cell types were more reliably responsive during the reward period than during licking movement alone. Only sessions/cell-types with significant reward-related responses were analyzed. Paired t-test, PV-INs day 1: $p = 1.6 \times 10^{-3}$, PV-IN day 7: $p = 1.6 \times 10^{-3}$ **(G)**, VIP-IN day 1: $p = 0.049$, VIP-IN day 7: $p = 0.014$ **(H)**, PN day 1 $p < 1 \times 10^{-3}$ **(I)**, SOM-IN day 7: $p = 0.030$ **(J)**.

PN: $n = 1029$ cells from 6 mice. PV-IN: $n = 316$ cells from 6 mice. VIP-IN: $n = 407$ cells from 4 mice. SOM-IN: $n = 189$ cells from 7. Error bars show SEM.

FIGURE SUPPLEMENT LEGENDS:

Figure 2 – Figure supplement 1. Z-scored ΔF traces from cells among each cell type on day 1 and day 7. Example calcium traces from 5 example cells (same mouse) during a representative ITI lick bout on day 1 (left), a trial on day 1 (middle) and a trial on day 7 (right). The same cells are tracked from day 1 to day 7. Blue dashed line indicates the time of the first lick in the lick bout (left). Grey shading represents time of the CS and the red dashed line indicates the time of the water reward delivery (middle, right).

Figure 2 – Figure Supplement 2. Z-Scored population activity of cells tracked from day 1 to day 7 for each cell type.

Z-scored activity of all the active neurons tracked across days during one representative trial on day 1 and day 7 for each cell type. Neurons were not sorted and the order of all cells were maintained for the day 1 and day 7 plot. Grey bar represents the timing of the CS. White line indicates the onset of water reward delivery. Black lines indicate licking.

Figure 4 – Figure supplement 1. Response properties of PV-IN and VIP-IN cells tracked from day 1 to 7 are consistent across days.

(A) Mean time from the CS onset to Ca^{2+} event onset for all PV-IN CS responses. Each line represents an individual mouse. The mean event onset time did not change from day 1 to day 7. Paired t-test, $p = 0.058$.

(B) Mean time from the first lick in response to reward delivery to Ca^{2+} event onset for all VIP-IN reward responses. Each line represents an individual mouse. The mean event onset time did not change from day 1 to day 7. Paired t-test, $p = 0.53$.

(C) Overall proportion of PV-INs that responded to the CS on day 1 did not change on day 7. Light red indicates the proportion of day 7 CS-responsive cells that also responded to the CS on day 1. Each line represents an individual mouse. Paired t-test, $p = 0.32$.

(D) PV-INs that responded to the CS on day 1, tracked to day 7. 97% of the cells that responded to CS on day 1 also responded to the CS on day 7. 3% of the cells that responded to the CS on day 1 became responsive to reward but not to the CS on day 7. Less than 1% responded to neither CS nor reward on day 7.

(E) Overall proportion of VIP-INs that responded to the reward stimuli on day 1 did not change on day 7. Light blue indicates the proportion of day 7 reward-responsive cells that also responded to reward on day 1. Each line represents an individual mouse. Paired t-test, $p = 0.32$.

(F) VIP-INs that responded to reward on day 1, tracked to day 7. 95% of the cells that responded to reward on day 1 also responded to reward on day 7. 4% of cells that responded to reward on day 1 became responsive to CS but not reward on day 7. Less than 1% responded to neither cue nor reward on day 7.

PV-IN: $n = 316$ cells from 6 mice. VIP-IN: $n = 407$ cells from 4 mice. Error bars show SEM.

Figure 4 – figure supplement 2. PV-IN and VIP-IN population plasticity is specific to associative learning.

(A) Trial structure for the no-reward control task. Gray shaded bar represents the response period analyzed for tone-responsive activity.

(B) Mean percent of tone-responsive PV-INs from the no-reward task. Violin plots show null distribution of percentage of responsive neurons made by re-sampling mice and shuffling the session 1,000 times (same as Figure 2). The circle represents the mean percentage of tone-responsive neurons. In the absence of reward, PV-INs did not show significant responses to tone. Bootstrap, day 1: $p = 0.11$, day 7: $p = 0.11$.

(C) The mean percent of tone-responsive PV-INs during no reward control do not change between day 1 and day 7. Paired t-test, n.s., non-significant, $p = 0.89$.

(D) Trial structure for the non-paired control task. Tone was followed by a randomly-varied 40-80s period before reward delivery and then a randomly varied ITI between 15-25s. First grey shaded bar shows the response period analyzed for tone-responsive activity (2.5s from tone onset). Second shaded bar shows the response period analyzed for reward-responsive activity (2.5s from onset of first lick after reward delivery).

(E) Mean lick rate during the tone from the non-paired task. Mice did not show increased anticipatory licking during the tone on day 7. One-way ANOVA, $p = 0.097$.

(F) Mean lick rate following the randomized water rewards from the non-paired task. Mice increased lick rate to consume the reward. The lick response was greater on day 7. One-way ANOVA, $p < 1 \times 10^{-3}$.

(G) Mean percent of tone-responsive PV-INs during the non-paired task. Violin plots show null distribution of percentage of responsive neurons made by re-sampling mice and shuffling the session 1,000 times. The circle represents the mean percentage of tone-responsive neurons. When tone was not paired with reward, PV-INs showed significant responses to tone on day 1 but not day 7. Bootstrap, Bonferroni correction for multiple comparisons, day 1 $p = 0.003$, day 7 $p = 0.069$.

(H) The mean percent of tone-responsive PV-INs during the non-paired task did not change between day 1 and day 7. Paired t-test, n.s., non-significant, $p = 0.43$.

(I) Mean reliability index of PV-INs that were responsive to the tone during the non-paired task. Cells were divided into High or Low Reliability Group based on the 50th percentile of tone reliability. The reliability of neither High nor Low Reliability PV-INs changed from day 1 to day 7. Paired t-test, n.s., non-significant, High: $p = 0.86$, Low: $p = 0.19$.

(J) Mean percent of reward-responsive PV-INs during the non-paired task. Violin plots show null distribution of percentage of responsive neurons made by re-sampling mice and shuffling the session 1,000 times. The circle represents the mean percentage of reward-responsive neurons. PV-INs responded significantly to randomly-timed rewards on day 1 and day 7. Bootstrap, Bonferroni correction for multiple comparisons, day 1: $p < 1 \times 10^{-3}$, day 7: $p < 1 \times 10^{-3}$.

(K) The mean percent of reward-responsive PV-INs during the non-paired task did not change between day 1 and day 7. Paired t-test, n.s., non-significant, $p = 0.51$

(L) Mean reliability index of PV-INs that were responsive to the reward. Cells were divided into High or Low Reliability Group based on the 50th percentile of reward reliability. The reliability of neither High nor Low Reliability PV-INs changed from day 1 to day 7. Paired t-test, n.s., non-significant, High: $p = 0.23$, Low: $p = 0.24$

(M) Mean percent of reward-responsive VIP-INs during the non-paired task. Violin plots show null distribution of percentage of responsive neurons made by re-sampling mice and shuffling the session 1,000 times. The circle represents the mean percentage of reward-responsive neurons. VIP-INs responded significantly to randomly timed rewards on day 1 and day 7. Bootstrap, day 1: $p < 1 \times 10^{-3}$, day 7: $p < 1 \times 10^{-3}$

(N) The mean percent of reward-responsive VIP-INs during the non-paired control did not change between day 1 and day 7. Paired t-test, n.s., non-significant, $p = 0.57$.

(O) Mean reliability index of VIP-INs that were responsive to the reward. Cells were divided into High or Low Reliability Group based on the 50th percentile of reward reliability. The reliability of neither High nor Low Reliability PV-INs changed from day 1 to day 7. Paired t-test, n.s., non-significant, High: $p = 0.53$, Low: $p = 0.22$.

No-water rewards: $n = 357$ PV-IN cells from 3 mice. Non-paired rewards: $n = 324$ PV-IN cells from 5 mice; $n = 208$ VIP-IN cells from 3 mice. Error bars show SEM.

Figure 7 – figure supplement 1. Lick cells were a small subset of active cells with stable mixed selectivity.

(A) The fraction of lick cells (in black) among all active cells tracked from day 1 to day 7 for each cell-type. Percentages indicate the percentage of cells that were non-lick cells. For all cell-types, the majority of cells were non-lick cells.

(B) Contingency table showing lick selectivity among all tracked active cells on day 1 and day 7. Individual cells were tracked across days to determine if licking-related activity changed across days. For all cell types, most cells were non-lick cells on day 1 and 7.

(C) Proportion of mixed selectivity for CS and reward among lick cells for all cell types on day 1 and 7.

(D) Percentage of lick cells with significant CS but not reward selectivity for all cell types. Each line represents an individual mouse.

(E) Percentage of lick cells with significant reward but not CS selectivity for all cell types. Each line represents an individual mouse.

(F) Percentage of lick cells with significant CS and reward selectivity for all cell types. Each line represents an individual mouse.

PN: n = 1029 cells from 6 mice. PV-IN: n = 316 cells from 6 mice. VIP-IN: n = 407 cells from 4 mice. SOM-IN: n = 189 cells from 7. Error bars show SEM.

Supplementary File 1

Summary of the number of mice, total ROIs, total active cells, and total active cells tracked from day 1 to day 7 in all experimental conditions.

ACKNOWLEDGEMENTS:

We thank the members of the Chen lab for discussions and providing feedback on the manuscript. This work was supported by grants for S.X.C. from Canada Research Chair (CRC) (Grant no. 950-231274) and Natural Sciences and Engineering Research Council of Canada (NSERC) (Grant no. 05308), and a grant for R. N. from NSERC (Grant no. 06972). E.H. was supported by a NSERC graduate scholarship. C.L. was supported by Ontario Graduate Scholarship and Queen Elizabeth II Graduate Scholarship.

COMPETING INTERESTS:

The authors declare no competing interests.

ETHICS:

All animal experiments were approved by the University of Ottawa Animal Care Committee (protocol #: CMM-2737) and in accordance with the Canadian Council on Animal Care guidelines.

MATERIALS AND METHODS

Mice

Experimental mice were group-housed in plastic cages with food and water ad libitum in a room with a reversed light cycle (12h - 12h). PV-Cre (008069), SOM-Cre (013044), VIP-Cre (010908), B6129SF1/J (101043) mouse lines were acquired from Jackson Laboratory (Bar Harbor, ME, USA). All mouse lines were homozygous and in C57BL/6 x 129S4 background. For all mouse lines, both male and females were used. Mice were between P40 and P60 at the time of surgery.

Surgery

Mice were deeply anesthetized under 1-2% isoflurane and given subcutaneous injections of Baytril (10 mg/kg) to prevent infection and buprenorphine (0.05 mg/kg) for analgesia. An incision was performed to remove a piece of the scalp and a custom head-plate was implanted onto the skull using instant glue (Krazy Glue) and dental cement (Lang Dental, Wheeling, IL, USA). A craniotomy of approximately 2 mm in diameter was performed over the right primary motor cortex. Virus (PNs: AAV1.CaMKII.GCaMP6f.WPRE.SV40; PV-IN, VIP-IN and SOM-IN: AAV1.Syn.Flex.GCaMP6f.WPRE.5v40) was diluted 1:5 in saline and injected at a depth of ~250 µm from the pia using a glass pipette. All virus was obtained from Addgene (Watertown, MA, USA). Injections were performed at 5 sites, centered on coordinates 1.5 mm lateral and 0.3 mm anterior to bregma. For PN groups, 20 nl per site was injected. For PV-IN, VIP-IN and SOM-IN, 40 nl per site was injected. All injections were performed at a rate of 10 nl/min and the pipette was left in place for 4 minutes following the injection to avoid backflow. A glass imaging window was then implanted over the craniotomy and sealed with dental cement. Following surgery, a subcutaneous injection of dexamethasone (2 mg/kg) and buprenorphine (0.1 mg/kg) was given. Mice were given a minimum of 1 week to recover prior to beginning water restriction.

Auditory Cued Reward Conditioning Behaviour

Mice were gradually water restricted down to ~1 ml per day and were maintained at ~80% of original body weight over two weeks prior to the start of imaging/behaviour sessions (Chen et al., 2015; Harvey, Coen, & Tank, 2012; Komiyama et al., 2010; O'Connor et al., 2013; Peters et al., 2014). Mice were then head-fixed for simultaneous two-photon imaging and exposed to the conditioned stimulus (a constant 9kHz auditory tone, 1s in duration) followed by a 1.5s delay period and a water reward (~ 10 µl). All lick times were measured by an infrared beam lick-o-meter and logged using the data acquisition software WaveSurfer (<https://wavesurfer.janelia.org/>). The inter-trial interval (ITI) between the previous water reward and subsequent CS onset was randomly varied between 60-120s. Each session was one hour in duration with 30-35 trials in total. Mice underwent one session per day for seven consecutive days. Two-photon calcium image was performed simultaneously on day 1 and day 7 of the behavioural task.

To assess licking behaviour, lick rate (number of licks per second, measured as infrared beam breaks) was calculated within 500ms bins, then averaged across all trials within a session for each mouse. Lick rate was then averaged across mice. Mean anticipatory lick rate was calculated as the mean lick rate from the time of the CS onset to the end of the delay period (2.5s in duration), not including the reward delivery. Mean ITI lick rate was calculated from the lick rate during the first 2.5s of self-initiated spontaneous lick bouts. ITI lick bouts were defined as licking events that followed the previous trial by at least 20s and preceded the subsequent trial by more than 2.5s. Mean reward lick rate was calculated from the lick rate from the time of reward delivery to 2.5s after.

All trials within a session were included in lick rate analysis in Figure 1. To ensure behavioural consistency across trials, only trials with at least 3 lick responses within 2.5s of the reward delivery time were included in all analysis of neural responses.

In the control experiments with water rewards omitted, mice were head-fixed and exposed to a constant 9kHz auditory tone, 1s in duration, followed by a randomly varied ITI between 15-25s. In the non-paired control experiments, mice were exposed to a constant 9kHz auditory tone, followed by a randomly varied delay period between 40-80s before delivery of a water reward (Krabbe et al., 2019). Water rewards were then followed by a 15-25s ITI. In both tasks, each session was 45min in duration with an average of 30 trials. Mice underwent one session per day for seven consecutive days.

Calcium Imaging and Analysis

In vivo imaging was performed using a commercial two-photon microscope (B-scope, Thorlabs, Newton, NJ, USA) and a 16x water immersion objective (Nikon) with excitation at 925 nm (InSight X3, Spectra-Physics, Milpitas, CA, USA) with a frame rate of 30 Hz. Images were taken at 512 x 512 pixels covering 755 by 650 μm .

Images were corrected for movement in the x and y plane using full-frame cross-correlation image alignment (Turboreg (Thévenaz, Ruttimann, & Unser, 1998) plug-in ImageJ). The entire session was visually inspected and regions of interests (ROIs) were manually drawn on neurons using a custom MATLAB program, described in Peters et al., 2014. The ROI template from day 1 was loaded onto day 7 and aligned along the x and y plane. Only ROIs that could be tracked from day 1 to day 7 were included in the dataset unless otherwise specified.

Fluorescence within an ROI was averaged across pixels. ΔF was calculated by subtracting a time-varying baseline fluorescence estimate (F_0) from the raw fluorescence trace. The calculation for baseline fluorescence (F_0) was calculated iteratively and based on inactive parts of the fluorescence trace as previously described (Chu et al., 2016; Kato et al., 2012; Peters et al., 2014; Peters et al., 2017)

We adapted a method by Driscoll et al. (Driscoll et al., 2017) to identify significant activity events for each neuron and then excluded ROIs with no significant activity events within the session, irrespective of the behaviour. For each neuron the ΔF trace was circularly shifted by a random integer 1,000 times and compared to the original trace. If the original ΔF trace was greater than the shifted data for at least 5 consecutive frames in at least 950 iterations, this was considered an active event. If a neuron did not have at least one active event in the entire session, irrespective of the behaviour, it was removed from the data set. This only accounted for a small proportion of ROIs as most of them are active on both day 1 and day 7, as shown in Figure 2D.

For all subsequent analyses, a modified z-score, adapted from Kato et al. (2015), was applied to ΔF . The z-score was calculated as $Z = (f(t) - \mu)/\sigma$, where $f(t)$ is the ΔF trace for a neuron, μ is the mean, and σ is the standard deviation of the neuron's ΔF during the baseline period. The baseline period was a concatenation of 2.5s preceding the CS onset (start of a trial) for all trials within a session.

Calculation of Tuning Coefficients

We quantified the tuning of individual neurons to the CS and reward stimuli delivered in our classical conditioning task using the non-parametric Spearman correlation ρ (scipy.stats.spearmanr) between the trial-averaged fluorescence and the timing of stimulus delivery

$$\rho_{m,n,s}^{(d)} = \text{corr} \left[\frac{1}{|\mathcal{T}_m^{(d)}|} \sum_{t \in \mathcal{T}_m^{(d)}} \mathbf{f}_{m,n,t-T_{\text{baseline}}:t+T_{\text{CS}}+T_{\text{delay}}+T_{\text{reward}}+T_{\text{post}}}^{(d)}, \mathbf{1}_s \right],$$

where t is the start time of a trial (defined as the start of the CS), $\mathcal{T}_m^{(d)}$ is the set of all trial start times from mouse m on day $d \in \{1, 7\}$, $\mathbf{f}_{m,n,t-T_{\text{baseline}}:t+T_{\text{CS}}+\dots}^{(d)}$ is the fluorescence trace of neuron n from mouse m during a single trial, $\mathbf{1}_s$ is an indicator function for stimulus $s \in \{\text{CS}, \text{reward}\}$, $|\mathcal{T}_m^{(d)}|$ is the number of trials, and ρ is the Spearman correlation coefficient. Analysis was carried out with $T_{\text{baseline}} = 2$ s and $T_{\text{post}} = 6$ s. We considered the “CS” period indicated by $\mathbf{1}_{\text{CS}}$ to range from the start of the CS at time t to the start of reward delivery at time $t + T_{\text{CS}} + T_{\text{delay}}$, and the “reward” period indicated by $\mathbf{1}_{\text{reward}}$ to be the first 2.5 seconds of reward delivery (see schematic in Figure 4A, 5A). We used the change in ρ from day 1 to day 7 as a cell-resolved measure of changes in tuning over the course of learning.

To summarize learning-associated changes in tuning, we calculated the mean change in the Spearman correlation for each cell type and trial component (CS or reward) from day 1 to day 7 as follows

$$\overline{\Delta\rho}_s = \frac{1}{|\mathcal{M}|} \sum_{m \in \mathcal{M}} \frac{1}{N_m} \sum_{n=1}^{N_m} (\rho_{m,n,s}^{(7)} - \rho_{m,n,s}^{(1)}),$$

where $\mathcal{M}$ is the set of mice used in the experiment, N_m is the number of neurons in mouse m , and $\rho_{m,n,s}^{(d)}$ is the Spearman correlation as defined above.

We used a non-parametric approach for statistical tests involving the mean change in Spearman correlation by scrambling trial times and bootstrapping mice to construct a null distribution for $\overline{\Delta\rho}_s$. Specifically, we first drew a random sample of $|\mathcal{M}|$ mice from $\mathcal{M}$ with replacement, then drew a random sample of $|\mathcal{T}_m^{(d)}|$ trial start times uniformly distributed between 0 and $T_{\text{session}}^{(d)} - (T_{\text{baseline}} + T_{\text{CS}} + T_{\text{delay}} + T_{\text{reward}} + T_{\text{post}})$ for each day d and randomly-selected mouse, and finally used these randomly-selected mice and scrambled trial start times to compute the change in tuning $\overline{\Delta\rho}_s$. This process was repeated 1000 times to approximate the distribution of $\overline{\Delta\rho}_s$ under the null hypothesis that changes in tuning are unrelated to the CS and reward delivery. We considered the observed changes in tuning $\overline{\Delta\rho}_s$ to be statistically significant at the * or ** level if they fell into the 5 % or 1 % tails of this distribution, respectively.

Activity Analysis

To identify neuron responses to the CS and reward, we applied a set threshold to each neuron. Neurons were defined as CS-responsive or reward-responsive on a trial-by-trial basis if they exceeded 1 z-score (excitation threshold used in Kato et al., 2015) for at least 5 consecutive frames within 2.5s of the CS onset or 2.5s of the reward delivery time, respectively. This was assessed for each trial with at least 3 lick responses within 2.5 s of the reward delivery time. We then took the median of the percent of responsive neurons across all trials in a session from one mouse, and the mean across mice. In the non-paired experiments, since the reward was not preceded by a CS, the mice required variable amounts of time to notice the water reward. Therefore, in the non-paired experiments, we calculated reward responses within 2.5s from the onset of the first lick following the water delivery. In the case that no licks were recorded before the subsequent tone, the trial was not included in the reward analysis.

We used a Monte-Carlo approach to validate the percent of CS and reward responsive neurons. The mean percentage of CS-responsive and reward-responsive neurons observed were compared to a null distribution made for each cell type on each day. We randomly sampled mice with replacement, then sampled the entire session, and then calculated the percentage of active cells (exceeding 1 z-score for at least 5 consecutive frames) during a randomly chosen 2.5s window. For each mouse, the number of samples was equal to the number of included trials (i.e., number of trials with at least 3 lick responses within 2.5s of reward delivery). We then took the median across the random samples and then took the mean across mice to obtain a mean percentage of responsive neurons during a randomly chosen time window. This was repeated 1000 times to generate a null distribution of mean percentage of active neurons. To assess whether the observed percentage of CS- and reward- responsive neurons was significantly different from the null distribution, the observed value was compared to the tails of the null distribution. This was done for each cell type on both day 1 and day 7. We considered the CS- or reward- responses to be statistically significant at the * or ** level if they fell into the 5% or 1% tails of this distribution, respectively and *** if there was no overlap with the distribution. Since this approach tests the null hypothesis that the observed neuronal responses are due to chance (in this case, baseline activity/noise), only cell types with a significantly higher percentage of responsive neurons for a given session were analyzed further.

The CS/reward reliability index was defined as the percentage of trials within a session where the neuron was CS/reward responsive. The reliability cumulative distribution was made by pooling the day 1 index values of all the neurons from a neuronal cell type (across mice). If a neuron's day 1 index value was lower or equal to the index value at the 50th percentile of the cumulative distribution (excluding non-responsive neurons with a reliability of 0) for that cell type, it was categorized into the Low Reliability group. If a neuron's day 1 index value exceeded the 50th percentile value, it was categorized into the High Reliability group. To assess changes in reliability at the population level, we took the mean reliability within each group on day 1 and day 7. To assess changes in reliability among individual Low reliability neurons, we used a $\Delta reliability$ measure where,
 $\Delta reliability_{CS} = (CS\ reliability_{Day\ 7} - CS\ reliability_{Day\ 1})$ and
 $\Delta reliability_{reward} = (reward\ reliability_{Day\ 7} - reward\ reliability_{Day\ 1})$. As a control, we randomly sampled the day 7 session matching the number of trials, and calculated the reliability to obtain a "random reliability" for each neuron. We then calculated a $\Delta reliability_{random}$ where,
 $\Delta reliability_{Random} = (reliability_{random} - reliability_{Day\ 1})$ for CS and reward reliabilities.

The onset time of neuronal activity following the CS was calculated as the time from the CS onset to the time of the first Ca^{2+} event (fifth frame above threshold) within the CS response period. For the onset time of reward-related neuronal activity, the time from the first lick (after reward delivery) to the time of the first Ca^{2+} event (fifth frame above threshold) within the reward response period was used. The latency for each cell was first calculated by taking the mean across all active trials for a single cell, then the median of all cells within a mouse was calculated. Only cells that were tracked between days were included.

To determine the proportion of PV-INs that responded to the CS on day 1 and day 7, we found the overall proportion of cells that responded to the CS out of total active cells on both day 1 and 7. To calculate the percentage of PV-INs that maintained CS responses across days, we found the proportion of day 7 cells that also had CS responses on day 1. We also found the proportion of day 1 cells with CS responses, that either maintained CS responses on day 7, became CS unresponsive but reward responsive, or became unresponsive to both CS and reward. The same analysis was performed on VIP-IN reward responses. Only active cells that were tracked from day 1 to 7 were included.

Licking-related Analysis

ITI lick bouts were defined as self-initiated licking events that occurred at least 20s after the preceding reward delivery time (trial end) and more than 2.5s prior to the subsequent CS onset (trial start). If individual licks were separated by 3s or more, they were considered to be a new lick bout. To remain consistent with CS and reward analyses, only the first 2.5s of a lick bout were analyzed for neural responses. ITI lick bout reliability indices were calculated as described above.

To determine lick cells, we found the first ITI lick bout in each ITI and calculated the mean z-scored ΔF during the first 2.5s of the lick bout. We then created a matrix of concatenated ITIs from the same session and randomly sampled the concatenated ITIs, and calculated the mean z-scored ΔF , matching the duration and number of ITI lick bouts. A paired t-test was used to compare the mean z-scored ΔF during ITI lick bouts and during the random samples. Cells with significantly higher ΔF during lick bouts were considered to be lick cells. We performed this analysis on day 1 and day 7 and tracked individual neurons to identify changes in selectivity on a cell-by-cell basis. To determine if cells have mixed selectivity, we performed the same analysis using the CS and reward response periods (2.5s from onset) and compared this activity to an equal number of random samples using a paired t-test.

Statistical Analysis

Statistical analysis for tuning coefficients were performed in Python and in R. All other statistical analyses were performed in Matlab using the Statistics and Machine Learning Toolbox. Two-way ANOVA was used to test for differences in anticipatory lick rate on day 1 and day 7. One-way ANOVA was used to test for differences in lick rate during ITI, CS and reward. One-way ANOVA was used to compare the percent of active cells across cell types on a single day. Monte Carlo (as described above) was used to test for significant percent of CS- and reward-responsive neurons, and for changes in tuning properties. Paired t-test was used to test for differences in the percentage of responsive cells and reliability index on day 1 and 7, and for differences in neuron reliability between ITI lick bouts, CS and reward. Paired t-test was used to determine mixed selectivity as described above. The Kolmogorov–Smirnov test was used to compare response reliability cumulative distributions and Δ reliability distributions. All values were reported as the mean and standard error of the mean unless otherwise specified. Power analysis was not performed to predetermine the sample size, and the experiments were not blinded.

Data Analysis and Code Availability

Tuning coefficient calculation and statistical tests were performed using Python 3.8 with the following libraries: NumPy, Pandas, h5py, and SQLAlchemy. Figures were prepared in Python using matplotlib and seaborn, and in R using ggplot2. Codes to reproduce the analysis for figures 1-2 and 4-7 are available at <https://github.com/clee162/Analysis-of-Cell-type-Specific-Responses-to-Associative-Learning-in-M1>. Codes to reproduce the analysis and figure 3 are available at <https://github.com/nauralcodinglab/interneuron-reward>. Data can be found on Dryad at <https://doi.org/10.5061/dryad.q573n5tjj>.

947 **REFERENCES**

- 948 Abe, M., Schambra, H., Wassermann, E. M., Luckenbaugh, D., Schweighofer, N., & Cohen, L. G. (2011).
 949 Reward improves long-term retention of a motor memory through induction of offline memory
 950 gains. *Curr Biol*, 21(7), 557-562. doi:10.1016/j.cub.2011.02.030
- 951 Atallah, B. V., Bruns, W., Carandini, M., & Scanziani, M. (2012). Parvalbumin-expressing interneurons
 952 linearly transform cortical responses to visual stimuli. *Neuron*, 73(1), 159-170.
 953 doi:10.1016/j.neuron.2011.12.013
- 954 Chen, S. X., Kim, A. N., Peters, A. J., & Komiyama, T. (2015). Subtype-specific plasticity of inhibitory
 955 circuits in motor cortex during motor learning. *Nat Neurosci*, 18(8), 1109-1115.
 956 doi:10.1038/nn.4049
- 957 Chu, M. W., Li, W. L., & Komiyama, T. (2016). Balancing the Robustness and Efficiency of Odor
 958 Representations during Learning. *Neuron*, 92(1), 174-186. doi:10.1016/j.neuron.2016.09.004
- 959 Cichon, J., & Gan, W. B. (2015). Branch-specific dendritic Ca(2+) spikes cause persistent synaptic
 960 plasticity. *Nature*, 520(7546), 180-185. doi:10.1038/nature14251
- 961 Donato, F., Rompani, S. B., & Caroni, P. (2013). Parvalbumin-expressing basket-cell network plasticity
 962 induced by experience regulates adult learning. *Nature*, 504(7479), 272-276.
 963 doi:10.1038/nature12866
- 964 Driscoll, L. N., Pettit, N. L., Minderer, M., Chettih, S. N., & Harvey, C. D. (2017). Dynamic Reorganization
 965 of Neuronal Activity Patterns in Parietal Cortex. *Cell*, 170(5), 986-999.e916.
 966 doi:10.1016/j.cell.2017.07.021
- 967 Duan, Z., Li, A., Gong, H., & Li, X. (2020). A Whole-brain Map of Long-range Inputs to GABAergic
 968 Interneurons in the Mouse Caudal Forelimb Area. *Neurosci Bull*, 36(5), 493-505.
 969 doi:10.1007/s12264-019-00458-6
- 970 Fishell, G., & Rudy, B. (2011). Mechanisms of inhibition within the telencephalon: "where the wild things
 971 are". *Annu Rev Neurosci*, 34, 535-567. doi:10.1146/annurev-neuro-061010-113717
- 972 Fu, Y., Tucciarone, J. M., Espinosa, J. S., Sheng, N., Darcy, D. P., Nicoll, R. A., . . . Stryker, M. P. (2014). A
 973 cortical circuit for gain control by behavioral state. *Cell*, 156(6), 1139-1152.
 974 doi:10.1016/j.cell.2014.01.050
- 975 Galea, J. M., Mallia, E., Rothwell, J., & Diedrichsen, J. (2015). The dissociable effects of punishment and
 976 reward on motor learning. *Nat Neurosci*, 18(4), 597-602. doi:10.1038/nn.3956
- 977 Gasselin C, Hohl B, Vernet A, Crochet S, Petersen CCH. Cell-type-specific nicotinic input disinhibits mouse
 978 barrel cortex during active sensing. *Neuron*. 2021 Mar 3;109(5):778-787.e3. doi:
 979 10.1016/j.neuron.2020.12.018.
- 980 Georgopoulos, A. P., Ashe, J., Smyrnis, N., & Taira, M. (1992). The motor cortex and the coding of force.
 981 *Science*, 256(5064), 1692-1695. doi:10.1126/science.256.5064.1692
- 982 Harvey, C. D., Coen, P., & Tank, D. W. (2012). Choice-specific sequences in parietal cortex during a
 983 virtual-navigation decision task. *Nature*, 484(7392), 62-68. doi:10.1038/nature10918
- 984 Jackson, J., Ayzenshtat, I., Karnani, M. M., & Yuste, R. (2016). VIP+ interneurons control neocortical
 985 activity across brain states. *J Neurophysiol*, 115(6), 3008-3017. doi:10.1152/jn.01124.2015
- 986 Kato, H. K., Chu, M. W., Isaacson, J. S., & Komiyama, T. (2012). Dynamic sensory representations in the
 987 olfactory bulb: modulation by wakefulness and experience. *Neuron*, 76(5), 962-975.
 988 doi:10.1016/j.neuron.2012.09.037
- 989 Kato, H. K., Gillet, S. N., & Isaacson, J. S. (2015). Flexible Sensory Representations in Auditory Cortex
 990 Driven by Behavioral Relevance. *Neuron*, 88(5), 1027-1039. doi:10.1016/j.neuron.2015.10.024

- Komiyama, T., Sato, T. R., O'Connor, D. H., Zhang, Y. X., Huber, D., Hooks, B. M., . . . Svoboda, K. (2010). Learning-related fine-scale specificity imaged in motor cortex circuits of behaving mice. *Nature*, 464(7292), 1182-1186. doi:10.1038/nature08897
- Krabbe, S., Paradiso, E., d'Aquin, S., Bitterman, Y., Courtin, J., Xu, C., . . . Lüthi, A. (2019). Adaptive disinhibitory gating by VIP interneurons permits associative learning. *Nat Neurosci*, 22(11), 1834-1843. doi:10.1038/s41593-019-0508-y
- Lee, C., Lavoie, A., Liu, J., Chen, S. X., & Liu, B. H. (2020). Light Up the Brain: The Application of Optogenetics in Cell-Type Specific Dissection of Mouse Brain Circuits. *Front Neural Circuits*, 14, 18. doi:10.3389/fncir.2020.00018
- Lee, K., Holley, S. M., Shobe, J. L., Chong, N. C., Cepeda, C., Levine, M. S., & Masmanidis, S. C. (2017). Parvalbumin Interneurons Modulate Striatal Output and Enhance Performance during Associative Learning. *Neuron*, 93(6), 1451-1463.e1454. doi:10.1016/j.neuron.2017.02.033
- Lee, S. H., Kwan, A. C., Zhang, S., Phoumthipphavong, V., Flannery, J. G., Masmanidis, S. C., . . . Dan, Y. (2012). Activation of specific interneurons improves V1 feature selectivity and visual perception. *Nature*, 488(7411), 379-383. doi:10.1038/nature11312
- Levy, S., Lavzin, M., Benisty, H., Ghanayim, A., Dubin, U., Achvat, S., . . . Schiller, J. (2020). Cell-Type-Specific Outcome Representation in the Primary Motor Cortex. *Neuron*, 107(5), 954-971.e959. doi:10.1016/j.neuron.2020.06.006
- Makino, H., & Komiyama, T. (2015). Learning enhances the relative impact of top-down processing in the visual cortex. *Nat Neurosci*, 18(8), 1116-1122. doi:10.1038/nn.4061
- Markram, H., Toledo-Rodriguez, M., Wang, Y., Gupta, A., Silberberg, G., & Wu, C. (2004). Interneurons of the neocortical inhibitory system. *Nat Rev Neurosci*, 5(10), 793-807. doi:10.1038/nnr1519
- Marsh, B. T., Tarigoppula, V. S., Chen, C., & Francis, J. T. (2015). Toward an autonomous brain machine interface: integrating sensorimotor reward modulation and reinforcement learning. *J Neurosci*, 35(19), 7374-7387. doi:10.1523/JNEUROSCI.1802-14.2015
- Moran, D. W., & Schwartz, A. B. (1999). Motor cortical representation of speed and direction during reaching. *J Neurophysiol*, 82(5), 2676-2692. doi:10.1152/jn.1999.82.5.2676
- Muñoz, W., Tremblay, R., Levenstein, D., & Rudy, B. (2017). Layer-specific modulation of neocortical dendritic inhibition during active wakefulness. *Science*, 355(6328), 954-959. doi:10.1126/science.aag2599
- Nikooyan, A. A., & Ahmed, A. A. (2015). Reward feedback accelerates motor learning. *J Neurophysiol*, 113(2), 633-646. doi:10.1152/jn.00032.2014
- O'Connor, D. H., Hires, S. A., Guo, Z. V., Li, N., Yu, J., Sun, Q. Q., . . . Svoboda, K. (2013). Neural coding during active somatosensation revealed using illusory touch. *Nat Neurosci*, 16(7), 958-965. doi:10.1038/nn.3419
- Peters, A. J., Chen, S. X., & Komiyama, T. (2014). Emergence of reproducible spatiotemporal activity during motor learning. *Nature*, 510(7504), 263-267. doi:10.1038/nature13235
- Peters, A. J., Lee, J., Hedrick, N. G., O'Neil, K., & Komiyama, T. (2017). Reorganization of corticospinal output during motor learning. *Nat Neurosci*, 20(8), 1133-1141. doi:10.1038/nn.4596
- Pfeffer, C. K., Xue, M., He, M., Huang, Z. J., & Scanziani, M. (2013). Inhibition of inhibition in visual cortex: the logic of connections between molecularly distinct interneurons. *Nat Neurosci*, 16(8), 1068-1076. doi:10.1038/nn.3446
- Pi, H. J., Hangya, B., Kvitsiani, D., Sanders, J. I., Huang, Z. J., & Kepecs, A. (2013). Cortical interneurons that specialize in disinhibitory control. *Nature*, 503(7477), 521-524. doi:10.1038/nature12676
- Poort J, Wilmes KA, Blot A, Chadwick A, Sahani M, Clopath C, Mrcic-Flogel TD, Hofer SB, Khan AG. Learning and attention increase visual response selectivity through distinct mechanisms. *Neuron*. 2021 Dec 9;S0896-6273(21)00954-5. doi: 10.1016/j.neuron.2021.11.016.

- Ramakrishnan, A., Byun, Y. W., Rand, K., Pedersen, C. E., Lebedev, M. A., & Nicolelis, M. A. L. (2017). Cortical neurons multiplex reward-related signals along with sensory and motor information. *Proc Natl Acad Sci U S A*, 114(24), E4841-E4850. doi:10.1073/pnas.1703668114
- Ramkumar, P., Dekleva, B., Cooler, S., Miller, L., & Kording, K. (2016). Premotor and Motor Cortices Encode Reward. *PLoS One*, 11(8), e0160851. doi:10.1371/journal.pone.0160851
- Richards, B. A., Lillicrap, T. P., Beaudoin, P., Bengio, Y., Bogacz, R., Christensen, A., . . . Kording, K. P. (2019). A deep learning framework for neuroscience. *Nat Neurosci*, 22(11), 1761-1770. doi:10.1038/s41593-019-0520-2
- Seybold, B. A., Phillips, E. A. K., Schreiner, C. E., & Hasenstaub, A. R. (2015). Inhibitory Actions Unified by Network Integration. *Neuron*, 87(6), 1181-1192. doi:10.1016/j.neuron.2015.09.013
- Thabit, M. N., Nakatsuka, M., Koganemaru, S., Fawi, G., Fukuyama, H., & Mima, T. (2011). Momentary reward induce changes in excitability of primary motor cortex. *Clin Neurophysiol*, 122(9), 1764-1770. doi:10.1016/j.clinph.2011.02.021
- Thévenaz, P., Ruttimann, U. E., & Unser, M. (1998). A pyramid approach to subpixel registration based on intensity. *IEEE Trans Image Process*, 7(1), 27-41. doi:10.1109/83.650848
- Turi, G. F., Li, W. K., Chavlis, S., Pandi, I., O'Hare, J., Priestley, J. B., . . . Losonczy, A. (2019). Vasoactive Intestinal Polypeptide-Expressing Interneurons in the Hippocampus Support Goal-Oriented Spatial Learning. *Neuron*, 101(6), 1150-1165.e1158. doi:10.1016/j.neuron.2019.01.009
- Wilson, N. R., Runyan, C. A., Wang, F. L., & Sur, M. (2012). Division and subtraction by distinct cortical inhibitory networks in vivo. *Nature*, 488(7411), 343-348. doi:10.1038/nature11347
- Wolff, S. B., Gründemann, J., Tovote, P., Krabbe, S., Jacobson, G. A., Müller, C., . . . Lüthi, A. (2014). Amygdala interneuron subtypes control fear learning through disinhibition. *Nature*, 509(7501), 453-458. doi:10.1038/nature13258
- Wood, K. C., Blackwell, J. M., & Geffen, M. N. (2017). Cortical inhibitory interneurons control sensory processing. *Curr Opin Neurobiol*, 46, 200-207. doi:10.1016/j.conb.2017.08.018
- Xu, T., Yu, X., Perlik, A. J., Tobin, W. F., Zweig, J. A., Tennant, K., . . . Zuo, Y. (2009). Rapid formation and selective stabilization of synapses for enduring motor memories. *Nature*, 462(7275), 915-919. doi:10.1038/nature08389
- Xu, X., & Callaway, E. M. (2009). Laminar specificity of functional input to distinct types of inhibitory cortical neurons. *J Neurosci*, 29(1), 70-85. doi:10.1523/JNEUROSCI.4104-08.2009
- Xue, M., Atallah, B. V., & Scanziani, M. (2014). Equalizing excitation-inhibition ratios across visual cortical neurons. *Nature*, 511(7511), 596-600. doi:10.1038/nature13321
- Zhang, S., Xu, M., Kamigaki, T., Hoang Do, J. P., Chang, W. C., Jenvay, S., . . . Dan, Y. (2014). Selective attention. Long-range and local circuits for top-down modulation of visual cortex processing. *Science*, 345(6197), 660-665. doi:10.1126/science.1254126

Lee et al., Figure 2

Lee et al., Figure 2 - figure supplement 1

Lee et al., Figure 2
- figure supplement 2

A

B

C

D

E

F