CELL-TYPE SPECIFIC RESPONSES TO REINFORCEMENT IN THE PRIMARY MOTOR CORTEX

CANDICE LEE

Thesis submitted to the University of Ottawa
in partial fulfillment of the requirements for the
Doctor of Philosophy in Neuroscience

September 29, 2022
Department of Cellular and Molecular Medicine
Faculty of Medicine
University of Ottawa

ABSTRACT

The primary motor cortex (M1) is an important site for learning new motor skills. While reward is known to both enhance and accelerate motor learning, the mechanism by which reward exerts these effects remains unclear. Previous studies in primates have demonstrated reward-related activity in M1, however, it is not known whether reward is represented among different neuronal cell types in M1, or if the representations change over the course of reward-based associative learning. We begin by reviewing advances in optogenetic methods that have enabled the dissection of cortical circuits underlying sensorimotor behaviours with a special focus on the functional roles of cell-type specific connections in governing sensorimotor information processing and learning and memory. We then used *in vivo*, two-photon calcium imaging to characterize reward and reward-related responses in pyramidal neurons (PNs), PV-INs, SST-INs and VIP-INs while mice simultaneously performed a head-fixed auditory classical conditioning task. We found that different cell types had distinct responses to the conditioned stimulus (CS) and to reward, and these responses underwent differential changes over the course of associative learning. Notably, VIP-INs preferentially represented reward and their reward responses increased with learning, while PV-INs preferentially represented the CS, and their CS responses increased with learning. Lastly, to identify which brain regions might provide reward-related input to VIP-INs, we performed cell-type specific monosynaptic rabies tracing and generated comparative brain-wide maps of input to VIP-INs, PV-INs, SST-INs and PNs in M1. We identified preferential input from the orbital frontal cortex (ORB) to VIP-INs compared to the other IN subtypes. These results suggest that ORB may convey reward-related input to VIP-INs and thereby disinhibit local MOP circuitry during reward-based learning. Together, these studies provide a potential mechanism for how reward modulates motor learning.

ACKNOWLEDGEMENTS

I would first like to thank my supervisor Dr. Simon Chen for years of attentive mentorship, guidance, and support. It has been a great privilege to be your first student and I am profoundly grateful for your supervision. Your mentorship has led me to grow as a scientist and as a person.

Thank you to my thesis advisory committee members Dr. Len Maler, Dr. Jean-Claude Béïque, Dr. Baptiste Lacoste and Dr. Richard Naud for invaluable feedback, discussion, and support. I have always been appreciative to have such a warm, approachable, and extremely knowledgeable TAC. A special thank you to Len for inviting me into your lab as an undergraduate student many years ago and thus fostering the beginning of my career in neuroscience research. A special thank you to Jean-Claude for your open, honest, and remarkably astute mentorship – your compassion and encouragement have energized and motivated me countless times.

A sincere thank you to each of my dear lab mates. To Irina Morozov for technical support throughout my project. To Dr. Sandrine Côté for being a superb collaborator and co-author. To Dr. Xuming Yin for collaboration and helpful discussion. To Liwen Cai for your careful dedication in managing our mouse colony. To Dr. Jungwoo Yang for always sharing your incredible expertise and never hesitating to help others. To Pablo Serrano for bringing so much joy and laughter to the lab. To Jan Sidiangco for a lovely new friendship and for all your help planning my transatlantic move. To Nima Raman for the many hours spent managing our lab, and for being a ray of sunshine – a constant source of warmth and vitality. To Nate Jones, for helpful discussion and a friendship that will always be dear to my heart. My labmates have become my family.

To my friends, Sébastien Maillé, Michael Lynn, and Emerson Harkin, thank you for your unconditional support, impeccable humour, and for readily sharing your extensive wealth of knowledge; I am constantly learning from each of you. To Marc Vani for encouraging me to try research and for helping me get my feet on the ground. To Dr. Tanya Foley for orienting me in the lab at the beginning of my PhD, for making late nights in the lab genuinely enjoyable, and for being a friend with whom I can candidly share this experience with. To Olanta Negeri for endless support and companionship; you are a true friend, and I am lucky to have you in my corner. Each of these wonderful friends/remarkable scientists fill me with admiration and gratitude.

To Dr. David Shulman, thank you for your kindness and patience in taking the time to teach a curious child what a neuron is.

To Phoebe, thank you for being my tireless study buddy, regardless of the hour. Your companionship means the world to me.

To my family, I thank my sister Shannon for her loving support. Finally, to my parents, I sincerely thank you for the innumerable sacrifices you've made to always prioritize my education. I express my deepest gratitude for your unconditional love, encouragement, patience, and support.

TABLE OF CONTENTS

LIST OF TABLES

Chapter 2

LIST OF FIGURES

Chapter 1

Chapter 2

Chapter 3

LIST OF ABBREVIATIONS

AAV	Adeno-associated virus
ACA	Anterior cingulate area
ACh	Acetylcholine
AI	Agranular insular area
ALM	Anterior lateral motor area
AP	Anterior-posterior
AM	Anteromedial Nucleus of thalamus
AUD	Auditory cortex
CFA	Caudal forelimb area
ChR2	Channelrhodopsin

CL	Central lateral nucleus of the thalamus
CLA	Claustrum
CM	Central medial nucleus of the thalamus
Contra	Contralateral
CT	Corticothalamic
DREADDs	Designer receptors exclusively activated by designer drugs
DV	Dorsal-ventral
EnvA	Envelope A
GABA	Gamma aminobutyric acid
IT	Intratelencephalic
Ipsi	Ipsilateral
L1	Layer 1
L2/3	Layer 2/3
L4	Layer 4
L5	Layer 5
L6	Layer 6
LD	Lateral dorsal nucleus of thalamus
LH	Lateral habenula
LTP	Long-term potentiation
M1	Primary motor cortex
M2	Secondary motor cortex
MD	Mediodorsal nucleus
ML	Medial-lateral

MOp	Primary motor cortex
MOs	Secondary motor cortex
NE	Norepinephrine
ORB	Orbital frontal cortex,
ORBl	Orbital frontal cortex, lateral part
ORBm	Orbital frontal cortex, medial part
ORBvl	Orbital frontal cortex, ventrolateral part
PCN	Paracentral nucleus
PF	Parafascicular nucleus
PL	Prelimbic area
PN	Pyramidal neuron
PO	Posterior complex of thalamus
POm	Medial posterior complex of thalamus
PT	Pyramidal Tract
PTLp	Posterior parietal association areas
PV-IN	Parvalbumin expressing interneuron
RSP	Retrosplenial cortex
RSPagl	Retrosplenial cortex, lateral agranular part
RSPd	Retrosplenial cortex, dorsal part
RSPv	Retrosplenial cortex, ventral part
RT	Reticular nucleus of the thalamuss
RV	Rabies virus
S1	Primary somatosensory cortex

SOM-IN	Somatostatin expressing interneuron
SS	Somatosensory cortex
SSp	Primary somatosensory cortex
SSs	Secondary somatosensory cortex
SST-IN	Somatostatin expressing interneuron
TC	Thalamocortical
TEa	Temporal association areas
VA	Ventral anterior nucleus of thalamus
VAL	Ventral anterior-lateral complex of the thalamus
VL	Ventral lateral nucleus of thalamus
VIP-IN	Vasoactive intestinal peptide expressing interneuron
VM	Ventral medial nucleus of the thalamus
vM1	Vibrissal area of primary motor cortex
VP	Ventral posterior complex of the thalamus
ΔG	Glycoprotein-deleted

GENERAL INTRODUCTION

1. The Primary Motor Cortex

Volitional movement is a means to interact with the world. Movement allows us to traverse through space, manipulate objects, sense your environment, and provide for your needs. Beyond our necessities, movement and the ability to learn new skilled movements is integral to our society; for example, a surgeon performing a highly dexterous and delicate surgery. Skilled movement also brings us many of our greatest joys, such as the rush of euphoria from a baseball player hitting a home run, or the sense of awe inspired by a sculptor completing the final exquisite details on a masterpiece.

The primary motor cortex (M1) is an essential node in the brain circuitry responsible for voluntary movement. M1 was first characterized in 1870 by Fritsch and Hitzig who observed that in dogs, not only could electrical stimulation of M1 elicit movement, but stimulating different regions within M1 evoked movement in different body parts (Fritsch & Hitzig, 1870). This revolutionary finding was shortly followed by the discovery of a similar body map organization in the motor cortex of monkeys (Ferrier, 1874). Since then, an extensive body of work has defined M1 as a mainly agranular cortical region (Brecht et al., 2004; Weiler et al., 2008) with topographic organization (Tennant et al., 2011). Neurons in M1 have intratelencephalic (IT) connectivity throughout all layers (Weiler et al., 2008; Hooks et al., 2011) alongside pyramidal tract projections that originate in the deep layers of cortex and terminate in brainstem and in the spinal cord (Muñoz-Castañeda et al., 2021. A key feature of motor cortex is its ability to evoke movement (Peters et al., 2017). Activation of pyramidal tract neurons ultimately evoke activation of motor neurons in the spinal cord, and thereby faithfully produce muscle contractions and movements.

Much of the early investigations of motor cortex were performed in primates and cats, however, due to the development of new experimental tools in mice, mice have recently become a powerful model for investigating M1 and motor learning (Ölveczky, 2011). For example, in mice trained on a forelimb reaching task, silencing M1 activity through optogenetic photostimulation of GABAergic inhibitory neurons in M1 prior to reaching movements blocked

movement initiation. When photostimulation was performed just after movement initiation, M1 silencing was sufficient to freeze the ongoing reach movement (Guo et al., 2015), demonstrating the importance of M1 in the execution of voluntary movement. Prior to the advent of optogenetics, studies mainly relied on performing lesions in M1 and resulted in impaired but intact movement, making the role of M1 in movement execution unclear (Walker & Fulton, 1938; Lawrence & Kuypers, 1968; Castro et al., 1972). However, lesions are permanent and result in compensatory mechanisms (Jones et al., 2017) which likely obscured the true role of M1 in movement execution. These mechanisms likely explain the stark difference in movement deficits following lesions compared to photoinhibition. Therefore, the advent of new experimental tools that are highly tractable in mice have now spurred a deeper understanding of motor cortex function. For this reason, the primary focus of this text will be on research performed in mice.

1.1 Topography

Topography within the motor cortex is conserved across mammals and has been reported in numerous model systems. In mice, Tennant et al. (2011), systematically applied low threshold brief intracortical stimulation across cortex to evoke movements and thereby mapped body part representations throughout cortex. The body map representation was found to span M1 but also extended into the secondary motor cortex (M2; also referred to as medial agranular cortex) and the anterior lateral motor area (ALM; this region is sometimes considered the frontal M2). The caudal forelimb, including the elbow, wrist, and digits, was shown to have the most prominent movement representation spanning 1.78 ± 0.09 mm^2. In the anterior-posterior axis, the caudal forelimb area (CFA) was roughly centered on bregma. Head, jaw, and neck representations were found anterior to the CFA and located mainly in ALM (Tennant et al., 2011). Other studies additionally report that tongue representations necessary for licking are also found in ALM (Komiyama et al., 2010; Mayrhofer et al., 2019). Vibrissa, hindlimb, and tail representations are located in M1 and are posterior to bregma (Tennant et al., 2011). Using optogenetic photostimulation, simple movements could be evoked by shorter light pulses, consistent with microstimulation studies. Evoked movements remain intact during the blockade of intracortical synaptic activity, demonstrating that descending projections are sufficient to evoke movement. Intriguingly, prolonged optogenetic photostimulation was capable of evoking complex, reproducible movements but these were dependent on intact intracortical connectivity (Harrison

et al., 2012). These findings are corroborated by studies in primates (Graziano et al., 2002). Moreover, the use of *in vivo* two-photon calcium imaging with cellular resolution was used to assess neuronal activity in M1 in behaving mice and demonstrated that neurons involved in different complex movements form distinct but overlapping clusters (Dombeck et al., 2009). These studies provide important insight on the organization and function of neurons in M1.

1.2 Laminar Organization

The elegant laminar connectivity within cortex has been intricately optimized to receive, integrate, process, and transmit numerous simultaneous inputs. Mammalian cortex is comprised of 6 layers (L1, L2/3, L4, L5a, L5b, L6a, L6b) and understanding the anatomical connectivity of cortical circuits and the functional flow of activity through this circuit are essential to understanding how cortex performs complex computations. A distinguishing feature of M1 (and other agranular cortical areas) compared to sensory cortex is that M1 is traditionally regarded as lacking a layer 4, defined by spiny stellate cells receiving ascending thalamic input (Weiler et al., 2008; Brecht et al., 2004). Recent findings, however, contest this notion and suggest that M1 does indeed possess a compressed layer 4 with thalamic input (Yamawaki et al., 2014; Muñoz-Castañeda et al., 2021), but this remains to be fully elucidated.

L1 contains apical dendrites of L2/3 and L5 neurons and is mostly devoid of cell bodies apart from some interneurons (Hooks et al., 2011; Tremblay et al., 2016; Schuman et al., 2019). L2/3 and L5a in M1 are smaller compared to the primary somatosensory cortex (S1; Hooks et al., 2011) and are comprised of IT neurons which project within cortex and to striatum (Muñoz-Castañeda et al., 2021). L5b and L6 in M1 are expanded compared to S1 (Hooks et al., 2011) and are comprised of both IT neurons and all the pyramidal tract (PT) neurons in M1 (neurons defined by projections to brainstem and spinal cord which can thereby initiate movement). PT neurons additionally have axon collaterals providing input to striatum and thalamus (Economo et al., 2018). L6 also contains corticothalamic (CT) neurons that project to thalamus (Muñoz-Castañeda et al., 2021).

Using mice, both Weiler et al. (2008) and Hooks et al. (2011) utilized cell-attached electrophysiological recordings with simultaneous laser scanning photostimulation and glutamate uncaging to map the connectivity within the caudal forelimb area and the vibrissal area of M1 (vM1), respectively. By systematically varying the cortical depth of electrophysiological

recordings and photostimulation, connectivity within M1 was evaluated across layers to construct a connectivity matrix. This approach revealed that the strongest connection within M1 was the descending pathway from L2/3 to L5. Conversely, the ascending pathway from L5 to L2/3 was relatively weaker, although it was the strongest ascending pathway in M1. Horizontal or intralaminar connectivity was strongest within L2 followed by L5 and was relatively weaker within L3 and L6, demonstrating that while all layers have recurrent connections, they are most dominant in L2 and L5. Furthermore, network modelling of M1 suggests that input targeted to L2/3 can trigger network-wide events, but input targeting the deep layers cannot. These findings support a framework where superficial layers of M1 receive input from various cortical and some thalamic regions and thereby integrate these multiple streams of information, acting as an amplifier to drive output from deep layer neurons in M1 (Weiler et al., 2008).

1.3 Long-range Input to M1

While long-range input to and output from M1 have been broadly characterized for many years, recent advances in tracing and imaging techniques have enabled higher granularity. Previous work has identified that the major input areas to M1 are S1 (Mao et al., 2011), motor thalamus including the anteromedial (AM), ventral anterior (VA) and ventrolateral (VL) nucleus (Tanaka et al., 2018), higher-order sensory thalamus including the posterior nucleus (PO; Deschênes et al., 1998; Casas-Torremocha et al., 2019), and the orbital frontal cortex (ORB; Hooks et al., 2013). However, these tracing experiments were not done with cell-type specific techniques, hence, they include inputs to all cell types in M1.

A major advance in neuronal tracing was the development of the modified rabies tracing system by Wickersham et al. (2007). This system utilizes a rabies virus (RV) that is modified in two key ways. First, it has been pseudotyped with envelope A (EnvA) and can therefore only infect neurons expressing the avian TVA receptor (a receptor that is not endogenous in mammals). Second, the rabies glycoprotein has been deleted (ΔG), which restricts the ability of the rabies virus to spread trans-synaptically. TVA receptors and rabies glycoprotein can be expressed in a specific population of neurons with either transgenic mice or through helper virus injections using the cre-lox recombination system. Then when rabies virus is injected, TVA receptors serve to restrict rabies infection to neurons expressing TVA receptors, and rabies glycoprotein serves to restrict trans-synaptic spread from the presynaptic infected neurons.

Thereby, this system enables efficient, monosynaptic retrograde tracing to a specific cell population (Wickersham et al., 2007).

In 2019, Luo et al. utilized RVΔG to map long-range input to L5 glutamatergic and GABAergic neurons in the forelimb area of M1 using Thy1-cre and VGAT-cre (vesicular GABA transporter) mice, respectively. They found that glutamatergic and GABAergic neurons received very similar input, and most of the inputs came from S1, S2, contralateral M1, and contralateral M2. Lesser input was found from the agranular insular cortex, ORB, claustrum (CLA), and thalamus, and less numerous input was found from various other regions (Luo et al., 2019). Duan et al. (2019) performed cell-type specific rabies tracing among three major subtypes of GABAergic inhibitory neurons (for more on cell types in M1, see section 2) using PV-Cre, SST-Cre and VIP-Cre mice. This allowed them to trace brain-wide input to specific GABAergic subtypes in M1. They found that different cell types received similar long-range input and confirmed previous findings that the major input region to M1 is from S1 with additional input regions such as basal ganglia, thalamus, and several other cortical regions (Duan et al., 2019).

1.4 Output from M1

M1 also has distinct and diverse long-range projections and overall, more than 110 regions have been identified as recipients of M1 output (Muñoz-Castañeda et al., 2021). IT neurons have been shown to project within cortex and have dense projections to ipsilateral S1, M2, and the contralateral M1. Corticothalamic neurons project primarily to thalamus but have axon collaterals in other brain regions, namely cortex. Finally, PT neurons project within cortex but additionally send axon collaterals to the striatum, thalamus, midbrain, pons, the medulla and spinal cord (Economo et al., 2018; Muñoz-Castañeda et al., 2021).

Some differences in the connectivity of descending pathways have been observed between species. Although both disynaptic and monosynaptic connectivity between the motor cortex and spinal motor neurons has been observed in primates (Witham et al., 2016), monosynaptic connectivity seems to be limited or non-existent in rodents (Yang & Lemon, 2003; Alstermark & Isa, 2004). In rodents, descending motor cortex projections primarily provide input to spinal cord motor neurons through connections with the brainstem pontine reticular formation, referred to as the reticulospinal tract, or through at minimum disynaptic connections within spinal cord, known as the corticospinal tract (Alstermark & Isa, 2004).

1.4 Neuronal Activity in M1

Successful execution of volitional movement must combine sensory input about the environment with a multitude of movement-related kinematics including direction, force, and speed. There is a comprehensive body of work investigating how M1 represents these movement-related parameters, including landmark findings from Georgopoulos et al. (1982) where the authors performed extracellular electrode recordings to record population activity of neurons in the motor cortex of rhesus monkeys performing a reaching task. In this task, the monkeys were trained to grasp a manipulandam and maneuver it outwards in one of eight directions toward the presented target. The authors found that approximately 75% of neurons in M1 demonstrated direction-tuning, meaning the neuron's discharge rate was highest for movement in a specific direction. As the movement direction diverged from a neuron's preferred direction, the discharge rate gradually declined such that the firing rate was a sinusoidal function of direction. Different neurons had different but overlapping tuning curves, and all eight directions of movement had a corresponding subset of tuned neurons. These changes in firing rate were observed as early as 80ms prior to movement initiation, recorded with electromyography. In some neurons, activity spanned parts of the movement or the entire movement (Georgopoulos et al., 1982). In a subsequent investigation, the authors utilized a similar task where the monkeys began the trial by first pressing a center button located directly in front of the monkey and then had to reach and press a second illuminated target button in one of eight possible equally distanced locations surrounding the center button. The authors then used multiple regression to represent the firing rate of each direction-tuned neuron during reaching movements in 3D space to estimate each neuron's preferred direction. This revealed a 3D continuum of broadly-tuned, directionally-selective neurons, while sharply-tuned neurons were not observed. Since individual neurons were only broadly tuned, the authors assessed population activity by expressing each neuron's activity as a vector and then summing the vectors across all recorded neurons. For each neuron, the vector's direction was based on the neuron's preferred direction, and its vectorial contribution was based on the magnitude of the neuron's change in firing rate. When the population of directionally-tuned neurons were summed, the authors were able to accurately resolve the direction of the actual movement, demonstrating that the direction of volitional movement is implemented through population code summation in M1 (Georgopolous et al., 1986).

In addition to movement direction, neurons in M1 have been shown to encode and control various other movement kinematics. These include neurons that represent load direction when movement must compensate for a force acting on the object being moved. These load-tuned neurons demonstrate similar tuning dynamics as direction-tuned neurons and activity vectors are also approximately linearly summed across the population (Kalaska et al., 1989). Neurons encoding target distance and movement speed prior to and during the movement have also been reported in M1 (Churchland et al., 2006).

Traditionally, investigations of activity in M1 have focused on movement planning and kinematics, however, M1 receives many long-range inputs from a multitude of brain regions with distinct functions (Luo et al., 2019; Duan et al., 2019), therefore, M1 is also modulated by non-movement-related inputs. M1 receives dense input from S1 including extensive innervation from the vibrissal area of S1, also known as the barrel cortex (Mao et al., 2011; Luo et al., 2019; Duan et al., 2019), and whisker deflections have been shown to reliably evoke sensation-related responses in M1 (Ferezou et al., 2007; Matyas et al., 2010; Mao et al., 2011). Moreover, during a whisker-based go/no-go task where the mice must discriminate two textures using their whiskers and lick for reward in response to the "go" texture, the proportion of M1-projecting S1 neurons with touch-related activity increased with learning (Chen et al., 2015a). Therefore, neurons in M1 integrate sensory input to dynamically modulate M1 activity based on features of the environment. In addition to S1, M1 receives a multitude of inputs from brain regions implicated in various behavioural and cognitive functions, therefore, it is likely that M1 neurons multiplex diverse information to guide movement. Although a broad delineation of long-range input to M1 has been established, a deeper understanding of the spatiotemporal dynamics and cooperativity of these inputs *in vivo* is still lacking. This will be critical to understanding how the brain develops and executes a motor plan, as well as how it dynamically updates movements based on a changing environment.

2. Cell Types in M1

Broadly, neurons within cortex can be categorized as excitatory (those that contribute to excitation in postsynaptic neurons), or inhibitory (those that contribute to blunting excitation in postsynaptic neurons). The main excitatory neurostransmitter in the brain is glutamate and the main inhibitory neurotransmitter is γ-Aminobutyric acid (GABA). A now outdated view once

held that computation occurred mainly within glutamatergic excitatory activity and inhibition was a passive reflection of excitation, serving mainly to maintain the excitation-inhibition balance in the brain. This erroneous assertion was in part due to limited experimental techniques for targeting specific cell populations. With the development of transgenic cre-driver mouse lines and optogenetics, selectively targeting inhibitory neurons (INs) has become highly tractable. In the past decade or so, the field has witnessed a major transformation in our understanding of IN function. Continually accumulating evidence has demonstrated that in addition to regulating excitation, inhibition serves as a powerful mechanism for actively shaping computation in the brain. Within rodent cortex, GABAergic INs are a small but powerful subset of neurons and comprise of 10-15% of the total neuronal population (Tremblay et al., 2016; Meyer et al., 2011).

GABAergic INs are a heterogeneous population, with diverse morphology, electrophysiological properties and connectivity. While INs can be categorized based on different properties, a dominant nomenclature has emerged for the major subtypes based on specific expression of molecular markers. The three major non-overlapping subtypes are the calcium-binding protein parvalbumin-expressing INs (PV-INs), the neuropeptide somatostatin-expressing INs (SST-INs), and the serotonin receptor 5HT3aR-expressing INs. Together, these subpopulations account for almost 100% of GABAergic INs in cortex. Vasoactive intestinal peptide expressing INs (VIP-INs) are a subpopulation of 5HT3aR-expressing INs (Kawaguchi and Kubota, 1997). These neurons are found intermingled throughout layers 2/3 and layers 5/6 of cortex (Kawaguchi & Kubota, 1997; Pfeffer et al., 2013). PV-INs, SST-INs and VIP-INs comprise approximately 80% of all GABAergic INs in cortex. Even within each of these subpopulations, transcriptome, morphology and electrophysiological properties vary (Tasic et al., 2016; Tremblay et al., 2016; Tasic et al., 2018). Despite this heterogeneity, cells within each subtype share key features that have enabled a more unified understanding of IN microcircuit function.

2.1 The Three Major GABAergic Subtypes

Within each of these IN subtypes, neurons can be further classified and characterized. PV-INs are the most abundant GABAergic IN subtype and account for about 40% of GABAergic neurons in cortex (Pfeffer et al., 2013; Tremblay et al., 2016). PV-INs include cells with two distinct morphologies – chandelier cells and basket cells. Chandelier cells are found in L2 and

L6, while basket cells can be found in L2-6. PV-INs have low input resistance, and fire action potentials with short duration, a high maximum firing frequency, and non-adapting patterns (Tremblay et al., 2016) and are therefore, often referred to as fast-spiking cells. PV-INs have also been shown to be electrically coupled to each other (Gibson et al., 1999). Basket cells primarily target the perisomatic region of PNs, while chandelier cells primarily target the axon initial segment (Kawaguchi & Kubota, 1997; Tremblay et al., 2016). The perisomatic location of PV-IN synapses on PNs, enable reliable and powerful control over PN excitation while the electrophysiological properties of PV-INs enable extremely fast control (Tremblay et al., 2016).

SST-INs comprise about 30% of GABAergic INs in cortex (Pfeffer et al., 2013; Tremblay et al., 2016) and can be further divided into Martinotti and non-Martinotti cells. Martinotti cells are found in L2-5 but are most abundant in L5 and have an axon plexus that ascends to L1 and synapses onto the apical dendrites of PNs (Kawaguchi & Kubota, 1996). Non-Martinotti cells are found in L4-5 and do not have an axon plexus in L1 (Tremblay et al., 2016). Martinotti cells have high-input resistance, a low-threshold for spiking, and a low maximum firing frequency (Kawaguchi & Kubota, 1997). Non-Martinotti cells also have low input resistance but have a higher maximum firing frequency than Martinotti cells. Unlike PV-INs, SST-INs primarily contact the apical dendrites of PNs (Kawaguchi & Kubota, 1996). Due to the cable properties of dendrites, inhibition exerts a more powerful effect on membrane potential when the inhibitory synapse is located closer to the soma (Koch et al., 1983). Therefore, SST-INs provide less powerful control over PN somatic spiking compared to PV-INs (Pfeffer et al., 2013). However, SST-INs still exert control over dendritic computations via shunting inhibition (Koch et al., 1983; Tremblay et al., 2016).

VIP-INs account for about 17% of GABAergic INs (Pfeffer et al., 2013) and have either bipolar or multipolar morphologies. VIP-INs are most dense in L2/3 of cortex but can be found in other layers (Prönneke et al., 2015). They have vertically running dendrites and can therefore receive input from multiple layers. A key feature of VIP-INs is their high input resistance, that leads to them being easily excitable. VIP-INs have been shown to have various spiking patterns including burst spiking and irregular spiking which refers to neurons that initially fire a burst followed by irregularly timed individual spikes (Kawaguchi & Kubolta, 1996; Kawaguchi & Kubolta, 1997; Tremblay et al., 2016).

2.2 A Common Inhibitory Microcircuit Motif in the Cortex

Perhaps the most striking property of PV-INs, SST-INs and VIP-INs are the conserved microcircuit motif they form throughout cortex and in some subcortical regions. Using either PV-Cre, SST-Cre or VIP-Cre transgenic mice lines and optogenetics, Pfeffer et al. (2013) performed systematic investigation of fine-scale connectivity between pairs of molecularly identified neurons in the primary visual cortex (V1). By expressing the excitatory opsin, channelrhodopsin (ChR2), in a single subtype and performing whole-cell patch-clamp recordings, the authors were able to selectively drive activity in a specific cell population and record post-synaptic responses. They then identified the post-synaptic cell type post-hoc, using single-cell reverse transcription PCR. By repeating this many times, the authors constructed a connectivity map between cell types in V1. Their findings revealed that PV-INs and SST-INs preferentially targeted PNs. 100% of recorded PNs received monosynaptic input from PV-INs and SST-INs. PV-INs provided the largest amplitude unitary inhibitory post-synaptic potentials (uIPSCs) onto PNs, followed by SST-INs. Connection probability of VIP-INs onto PNs was found to be approximately 12.5% and uIPSC amplitude was much smaller than PV-INs and SST-INs. Among the INs, PV-INs primarily inhibited other PV-INs, provided a small amount of inhibition to VIP-INs, and almost no inhibition to SST-INs. SST-INs provided input to VIP-INs, relatively weaker input to PV-INs, and almost no input onto other SST-INs. Finally, VIP-INs primarily targeted SST-INs and provided little input to other IN subtypes. These findings characterize a highly specific microcircuit motif where PV-INs and SST-INs are highly connected to PNs. PV-INs provide strong input onto PNs and other PV-INs. SST-INs provide equally as extensive but relatively weaker input to PNs. Finally, VIP-INs primarily target other SST-INs, resulting in disinhibition of pyramidal neurons (Pfeffer et al., 2013). A similar motif has been reported in multiple brain regions including the auditory, prefrontal (Pi et al., 2013) and somatosensory cortices (Lee et al., 2013) and the basolateral amygdala (Krabbe et al., 2019).

2.3 The Role of GABAergic Inhibitory Neurons in Sensory Processing

In other cortical and subcortical brain regions, investigations of PV-INs, SST-INs and VIP-INs have revealed complex roles in integrating a range of signals including sensory input, attention, novelty, and aversion (Wilson et al., 2012; Pi et al., 2013; Zhang et al., 2014; Makino et al., 2015; Kuchibolta et al., 2017; Muñoz et al., 2017; Krabbe et al., 2019; Garrett et al., 2020).

Much of the seminal work on the functional roles of INs was performed in primary sensory cortices. In the primary visual cortex (V1), PV-INs and SST-INs have been shown to be broadly tuned to visual stimuli properties (Hofer et al., 2011; Cottam et al., 2013) and have distinct effects on PN orientation tuning. PV-INs have been shown to tightly regulate PN firing in response to visual stimulation and linearly modulate PN activity while preserving orientation tuning (Atallah et al., 2012; Wilson et al., 2012). SST-IN activation instead, subtracted from PN tuning curves and sharpened PN tuning (Wilson et al., 2012). Moreover, SST-IN activation has been shown to also perform subtraction on PV-IN tuning curves, thereby sharpening the orientation tuned response from the PV-IN population (Cottam et al., 2013). Top-down modulation from the cingulate area activated all three subtypes in V1, and activation of SST-INs and PV-INs broadly suppressed V1 neurons, while VIP-IN activation selectively activated a population of nearby neurons. Therefore, long-range input may recruit this inhibitory microcircuit to control the spatial specificity of visual responses in V1, (Zhang et al., 2014; Ayzenshtat et al., 2016). This effect is mediated through VIP-IN inhibition of SST-INs, thereby removing inhibition from select PNs and providing a mechanism for disinhibition of visual responses, potentially based on attention (Karnani et al., 2016). VIP-IN activity has also been shown to facilitate states of high activity in V1 across multiple behavioural states and provide a means for local gain control (Fu et al., 2014; Jackson et al., 2016). VIP-INs in V1 also respond to novelty (Garett et al., 2020). In various cortical regions, VIP-INs, unlike SST-INs, PV-INs or PNs, are selectively depolarized by ionotropic nicotinic input (Alitto & Dan, 2012; Pi et al., 2013; Gasselin et al., 2021; Ren et al., 2022). Therefore, one possibility is that during states of high attention, cholinergic input depolarizes VIP-INs, thereby increasing cortical gain through disinhibition.

In the barrel cortex, INs have similarly been shown to integrate sensorimotor input and regulate local activity accordingly. *In vitro* whole-cell recordings in voltage-clamp configuration demonstrated that input from vM1 strongly recruited VIP-INs when compared to PNs, PV-INs and SST-INs which in turn inhibited SST-INs. *In vivo*, M1 input to S1 recruits VIP-INs during whisking (Lee et al., 2013). While some SST-INs were inhibited by VIP-INs during active whisking, other SST-INs were active and presumably received little VIP-IN input, as their activity was not altered by VIP-IN manipulation. Instead, these whisking-related responses appear to be driven by cholinergic activation of muscarinic receptors on SST-INs and can be

blocked with a muscarinic antagonist (Muñoz et al., 2017). This demonstrates that distinct modules of INs in cortex are regulated by attention and active sensing. Moreover, vM1 input to S1 has been shown to evoke nonlinear dendritic events in L5 PNs during whisker touch events (Xu et al., 2012). These nonlinear dendritic events are thought to be calcium plateau potentials and when combined with bottom-up input, can induce burst spiking (Larkum et al., 1999) thought to correlate with perception (Takahashi et al., 2016; Doron et al., 2020). Optogenetic activation of SST-INs was sufficient to reduce the probability of perception in behaving mice, while optogenetic activation of L5 PN apical dendrites was sufficient to increase the probability of perception (Takahashi et al., 2016). The emerging picture is that during active sensing, VIP-INs and SST-INs are differentially driven by cholinergic input, which recruits nicotinic and muscarinic receptors on VIP-INs and SST-INs respectively (Alitto & Dan, 2012; Muñoz et al., 2017; Gasselin et al., 2021). Thus, long-range input from vM1 and possibly other regions recruit VIP-INs, which in turn, inhibit a subpopulation of SST-INs. SST-INs are thereby positioned to gate non-linear dendritic activity and burst spiking in L5 PNs. In this way, this microcircuit modulates the neural correlates of perception in S1. However, how other long-range inputs and behavioural states interact to drive this circuit and thereby modulate perception remain an open question.

2.4 The Role of GABAergic Inhibitory Neurons in Reinforcement

The ability to relate a stimulus to its subsequent outcome is critical to survival for all animals. Fear conditioning is a common model for associative learning. In fear learning, a neutral auditory tone, known as the conditioned stimulus (CS) is paired with a subsequent aversive stimulus, known as the unconditioned stimulus (US), such as a foot shock. When this process is repeated multiple times, the animal learns to associate the once neutral tone with the shock, resulting in freezing behaviour in anticipation of the shock, demonstrating the formation of a fear memory (Lucas & Clem, 2018). In the auditory cortex, an unidentified layer 1 disinhibitory IN has been shown to be recruited by cholinergic input during auditory fear learning. This disinhibition is necessary for learning the association between the CS and US and can be blocked with application of a nicotinic receptor antagonist (Letzkus et al., 2011). Subsequent work has shown that VIP-INs in the auditory cortex are strongly recruited by reward and punishment (Pi et al., 2013) and are also depolarized by nicotine (Askew et al., 2019). Since there is a small population of VIP-INs located in L1 (Schuman et al., 2019), it is possible that VIP-INs may be the L1

disinhibitory INs shown to mediate fear-learning in the auditory cortex (Letzkus et al., 2011). However, this is not definitive since other disinhibitory INs are present in L1 of the auditory cortex and also have been shown to have a role in fear learning (Pardi et al., 2020).

In the basolateral amygdala where the SST-IN, PV-IN and VIP-IN microcircuit is conserved, VIP-INs are strongly recruited by foot shocks during an auditory fear learning task. VIP-IN responses are highest at the beginning of learning and gradually decline with repeated CS-US pairings. VIP-IN activation then disinhibits projection neurons in the BLA through inhibition of SST-INs and to a lesser degree, PV-INs. Intriguingly, VIP-IN silencing during the US impairs fear learning (Krabbe et al., 2019). These results suggest a general principal, where long-range input and potentially cholinergic input recruit VIP-INs during aversive stimuli, and through disinhibition of PN dendrites, gate plasticity and learning.

3. Circuit Mechanisms of Motor Learning

Humans are born with a very limited repertoire of innate motor skills; the ability to reach and grasp, to walk, and to ultimately perform complex and highly skilled movements are all learned through repetitive practice. Motor learning or the formation of motor memories is a form of procedural learning where performance in a motor behaviour undergoes permanent improvement as a result of repetitive practice (Brem et al., 2013). If one considers motor control as the process of using sensory input to effectively guide motor output, then motor learning can be viewed as the process of mastering this transformation through practice (Wolpert & Flanagan, 2001). Hallmarks of motor learning include enhanced speed, accuracy, or reproducibility of a movement (Papale & Hooks, 2018). This section will survey the extensive evidence implicating M1 as an important site for motor learning.

3.1 Pyramidal Neuron Plasticity

Early investigations of the substrate for motor learning found that in rats trained on a forepaw reaching task, the apical dendrites of L5 PNs underwent morphological expansion in the contralateral motor-sensory forelimb cortex. This dendritic expansion included an increase in dendritic field size, total dendritic length, number of branches and the length of the terminal branches (Greenough et al., 1985). The changes were later shown to be biphasic, as they are followed by a subsequent refinement in dendritic length and arborization, thus returning to the

baseline size but leaving the learned motor skill remains intact (Gloor et al., 2015). During motor learning using rotarod, *in vivo* electrode recordings showed that on the first day of learning, more L5 PNs were recruited and became active as the session continued. In contrast, on subsequent days, less neurons were recruited, and the number of recruited neurons were similar to the number of dismissed neurons (Costa et al., 2004). The recruitment of additional L5 PNs early in learning is complimentary to the expansion seen in L5 dendrites. Withers & Greenough also reported structural remodeling in the basal dendrites of subpopulation of L2/3 neurons in M1 (Withers & Greenough, 1989). These findings suggest that neurons within different layers of M1 receive increased input during the initial phases of motor learning which drives structural remodeling. Additionally, *in vitro*, contralateral M1 had higher amplitude local field potentials and reduced capacity for long-term potentiation induction (LTP) after learning occurred (Rioult-Pedotti et al., 1998), suggesting motor learning is accompanied by LTP and increased excitability in M1 which occludes further LTP induction. Supporting this notion are findings that pharmacologically blocking protein synthesis in M1, but not in the cerebellum or parietal cortex, impaired acquisition in a motor learning task in rats (Luft et al., 2004).

These early studies were extremely influential in building a foundational understanding of M1's role in motor learning, however, the scope of these investigations was limited by the lack of experimental tools that could be applied *in vivo* to assess brain structure and function. Early studies relied on lesions and pharmacological manipulation or extracellular recordings to assess neuronal activity *in vivo*, all of which provide limited cell type specificity. To investigate brain structure or synaptic physiology, animals had to be trained on a motor learning task and then sacrificed for the brain to be examined *in vitro*. Although these approaches enabled many important findings, they were limited in the ability to understand the dynamism that occurs throughout learning. Advances in *in vivo* imaging techniques such as two-photon imaging have greatly improved our understanding of phasic changes that occur throughout the process of learning. One particularly influential study was performed by Xu et al. in 2009, where the authors used transgenic mice that sparsely expressed yellow fluorescent protein (YFP) in L5 PNs to track the dendrites of these neurons throughout motor learning. By thinning the skull and implanting a headplate, the authors could then head-fix the anesthetized mouse and perform transcranial, chronic, *in vivo* two-photon imaging. Mice were trained on a forelimb reaching task and L5 apical dendrites were imaged within 1 hour after the training session. They acquired z-stacks and

created 3D reconstructions of the apical dendrites and used this technique to track the same apical dendrites across days. They found that after the first session with successful reaches, spinogenesis significantly increased, and this effect was not observed during shaping or in control mice that received similar handling. Even more intriguing, the proportion of newly added spines was highly correlated with the number of successful reaches during the training session. When the mice were trained for several days, increased spine formation during the initial phase of training was then followed by increased spine elimination, which ultimately returned spine density to baseline levels after the transient increase. Initially, the newly formed spines were more likely to be eliminated, suggesting many of these newly formed spines were not stable, however, a subset of new spines persisted to the end of learning. A hallmark of motor learning is long-lasting motor memory and the ability to perform the motor skill after an extended period of time without practice. Indeed, when mice were subjected to the task 4 months later, they maintained their skilled reaching and did not undergo an increase in spinogenesis. However, when these same mice were trained on a different reaching task, spinogenesis once again occurred in M1, but the spines from the first task remained stable, suggesting distinct sets of synapses correspond to different tasks (Xu et al., 2009). Similar dendritic spine dynamics have been observed in mice learning rotarod (Yang et al., 2010). Together, these findings demonstrate that during the acquisition of a new motor skill, new synapses are formed in M1, evidenced by the increase in dendritic spines. Following this expansion, the circuit undergoes a selective pruning of synapses, preserving those synapses that facilitate the newly learned skill and eliminating those that do not.

In the years following these findings, *in vivo* two-photon calcium (Ca^{2+}) imaging underwent a number of advancements (Komiyama et al., 2010; Akerboom et al., 2011; Chen et al., 2013; Guo et al., 2014) enabling high-speed, reliable measurements of neuronal activity in the same population of neurons across days in awake and behaving head-fixed animals. Genetically encoded calcium indicators such as GCaMP fluoresce in the presence of calcium and have rise and decay times on the order of milliseconds (Akerboom et al., 2011; Chen et al., 2013). When combined with high-speed *in vivo* two-photon imaging, the change in fluorescence can be calculated for each neuron, providing a non-linear approximation of neuronal activity in an awake and behaving animal. Although electrophysiological recordings are faster than calcium imaging, the same cells cannot be reliably tracked across days, and it is limited in identifying cell

types. With calcium imaging, the same population of cells can be reliably tracked across days, and when employed in combination with the cre-lox system and cre-driver transgenic mouse lines, GCaMP expression can be limited to genetically-defined populations of cells. For these reasons, the development of *in vivo* two-photon calcium imaging enabled significant advancements in our understanding of the mechanisms of learning since the activity of genetically identified cells can now be monitored throughout the entire process of motor learning.

By expressing GCaMP5G in M1 and implanting a headplate and a chronic imaging window over M1, Peters et al. (2014) used two-photon calcium imaging to visualize the activity of L2/3 PNs in M1 across days while mice simultaneously learned a head-fixed forelimb lever-press task. The authors found that initially, the movement trajectory of individual lever-press attempts had low correlation but as mice learned the task, the correlation between lever-presses increased and eventually became highly reproducible and stereotyped. Concurrently, in naïve mice first learning the task, the fraction of movement-related neurons with activity during lever presses transiently increased and then subsequently decreased as the mice learned the task across sessions. Additionally, the trial-to-trial correlation of the neural activity increased with learning, resulting in a stable a spatiotemporally reproducible activity pattern that accompanied the learned movement in expert mice. Analogously, the authors demonstrated that the dendritic spines of L2/3 PNs undergo a similar increase in dendritic spine formation followed by increased elimination (Peters et al., 2014). Together, these findings demonstrate that early in motor learning when movements are highly uncorrelated, there was an expansion in the movement-related population in M1, evidenced by the increase in the fraction of movement-related neurons and the increase in dendritic spines. As learning progressed and performance improved, the population of movement-related neurons was gradually refined and dendritic spines were selectively eliminated. Finally, the dendritic spines that are presumably involved in evoking the expert movement were stabilized, and these plastic changes enabled spatiotemporally reproducible neural activity in M1 that ultimately produce a highly correlated, skilled movement. Disrupting remodeling of dendritic spines using the synaptic optoprobe Activated Synapse Targeting Photoactivatable Rac1 (AS-PaRac1), which enables selective optical shrinkage of recently potentiated spines, impaired the acquisition of motor learning in mice. This study demonstrated for the first time that potentiation and stabilization of task-related spines are necessary for learning (Hayashi-Takagi et al., 2015).

Structural and functional remodeling in PNs in M1 have been rigorously shown to be a key mechanism for the acquisition, consolidation and long-term retention of motor skills. To fully understand the mechanisms that spatiotemporally control these plastic changes, a holistic view of M1 circuitry is required. Since GABAergic INs can inhibit PN activity, and plasticity is activity-dependent, it is logical to hypothesize that GABAergic INs can regulate PN plasticity.

3.2 The Role of GABAergic Inhibitory Neurons in M1 During Motor Learning

Mounting evidence has demonstrated that GABAergic INs not only balance the excitation to inhibition (E:I) ratio in M1, but are also indispensable to regulating plastic changes in PNs. Numerous studies have implicated plastic changes in the dendrites of L2/3 and L5 PNs in M1 as an important neural correlate of motor learning (Greenough et al., 1985; Withers & Greenough, 1989; Xu et al., 2009; Yang et al., 2010; Peters et al., 2014; Hayashi-Takagi et al., 2015; Gloor et al., 2015). SST-INs are key regulators of apical dendritic activity in PNs (Kawaguchi & Kubota, 1996); therefore, they are well positioned to regulate the synaptic remodeling of dendritic input during motor learning.

Motor learning induced a transient increase in spine density in distal dendrites of PNs (Xu et al., 2009; Yang et al., 2009; Peters et al., 2014). Intriguingly, no change was seen on the perisomatic dendrites, demonstrating that L2/3 PNs undergo compartment-specific changes during motor learning (Chen et al., 2015b). Since the perisomatic region and the distal dendrites receive compartment-specific inhibition from PV-INs and SST-INs respectively, the authors investigated whether SST-INs might be involved in distal dendrite spine formation. Chronic imaging of the axonal bouton density of SST-INs and PV-INs in M1 as mice learned a lever-press task revealed distinct structural changes in both IN populations. Motor learning induced a transient increase in PV-IN bouton density, while SST-INs underwent a persistent decrease. Moreover, optogenetic inhibition of SST-INs during training led to an increase in the density of dendritic spines, while activation of SST-INs resulted in decreased spine density. This suggests that SST-IN mediated inhibition regulates spine stabilization and elimination in L2/3 PNs, consistent with previous work showing that increased inhibition onto dendrites blocks the potentiation of new spines (Chiu et al., 2018). In addition, it has been shown that successful motor learning induces a reduction in the fraction of neurons with movement-related responses and is accompanied by the emergence of a spatiotemporally reproducible activity sequence in M1

(Peters et al., 2014). Hence, SST-INs are likely to be key to regulating which inputs are potentiated, while PV-INs may have a role in homeostatic plasticity and maintaining the E:I ratio.

Another study investigated the role of SST-INs in PN plasticity in M1 during a treadmill running task for mice. Using *in vivo* two-photon calcium imaging, Cichon & Gan (2015) imaged the activity in apical tuft branches of sparsely labelled L5 PNs while mice were trained on a task-switching, forward and backward running task. Running forwards compared to backwards induced events on different branches, suggesting different tasks evoke branch-specific input to L5 PN apical dendrites. When SST-INs were ablated or silenced, branch specificity was lost in L5 PNs, and forward and backward running evoked events on the same dendritic branch. SST-IN ablation also impaired running behaviour following a switch in running direction, suggesting SST-INs regulate the specificity of neuronal activity and synapse formation in PN dendrites during different tasks (Cichon & Gan, 2015).

In a subsequent study, forward running on a head-fixed treadmill was shown to be accompanied by a sequential activity pattern in L2/3 PNs. However, if SST-INs were chemogenetically silenced or if VIP-INs were activated, which would presumably inhibit SST-INs, sequential activity in L2/3 PNs was disrupted and motor learning was impaired. Moreover, calcium imaging in SST-INs during forward running revealed a range of activity with some SST-INs increasing their activity and some decreasing. When mice then had to run backwards, approximately 75% of SST-INs reversed their activity levels, suggesting SST-INs have task-specific activity and are involved in regulating the recruitment of PNs during movement (Adler al., 2019).

Since VIP-INs primarily inhibit SST-INs, they are well-positioned to regulate PN plasticity through SOM-IN mediated-inhibition. Ren et al. (2022) have shown that VIP-IN activity increased early in learning during a forelimb lever-press task but as learning progressed over days, VIP-IN activity decreased while SOM-IN activity inversely increased. Calcium imaging of SST-INs while VIP-INs were simultaneously silenced using DREADDs resulted in increased SST-IN activating, thus confirming that VIP-INs typically exert inhibitory control over SST-INs at the beginning of learning. Silencing VIP-INs in M1 also impaired motor learning, evidenced by reduced performance in the current and subsequent training session. Together, Ren

et al. demonstrate that precise VIP-IN and SST-IN activity are critical for motor learning (Ren et al., 2022).

These findings convey a cohesive story: early phases of motor learning are accompanied by an expansion in the M1 ensemble with movement-related activity. During this phase of expansion, PNs undergo synaptic remodeling and select synapses on the apical dendrites of PNs are potentiated. This phase is gated by VIP-IN activity which inhibits SST-INs, such that PN apical dendrites are more disinhibited. Following this initial expansion, VIP-INs reduce their activity and SST-INs become more active, thereby exerting increased inhibition onto PN apical dendrites. M1 circuitry then undergoes a refinement process, eliminating synapses that are not conducive to the newly learned motor skill and this process is regulated by SST-INs. Therefore, not only do GABAergic INs have active roles in the formation and consolidation of motor memories, but distinct subtypes have distinct roles.

3.3 The Role of Long-Range Input During Motor Learning

Substantial progress has been made in our understanding of local plasticity in M1 during motor learning, however, the rules that initiate and orchestrate this process in M1 are still unclear. Plasticity in M1 is likely the result of concerted inputs, including long-range glutamatergic input from brain regions involved in decision making and reinforcement learning, neuromodulators involved in learning and reinforcement, and the interplay of these inputs with different cell types in local M1 circuitry (Richards et al, 2019).

M1 is highly interconnected (Luo et al., 2019; Muñoz-Castañeda et al., 2021), therefore, to fully understand M1 plasticity, activity in other brain regions must be considered. Motor thalamus is a major input to M1 and includes the ventral anterior, medial and lateral nuclei of thalamus (VA, VM and VL respectively; Sieveritz et al., 2018). High frequency tetanic stimulation in S1 but not VL drove LTP in M1. However, simultaneous high-frequency stimulation of S1 and VL drove LTP in VL input to M1 (Iriki et al., 1989). These findings demonstrate how neurons in M1 integrate cooperative streams of input from distinct long-range inputs, resulting in long term plasticity. Following a forelimb reaching task, VA/VL input was potentiated onto corticospinal neurons in M1 that project to the forelimb area of the spinal cord but no potentiation was observed for neurons projecting to the hindlimb area of the spinal cord (Biane et al., 2016). This demonstrates that in addition to plasticity within M1, thalamic input to

M1 is also plastic during motor learning and these changes are highly specific to task-related neurons. Another study that performed axonal bouton imaging of axons from VA/VL and VM in M1 during a lever-pull task demonstrated that the thalamocortical boutons show task-related activity and that this activity is plastic during motor learning. Using a linear multiple regression model to predict lever trajectory from thalamocortical bouton activity, the authors found that the prediction accuracy improved with learning. This implies that the activity of thalamocortical axons had a stronger representation of the movement after learning. Additionally, the responses became more time-locked to the movement (Tanaka et al., 2018).

There is also evidence for plastic changes within cortico-cortical input during motor learning. Using wide-field calcium imaging, Makino et al. (2017) tracked the mesoscale flow of activity throughout cortex across days as mice learned a lever-press task. Numerous cortical regions showed movement-related activity. Retrosplenial cortex was consistently one of the first nodes to respond. Interestingly, frontal M2 (M2 in this study includes with ALM), shifted to earlier activation times with learning and the entire response sequence across cortex became more compressed and less variable (Makino et al., 2017). These cortex-wide transformations during motor learning provide further evidence that an in-depth and comprehensive understanding of motor learning necessitates examination of activity changes in other regions that interact with M1.

While investigations of long-range input to M1 during motor learning are currently limited, there is strong evidence that these inputs are key drivers in M1 plasticity and motor learning. Dissecting the dynamics and any potential plastic changes in different long-range inputs during motor learning will be key to understanding how M1 plasticity is initiated and regulated.

4. The Effects of Reward on Motor Learning

Reward is a fundamentally important and motivating for animals. The brain has evolved to have dedicated circuitry for detecting, predicting, and seeking reward. Reward also serves as a positive reinforcer, meaning it increases the frequency of the behaviour linked to the reward (Schultz, 2000). Animals are capable of multiple forms of reward-based learning, the simplest of which is classical conditioning. In this form of associative learning, a neutral stimulus known as the conditioned stimulus (CS) is paired with a reinforcer (this may be reward or punishment), known

as the unconditioned stimulus (US). Through repeated pairing, the animal learns the association, and upon CS presentation, can predict the US contingency (Bouton & Moody, 2004). Neurons in various brain regions have been shown to be sensitive to reward, evidenced by reward-prediction, reward-prediction errors (a discrepancy between predicted and received reward) and received reward. For example, in the orbital frontal cortex (ORB) of monkeys, not only do neurons respond to reward, but some neurons actually discriminate between rewards by firing more for the animal's preferred reward (Tremblay & Schultz, 1999) demonstrating that the brain encodes the presence and relative value of reward.

An important caveat that can often be difficult to distinguish in reward-based learning is whether reward is truly enhancing learning, or if it is simply enhancing motivation. Through proper controls however, differences in motivation levels can be accounted for. One study utilized a motor learning task where human subjects were instructed to manipulate a force transducer using their right thumb and index finger to control a cursor on a computer screen with the goal of maintaining it inside a target. The subjects were randomly divided into three groups: a reward group, a punishment group and a neutral group. To test the effects of different types of reinforcement, the reward group was told they began the experiment with $0 but would earn money for time spent on the target, the punishment group was told they began with $72 but would lose money for time spent off target, finally the neutral group was told they would receive $40 at the end of the study. All three groups were given four training blocks with 20 trials each. Performance was tested before, immediately after, 6 hours, 24 hours, and 30 days after training. Immediately after training, all three groups had comparable improvement in the task suggesting acquisition was unaffected by the mode of reinforcement. Intriguingly, at the 6h time point the reward group maintained their performance level while both the punishment and neutral groups experienced significant forgetting. By 24h, the reward group demonstrated enhanced performance indicating they experienced offline improvements consistent with memory consolidation (Censor et al., 2012) while the other groups continued to forget. Finally, after 30 days, the rewarded group maintained the improvements seen at the 24h mark, while the other groups showed significant forgetting (Abe et al., 2011). These findings are striking because firstly, it is well-established that humans are more sensitive to potential losses than to potential gains, a concept known as loss aversion (Tom et al, 2007), therefore, the inclusion of a punishment group appropriately controls for motivation levels during the task. Second, this study

implies that the initial acquisition phase of learning is not affected, but rather reward can enhance offline memory consolidation. Additionally, other studies using human subjects have reported that reward enhanced and accelerated learning (Wächter et al., 2009; Nikooyan & Ahmed, 2015) including implicit reward that was not disclosed until after the action was performed (Palminteri et al., 2011). How then does the brain utilize reward to enhance motor learning? The motor cortex receives the majority of its input from cortex, thalamus and neuromodulatory systems in brainstem (Papale & Hooks, 2018) and it remains unclear if and how these different inputs orchestrate reward-based enhancement of learning? Where in the brain do reward-related computations occur and what circuit mechanisms are responsible for integrating reward signals with motor activity to enhance learning and consolidation? While humans are ideal subjects for behavioural assessment of the effects of reward on learning, detailed understanding of the neural mechanisms enabling this effect must be performed in animal models.

5. Reward Representations in M1

M1, being a primary cortical region, has traditionally been seen as a command center for movement. Simulations of M1 plasticity have been able to explain experimentally observed network reorganization seen in M1 using a simple Hebbian learning rule – however, implementation of this rule relied on correlation between presynaptic activity and a global reward signal capable of updating synaptic weights (Legenstein et al., 2010). This suggests reinforcement could be a biologically plausible signal for guiding plasticity in M1 during learning. Using extracellular electrode recordings in monkeys, multiple studies have reported reward-related activity in M1. Marsh et al. (2015) recorded from M1 in macaque monkeys performing a reaching task where they had to move a cursor to a target. In this task, the colour of the target predicted whether or not the movement would be rewarded. 44% of neurons increased their firing rate in anticipation of reward, 24.9% increased their firing rate when no reward was expected and 31.2% had no reward-related response. Interestingly, when the monkeys performed a separate, purely observational version of the task where they simply watched the cursor on the screen and still received reward, about 47% of neurons in M1 responded higher for reward, 40% responded higher for no reward and the remainder had no reward response. Furthermore, a linear classifier was able to accurately predict reward from the activity of 20-35% of neurons during the post-reward period (Marsh et al., 2015). A subsequent study where monkeys were trained to

manipulate a manipuladum, found a small population of neurons that responded categorically to reward. Intriguingly, more neurons in this study signaled the absence of reward than the presence (Ramkumar et al., 2016). Neurons with reward-related activity in monkeys have also been shown to multiplex information related to the task and the movement (Ramakrishnan et al., 2017). These studies all report reward-representation and reward-prediction errors in M1 neurons in well trained animals. However, because they all employed extracellular electrode recordings, the cell types could not be identified. Since PNs greatly outnumber INs in M1 (Tremblay et al., 2016), it is likely that most of the recorded neurons were PNs. Therefore, it is unclear how different cell types in M1 represent reward and how the representations might evolve with learning.

In humans, studies using transcranial magnetic stimulation (TMS) have demonstrated that M1 activity is modulated by reward. In one study, TMS was used to test intracortical inhibition during a slot machine simulation delivering variable amounts of monetary reward. Reward expectation as well as previous rewards were associated with an increased paired-pulse inhibition ratio, suggesting that reward modulates excitability in M1 (Kapogiannis et al., 2008). A separate study applying TMS to M1 during a behavioral task involving pseudorandomly delivered monetary rewards also found that reward modulated M1 excitability, likely through a $GABA_A$-dependent mechanism (Thabit et al., 2011).

Intriguingly in expert mice performing a forelimb pellet reaching task, distinct populations of L2/3 PNs reported the completion of successful vs unsuccessful reaching movements. When the food pellets were replaced with plastic pellets, the responses of L2/3 PNs to successful vs failed reaches remained intact, therefore, these neurons encoded movement outcome but not appetitive reward (Levy et al., 2020).

Thus far, research on reward representations in M1 leaves many open questions. Investigations in monkeys suggest that M1 encodes reward, reward expectation and lack of reward even in the absence of reward expectation (Marsh et al., 2015; Ramkumar et al., 2016). Moreover, single neurons multiplex reward representation with movement-related activity (Ramkrishnan et al., 2017). In humans, there is evidence that reward transiently modulates inhibition in M1 (Kapogiannis et al., 2008; Thabit et al., 2011). Finally in mice trained to expert level on a motor task, reward representations were not observed in L2/3 PNs (Levy et al., 2020). Therefore, it remains unclear if reward representations are present in M1 of mice. If reward is

indeed represented in M1 of mice, it is unclear how reward might modulate the activity of different cell types in M1, which long-range inputs to M1 drive putative reward representations and if reward representations change with learning.

RATIONALE

In this section, I briefly summarize the key literature findings described in the introduction that motivated this thesis. First, M1 is a highly interconnected brain region that integrates input from numerous brain regions, and performs local computations to multiplex sensory, reward, and movement-related information. Moreover, M1 activity is highly plastic and the coordinated activity of the GABAergic INs, most notably SST-INs, have key roles in regulating local plasticity in M1 and are necessary for successful motor learning. Since VIP-INs primarily inhibit SST-INs, VIP-INs are a good candidate for gating M1 plasticity. Since reward is known to promote and enhance learning, it is possible that reward may serve as a teaching signal in M1. Indeed, in monkeys and humans, it has been shown that reward modulates M1 activity. However, a mechanistic understanding of which brain regions drive reward-related input in M1, which cells represent reward in M1, and how this reward-related input might regulate M1 activity and plasticity to promote motor learning have not been elucidated.

HYPOTHESIS

Following reward, long-range reward-related input drive VIP-INs in M1 to disinhibit PNs, and thereby, VIP-INs serve as a disinhibitory gate to open a window of plasticity in M1.

OBJECTIVES

(1) Identify whether reward is represented among the major neuronal cell types in M1 of mice.

(2) Characterize if and how these reward representations among different cell types change with associative learning.

(3) Uncover whether different GABAergic IN subtypes in M1 receive distinct long-range input.

(4) Identify candidate reward-related brain regions that provide biased input to VIP-INs

CHAPTER 1 – Light Up the Brain: The Application of Optogenetics in Cell-Type Specific Dissection of Mouse Brain Circuits

Lee C.[1], Lavoie A.[2,3], Liu J.[2,3], Chen S.X.[1,4,5], Liu B[2,3].

[1]Department of Cellular and Molecular Medicine, University of Ottawa, Ottawa, ON, Canada

[2]Department of Biology, University of Toronto Mississauga, Mississauga, ON, Canada

[3]Department of Cell and Systems Biology, University of Toronto, Toronto, ON, Canada

[4]Brain and Mind Research Institute, University of Ottawa, Ottawa, ON, Canada

[5]Center for Neural Dynamics, University of Ottawa, Ottawa, ON, Canada

Front Neural Circuits 14 (2020):18. DOI: 10.3389/fncir.2020.00018

Author Contributions

All authors contributed to writing the review.

Abstract

The exquisite intricacies of neural circuits are fundamental to an animal's diverse and complex repertoire of sensory and motor functions. The ability to precisely map neural circuits and to selectively manipulate neural activity is critical to understanding brain function and has, therefore been a long-standing goal for neuroscientists. The recent development of optogenetic tools, combined with transgenic mouse lines, has endowed us with unprecedented spatiotemporal precision in circuit analysis. These advances greatly expand the scope of tractable experimental investigations. Here, in the first half of the review, we will present applications of optogenetics in identifying connectivity between different local neuronal cell types and of long-range projections with both *in vitro* and *in vivo* methods. We will then discuss how these tools can be used to reveal the functional roles of these cell-type specific connections in governing sensory information processing, and learning and memory in the visual cortex, somatosensory cortex, and motor cortex. Finally, we will discuss the prospect of new optogenetic tools and how their application can further advance modern neuroscience. In summary, this review serves as a primer to exemplify how optogenetics can be used in sophisticated modern circuit analyses at the levels of synapses, cells, network connectivity and behaviors.

Main Text

Introduction

In the past decades, numerous newly developed techniques have greatly assisted in dissecting connectivity and function of the brain. However, only a handful of them have influenced and advanced modern neuroscience as heavily as optogenetics. This state-of-the-art technique utilizes light-sensitive channels or pumps, known as opsins, to manipulate the activity of neurons. In addition, when it is combined with the Cre-Lox recombinase system, it provides a spatiotemporally precise method to reversibly turn on and off the activity of genetically defined or projection-specific groups of neurons. In this review, we will first highlight the use of optogenetics in the investigation of neural connectivity, both within and between brain regions, and then its applications in identifying the functional roles of specific neural circuit components in behavior and physiology. Finally, we will discuss some of the limitations and future directions of optogenetics. Although most of the examples in this review come from studies of sensory and motor systems, their diverse experimental designs and underlying principles are potentially useful for advancing our understanding of the structure and function of other brain circuits.

Opsins used in optogenetics were first discovered in microbes (Soliman and Trüper, 1982; Mukohata et al., 1988) and later cloned and introduced into neurons (Boyden et al., 2005). These microbial opsins can be divided into excitatory opsins and inhibitory opsins. The most commonly used excitatory opsin is channelrhodopsin (ChR2), a cation channel that opens in the presence of blue light (~470 nm) to depolarize neurons (Nagel et al., 2003). In contrast, inhibitory opsins, such as the chloride pump halorhodopsin (eNpHR) and the proton pump Archaerhodospin (Arch), mediate hyperpolarizing currents which impede action potentials upon yellow light illumination (~580 nm) (Zhang et al., 2007; Chow et al., 2010). In this review, we will primarily focus on the applications of commonly used opsins, rather than covering all the different variants. However, we will introduce a few recently developed opsins in order to illustrate the diverse properties of optogenetics and its unique applications.

Connectivity

Neural circuits consist of heterogeneous cell-types receiving distinct inputs from both local and long-range sources. Dissecting the intricate connections of these neural circuits has been a long-standing challenge for neuroscientists, largely due to technical limitations in identifying and

targeting specific neuronal cell-types. Although traditional methods of circuit analysis have been useful in gaining a gross understanding of macroscale and mesoscale features of brain connectivity, these techniques are limited. For example, anatomical circuit tracing with anterograde or retrograde reagents can only suggest potential innervations, without confirming the presence of functional synaptic contacts (Zeng, 2018); electron microscopy, despite its capability of identifying synapses, is labor and time intensive and cannot reveal the type or function of the synapse (Burette et al., 2015); electrical stimulation of axonal tracts, which is used to reveal functional connectivity, indiscriminately activates all fibers passing the stimulation site (Klauer et al., 1990); pairwise whole-cell recording, a gold standard for establishing connectivity, is technically challenging and time consuming and suffers from a small yield (Xue et al., 2014). Mitigating all the above issues, optogenetics, combined with the Cre-Lox recombinase system, provides a cell-type specific and high-throughput method for dissecting circuit connectivity.

Local Connectivity

Neural circuits are characterized by entangled connections between various types of neurons within the network, making detailed dissection of local circuit connectivity extremely difficult. Optogenetics has simplified experimental designs for analysis of local connectivity and has greatly boosted the efficiency of data collection. For example, by selectively expressing ChR2 in a specific population of neurons, one can study the connectivity from those ChR2-expressing neurons onto other non-ChR2 expressing neurons with ease and speed (Adesnik & Scanziani, 2010; Adesnik et al., 2012). In this experimental design, optogenetic stimulation replaces the electrical stimulation in paired whole-cell clamp recordings and greatly increases the yield and the chance of detecting connectivity, since multiple presynaptic neurons can be activated simultaneously by the light-evoked current (Seeman et al., 2018) and the spatiotemporal pattern of optogenetic stimulation can be flexibly readjusted (Adesnik & Scanziani, 2010; Adesnik et al., 2012). Notably, when this optogenetic method of local circuit analysis was compared to traditional pairwise patch clamp methods, both gave rise to similar connection probabilities, validating the utility of optogenetics in the study of local circuit connectivity (Seeman et al., 2018).

When combined with cell-type specific Cre mouse lines, optogenetics can also be used to study the connectivity of genetically-defined populations within local circuits. For instance, in the

visual cortex, Pfeffer et al. (2013) examined the pattern of connectivity between three major subtypes of GABAergic inhibitory interneurons, parvalbumin (Pvalb), somatostatin (Sst), and vasoactive intestinal peptide (VIP) expressing interneurons (Figure 1A). Until recently, these different GABAergic subtypes were poorly characterized, as they are intermingled in the cerebral cortex and could not be specifically targeted with electrical or pharmacological manipulations. By expressing ChR2 in one population at a time and recording from different interneuron subtypes identified by single-cell reverse-transcription PCR, the authors were able to elucidate microcircuit motifs. They found that Pvalb interneurons preferentially inhibit pyramidal neurons and other Pvalb interneurons; Sst interneurons preferentially inhibit pyramidal neurons and all other interneuron types except themselves; and VIP interneurons preferentially inhibit Sst interneurons (Pfeffer et al., 2013). The sample sizes required to deduce connection probabilities and circuit motifs are difficult to achieve using paired recordings. But with a high-throughput optogenetic design, as the above experiment, cell-type specific connectivity analysis becomes surmountable.

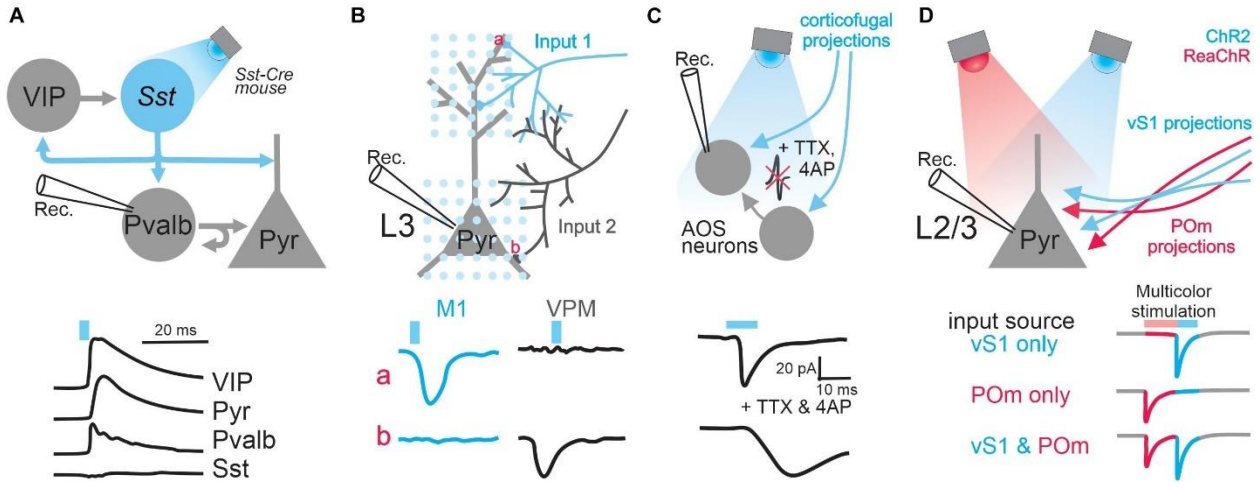

Figure 1. Applications of optogenetics in the analysis of circuit connectivity. (A) The examination of the connection from Sst positive GABAergic interneurons to excitatory pyramidal neurons and various types of GABAergic interneurons by photo-stimulating Sst neurons expressing ChR2. Top, schematic of experimental design; bottom traces, example IPSCs received by different types of cells upon Sst activation. **(B)** Subcellular ChR2-assisted circuit mapping (sCRACM) to reveal subcellular organization of two different long-range inputs impinged onto L3 pyramidal neurons. Bottom traces, example EPSCs when stimulating M1 or VPM inputs at particular domains (a, the tuft branches; b, basal dendrites). **(C)** The examination of corticofugal innervation onto brainstem AOS neurons by pharmacologically blocking action potentials.

Bottom traces: example EPSCs before and after blocking action potentials. **(D)** Dual-channel mapping to resolve two different inputs received by L2/3 pyramidal neurons. ChR2 and ReaChR are expressed in axons coming from vS1 and POm, respectively. Bottom traces, the colored segments represent the EPSC components mediated by activating POm input or vS1 input. The blue or pink bars above traces represent light stimulation.

Local connectivity can also be investigated between layers of the cortex. For instance, Bortone et al. (2014) used optogenetics to investigate how layer 6 (L6) excitatory neurons can regulate the strength of cortical responses throughout cortical depth. With the help of the L6 specific Cre mouse line, neurotensin receptor 1 (NTSR1)-Cre, they optogenetically stimulated L6 neurons and identified the recruitment of unique L6 fast-spiking interneurons with massive translaminar axons whose activation suppresses neurons across laminar layers.

Projection-Specific Connectivity

Prior to the development of optogenetics, it was practically impossible to activate only a specific group of axonal projections. Electrical stimulation of axon tracts indiscriminately activates all fibers that pass through the stimulated area, including axons originating from different brain structures and ones projecting to different areas (Schwarz et al., 2015). Therefore, electrical stimulation will activate several pathways in parallel, which complicates data interpretation. These off-target effects can now be mitigated through the use of optogenetics. Optogenetics enables specificity at several levels: the injection location of the opsin-expressing virus provides a degree of spatial specificity for presynaptic source; the use of the Cre-lox system enables cell-type specificity; and the range of the light illumination provides a final level of specificity. For example, this method was successfully applied to explore the differential connectivity of thalamocortical and corticothalamic pathways that are entwined with each other (Cruikshank et al., 2010). This approach was also applied to dissect specific basolateral amygdala projections to the central nucleus of the amygdala (Tye et al., 2011).

The power of optogenetic-based projection analysis is exemplified by a study that dissected the laminar organization of long-range callosal projections linking the barrel cortices of the two hemispheres in slice preparation (Petreanu et al., 2007). Layer 2/3 of the barrel cortex receives input from several structures such as the thalamus, other cortical areas including the contralateral barrel cortex, and local circuits (Lubke and Feldmeyer, 2007). Since the axons of

these inputs are intermingled in the barrel cortex, it is impossible to electrically stimulate only axons coming from the contralateral barrel cortex, namely the callosal axons, in order to investigate the laminar organization of their innervations. Therefore to address this question, Petreanu et al. (2007) unilaterally expressed ChR2 in layer 2/3 of the barrel cortex contralateral to the recording sites (Figure 1B). Because ChR2 was expressed throughout the neurons, including their axons projecting to the recording site of the barrel cortex, blue light illumination over the recording sites activated ChR2-expressing axons directly and thus stimulated only callosal input. Using this so-called ChR2-assisted circuit mapping (CRACM), the authors (Petreanu et al., 2007) systematically examined the strength of long-range callosal innervation received by neurons in individual layers of the barrel cortex and found that laminar specificity of this long-range cortical innervation is identical to local innervation (Petreanu et al., 2007). This study demonstrated that the CRACM method can reliably drive projection-specific inputs without the need to preserve their tracts in slices. However, one should be cautious of some limitations of this technique: its validation requires knowledge of the anatomy and cell types of the circuits under investigation; and severed axons in slice preparation have a limited supply of synaptic vesicles, which can be quickly depleted if one uses prolonged or particularly strong stimulation (Hass and Glickfeld, 2016).

Optogenetics can also be used to locate the synaptic innervation of long-range projections tagged by the expression of ChR2. In a subsequent study, Petreanu et al. (2009) slightly modified their CRACM protocol (Figure 1B): a blue laser beam was restricted to a small spot and raster scanned the area containing the dendritic tree of a cortical pyramidal neuron; direct activation of presynaptic terminals was achieved by pharmacologically blocking the propagation of optogenetically evoked action potentials. By recording from a postsynaptic neuron while systematically photostimulating its inputs at different locations, one can generate a 2D map of the long-range connections, a method named subcellular ChR2-assisted circuit mapping (sCRACM). With this elegant design, Petreanu et al. examined the spatial distribution of synaptic inputs onto the dendritic arborization of layer 3 and layer 5 pyramidal neurons in the barrel cortex (Petreanu et al., 2009). They found that different inputs target different apical or basal domains of those neurons.

Establishing Polysynaptic vs. Monosynaptic Connections

Optogenetic stimulation of long-range projections can induce responses in recorded neurons via direct synaptic contacts (monosynaptic input) and/or indirectly via recurrent connections from other neurons in the network (polysynaptic input). To establish monosynaptic connectivity, one can pair ChR2 assisted optogenetic stimulation with pharmacology. By blocking voltage-gated sodium channels with tetrodotoxin (TTX) and potassium channels with 4-Aminopyridine (4AP), action potentials and thereby polysynaptic inputs will be prevented. Consequently, any remaining light-evoked response in recorded neurons must come directly from the activation of axon terminals that express ChR2. For instance, the long-range connectivity between the visual cortex and the brainstem accessory optic nuclei (AOS) was investigated by Liu et al. (2016; Figure 1C). The authors expressed ChR2 in the visual cortex and then photostimulated terminals of corticofugal axons while performing whole-cell recordings from AOS neurons. By suppressing action potentials with two different cocktails of drugs, they definitively confirmed the presence of monosynaptic connections from the visual cortex to the AOS. Without optogenetics, this experiment would not be possible, as the corticofugal axons from the visual cortex to AOS do not form a single nerve bundle, which is required for effective electrical stimulation, but instead intermingle with axons of other types of inputs. A similar design was also used to examine monosynaptic connectivity from burst-firing neurons in the subiculum to the neurons in the entorhinal cortex (Wozny et al., 2018). ChR2 was selectively expressed in burst-firing neurons by injecting the Cre-dependent ChR2 virus into the subiculum of VGLUT2-ires-Cre mice, where Cre exists only in burst-firing neurons. In the presence of TTX, the authors elucidated monosynaptic connections from the axons of those burst-firing subiculum neurons with slice electrophysiology.

Characterization of the Synapse

Beyond the identification of connectivity, when combined with pharmacology, optogenetics can be used to probe the properties of synapses. It was thought that dopaminergic projections might co-release glutamate and dopamine (Stuber et al., 2010). However, evidence supporting this idea came from studies where electrical stimulation was used to activate dopaminergic neurons, which could activate glutamatergic neurons in the neighborhood (Hnasko et al., 2010). The non-specific

activation makes it difficult to discern whether the release of glutamate and dopamine indeed occur from the same terminal or two different termimals (Gu, 2010). To solve this issue, Tritsch et al. (2012) expressed ChR2, using the Slc6a3-IRES-Cre mouse line, in dopaminergic neurons to specifically activate dopaminergic neurons in the substantia nigra pars compacta (SNc) and pharmacologically isolated currents mediated by different neurostransmitters. Surprisingly, when they activated these neurons in the SNc and recorded from neurons in the dorsal striatum, they found both glutamatergic excitatory post-synaptic current and GABAergic inhibitory post-synaptic current, as well as amperometric dopamine transients. By combining cell-type specific optogenetics with pharmacology, they were able to definitively identify co-release of dopamine, glutamate and GABA (Tritsch et al., 2012). A similar method was also used to identify the co-release of both glutamate and dopamine by dopaminergic axonal terminals coming from the ventral tagmental area (VTA) to the prefrontal cortex (Perez-Lopez et al., 2018).

The ability to selectively activate specific projections using optogenetics also enables the characterization of synapses of local versus long-range neuronal populations. The VTA receives long-range inhibitory input from GABAergic neurons in the rostromedial tegmental area. Using ChR2, Polter et al. (2018) activated this inhitibitory input and observed that glycine receptor blocker, strychnine, significantly reduced the amplitude of inhibitory postsynaptic currents (IPSCs) in VTA neurons. They then added bicuculline, a GABAA receptor blocker, along with strychnine and all residual IPSCs were abolished, suggesting this inhibitory projection co-releases GABA and glycine. They also performed the same recording with ChR2 expressed in VTA GABAergic interneurons to stimulate local inhibition, but did not find co-release of glycine (Polter et al., 2018). They proceeded to perform further characterization of these two inhibitory inputs received by VTA neurons. For example, they utilized light stimulation to compare paired-pulse ratios of those two types of inhibitory synapses, demonstrating how optogenetics can be used to characterize synaptic properties of distinct neuronal populations.

Dual-Channel Mapping

A single neuron might receive multiple inputs coming from different regions, representing multiplexed streams of information transmission (Petreanu et al., 2009; Oh et al., 2014). Understanding this input convergence is a fundamental but difficult task since it requires the technical capability of individually manipulating different types of inputs. This independent

control of different inputs is impossible with electrical stimulation when the axons of those inputs are intermingled. This technical challenge was solved through the use of dual-color optogenetics when Hooks et al. studied the convergence of two different intermingled inputs in the primary motor cortex (Hooks et al., 2015; Figure 1D). In order to achieve this, they chose two opsin variants that prefer different wavelengths of light: blue light sensitive ChR2 was used to excite axons originating from the barrel cortex and ReaChR, an opsin activated by orange light (Lin et al., 2013), to perturb input from the posterior medial thalamic nucleus. Despite the distinct optimal excitation wavelengths of the two opsins, blue light can in fact excite ReaChR as well as ChR2. Therefore, when illuminated by blue light, both pathways will be stimulated, complicating data interpretation. To solve this complication, the authors created a clever protocol that reversibly inactivates ReaChR prior to activating ChR2 (Hooks et al., 2015), allowing complete separation of the two inputs. Indeed, with whole-cell recording they found that the input from the somatosensory cortex and that from the thalamus converge on the same layer 2/3 neurons in the motor cortex. Since this optogenetic dual-channel stimulation does not require the spatial segregation of axons of different inputs – a condition demanded by electrical stimulation – this mapping method can in principle be generalized on any convergent circuit system and is especially indispensable in the case where axons of distinct origins intermingle.

In vitro slice electrophysiology is a gold standard to determine connectivity between pairs of neurons. Although care is taken to preserve the integrity of the circuit during slice preparation and to maintain similar physiological conditions to live animals during recording, some damage and cell death is inevitable, and impedes faithful quantification of neuronal connectivity. To address this issue, efforts had been made to examine connectivity in a physiologically pristine environment with the help of optogenetics. For example, Pala and Petersen (2015) performed *in vivo* whole-cell recordings from GABAergic interneurons in L2/3 of the barrel cortex to examine the connection from excitatory pyramidal neurons. To precisely measure unitary excitatory postsynaptic potentials (uEPSPs) in GABAergic interneurons, these researchers introduced plasmid DNA encoding ChR2 into a single L2/3 pyramidal cell with electroporation and elicited one action potential per stimulus with very brief light pulses. In particular, they compared the uEPSPs response between Pvalb and Sst inhibitory neurons, finding that each inhibitory population differed in the probability, time course, strength, reliability, and short-term synaptic

plasticity of their response to excitatory stimulation. This finding largely agrees with previous *in vitro* results (Pala and Petersen, 2015).

Functional Dissection

Thus far, we have discussed how optogenetic strategies can aid the dissection of circuit connectivity. Beyond characterizing connectivity, optogenetics is also an extremely powerful tool when investigating the functional roles of specific neural circuits in animal behavior and physiology. *In vivo* investigations present a unique set of demands and limitations. For many years, researchers have relied on extracellular recordings to investigate neural activity *in vivo*. Although this method allows monitoring activity from large populations of neurons, the type of neuron recorded cannot be identified in most cases, limiting its applications in the study of cell-type specific functions. Optogenetic tagging provides a feasible solution to this problem. For example, Lima et al. (2009) restricted the expression of ChR2 to Pvalb interneurons in the auditory cortex and inserted a recording electrode into this cortical area. When they illuminated this area with brief pulses of blue light, short latency spikes that are precisely synchronized with the light pulse were reliably elicited in a population of neurons, distinguishing them as Pvalb interneurons expressing ChR2. These authors also used this approach to identify auditory cortical neurons projecting to the contralateral hemisphere, namely callosal projection neurons, for which a retrograde herpes simplex virus-1 (HSV-1) encoding ChR2 injected to the auditory cortex in one hemisphere was used to tag callosal projection neurons in the other hemisphere (Lima et al., 2009). These two examples demonstrate the utility of this optogenetic tagging technique for both cell-type and projection-type specific functional circuit analyses.

In addition to identifying neuron type *in vivo*, optogenetics can also be used to manipulate neuronal activity during behavior. The canonical experiment to define a neural correlate of a behavior is to activate or silence the putative correlate *in vivo* and assess whether and how the behavior is altered. For decades, the field had relied on electrical stimulation/lesions and pharmacological activation/silencing, however, these techniques of circuit manipulation have intrinsic problems. They offer little to no temporal precision, nor cell-type specificity; furthermore, lesions are permanent and can result in unpredicted plasticity and compensatory mechanisms that confound results (Whishaw, 2000; Murphy and Corbett, 2009). The advent of optogenetics enables immediately reversible manipulation and allows trial-by-trial and within-animal comparisons in a single session. The remainder of the review will examine how the use of

optogenetics has enabled cell-type and projection-specific circuit dissection of visual, somatosensory and motor function *in vivo*.

Visual Cortex

Since Hubel and Weisel first discovered the fundamentals of visual processing in the primary visual cortex (Hubel and Wiesel, 1959, 1963, 1968), numerous studies have extensively examined and characterized visual processing in different model systems. These studies revealed that stimulus features such as orientation and direction are encoded in the mammalian visual cortex through tuned neural responses at both the level of single neurons and cortical columns. However, due to the prevalence of pyramidal neurons over GABAergic inhibitory neurons in the cortex, the "blind" electrophysiological recording methods used in these studies primarily characterized pyramidal neurons and could not clearly elucidate how other cell-types might contribute to this neural code.

A long-standing question in the field has been to understand how orientation selectivity emerges in the visual cortex. For many years it had been speculated that inhibition from GABAergic interneurons may shape tuning of pyramidal neurons (Sillito et al., 1980; Sato et al., 1996; Ringach et al., 2003; Liu et al., 2011), although findings have been inconclusive until recently since this hypothesis could not be directly examined before the invention of optogenetics and inhibitory neuron-specific mouse lines. Using high-speed calcium imaging or *in vivo* cell-attached recordings paired with subtype-specific optogenetic activation, Wilson et al. (2012) demonstrated that different subtypes of GABAergic interneurons indeed have distinct roles in shaping pyramidal neuron responses to visual stimuli. The authors expressed ChR2 in either Pvalb or Sst interneurons in the primary visual cortex in order to activate each subtype independently while showing the mouse oriented drifting gratings (Figure 2A). They also used simultaneous calcium imaging or cell-attached recordings to characterize the orientation tuning of pyramidal neurons. By comparing pyramidal neuron tuning curves with and without optogenetic activation of Sst or Pvalb interneurons, they found that Pvalb and Sst populations exert distinct computational control on the responses of pyramidal neurons. Sst interneuron activation resulted in uniform subtractive inhibition in pyramidal neurons across all directions of drifting gratings. Pvalb interneuron activation, on the other hand, led to divisive inhibition in pyramidal cells, where the effects of inhibition were strongest when the grating was at the neuron's preferred orientation. In this way, Sst interneurons sharpen stimulus selectivity in the visual cortex, while

Pvalb interneurons modulate response gain but preserve stimulus selectivity (Wilson et al., 2012). Moreover, bidirectional manipulation of Pvalb interneurons in another study using ChR2 or Arch to activate or silence Pvalb interneurons respectively, further supports the role of Pvalb interneurons in modulating the gain of pyramidal neuron responses without strongly altering tuning properties (Atallah et al., 2012). Importantly, without this cell-type specific optogenetics, it would be impossible to target or manipulate different types of neurons.

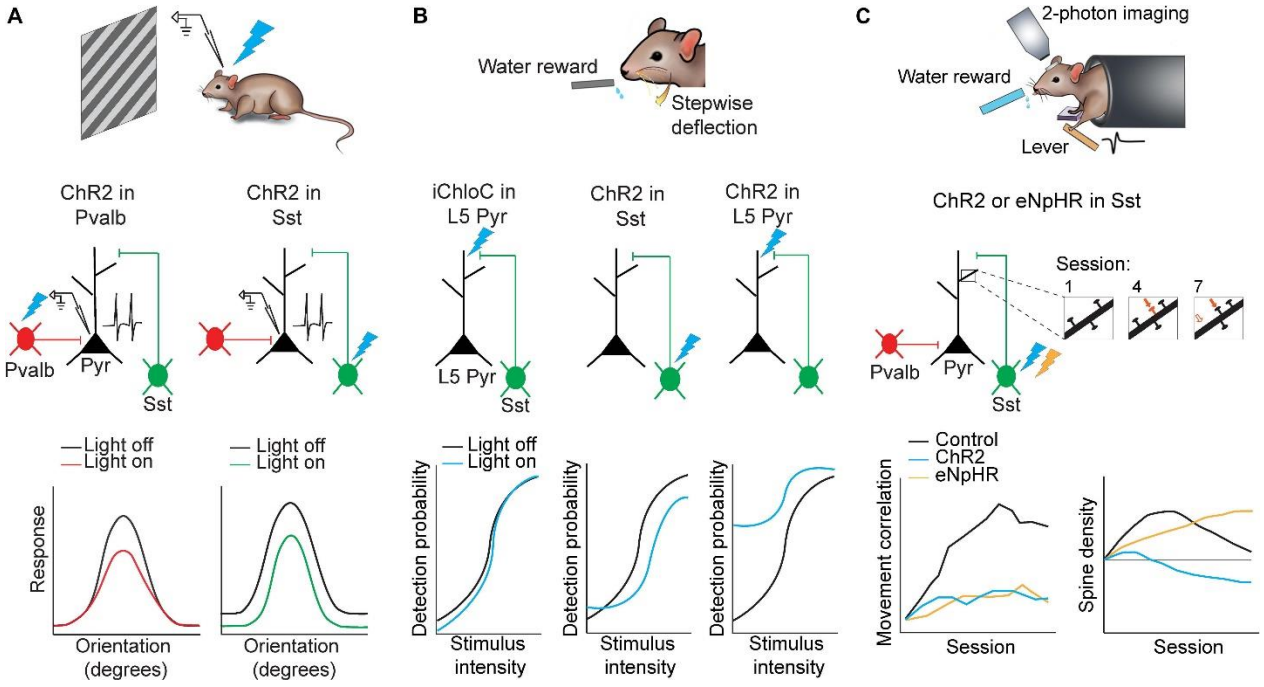

Figure 2. Applications of optogenetics in the analysis of circuit functions. (A) Mice were presented with oriented drifting gratings while *in vivo* cell-attached recordings of putative pyramidal neurons in the visual cortex were performed. Either Pvalb or Sst interneurons were photoactivated using ChR2. Pvalb interneuron photoactivation (left) resulted in divisive inhibition of pyramidal neurons where suppression was stronger at orientations where pyramidal neuron responses were also stronger. Sst interneuron photoactivation (right) resulted in subtractive inhibition where suppression was uniform along all orientations. Adapted from Wilson et al. (2012). **(B)** Mice were trained to lick for a water reward in response to whisker stimulation of varying intensities. Perceptual detection as a function of stimulus intensity formed a sigmoid curve. When L5 pyramidal neuron apical dendrites were photoinhibited (iChloC) or when Sst interneurons were photoactivated (ChR2), the curve shifted to lower detection probabilities. When L5 pyramidal neuron apical dendrites were photoactivated (ChR2), the curve shifted to higher detection probabilities. Adapted from Takahashi et al. (2016). **(C)** Mice were trained to press a lever following an auditory cue to obtain a water reward. Two-photon imaging was performed in M1 throughout learning to track dendritic spine dynamics in pyramidal neurons. Control animals developed a stereotyped lever-press movement with learning but this was impaired if Sst interneurons were photoactivated (ChR2) or photoinhibited (eNpHR) (bottom

left). Furthermore Sst interneuron photoinhibition resulted in increased stabilization of dendritic spines while photoactivation resulted in increased elimination. (bottom right). Adapted from Chen et al. (2015).

In addition to its applications in local microcircuitry, optogenetics has also been used to investigate the role of long-range inputs in visual processing. To parse the role of specific inputs to the visual cortex, Zhang et al. (2014) used a virus to express ChR2 in the cingulate cortex and then shone blue light on the visual cortex, thus exciting only the axons projecting from the cingulate cortex to the visual cortex. Activation of this top-down projection sharpened tuning of pyramidal neurons in the visual cortex and improved behavioral performance in a visual discrimination task (Zhang et al., 2014). The authors further examined which cell-types in the visual cortex might be receiving this input from the cingulate cortex. Tracing experiments revealed that Pvalb, Sst and VIP interneurons in the visual cortex all receive long-range top-down input from the cingulate cortex. Hence, to delineate subtype-specific roles in this projection, the authors again, expressed ChR2 in the cingulate cortex, but they additionally expressed the inhibitory opsin, eNpHR, in either Pvalb, Sst or VIP interneurons in the visual cortex. By patching pyramidal neurons in the visual cortex and either (i) silencing one subtype of local inhibitory neurons, (ii) activating cingulate axons, or (iii) doing both simultaneously, the effect of cingulate cortex projections on the different inhibitory populations was dissected. By taking advantage of projection-based and cell-type based optogenetics, the researchers were able to demonstrate that cingulate cortex input activates all three types of inhibitory neurons. However, Sst and Pvalb interneuron activation inhibited a broad cortical area, while VIP interneuron activation selectively enhances responses in the center region. Therefore, VIP interneurons may have a unique role of disinhibiting the center, while Sst interneurons inhibit the surround, thus explaining a facilitatory center and a suppressive surround of top-down modulation during visual processing.

Somatosensory Cortex

Optogenetics has also been employed in the somatosensory cortex to dissect the function of specific projections during perception of sensory stimuli. By expressing the inhibitory opsin Arch in the secondary motor cortex (M2), Manita et al. (2015) used amber light in the primary somatosensory cortex (S1) to inactivate projections from M2 to S1, while simultaneously

recording from S1 with either multi-unit electrodes or patch-clamp electrode. Using this preparation, the authors compared the neural response in S1 following hindpaw stimulation with and without perturbation of M2 innervation. They found that silencing M2 projections did not alter fast, putative bottom-up responses to sensory stimuli but did lead to a significant reduction in the slower second wave of top-down activity, suggesting M2 provides significant top-down modulation on S1 activity following sensory stimuli. Furthermore, using a miniature wireless LED device that can be implanted onto a freely moving, behaving mouse, they found that silencing M2 projections to S1 impaired tactile discrimination, revealing that M2 projections to S1 contribute to this type of sensory processing (Manita et al., 2015). This study demonstrates the importance of *in vivo* optogenetic approaches, since large scale network activity in response to sensory stimuli cannot be studied *in vitro*.

Another study used optogenetics in S1 to reveal the importance of dendritic activity in layer 5 (L5) pyramidal neurons in sensory perception (Takahashi et al., 2016). The authors first trained mice to lick for a reward following whisker deflection until an 80% success rate was achieved. They then constructed a psychometric curve by varying the intensity of the whisker deflection to uncover the intensity threshold for perceptual detection (Figure 2B). A multi-pronged approach was then employed to dissect the neural correlate of sensory perception in S1. Remarkably, using three distinct optogenetic approaches, the animal's detection probability curve was shifted bidirectionally. First, the authors used iChloC, an inhibitory chloride-conducting channelrhodopsin (Wietek et al., 2015), expressed in L5 pyramidal neurons to silence apical dendrites of L5 pyramidal neurons in behaving mice, which was achieved by calibrating light intensity *ex vivo* to target superficial layers of S1 without affecting activity at the soma of L5 pyramidal neurons located deeper in the cortex. Silencing the L5 pyramidal neuron dendrites shifted the detection probability curve to higher intensity values when compared to light-off trials, implying mice had impaired detection ability. Second, to further understand the local microcircuit, the authors expressed ChR2 in Sst interneurons, which preferentially inhibit apical dendrites of pyramidal neurons. Activating Sst interneurons also significantly shifted the detection probability to higher intensities. Finally, activation of ChR2 expressed in L5 pyramidal neurons shifted the detection probability to significantly lower detection probabilities and caused a substantial increase in false detection rates. The use of optogenetics enables a reversible approach for within animal and within session comparisons between light on and light off trials.

Motor Cortex

The motor cortex is unique among primary cortical areas in that its primary function is to activate muscles and execute movements. Inactivation techniques such as pharmacological inactivation or lesions lack the temporal resolution needed to dissect circuit contributions to specific aspects or components of a movement, such as an arm movement versus digit movement in a reaching task. Additionally, motor cortex lesions impair movement but do not result in the complete loss of movement (Castro, 1972; Alaverdashvili and Whishaw, 2008) due to compensatory mechanisms that change motor function over time (Whishaw, 2000), further obscuring motor cortex functions under normal conditions.

The ultra-fine temporal precision of optogenetics makes it an ideal tool for addressing these open questions regarding the motor cortex. In recent years, the use of optogenetics in the motor cortex has elicited a range of results, some of which do not support previous findings in lesion studies, further demonstrating how lesion and pharmacological studies can be limiting when probing complex systems conveying multiplexed information. To test the role of the motor cortex in learned behaviors, Guo et al. (2015) used a transgenic mouse line expressing ChR2 in all GABAergic neurons and trained head-fixed mice on a forelimb pellet reaching task until the mice achieved expert level. Following training, they used optogenetic activation of GABAergic inhibitory neurons to silence the central forelimb area of the motor cortex. Remarkably, they demonstrated that silencing the motor cortex before movement blocked the initiation of reaching, while silencing the same area following movement initiation caused the mouse to pause with its limb protracted. The other limbs and any untrained movements such as grooming were unaffected. In fact, turning off the motor cortex and then releasing the inhibition was sufficient to initiate reaching movement outside of trials when there was no task-related cue or food reward present, indicating that the release of inhibition could trigger the full activation sequence for the trained reach movement. Previous lesion studies have demonstrated impairments in behaviors but failed to reveal such a robust effect (Castro, 1972; Alaverdashvili and Whishaw, 2008). This can perhaps be explained by compensatory mechanisms that can contribute to recovery of movement when the motor cortex malfunctions for long periods of time (Murphy and Corbett, 2009), demonstrating the advantage of fast and reversible optogenetics.

Furthermore, learning a new movement has been shown to be associated with the emergence of a reproducible neuronal activation pattern in the motor cortex that is specific to the

learned movement (Peters et al., 2014). Subcellularly, motor learning also induces the addition and selective elimination of dendritic spines on pyramidal neurons to reorganize synaptic connections in the motor cortex (Xu et al., 2009; Peters et al., 2014; Chen et al., 2015b). However, until recently, the role of GABAergic interneurons in motor learning was not clear. Using a combination of two-photon imaging and optogenetics, Chen et al. (2015) demonstrated that Pvalb and Sst interneurons have distinct roles in motor learning (Figure 2C). The authors trained mice on a cued forelimb lever press task. Mice demonstrated increased success rate and highly correlated, stereotyped press movements over the course of training, hallmarks of learning. The authors then used ChR2 or eNpHR to activate or inhibit Sst interneurons respectively during each session throughout learning, while assessing learning-induced spine reorganization in pyramidal neurons in the motor cortex. Activation of Sst interneurons during motor learning led to increased elimination of newly-formed spines, thus abolishing learning-induced spine reorganization. In contrast, inhibiting Sst interneurons led to decreased elimination of newly-formed spines and hyper-stabilized the spines, thus disrupting the reorganization process. Both optogenetic activation or inhibition of Sst interneurons impaired learning and impaired the ability to form stereotyped press movements across trials. Lastly, when Sst interneurons were manipulated during the task in mice that had already been trained, the movement was unaffected, indicating that the activity of Sst interneurons may be involved in regulating structural remodeling in pyramidal neurons only during motor learning and is, therefore, most critical for the acquisition of motor memories but not the execution of previously learned movements (Chen et al., 2015b).

The contribution of long-range input to motor function has also been examined using optogenetics. To understand the role of thalamus and the anterior motor lateral area (ALM), two main sources of inputs to the primary motor cortex, in movement, Guo et al. (2017), utilized potent circuit silencing strategies that activate GABAergic inhibitory neurons innervating those two brain regions. The authors used ChR2 to optogenetically stimulate GABAergic axonal terminals coming from the thalamic reticular nucleus to remove the contribution from thalamic input or, they optogenetically activated inhibitory interneurons in ALM to examine the contribution from ALM. Along with silencing either thalamic activity or ALM activity, they also conducted simultaneous multi-unit recordings in ALM or thalamus while mice performed a lick – no lick task. The task made use of a delay period between the go cue and the movement, which

allowed for the assessment of preparatory activity in these regions preceding movement. Intriguingly, they found that silencing thalamus during the delay period almost entirely abolished ALM spiking. Similarly, silencing ALM strongly reduced thalamic spiking. This suggests that ALM and thalamus are connected in a recurrent loop resulting in persistent spiking. Lastly, silencing either area significantly impaired the mice's performance in the task when contralateral movements were required but ipsilateral movements were unaffected (Guo et al., 2017), demonstrating that this loop is critical for movement preparation. Optogenetics is essential when dissecting movements that operate on millisecond to second timescales, as this type of detailed movement deconstruction and delineation is not possible with slower approaches such as pharmacology or DREADDs.

Optogenetic activation of GABAergic inhibitory neurons in the visual cortex or in the thalamic reticular nucleus has also been used to shut down local recurrent excitation or long-range thalamocortical excitation respectively, in order to elucidate their contribution to cortical visual processing (Li et al., 2013; Lien and Scanziani, 2013; Reinhold et al., 2015).

Optogenetics in Diverse Animal Models

This review has primarily focused on the use of optogenetics in mice because of the extensive availability of transgenic mouse lines that permit cell-type specific circuit dissection. However, it is important to note that optogenetics can also be used in other animal models. In particular, rats present a desirable animal model, because they can learn more complex tasks, and because their larger size, relative to mice, enables them to tolerate implants more easily. Recently, more and more transgenic rat lines have become available, which makes the aforementioned cell-type specific experimental designs feasible for this animal model. For example, Witten et al. (2011) used Tyrosine hydroxylase (Th):Cre transgenic rats to drive the expression of the Cre recombinase under the dopaminergic specific Th promoter. By injecting Cre-dependent ChR2 virus in the VTA, the expression of ChR2 was restricted to dopaminergic neurons in the VTA. With this model, they investigated cell-type specific circuitry driving positive reinforcement during intracranial self-stimulation (ICSS). During a freely moving ICSS task, optogenetic stimulation of dopaminergic neurons in the VTA occurred when rats nosepoked one of the two identical ports, but not when they poked the other one. This simple optogenetic activation of dopaminergic system during behavior strongly biased rats to the port associated with optogenetic

stimulation, but not the other port, thus demonstrating that activation of dopaminergic cells in the VTA is sufficient to drive ICSS.

Optogenetics has also been used in rats to investigate projection-specific contributions to motivation. For instance, Warden et al. (2012) targeted medial prefrontal cortex (mPFC) neurons that project to the dorsal raphe nucleus (DRN) by virally expressing ChR2 in mPFC neurons and photostimulating their axons in the DRN. The authors used a protocol of alternating 2 min light-on epochs followed by 2 min light-off epochs for five cycles during the forced swim test, allowing multiple within session comparisons. They observed a robust increase in kick frequency during light-on epochs relative to light-off epochs. In contrast, when they activated the mPFC-lateral habenula pathway in the same way, kick frequency during light-on epochs was robustly decreased. These projection-specific, opposing behavioral phenomena were not seen when all mPFC neurons were non-selectively photostimulated. These results demonstrate that optogenetic activation of projection-specific pathways arising from the same brain area can result in opposite behavioral effects and further corroborate the need for circuit specific perturbation in behavioral studies.

Optogenetics has also been used in higher mammals, such as Rhesus monkeys, enabling the causal study of circuits involved in complex behaviors and neurological diseases (Han et al., 2009; Jazayeri et al., 2012; Shewcraft et al., 2020). In the first study that validated the use of optogenetics in non-human primates, Han et al. used lentivirus to express ChR2 in the frontal cortex of macaques, and established that viral expression of opsins is safe and can support long-term experiments (Han et al., 2009). Even though there is a lack of genetically modified Cre-expressing monkey lines, cell-type specific promoters can be used to target unique neuronal populations. For instance, Stauffer et al. (2016) used virus with a dopaminergic cell-specific Th promoter to express Cre recombinase alongside a cre-dependent ChR2 virus, thus restricting ChR2 expression to dopaminergic neurons in the midbrain. They then demonstrated that the activation of this dopaminergic population promotes reward-related learning at both cellular and behavioral levels. Similarly, the Purkinje cell specific promotor L7/Pcp2 has been used to selectively drive ChR2 expression in Purkinje cells in the cerebellum occulomotor vermis (El-Shamayleh et al., 2017), and the CaMKII promotor has been used to express ChR2 in pyramidal neurons in the frontal cortex (Han et al., 2009). In addition to cell-type specificity, projection-specific optogenetics has also been achieved in Macaque monkeys. Galvan et al. expressed ChR2

in the primary motor cortex of macaque monkeys and optogenetically activated the terminals of the corticothalamic projection from the motor cortex, revealing that this projection pathway plays a modulatory role in the thalamus (Galvan et al., 2016). These studies in non-human primates demonstrate the wider applicability of optogenetics for circuit and behavior analysis with both projection (Inoue et al., 2015; Galvan et al., 2016) and cell-type specificity (Han et al., 2009; Jazayeri et al., 2012; Klein et al., 2016; Nakamichi et al., 2019).

Future Directions

New Variants of Opsins

Over the past decade, several newly developed variants of opsins have enabled the use of optogenetics in investigations that were not technically possible with earlier variants. Although this review is not intended to provide a comprehensive overview of different opsins, this section will highlight some experimental designs that leveraged unique properties of engineered opsin variants.

When first introduced into neurons, ChR2 was demonstrated to drive spike trains ranging from 5 to 30 Hz (Boyden et al., 2005). However many neurons are capable of firing far beyond 30 Hz (O'Connor et al., 2010). High-speed opsins are indispensable to investigate fast spiking neuron types such as Pvalb interneurons (Chrobak and Buzsaki, 1998; Hajos et al., 2004); hence, oChIEF was developed in 2009 (Lin et al., 2009), ChETA developed in 2010 (Gunaydin et al., 2010), and Chronos developed in 2014 (Klapoetke et al., 2014). These ChR2 variants can reliably evoke ultra fast spiking in neurons, up to 100 Hz.

One advantage of these ultra fast opsins was demonstrated when studying the importance of long-term potentiation (LTP) for memory formation (Bliss and Lomo, 1973). Although this hypothesis had existed for many years, it was difficult to test it directly *in vivo*, because of the lack of tools to induce LTP in live animals. A classical LTP protocol for slice preparation involves high-frequency tetanic electrical stimulation (a train of brief pulses at 50–100 Hz) of the tract of presynaptic axons, which is not feasible for *in vivo* experiments. However, the ultra fast opsins which support reliable optogenetic stimulation up to 100 Hz provide an opportunity. Nabavi et al. (2014) used AAV to express oChIEF in the medial geniculate nucleus and the auditory cortex of rats and implanted a cannula in the lateral amygdala to deliver light to the

oChIEF expressing axon terminals originating from the medial geniculate nucleus and auditory cortex. The authors then delivered five trains of brief light pulses at 100 Hz and used *in vivo* field recordings to confirm the induction of LTP. When this optical stimulus used to induce LTP, was paired with a foot shock in a cued fear conditioning paradigm, the rats performed the conditioned fear response. In contrast, when the optical stimulus and the shock were unpaired, the rats did not perform the conditioned fear response, suggesting they had not formed the cued-fear memory. In this experiment, the reliable high frequency optogenetic stimulation of axon terminals depends on several desirable properties of oChIEF, such as fast opening and closing rates, high light sensitivity and the low channel inactivation (Lin, 2011; Hass and Glickfeld, 2016).

In addition to ChR2 variants with faster kinetics, step function opsin (SFO) (Berndt et al., 2009) and stabilized step function opsin (SSFO) (Yizhar et al., 2011) have been engineered to produce very slow kinetics, which confer a unique set of experimental advantages. These opsins have mutations at the C128 position in ChR2 that substantially extend the period of activation. They are also capable of bistable switching, in which they can be activated with blue light to produce prolonged depolarization that can be rapidly terminated with amber light. While SFO can support stable depolarization for minutes (Berndt et al., 2009), SSFO mediated depolarization is stable over the 30 min time scale (Yizhar et al., 2011). This unique property makes SSFO useful for *in vivo* investigations since complex animal behaviors typically occur on the timescale of minutes. In addition, SSFO can also be useful when optogenetics is paired with two-photon calcium imaging, since it solves the issue of visible light from optogenetic stimulation interfering with GCaMP based calcium imaging. For example, Makino and Komiyama (2015) used a virus carrying a Cre-dependent SSFO gene to conditionally express SSFO in Sst inhibitory neurons of Sst-Cre mice, and expressed GCaMP6f in pyramidal neurons. With this design, the researchers optogenetically turned Sst interneuron activity on (using blue light) and off (using amber light), while performing calcium imaging of mostly pyramidal neurons. The use of SSFO in this study enabled cell-type specific optogenetic manipulation throughout the entire course of a complex behavioral task with simultaneous calcium imaging. This study also exemplifies the flexibility of optogenetics, enabling within animal and within session controls.

The level of control achieved through optogenetics can be extended even further using the opto-XR family of opsins. This family of opsins was engineered by replacing the intracellular loops of rhodopsin with those from specific G-protein coupled receptors, allowing

photoactivation with light to initiate intracellular signaling pathways (Airan et al., 2009). One example is opto-α1-adrenergic receptor (AR), which when photostimulated, led to a significant increase in IP3 signaling in HEK cells. When this chimeric opsin-receptor was expressed in the nucleus accumbens of mice and photoactivated during a place preference task, mice spent significantly more time in the conditioned area in the subsequent session than control mice without the optical stimulation, indicating that opto-α1-AR can alter cell signaling *in vivo* and affect animal behavior (Airan et al., 2009). This experiment demonstrates the applicability of optogenetics, beyond ionic manipulations, to study effects of metabotropic cell signaling on behavior.

Single Cell Optogenetics

Neurons in the cortex, even of the same cell type, fire at different times, and this asynchronous firing across space and time is essential for information coding (Goard and Dan, 2009). However, so far, most of the optogenetic applications took reductionist approaches in which all opsin-expressing neurons or axonal projections were activated or inhibited at the same time by uniform light illumination. Because these approaches cannot recapitulate the natural firing patterns of groups of neurons, they limit the understanding of the physiology of neural circuits. Moreover, the dependence of cell-type specific optogenetics on Cre mouse lines limits the freedom of experimental designs. To solve these issues, new techniques which enable selective manipulation of arbitrarily defined neuronal ensembles at a single-cell spatial resolution and at a sub-millisecond temporal resolution are required. In recent years, advances in two directions have been improving fine spatiotemporal optogenetic control: single-cell photostimulation and new opsins of fast kinetics, high conductance, and subcellularly restricted expression. In an early attempt to achieve photostimulation at a single-cell resolution, a small, focal laser spot of blue light was used to raster scan brain tissue in which some neurons expressed ChR2. (Wang et al., 2007). However, this one-photon excitation method suffers from low penetration depth, lack of optical sectioning, and degraded lateral resolution due to light scattering, resulting in unwanted activation of the axons of passage and out of focus somata. To mitigate those issues, the non-linearity of two-photon excitation was used (Rickgauer and Tank, 2009; Prakash et al., 2012). However, traditional two-photon excitation methods are limited by poor temporal precision, as a focused infrared laser must raster scan the surface of opsin-expressing somata to induce enough

photocurrent. To surmount the limitation of slow scanning, two parallel illumination methods were invented – computer generated holography (CGH) (Begue et al., 2013; Dal Maschio et al., 2017; Forster et al., 2017; Ronzitti et al., 2017) and generalized phase contrast (GPC) (Papagiakoumou et al., 2010). These techniques rely on phase modulation to make a pattern of infrared illumination precisely matching the three-dimensional profiles of neurons of interest, therefore allowing simultaneous two-photon activation of multiple neuronal targets at once (Gerchberg and Saxton, 1972; Sinclair et al., 2004). When combined with temporal focusing, these parallel illumination methods can reach micrometer resolution even at a depth of hundreds of micrometers in scattering brain tissues. Advances in these state-of-the-art approaches provides a new avenue to systematically map the connectivity of neuronal networks at a cellular resolution.

In addition to the development of single-cell photostimulation, the engineering of new opsins also contributes to single-cell optogenetics. For example, neurons in most of the brain areas are surrounded by numerous neurites of nearby cells, and most methods for expressing opsins will lead to opsin expression in neighboring neurites as well. As a consequence, even with single-cell photostimulation methods mentioned above, the stimulation light will inevitably stimulate dendrites and axons which pass by the targeted neurons. To address this issue, a high-performance somatic opsin, soCoChR, was generated (Shemesh et al., 2017). The expression of the high-photocurrent channelrhodopsin (CoChR) was restricted primarily to the cell bodies of cortical neurons by fusing a short amino-terminal segment of the kainate receptor KA2 subunit to this opsin. With the help of two-photon CGH, soCoChR allows the activation of individual neurons at cellular resolution with sub-millisecond temporal precision. Future work employing both single-cell stimulation and soCoChR will enable more reliable mapping of neuronal connectivity in intact brain circuits.

Moreover, sub-millisecond optogenetics can also be helpful for exploring the connectivity of fast-spiking neurons, such as interneurons (Hajos et al., 2004). Chronos is a fast opsin that can be combined with two-photon parallel illumination to drive spiking up to 100 Hz with sub-millisecond onset precision and cellular resolution. The high spatiotemporal resolution and high rate of optogenetic stimulation exemplified in the above will be essential for faithfully replicating the asynchronous activity in neuronal networks *in vivo*.

Limitations And Alternative Strategies

Optogenetics is, however, not without limitations. In fact, in some circumstances, it can result in paradoxical and undesirable effects. Inhibitory chloride pumps such as eNpHR can lead to a change in inhibitory signaling that outlasts the photostimulation (Raimondo et al., 2012; Mahn et al., 2016). Long duration photoinhibition through chloride pumps can alter the chloride ion gradient and therefore change GABAA receptor reversal potential, resulting in synaptically evoked spiking following photoinhibition (Raimondo et al., 2012). A similar phenomenon can be seen when using light-driven outward proton pumps such as Arch. Prolonged eNpHR activation overwhelms the neuron's mechanisms for chloride homeostasis (Raimondo et al., 2012), while prolonged Arch activation can overwhelm the cell's mechanisms for pH regulation (Mahn et al., 2016). These effects can be mitigated through the use of a pulsing or sinusoidal light stimulus instead of a constant stimulus, by minimizing the intensity and duration of photoinhibition (Raimondo et al., 2012) and by using ramp termination of photostimulation instead of step termination (Mahn et al., 2016).

Another option for photoinhibition is to use ChR2 to activate GABAergic inhibitory neurons. Although this method is more effective than direct photoinhibition by the proton pumps Arch and Jaws (Chuong et al., 2014; Li et al., 2019), it also has its own limitations. For instance, optogenetically activating all GABAergic inhibitory neurons in the cortex will drastically shut down the activity of all principal cells, including various types of cortical projection neurons (e.g., corticocortico, corticothalamic, and corticofugal projections) (Babl et al., 2019). This non-specifical circuit manipulation introduces a confounding factor: we cannot determine if experimental effects result from the direct suppression of the specific pathway of interest or the concurrent suppression of multiple output pathways (Babl et al., 2019). Furthermore, another confounding factor is the rebound activity caused by abruptly terminating the photostimulation of inhibitory neurons. Those synchronized spiking events cross a large number of neurons and can last several seconds, disrupting the functionality of cortical circuits. Therefore, caution should be used when choosing an opsin and when designing a photostimulation method (Liu et al., 2016; Guo et al., 2017).

Moreover, the methods of expression should also be carefully considered. Transgenic lines result in global opsin expression, which leads to at least two complications: first, photostimulation may activate neurons that are located outside of the target region but have axons

passing through this region; second, one-photon photostimulation, which has poor spatial resolution, may activate opsin expressing neurons which neighbor the target region. While the use of virus can resolve the issues associated with transgenenic lines by spatially confining opsin expression, some caveats associated with this expression method should also be considered. First, Jackman et al. (2014) reported artificial synaptic depression in hippocampal neurons when ChR2 was expressed using specific AAV virus serotypes, including AAV1, AAV5, and AAV8. But this artifact did not occur when ChR2 was expressed transgenically or with AVV9. Second, some types of viral expression systems are cytotoxic. For instance, the expression of target genes mediated by rabies virus leads to rapid cell death starting 1 week after viral injection (Wickersham et al., 2007), preventing its use for long-term functional experiments. Third, virus-based expression can also be limited by viral tropism in certain brain regions. For example, canine adenovirus and herpes simplex virus 1, two retrograde viruses, labeled largely non-overlapping populations of the basolateral amygdala that project into the prefrontal cortex (Senn et al., 2014). Therefore, one should carefully design experiments, considering the trade-off between transgenic lines and viral methods.

Furthermore, *in vivo* optogenetics is heavily limited by the level of invasiveness. Implementing optogenetics *in vivo* requires either a cranial window for superficial brain areas or an optical fiber implant to deliver light. For deep brain regions, implants also require aspiration of tissue superficial to the target site. Chemogenetics is therefore an appealing alternative as it is effective, minimally invasive, reversible and highly specific. Chemogenetic approaches include Pharmacologically Selective Actuator Modules (PSAMs) and their corresponding Pharmacologically Selective Effector Modules (PSEMs) (Magnus et al., 2011). PSAMs are a toolbox of engineered chimeric ligand-gated ion channels that selectively bind the corresponding, engineered PSEM. For example, combining a pharmacologically selective ion binding domain with the serotonin 3a receptor or with glycine receptors, results in a chimeric channel that selectively binds PSEM and gates depolarizing cation current or hyperpolarizing chloride current respectively. These channels can be expressed *in vivo* and PSEM can be delivered through an intraperitoneal (i.p.) injection (Magnus et al., 2011). Designer Receptors Exclusively Activated by Designer Drugs (DREADDs) are another common chemogenetic toolbox that utilize engineered G-protein coupled receptors to regulate neuronal activity. Two of the most commonly used DREADDs are hM3Dq, an engineered Gq-coupled receptor and hM4di, an engineered Gi-

coupled receptor (Armbruster et al., 2007). hM3Dq and hM4Di, which effectively mediate neuronal activation or silencing, respectively, when exposed to their pharmacologically inert ligand, clozapine-N-oxide (CNO) (Armbruster et al., 2007; Alexander et al., 2009; Zhu et al., 2014). For *in vivo* experiments, CNO can be delivered via an i.p. injection to perturb the activity of all neurons which express DREADDs. Alternatively, when DREADDs are expressed in specific populations of projection neurons and CNO is be locally infused in a post-synaptic region, projection-specific activation (Vazey and Aston-Jones, 2014) or silencing (Mahler et al., 2014; Stachniak et al., 2014; Zhu et al., 2016) of axon terminals can be achieved. However, it is worth noting that CNO has dose-dependent side effects: excessive CNO is converted to clozapine and could alter animal behavior independent of DREADDs (Manvich et al., 2018). These side effects can be reduced by using an appropriate dosage of CNO. Overall, chemogenetics offers many of the same advantages as optogenetics but has the additional advantage of being minimally-invasive, making it an ideal technique for targeting deep brain structures. When combined with an intersectional approach, chemogenetics can also be used for cell-type and projection-specific dissection. Chemogenetics however, lacks the fast kinetics of optogenetics, therefore limiting its applications in reversible circuit manipulation.

Other less prevalent solutions to the invasiveness of optogenetics also exist. Transcranial near-infrared optogenetics utilizes synthesized Lanthanide-doped upconversion nanoparticles (UCNPs), which convert near infrared photons into high energy visible light sufficient to activate opsins (Chen et al., 2018). For example, the authors expressed ChR2 in the VTA, and injected UCNPs into the VTA. Afterward, they delivered near infrared light outside of the skull, which was sufficient to evoke firing in ChR2 expressing neurons in the VTA (Chen et al., 2018). Additionally, magnetic manipulation of neural activity, dubbed Magneto2.0 (Wheeler et al., 2016), also offers a non-invasive solution; however, it remains unclear whether this approach can reliably manipulate neuronal circuits in various brain structures, as conventional optogenetic methods have (Meister, 2016; Kole et al., 2019; Wang et al., 2019; Xu et al., 2019). In conclusion, experimentalists now have numerous techniques at their finger tips for flexible manipulations of neuron activity both *in vitro* and *in vivo*. Among these techniques optogenetics provides numerous powerful advantages for modern circuit dissection.

Funding

This work was supported by grants from Natural Sciences and Engineering Research Council of Canada (NSERC).

Conflict of Interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest

CHAPTER 2 - Cell-type-specific responses to associative learning in the primary motor cortex

Lee C.[1], Harkin E.F.[1], Yin X.[1], Naud R.[1-4], Chen S.X.[1,3,4]

[1]Department of Cellular and Molecular Medicine, University of Ottawa, Ottawa, ON, Canada;

[2]Department of Physics, STEM Complex, University of Ottawa, Ottawa, ON, Canada;

[3]Brain and Mind Research Institute, University of Ottawa, Ottawa, ON, Canada;

[4]Center for Neural Dynamics, University of Ottawa, Ottawa, ON, Canada

eLife; 2022;11:e72549. DOI: 10.7554/eLife.72549

Rights to Reproduce

Author Contributions

Lee C.: Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Supervision, Validation, Visualization, Writing - original draft, Writing - review and editing.

Harkin E.: Formal analysis, performed the analyses for Figure 3.

Yin X.: Data curation.

Naud R.: Formal analysis, Supervision, performed the analyses for Figure 3

Chen S.X.: Conceptualization, Funding acquisition, Supervision, Writing - original draft, Writing - review and editing

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. Pyramidal neurons 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 and 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 and 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 and 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 nonoverlapping 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, 1-s duration), followed by a 1.5-s 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 intertrial intervals (ITIs) between 60 and 120 s (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., 2015b; Cichon and 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., 2015b; 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 days 1 and 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

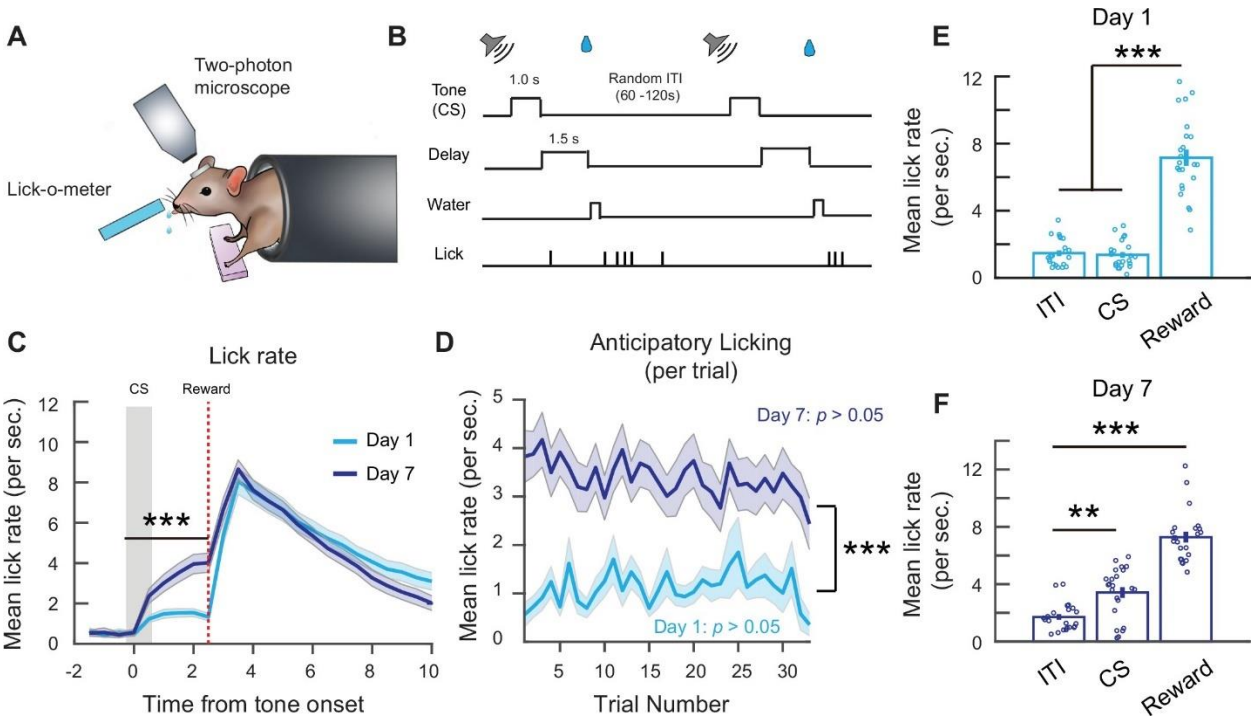

Figure 1. Associative learning during a head-fixed classical conditioning task. (A) Schematic of head-fixed classical conditioning task. **(B)** Trial structure. **(C)** Mean lick rate per second on days 1 and 7. Binned over 0.5-s intervals. Lick rate following cue stimulus (CS) onset up to reward delivery time is higher on day 7. Two-way analysis of variance (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 days 1 and 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 2.5 s of intertrial interval (ITI) lick bouts, 2.5 s following CS onset and 2.5 s 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 2.5 s of ITI lick bouts, 2.5 s following CS onset and 2.5 s 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. **$p < 0.01$, ***$p < 0.001$. Error bars show standard error of the mean (SEM).

(in the absence of the CS or reward). To be consistent with the 2.5-s analysis window for CS responses, we analyzed the first 2.5 s of ITI lick bouts and 2.5 s following reward delivery. On day 1, the mean lick rate during ITI lick bouts (1.47 ± 0.17/s) and the CS (1.37 ± 0.17/s) were similar, while the lick rate following reward delivery was significantly higher (7.16 ± 0.48/s). In contrast, on day 7 following associative learning, the lick rate during the CS period (3.42 ±

0.37/s) was significantly higher than during ITI lick bouts (1.71 ± 0.2/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 adenoassociated virus (AAV) carrying a Ca^{2+} indicator (GCaMP6f) driven by the CaMKII promoter (AAV1.CaMKIIa.GCaMP6f) into M1 of wild-type 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 days 1 and 7 (Figure 2A). We identified all the active neurons within a session, irrespective of the behavioral 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 days 1 and 7 after associating learning (Figure 2—figure supplements 1 and 2). 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 behavioral task) on days 1 and 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 behavioral task, and then calculating the mean percentage of active neurons during the sampled period. By repeating this 1000 times for each cell type on days 1 and 7, we created a distribution of the percentage of active neurons that were present at baseline

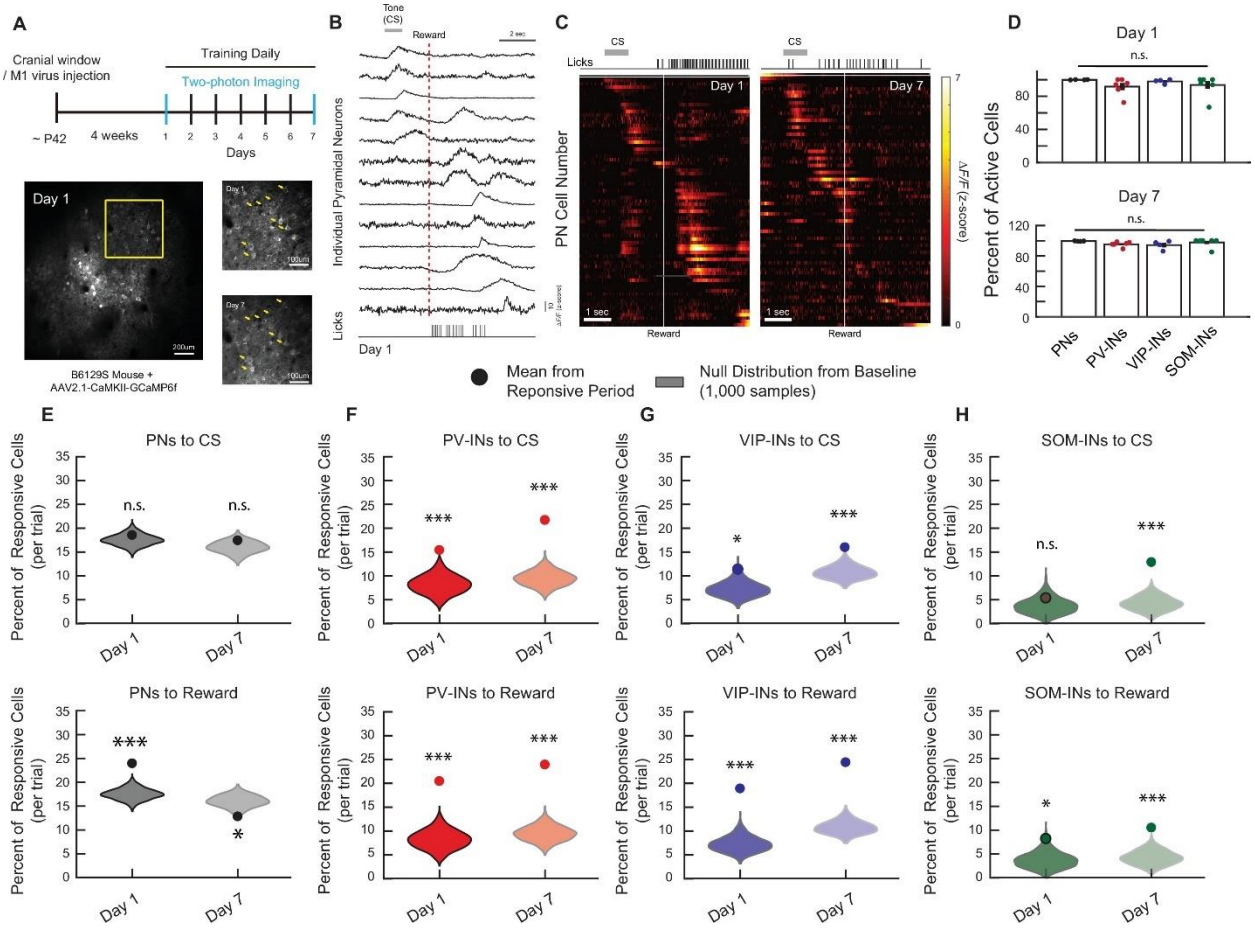

Figure 2. Longitudinal *in vivo* Ca2+ imaging of neuronal responses in M1 during a classical conditioning task. (A) Experimental timeline (top). *In vivo* two-photon imaging of L2/3 pyramidal neurons (PNs) expressing GCaMP6f in M1 (bottom left). The same population can be tracked from days 1 to 7 (bottom 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. Gray bar represents the timing of the cue stimulus (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 days 1 and 7, sorted by timing of maximum activity following the CS onset. Gray 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 behavioral task, for PNs, PV-INs, VIP-INs, and SOM-INs on days 1 and 7. All cell types showed a similar percentage of active neurons. One-way analysis of variance (ANOVA), n.s., nonsignificant, day 1: $p = 0.38$, day 7: $p = 0.22$. Error bars show standard error of the mean (SEM). **(E–H)** Mean percent of responsive neurons to CS (top) and reward (bottom) within 2.5 s of CS/reward onset for each cell type. Violin plots show null distribution of percentage of responsive neurons made by randomly resampling mice and shuffling the session, 1000 times (see Methods). The circle represents the mean percentage of tone- or reward-responsive neurons. Monte-Carlo with Bonferroni correction, n.s., nonsignificant, *$p < 0.05$, ***$p < 1 \times 10{-}3$. 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

six 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 six 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 four 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 seven mice (H).

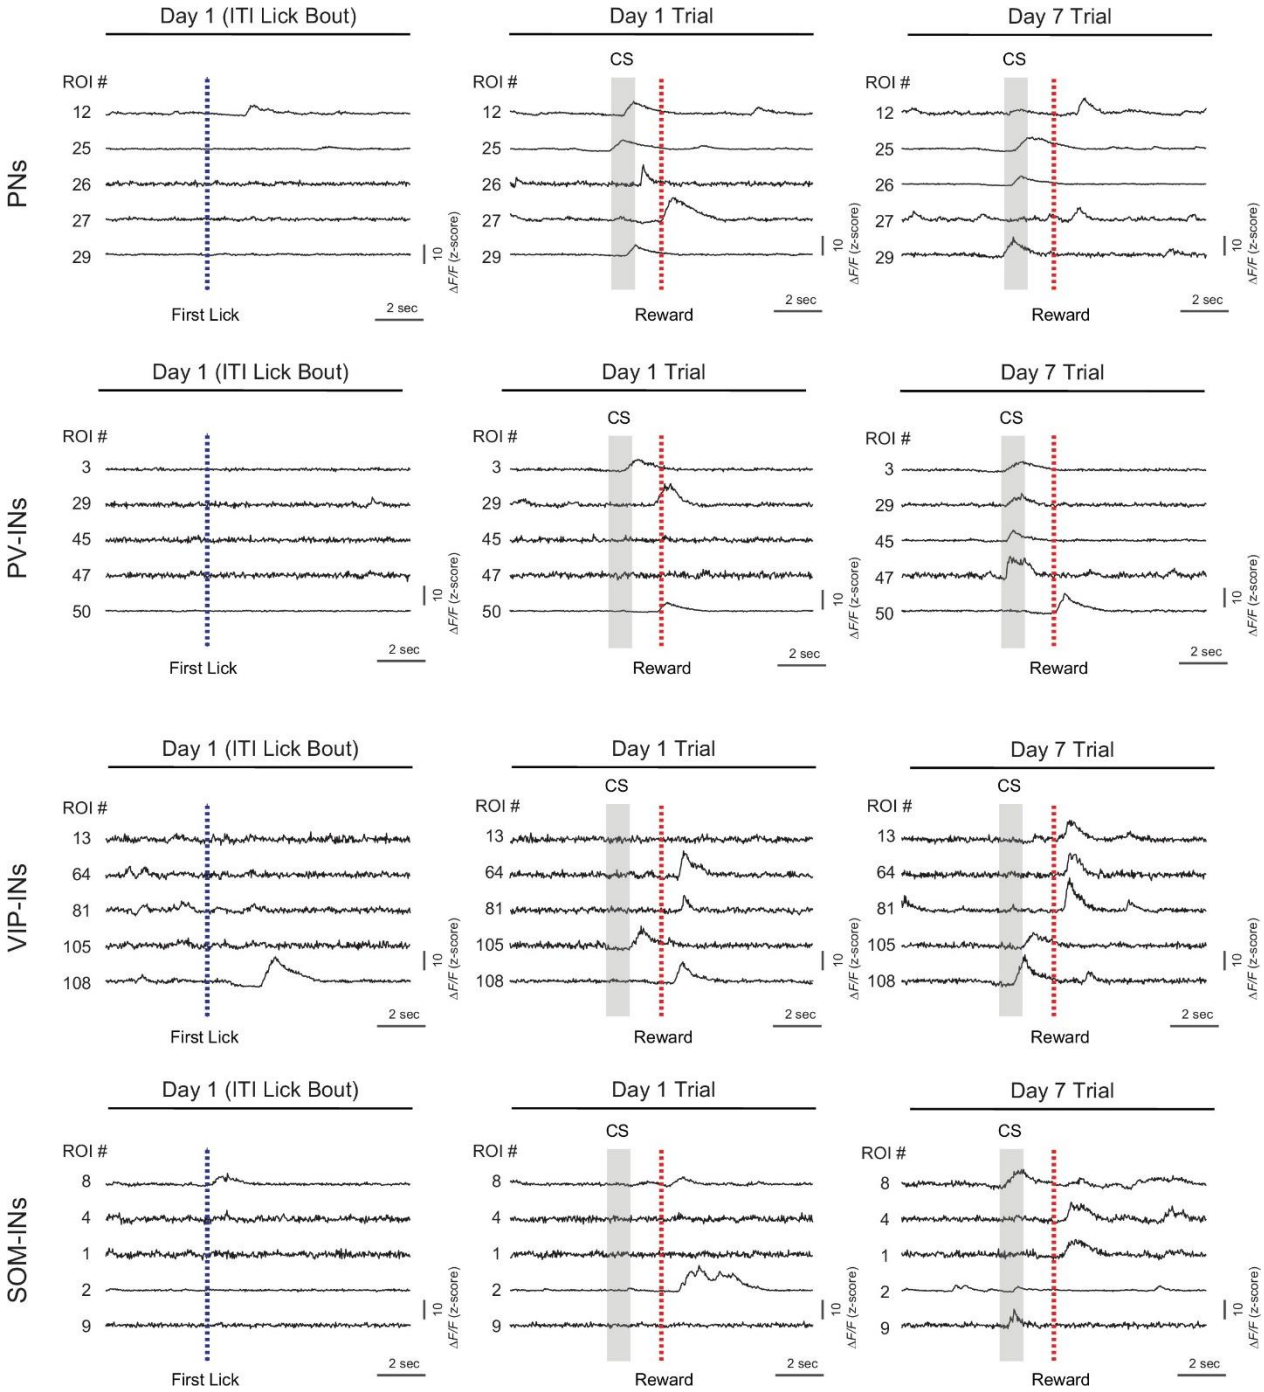

Figure 2 – figure supplement 1. Z-Scored ΔF traces from cells among each cell type on days 1 and 7. Example calcium traces from five example cells (same mouse) during a representative intertrial interval (ITI) lick bout on day 1 (left), a trial on day 1 (middle) and a trial on day 7 (right). The same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

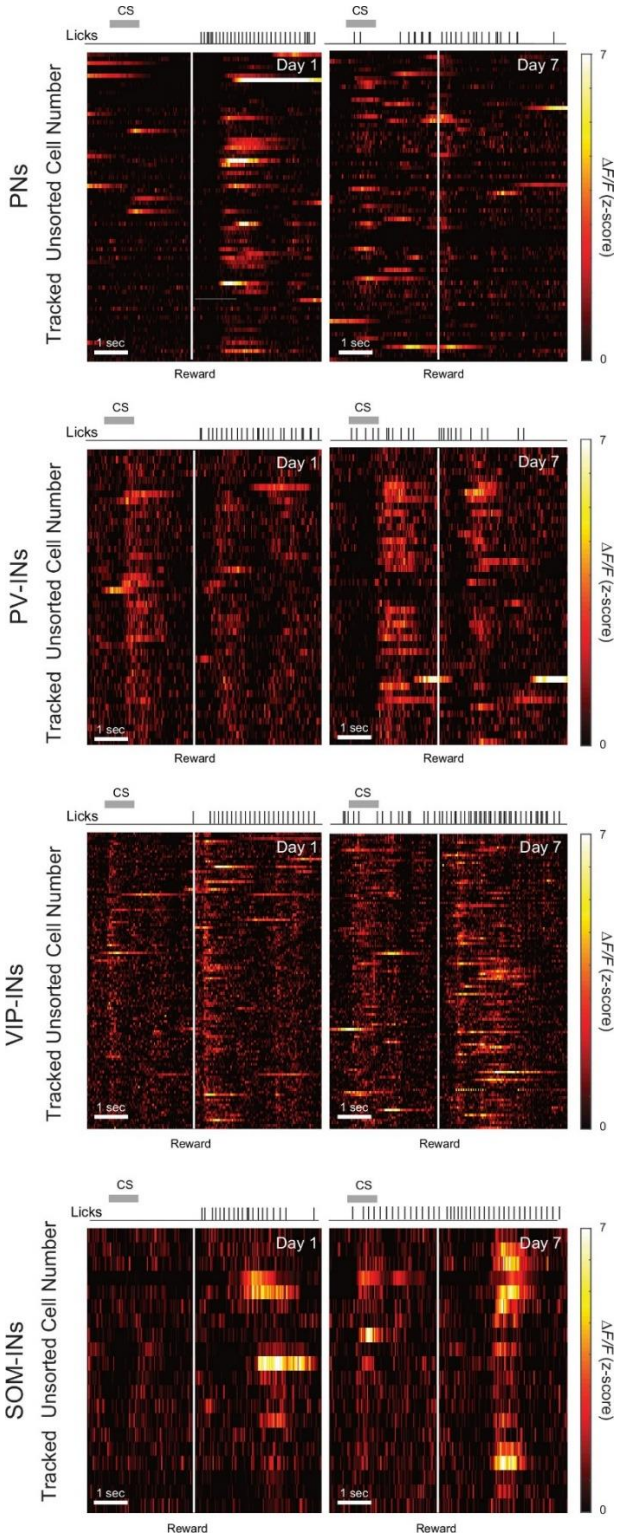

Figure 2 – figure supplement 2. Z-Scored population activity of cells tracked from days 1 to 7 for each cell type. Neurons were not sorted and the order of all cells were maintained for the days 1 and 7 plot. Gray bar represents the timing of the cue stimulus (CS). White line indicates the onset of water reward delivery. Black lines indicate licking

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 days 1 and 7 (PV-IN CS: day 1: 15.26% ± 2.11%, day 7: 24.09% ± 2.98%; PV-IN reward: day 1: 20.17% ± 3.27%, day 7: 27.47% ± 3.66%; VIP-IN CS: day 1: 11.29% ± 3.23%, day 7: 16.59% ± 2.01%, VIP-IN reward: day 1: 18.65% ± 5.91%, day 7: 26.3% ± 6.04%; Figure 2F, G). PN responses to the CS were not different from the null distribution on both days 1 and 7 (day 1: 18.42% ± 1.05%, day 7: 16.41% ± 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% ± 2.49%, day 7: 12.12% ± 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% ± 2.7%, day 7: 13.56% ± 3.17%; Reward: day 1: 9% ± 3.66%, day 7: 12.15% ± 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 nonsignificant percent of active cells during the stimulus period cannot be readily distinguished from nontask-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 days 1 and 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 (2.5-s window) using the nonparametric 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 days 1 and 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

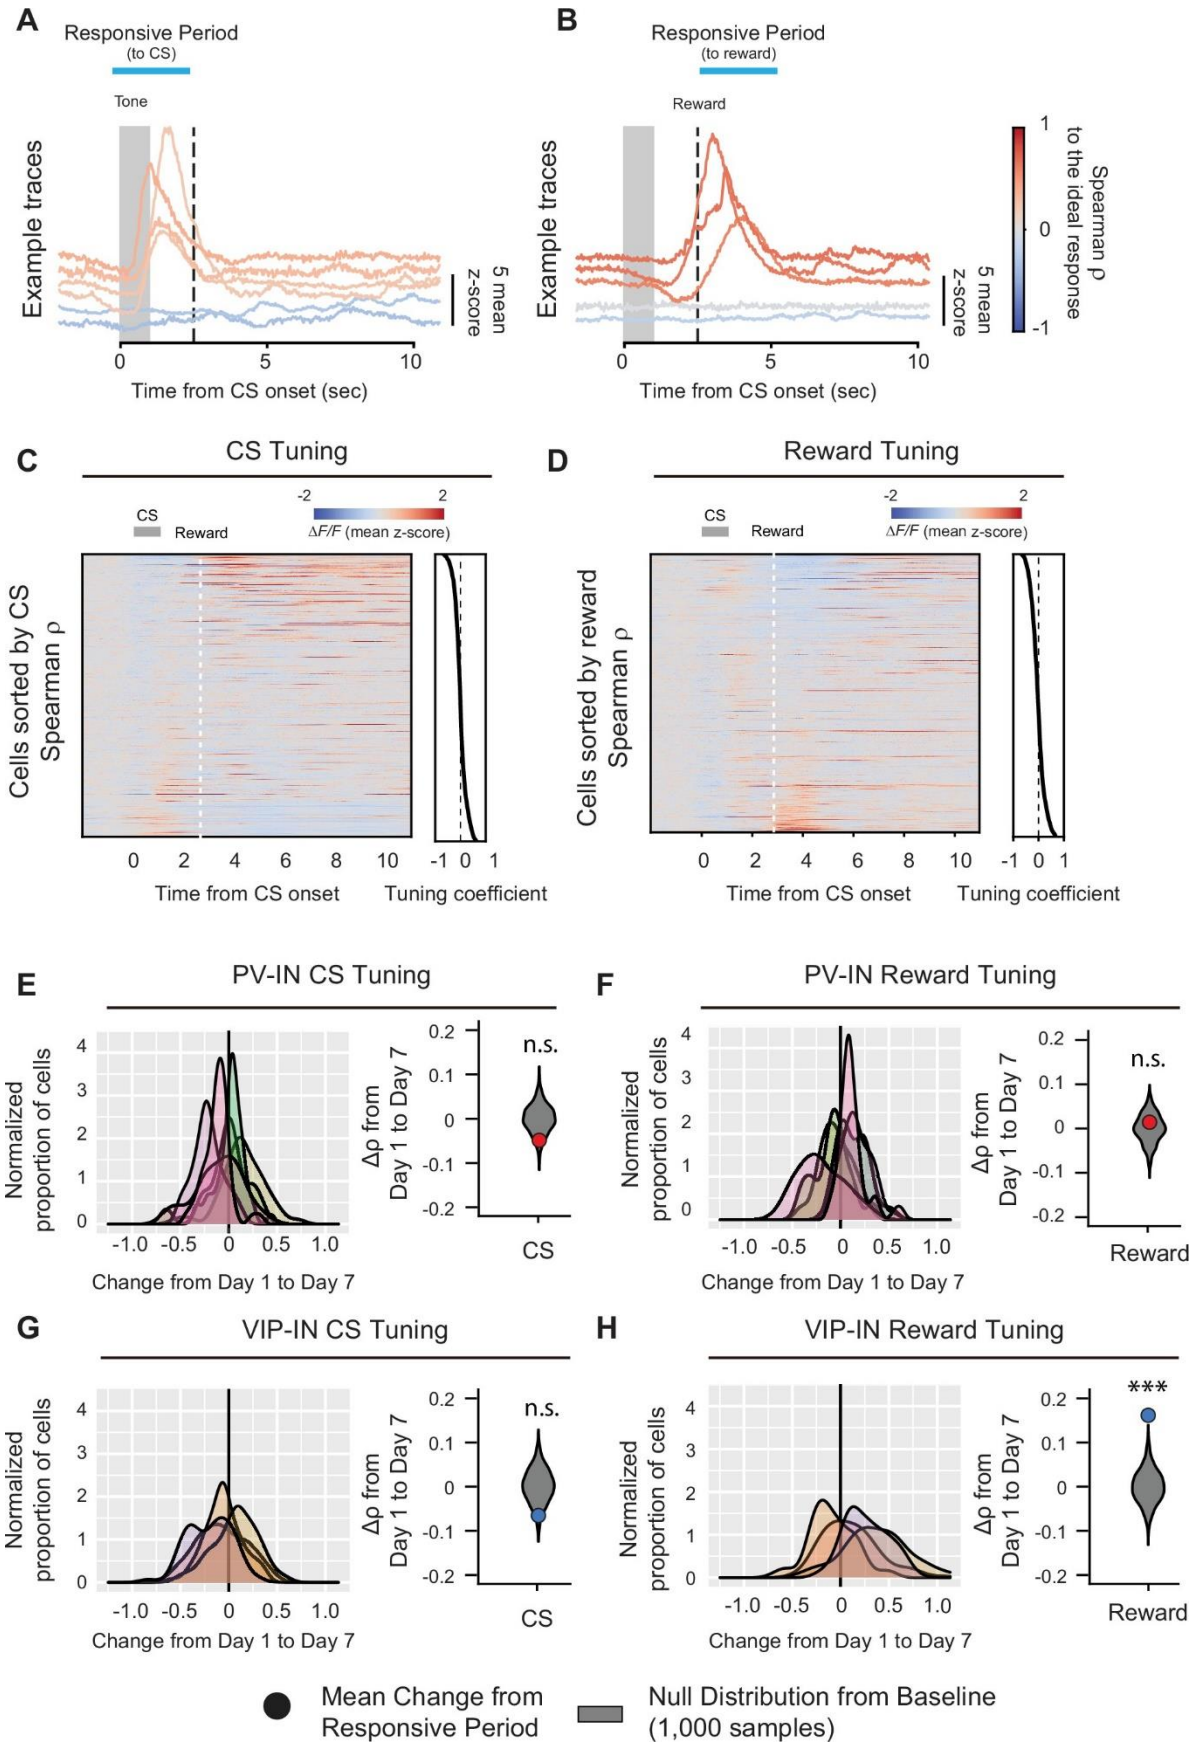
A
Responsive Period
(to CS)
Tone
Example traces
5 mean z-score
0
5
10
Time from CS onset (sec)
B
Responsive Period
(to reward)
Reward
Example traces
5 mean z-score
1
0
-1
Spearman ρ to the ideal response
0
5
10
Time from CS onset (sec)
C
CS Tuning
CS
Reward
-2
2
ΔF/F (mean z-score)
Cells sorted by CS Spearman ρ
0 2 4 6 8 10
Time from CS onset
-1 0 1
Tuning coefficient
D
Reward Tuning
CS
Reward
-2
2
ΔF/F (mean z-score)
Cells sorted by reward Spearman ρ
0 2 4 6 8 10
Time from CS onset
-1 0 1
Tuning coefficient
E
PV-IN CS Tuning
Normalized proportion of cells
Change from Day 1 to Day 7
Δρ from Day 1 to Day 7
n.s.
CS
F
PV-IN Reward Tuning
Normalized proportion of cells
Change from Day 1 to Day 7
Δρ from Day 1 to Day 7
n.s.
Reward
G
VIP-IN CS Tuning
Normalized proportion of cells
Change from Day 1 to Day 7
Δρ from Day 1 to Day 7
n.s.
CS
H
VIP-IN Reward Tuning
Normalized proportion of cells
Change from Day 1 to Day 7
Δρ from Day 1 to Day 7
***
Reward
Mean Change from Responsive Period
Null Distribution from Baseline (1,000 samples)

Figure 3. Learning-associated changes in single-neuron tuning properties in M1. Example fluorescence traces, color-coded based on nonparametric Spearman correlation with cue stimulus (CS) **(A)** or reward **(B)**. Each trace is a trial-averaged response from different neurons on day 7. Trial-averaged fluorescence of all the neurons recorded on day 7 and sorted to the value of the Spearman correlation (−1−1) to CS **(C)** or reward **(D)**. Active neurons during the CS- or reward-responsive period showed higher tuning coefficient. Left, distribution of changes in Spearman correlation Δρ 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 Δρ 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 resampling each mouse and shuffling trials 1000 times (see description of calculation of tuning coefficients in Methods). VIP-IN reward tuning significantly increased with associative learning. Monte-Carlo, ***p < 1 × 10−3, n.s., nonsignificant, 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 × 10−3 **(H)** PV-IN: n = 316 cells from six mice. VIP-IN: n = 407 cells from four mice.

any significant changes in either CS or reward tuning between days 1 and 7 ($\overline{\Delta\rho}_{CS} = -0.049 \pm 0.046$, $\overline{\Delta\rho}_{reward} = 0.014 \pm 0.054$; 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 days 1 and 7 $\overline{\Delta\rho}_{CS} = -0.065 \pm 0.048$, but VIP-INs' reward tuning significantly increased on day 7 ($\overline{\Delta\rho}_{reward} = 0.161 \pm 0.086$; 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 days 1 and 7 (Figures 4B and 5B). Hence, by identifying and tracking the same neurons from days 1 to 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 days 1 and 7, we found that the average percent of CS-responsive PV-INs during a trial increased significantly by day 7 (day 1: 15.26% ± 2.11%, day 7: 24.09% ± 2.98%; Figure 4C), while the percent of CS-responsive VIP-INs did not change (day 1: 11.29% ± 3.23%, day 7: 16.59% ± 2.01%; Figure 4D), demonstrating that more PV-INs became responsive to the CS after associative learning. We then assessed the reliability of the

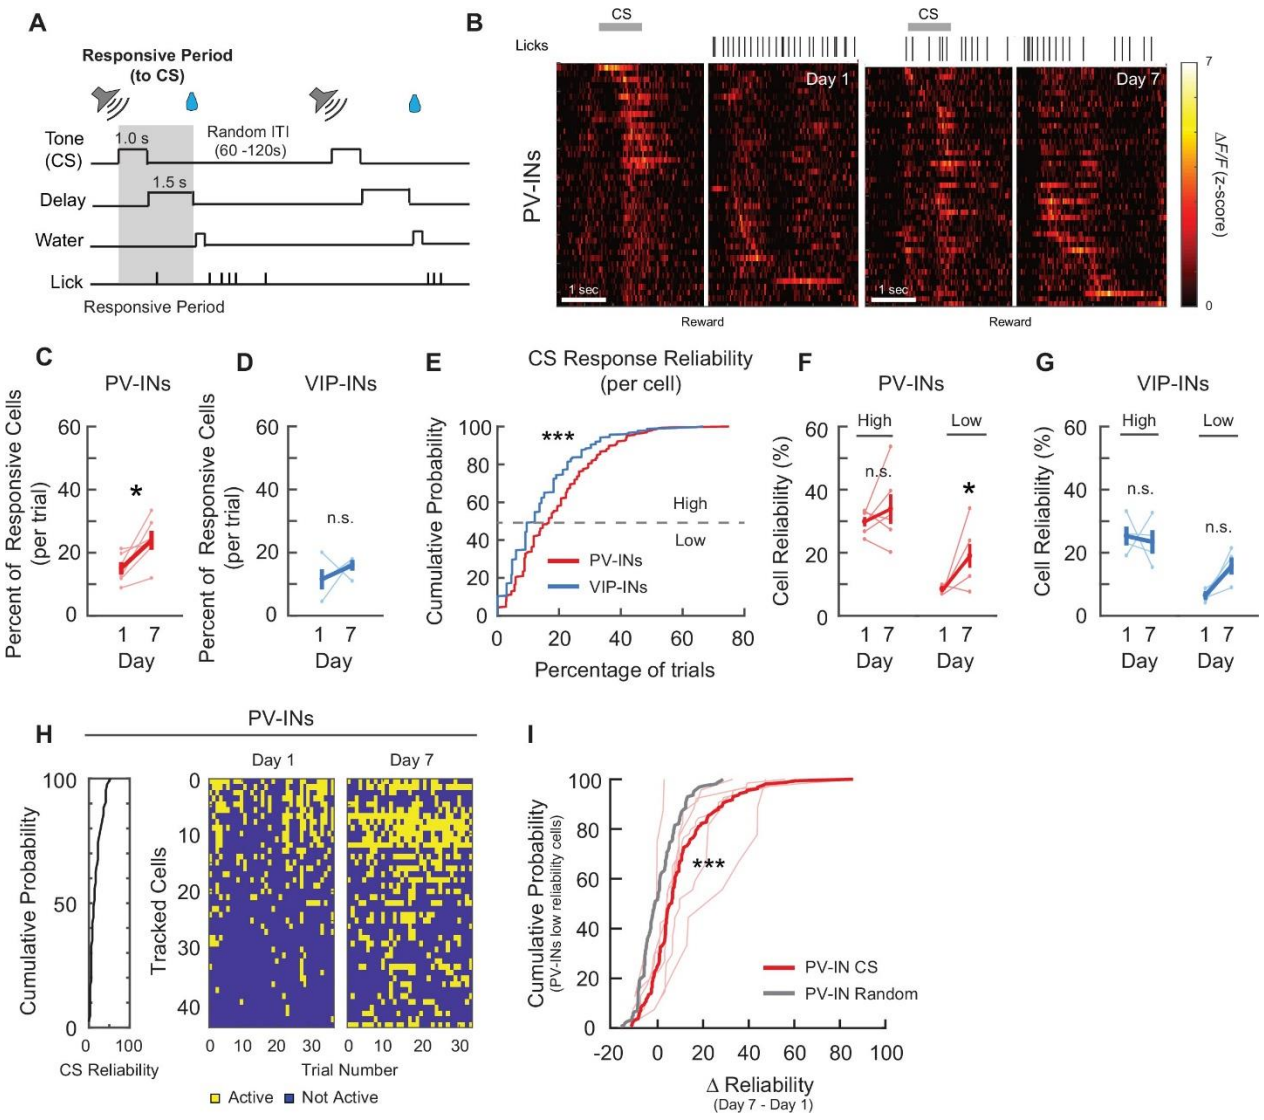

Figure 4. PV-IN and VIP-IN cue stimulus (CS)-related responses before and after associative learning. (A) Trial structure. Gray shaded bar represents the response period analyzed for CS-responsive activity. **(B)** Z-Scored activity of all the active PV-INs from an example mouse during one representative trial on days 1 and 7, sorted by timing of maximum activity following the CS onset. Gray bar represents the timing of the CS. White line indicates the onset of water reward delivery. 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., nonsignificant, 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, $p < 1 \times 10{-}3$. 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 days 1 and 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 days 1 and 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 days 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 (gray, $\Delta reliability_{random}$from the same population of PV-INs). Kolmogorov–Smirnov test, ***p < 1 × 10−3 PV-IN: n = 316 cells from six mice. VIP-IN: n = 407 cells from four mice. Error bars show standard error of the mean (SEM).

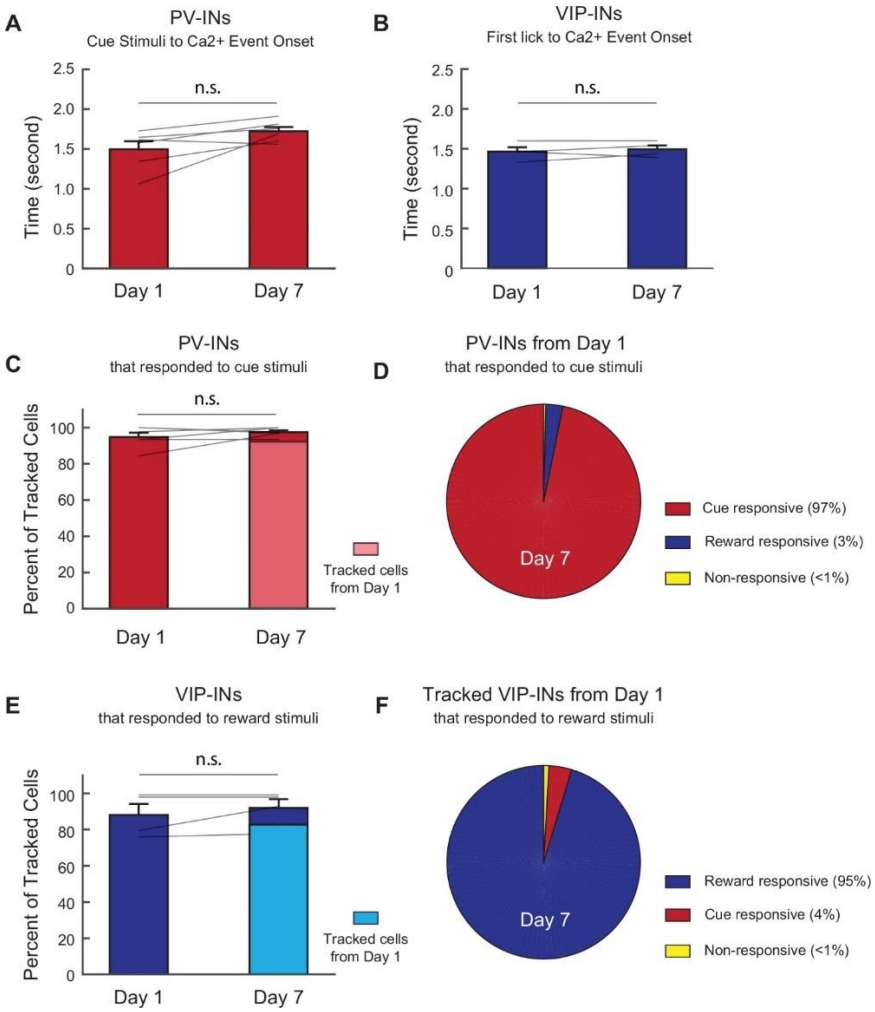

Figure 4-figure supplement 1. Response properties of PV-IN and VIP-IN cells tracked from days 1 to 7 are consistent across days. (A) Mean time from the cue stimulus (CS) onset to Ca2+ event onset for all PV-IN CS responses. Each line represents an individual mouse. The mean event onset time did not change from days 1 to 7. Paired t-test, p = 0.058. **(B)** Mean time from the first lick in response to reward delivery to Ca2+ event onset for all VIP-IN reward responses. Each line represents an individual mouse. The mean event onset time did not change from days 1 to 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 six mice. VIP-IN: n = 407 cells from four mice. Error bars show standard error of the mean (SEM).

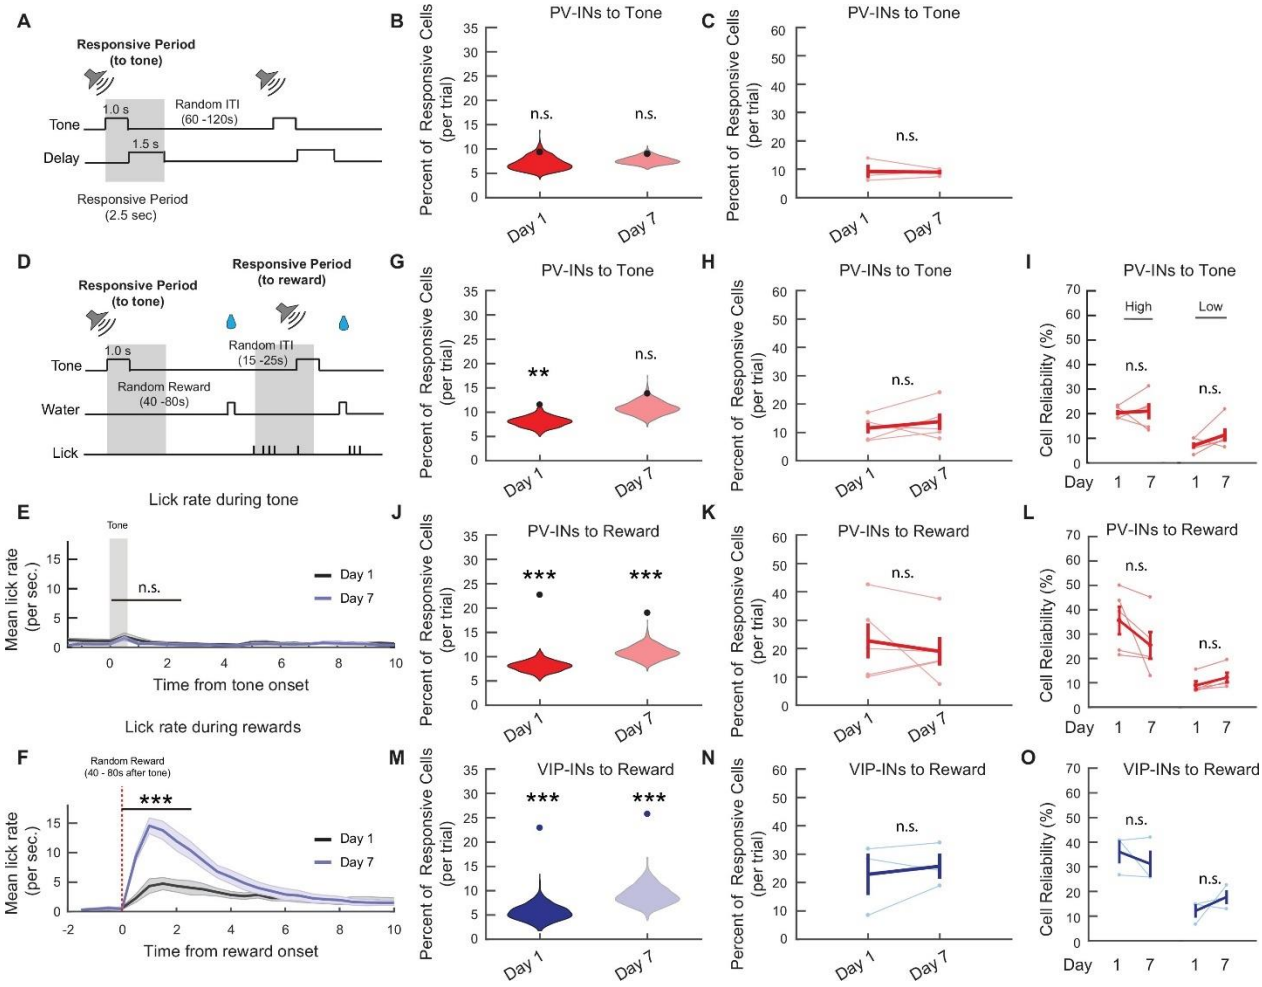

Figure 4-figure supplement 2. PV-IN and VIP-IN population plasticity is specific to associative learning. (A) Trial structure for the no-reward control task. Gray shaded bar represents the response period analyzed for tone-responsive activity. **(B)** Mean percent of tone-responsive PV-INs from the no-reward task. Violin plots show null distribution of percentage of responsive neurons made by resampling mice and shuffling the session 1000 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 days 1 and 7. Paired t-test, n.s., nonsignificant, p = 0.89. **(D)** Trial structure for the nonpaired control task. Tone was followed by a randomly varied 40- to 80-s period before reward delivery and then a randomly varied intertrial interval (ITI) between 15 and 25 s. First gray

shaded bar shows the response period analyzed for tone-responsive activity (2.5 s from tone onset). Second shaded bar shows the response period analyzed for reward-responsive activity (2.5 s from onset of first lick after reward delivery). **(E)** Mean lick rate during the tone from the nonpaired task. Mice did not show increased anticipatory licking during the tone on day 7. One-way analysis of variance (ANOVA), $p = 0.097$. **(F)** Mean lick rate following the randomized water rewards from the nonpaired 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 nonpaired task. Violin plots show null distribution of percentage of responsive neurons made by resampling mice and shuffling the session 1000 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 nonpaired task did not change between days 1 and 7. Paired t-test, n.s., nonsignificant, $p = 0.43$. **(I)** Mean reliability index of PV-INs that were responsive to the tone during the nonpaired 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 days 1 to 7. Paired t-test, n.s., nonsignificant, High: $p = 0.86$, Low: $p = 0.19$. **(J)** Mean percent of reward-responsive PV-INs during the nonpaired task. Violin plots show null distribution of percentage of responsive neurons made by re-sampling mice and shuffling the session 1000 times. The circle represents the mean percentage of reward-responsive neurons. PV-INs responded significantly to randomly timed rewards on days 1 and 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 nonpaired task did not change between days 1 and 7. Paired t-test, n.s., nonsignificant, $p = 0.51$. **(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 days 1 to 7. Paired t-test, n.s., nonsignificant, High: $p = 0.23$, Low: $p = 0.24$. **(M)** Mean percent of reward-responsive VIP-INs during the nonpaired task. Violin plots show null distribution of percentage of responsive neurons made by resampling mice and shuffling the session 1000 times. The circle represents the mean percentage of reward-responsive neurons. VIP-INs responded significantly to randomly timed rewards on days 1 and 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 nonpaired control did not change between days 1 and 7. Paired t-test, n.s., nonsignificant, $p = 0.57$. **(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 days 1 to 7. Paired t-test, n.s., nonsignificant, High: $p = 0.53$, Low: $p = 0.22$. $^{**}p < 0.01$, $^{***}p < 0.001$. No-water rewards: n = 357 PV-IN cells from three mice. Nonpaired rewards: n = 324 PV-IN cells from five mice; n = 208 VIP-IN cells from three mice. Error bars show standard error of the mean (SEM).

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% ± 1.51%, day 7: 33.87% ± 4.72%), while PV-INs that began as Low Reliability became significantly more reliable (8.47% ± 0.46%, day 7: 18.99% ± 3.76%; Figure 4F). In contrast, the reliability of both High and Low VIP-INs did not change (High Reliability: day 1: 26.55% ± 2.62%, day 7: 25.93% ± 3.81%; Low Reliability: day 1: 6.32% ± 0.76%, day 7: 14.24% ± 2.6%; Figure 4G). We then followed individual Low Reliability PV-INs and calculated the change in reliability to the CS ($\Delta reliability_{CS}$) from days 1 to 7. As a control, we randomly sampled the day 7 session irrespective of the behavioral 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 ($\Delta reliability_{random}$) for each Low Reliability neuron (see Methods). When we compared the two distributions, we found that among the Low Reliability PV-INs, $\Delta reliability_{CS}$ was significantly greater than the $\Delta reliability_{random}$ 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 nonpaired randomly timed rewards. In the first experiment, water-

restricted PV-Cre mice were exposed to the same auditory tone used as the CS (1-s duration) but all water rewards were omitted (randomly varied ITI between 60 and 120 s; ~30 trials/session; 1 session/day for 7 days). We imaged PV-IN activity on both days 1 and 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 7 (Figure 4—figure supplement 2B); hence, the average percent of tone-responsive PV-INs per trial also did not increase from days 1 to 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 'nonpaired' water reward that was given at randomly varied time intervals (40–80 s). Mice were also trained for ~30 trials/session (1 session/day for 7 days) with a randomly varied ITI between 15 and 25 s (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 nonpaired 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 days 1 to 7 in the nonpaired 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 days 1 to 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 2.5 s of the reward delivery time (Figures 4B and 5A, B). We tracked the same neurons from days 1 to 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

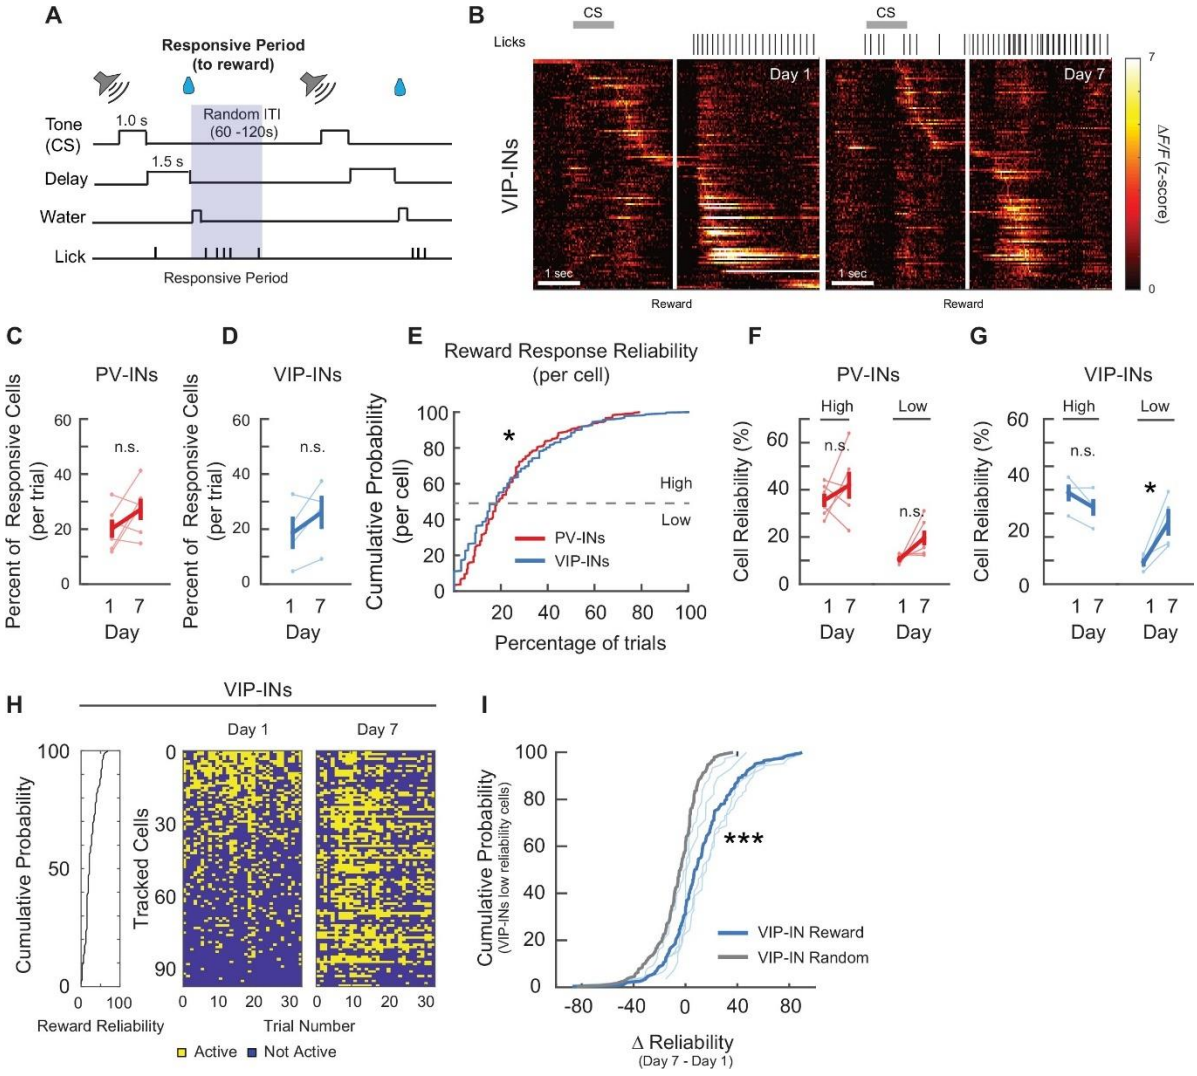

Figure 5. PV-IN and VIP-IN reward-related responses before and after associative learning. (A) Trial structure. Gray shaded bar represents the response period analyzed for reward-responsive activity. **(B)** Z-Scored activity of all the active VIP-INs from an example mouse during one representative trial on days 1 and 7, sorted by timing of maximum activity following the cue stimulus (CS) onset. Gray bar represents the timing of the CS. White line indicates the onset of water reward delivery. 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., nonsignificant, 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, $*p < 0.05$, $p = 5.2 \times 10{-}3$. 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 days 1 and 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 days 1 and 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 days 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 (Gray, $\Delta reliability_{random}$from the same population of VIP-INs). Kolmogorov–Smirnov test, ***p < 1 × 10−3. PV-IN: n = 316 cells from six mice. VIP-IN: n = 4 07 cells from four mice. Error bars show standard error of the mean (SEM).

cells per trial (PV-IN: day 1: 20.17% ± 3.27%, day 7: 27.47% ± 3.66%; VIP-IN: day 1: 18.65% ± 5.91%, day 7: 26.3% ± 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% ± 2.71%, day 7: 35.88% ± 4.32%), and the Low Reliability VIP-INs became significantly more reliable on day 7 (VIP-IN Low Reliability: day 1: 10.25% ± 1.67%, day 7: 24.06% ± 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% ± 2.81%, day 7: 41.94% ± 5.7%; Low Reliability: day 1: 10.58% ± 0.74%, day 7: 19.59% ± 3.08%; Figure 5F). We also followed individual Low Reliability VIP-INs and calculated the change in reliability to reward ($\Delta reliability_{reward}$) from days 1 to 7. As a control, we randomly sampled the day 7 session irrespective of the behavioral 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 nonpaired 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 (2.5 s from the first lick after reward delivery) compared to the chance level on both days 1 and 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 days 1 to 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 days 1 and 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 days 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 ($\Delta\rho_{tone} = 0.059 \pm 0.031$, Figure 6C), suggesting their responsivity did not change with learning. Interestingly, when we assessed the cumulative distribution of CS response reliability on days 1 and 7, the cumulative distribution function shifted to significantly higher

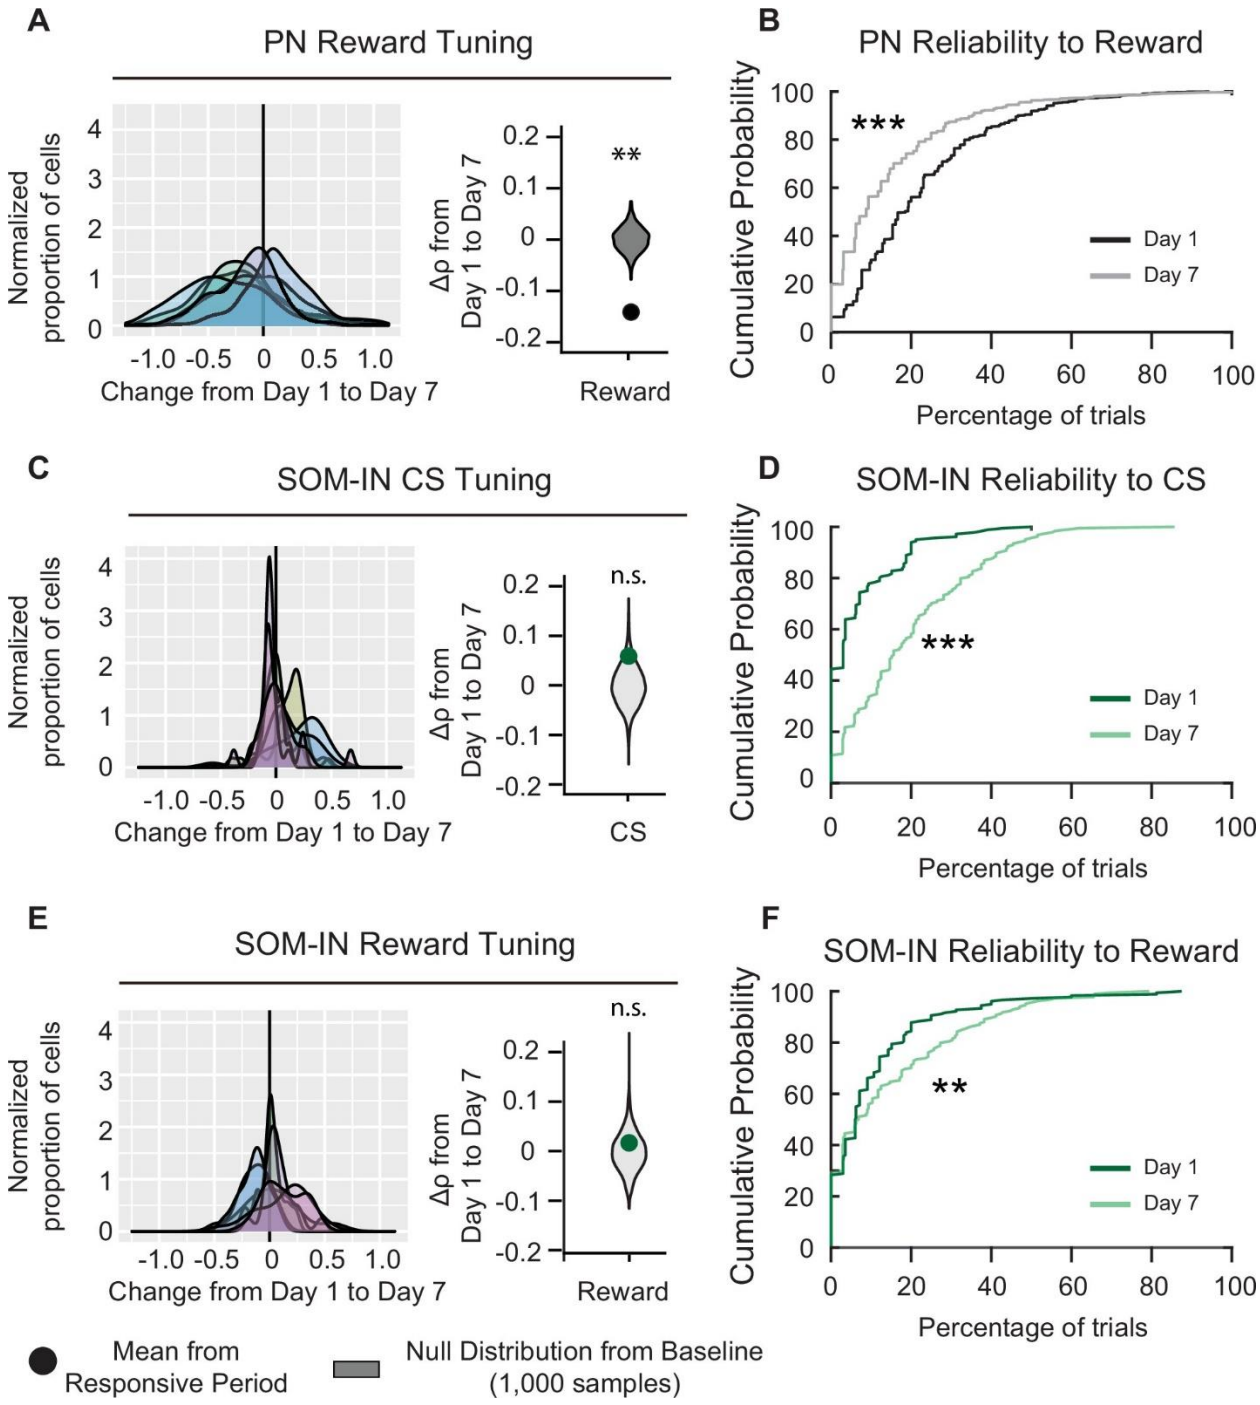

Figure 6. Pyramidal neuron (PN) and SOM-IN reliability is altered after associative learning. (A) Left, distribution of changes in PN Spearman correlation $\Delta\rho$ for reward. Each curve represents a Gaussian kernel density estimate of the distribution of $\Delta\rho$ in a single mouse. Right, mean change in Spearman correlation $\overline{\Delta\rho}$. Null distributions (gray) were estimated by resampling each mouse and shuffling trials 1000 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 days 1 and 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 cue stimulus (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., nonsignificant, p = 0.128. **(D)** Cumulative probability plots showing the percent of trials that each neuron responded to the CS for SOM-INs on days 1 and 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., nonsignificant, p = 0.598. (**F**) Cumulative probability plots showing the percent of trials that each neuron responded to the reward for SOM-INs on days 1 and 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 six mice. SOM-INs: *n* = 189 cells from seven mice.

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 days 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 days 1 and 7 $\Delta\rho_{reward} = 0.017 \pm 0.040$, Figure 6E). However, SOM-IN reliabilities also shifted to higher values on day 7 (Figure 6F). Altogether, these results suggest that in naive 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

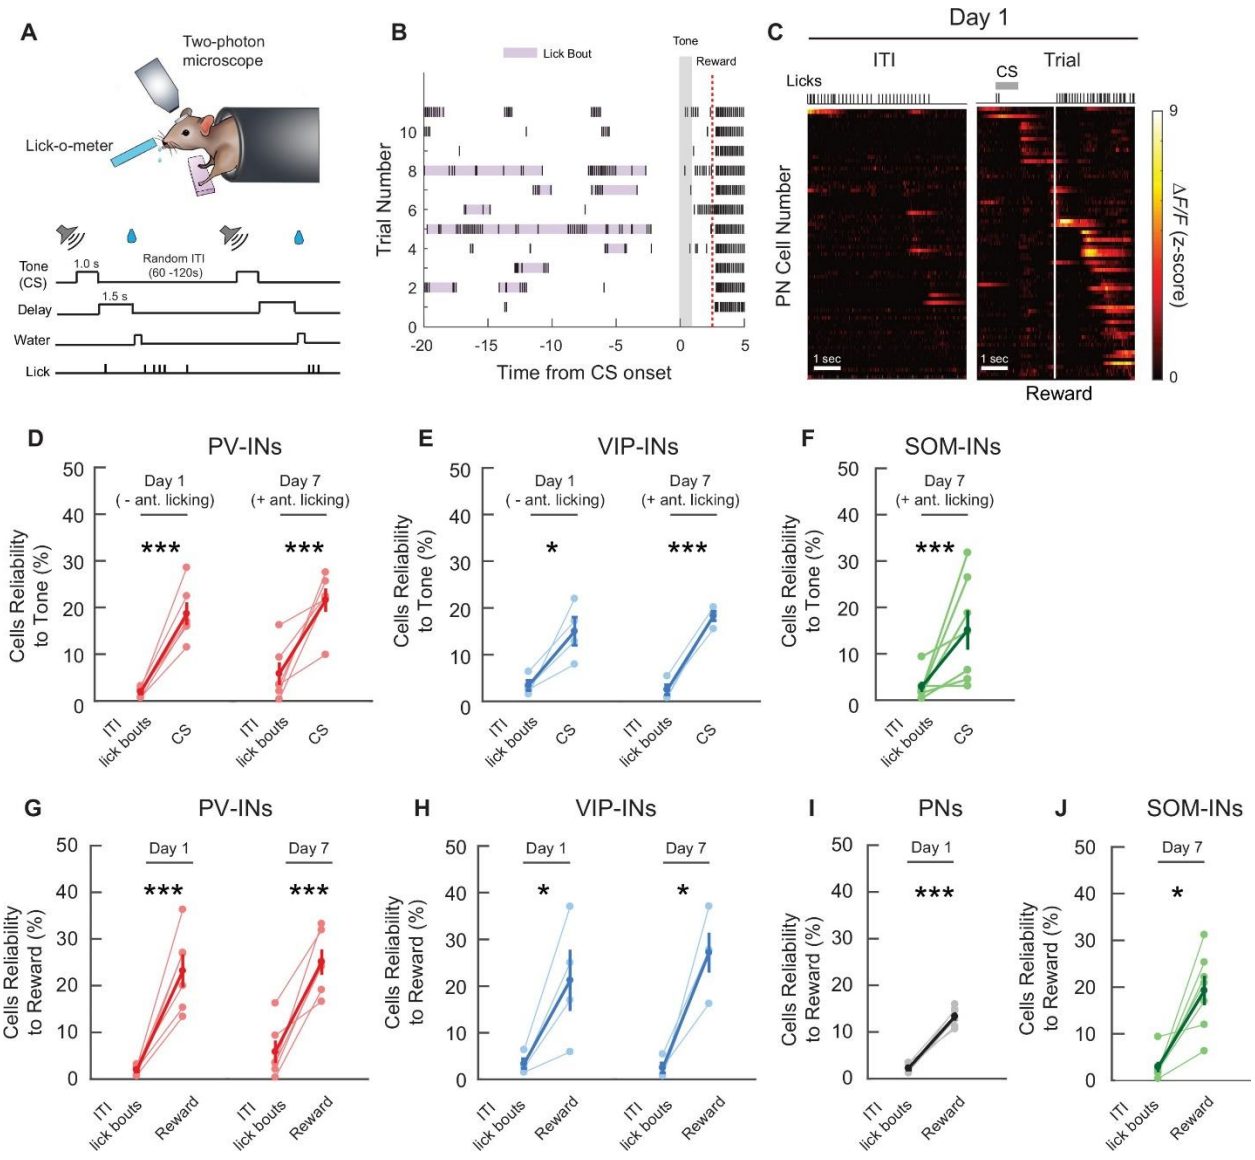

Figure 7. Cell-type-specific cue and reward activity are not due to licking movements. (A) Schematic of the cued reward conditioning task (top) and the trial structure (bottom). **(B)** Example licking behavior during the intertrial intervals (ITIs) from one mouse on day 1. Purple shading shows licks that were considered to be an individual lick bout. Gray shaded bar shows the cue stimulus (CS) timing and the red dotted line shows the reward timing. **(C)** Z-Scored activity of all the active pyramidal neurons (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. Gray 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, *$p < 0.05$, ***$p < 0.001$. 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 six mice. PV-IN: n = 316 cells from six mice. VIP-IN: n = 407 cells from four mice. SOM-IN: n = 189 cells from seven mice. Error bars show standard error of the mean (SEM).

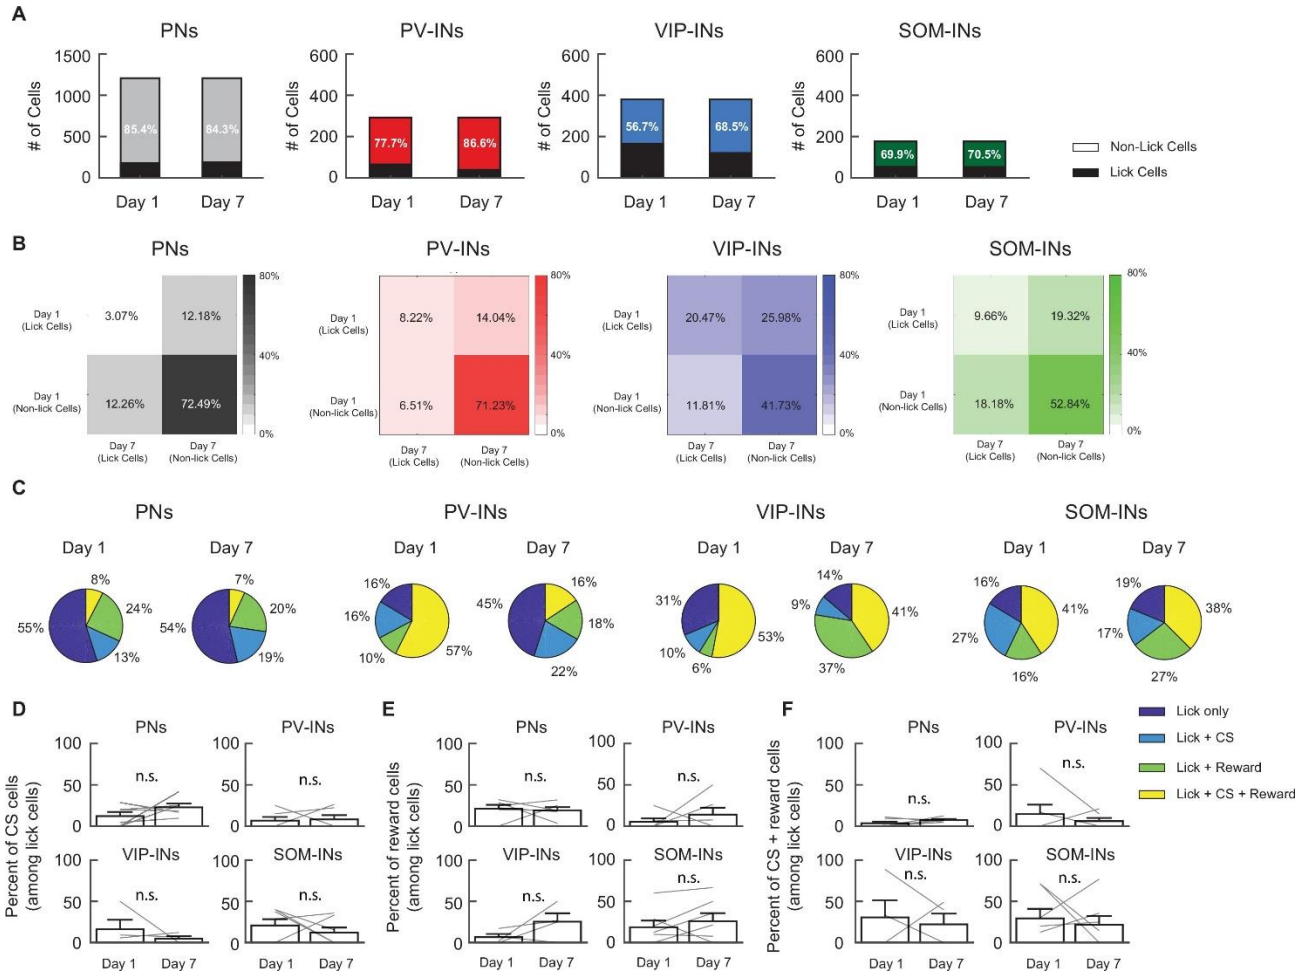

Figure 7-figure supplement 1. Lick cells were a small subset of active cells with stable mixed selectivity. (A) The fraction of lick cells (in black) among all active cells tracked from days 1 to 7 for each cell type. Percentages indicate the percentage of cells that were nonlick cells. For all cell types, the majority of cells were nonlick cells. **(B)** Contingency table showing lick selectivity among all tracked active cells on days 1 and 7. Individual cells were tracked across days to determine if licking-related activity changed across days. For all cell types, most cells were nonlick cells on days 1 and 7. **(C)** Proportion of mixed selectivity for cue stimulus (CS) and reward among lick cells for all cell types on days 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 six mice. PV-IN: n = 316 cells from six mice. VIP-IN: n = 407 cells from four mice SOM-IN: n = 189 cells from seven mice. Error bars show standard error of the mean (SEM).

each ITI on days 1 and 7. We observed that in each cell type, the majority of the neurons were nonlick neurons on both days 1 and 7 (Figure 7—figure supplement 1A). We then tracked individual lick and nonlick neurons and examined if they shifted their responses after associative learning. We found that most of the neurons maintained the same responses, and nonlick 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 days 1 and 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 days 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 behavioral 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 naive 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 and 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 naive stage (Levy et al., 2020). Hence, one possibility is that L2/3 PNs in M1 encode reward signals during the naive 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 and 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 7, while mice exposed to a nonpaired 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 behaviorally relevant cues such as those that predict reward. Finally, while VIP-INs remained responsive to reward in the nonpaired 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 and Rudy, 2011; Xu and 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 behavior while PV-IN suppression reduced freezing, indicating bidirectional control of the conditioned response (Wolff et al., 2014). Our results demonstrate that in a naive 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; Gasselin 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 behaviorally 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.

Materials and Methods

All animal experiments were approved by the University of Ottawa Animal Care Committee and in accordance with the Canadian Council on Animal Care guidelines. 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, Kim, Peters, & Komiyama, 2015; Harvey, Coen, & Tank, 2012; Komiyama et al., 2010; O'Connor et al., 2013; Peters, Chen, & Komiyama, 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 et al., 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

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

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 $d \in \{1,7\}$, $\mathbf{f}_{m,n,t-T_{\mathrm{baseline}}:t+T_{\mathrm{CS}}+\cdots}^{(d)}$ is the fluorescence trace of neuron n from mouse m during a single trial, $\mathbf{1}_s$ is an indicator function for stimulus $s \in$ Error! Bookmark not defined., $|\mathcal{T}_m^{(d)}|$ is the number of trials, and ρ is the Spearman correlation coefficient. Analysis was carried out with T_{baseline} =2 s and T_{post} =6 s. We considered the "CS" period indicated by $\mathbf{1}_{\mathrm{CS}}$ to range from the start of the CS at time t to the start of reward delivery at time $t + T_{\mathrm{CS}} + T_{\mathrm{delay}}$, and the "reward" period indicated by $\mathbf{1}_{\mathrm{reward}}$ 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

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

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 50^{th} 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 50^{th} 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.

Ca^{2+} event timing after the CS was calculated for each cell as the time from the cue onset to the time of the fifth frame above threshold within the CS response period (2.5s). Since reaction time and therefore, reward consumption time varies, Ca^{2+} event timing after the reward was calculated as the time from the first lick within the reward response period to the time of the fifth frame above threshold within 2.5s of the first lick. The mean latency for each cell was first calculated by taking the mean across all trials for a single cell, then the median of all cells within a mouse was calculated. Only cells that could be tracked between days were used.

To determine the proportion of PV-INs that respond to the CS on day 1 and day 7, we found the overall proportion of cells that responded to the CS 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 Δreliability 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.

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.

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.

Conflict of Interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest

Supplementary Information

Experimental Condition	Cell type	Total ROIs on Day 1	Total Active Cells on Day 1	Total Active Cells Tracked From Day 1 to 7
Paired CS-Reward	PN	211	211	211
		73	73	73
		299	299	299
		184	183	183
		256	255	255
		186	186	186
	PV-IN	102	102	101
		45	43	43
		40	40	40
		65	61	61
		47	34	32
		17	15	15
	VIP-IN	120	113	112
		118	118	117
		104	103	98
		65	64	54
	SOM-IN	21	20	20
		32	30	26
		22	22	22
		40	40	40
		18	12	12
		20	20	20
		36	36	36
No rewards	PV-IN	144	141	136
		109	107	107
		104	101	101
Non-paired rewards	PV-IN	52	51	51
		59	55	51
		57	40	37
		47	47	46
		109	99	91
	VIP-IN	76	74	58
		62	59	59
		70	69	69

Supplementary table 1. Summary.of the number of mice, total regions of interests (ROIs), total active cells, and total active cells tracked from days 1 to 7 in all experimental conditions.

CHAPTER 3 – Whole-brain mapping of long-range inputs to the VIP-expressing inhibitory neurons in the primary motor cortex

Candice Lee[1†], Sandrine L. Côté[1†], Nima Raman[1], Hritvic Chaudhary[1], Simon X Chen[1,2,3]

[1]Department of Cellular and Molecular Medicine, University of Ottawa, Ottawa, ON, Canada

[2]Brain and Mind Research Institute, University of Ottawa, Ottawa, ON, Canada

[3]Center for Neural Dynamics, University of Ottawa, Ottawa, ON, Canada

[†] These authors contributed equally

Author Contributions

C.L., S.L.C. and S.X.C. conceived the project. C.L. and S.L.C. conducted the experiments. C.L., S.L.C., N.R., and H.C. performed image registration and cell quantifications. C.L and S.L.C. analysed the data under the supervision of S.X.C. C.L., S.L.C., and S.X.C. wrote the manuscript.

Abstract

The primary motor cortex (MOp) is an important site for motor skill learning. Interestingly, neurons in MOp possess reward-related activity, presumably to facilitate reward-based motor learning. While pyramidal neurons (PNs) and different subtypes of GABAergic inhibitory interneurons (INs) in MOp all undergo cell-type specific plastic changes during motor learning, the vasoactive intestinal peptide-expressing interneurons (VIP-INs) in MOp have been shown to preferentially respond to reward and play a critical role in the early phases of motor learning by triggering local circuit plasticity. To understand how VIP-INs might integrate various streams of information, such as sensory, pre-motor, and reward-related inputs, to regulate local plasticity in MOp, we performed monosynaptic rabies tracing experiments and employed an automated cell counting pipeline to generate a comprehensive map of brain-wide inputs to VIP-INs in MOp. We then compared this input profile to the brain-wide inputs of somatostatin-expressing inhibitory interneurons (SST-INs), parvalbumin-expressing inhibitory interneurons (PV-INs) and PNs in MOp. We found that while all cell types received major inputs from sensory, motor, and prefrontal cortical regions, as well as from various thalamic nuclei, VIP-INs received more inputs from the orbital frontal cortex (ORB) - a region associated with reinforcement learning and value predictions. Our findings provide insight on how the brain leverages microcircuit motifs by both integrating and partitioning different streams of long-range input to modulate local circuit activity and plasticity.

Introduction

The primary motor cortex (MOp) has a well-established role in the execution of voluntary movement (Guo et al., 2015), and recent studies have also identified it as a critical site for motor learning (Xu et al., 2009; Peters et al., 2014; Chen et al., 2015b; Kawai et al., 2015). Like other cortical areas, MOp is primarily composed of glutamatergic pyramidal neurons (PNs) and different subtypes of GABAergic inhibitory interneurons (INs) with distinctive patterns of local connectivity. In particular, parvalbumin-expressing INs (PV-INs) primarily inhibit the perisomatic region of PNs, somatostatin-expressing INs (SST-INs) primarily inhibit the apical dendrites of PNs, while vasoactive intestinal peptide-expressing interneurons (VIP-INs) mainly inhibit SST-INs and thereby disinhibit PNs (Pfeffer et al., 2013; Pi et al., 2013; Tremblay et al., 2016). During motor learning, both PNs and INs in MOp undergo structural and functional plastic changes (Xu et al., 2009; Peters et al., 2014; Chen et al., 2015b; Ren et al., 2022). Interestingly, INs also exhibit subtype-specific modifications, in which PV-INs undergo a transient increase in axonal bouton density, while SST-IN axonal bouton density undergoes a persistent decrease. Optogenetically activating or inhibiting SST-INs, but not PV-INs, disrupts PN plasticity and impairs motor learning (Chen et al., 2015b; Cichon and Gan, 2015; Adler et al., 2019; Yang et al., 2022). Furthermore, Ren et al. (2022) recently discovered that during the early phase of motor learning, the VIP-IN population is highly active and its activity decreases with learning, while SST-INs conversely become more active with learning. When VIP-INs in MOp were chemogenetically inactivated during the early stage of learning, SST-IN activity increased and motor learning was impaired, demonstrating that VIP-IN mediated inhibition of SST-INs during the early phase of learning is critical for the acquisition of new motor skills. Altogether, these subtype-specific plastic changes allude to unique functions for different IN subtypes during motor learning.

Interestingly, in addition to motor-related activity in MOp, studies in macaque monkeys also showed widespread representations of reward anticipation and reward-related activity among cortical neurons in MOp (Marsh et al., 2015; Ramkumar et al., 2016; Ramakrishnan et al., 2017). Moreover, a recent study that imaged the activity of different neuronal subtypes in MOp also revealed varying degrees of responsiveness to reward and to reward-predicting cues when animals were engaged in a classical conditioning task. Notably, VIP-INs in MOp showed

preferential responses to reward compared to PNs, PV-INs and SST-INs, and after associative learning, VIP-INs became more reliably responsive to reward (Lee et al., 2022). This is consistent with other findings throughout the cortex (Pi et al., 2013; Ren et al., 2022; Szadai et al., 2022), the amygdala (Krabbe et al., 2019), and the hippocampus (Turi et al., 2019), in which VIP-INs respond robustly to reinforcement signals, and in turn, disinhibit local microcircuitry and increase local neural gain (Fu et al., 2012; Jackson et al., 2016). Thus, VIP-INs could serve as a key regulator of cortical activity, and VIP-IN activation following positive reinforcement may gate the window of plasticity during reward-based motor skill learning.

Although the plastic changes attributed to VIP-INs in MOp underscore the importance of its function during motor learning, it is unclear where these learning-related inputs to VIP-INs arise from. While MOp is a critical site for motor learning, several other brain regions known to provide input to MOp also undergo plastic changes during motor learning. These include cortical regions such as the secondary motor cortex (MOs; Cao et al., 2015), anterior lateral motor area (ALM; Chabrol et al., 2019), restrosplenial cortex (RSP; Makino et al., 2017), and subcortical regions such as thalamus (Biane et al., 2016; Tanaka et al., 2018). Moreover, many input regions have also been shown to undergo plastic changes after reward-based associative learning such as the primary somatosensory cortex (Chen et al., 2015a), auditory cortex (Kisley and Gerstein, 2001; Lee and Rothschild, 2021), RSP (Hattori et al., 2019), ALM and MOs (Komiyama et al., 2010). However, whether these brain regions provide preferential input to a specific neuron subtype in MOp or if they provide reinforcement-related signals to VIP-INs remains unexplored.

Here, we utilized a monosynaptic rabies tracing strategy and performed brain-wide mapping of long-range inputs to the four major cell types in MOp (VIP-INs, PV-INs, SST-INs, and PNs). By systematically comparing the proportion of inputs from different brain regions to VIP-INs with the inputs to PV-INs and SST-INs, we found that VIP-INs received significantly more inputs from the orbital frontal cortex (ORB). Considering that ORB has been shown to encode value and reward (Baltz et al., 2018; Namboodiri et al., 2019; Zhou et al., 2019; Wang et al., 2020), our results point toward ORB serving as an important node in a reward-related input stream projecting to VIP-INs in MOp. In contrast, SST-INs received more input from the RSP, demonstrating that different IN subtypes receive preferential long-range input from specific brain regions. Taken together, our comprehensive whole-brain mapping uncovers input from ORB that

could be responsible for activating VIP-INs in MOp in response to reward and may thereby gate local circuit plasticity during reward-based motor learning.

Results

To identify long-range input regions that are specific to VIP-INs, PV-INs, and SST-INs in the caudal forelimb area of MOp, we utilized the monosynaptic rabies virus (RV) tracing system (Wickersham et al., 2007; Callaway and Luo, 2015; Wall et al., 2016). Helper virus (AAV1-EF1a-DIO-TVA950-T2A-CVS11G or AAV2/DJ-hSyn-FLEX-TVA-P2A-eGFP-2A-oG) was first injected into the right MOp forelimb area of *VIP-Cre*, *PV-Cre*, or *SST-Cre* mice to express avian TVA receptors, rabies glycoprotein (G) and GFP in each cell type, respectively. After 3 weeks, we injected the pseudotyped G-deleted rabies virus (EnvA-RVdG-mCherry) into the same site (**Figure 1A)**. Animals were perfused 1 week later, and coronal sections were imaged at a 120 μm increment across the entire brain (see Methods). To automatically and unbiasedly quantify RV-labelled cells throughout the brain, we employed the software 'Wholebrain' (Fürth et al., 2018), which enables automated detection and quantification of labeled neurons. Most importantly, Wholebrain enables scale-invariant registration of brain sections to the Allen Mouse Brain Atlas (**Figure 1B - C**), thus providing a method to identify all the brain regions from different brain samples within a standardized framework. By applying the Wholebrain software to sections from 2.945mm anterior to bregma through to 5.055mm posterior to bregma (cells were not found outside of these coordinates), we generated comprehensive and comparative maps of whole-brain input to VIP-INs, PV-INs, and SST-INs (**Figure 1D - F**). We first quantified the number of input cells ($mCherry^+$ cells outside of MOp) within each region identified in the Allen Brain Atlas, including subdivisions and cortical layers. We found that the total number of labelled input cells varied between each IN subtype (VIP-INs: 3887 ± 2021 cells, PV-INs: 1593 ± 234 cells, SST-INs: 1253 ± 344 cells). Since the total number of input cells could be dependent on the number of starter cells (colocalized GFP^+ and $mCherry^+$ cells within MOp) that were labelled from the viral injection, we also calculated an approximate ratio of starter cells to the total number of input cells in the entire brain. We found that the ratio of starter-to-input cells were very similar between the IN subtypes (VIP-INs: 1:15, PV-INs: 1:19, SST-INs: 1:18).

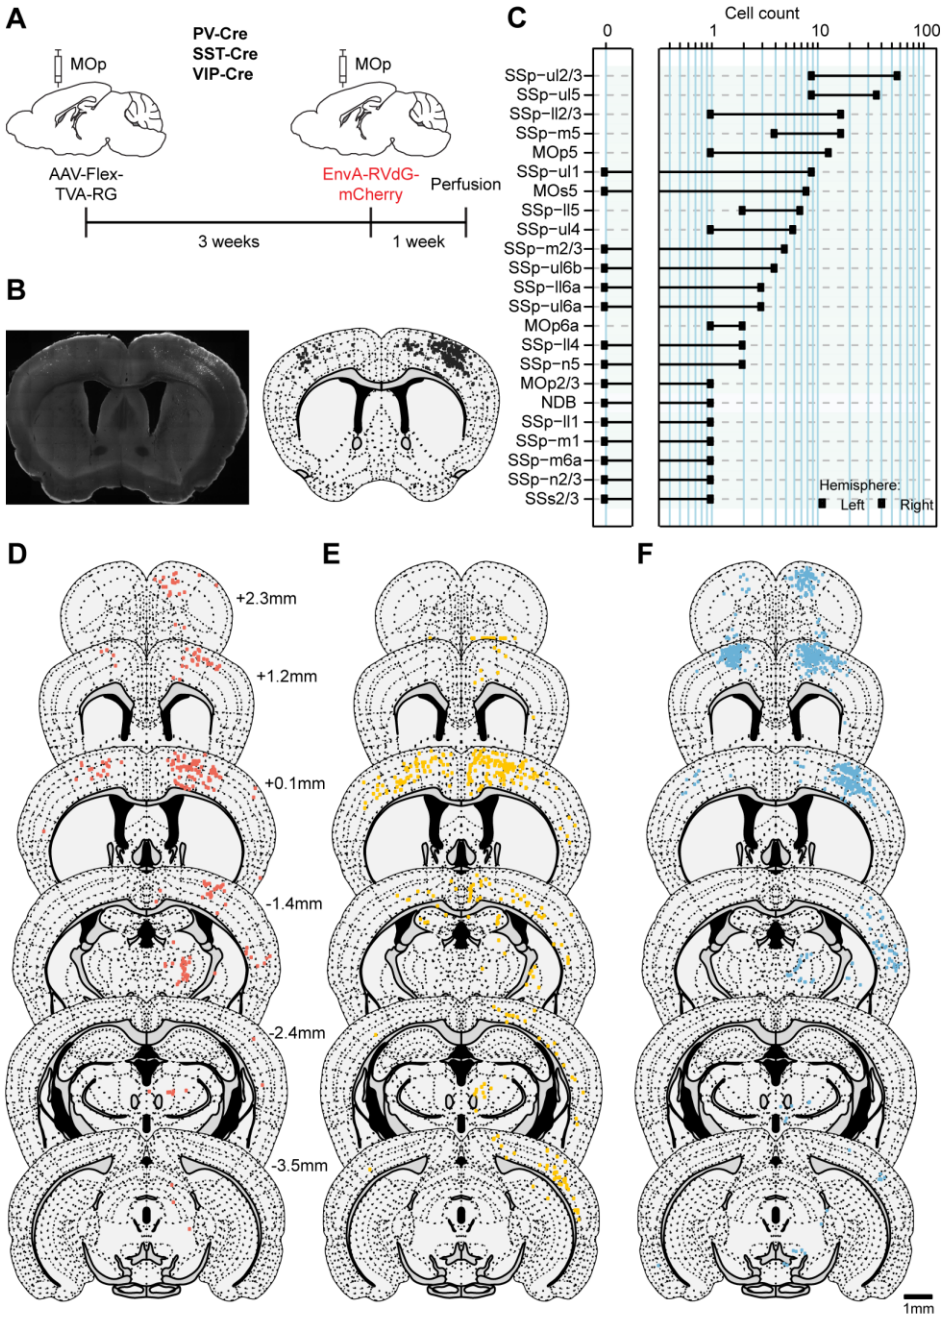
A
PV-Cre
SST-Cre
VIP-Cre
MOp
MOp
AAV-Flex-
TVA-RG
EnvA-RVdG-
mCherry
Perfusion
3 weeks
1 week
B
C
Cell count
0
1
10
100
SSp-ul2/3
SSp-ul5
SSp-ll2/3
SSp-m5
MOp5
SSp-ul1
MOs5
SSp-ll5
SSp-ul4
SSp-m2/3
SSp-ul6b
SSp-ll6a
SSp-ul6a
MOp6a
SSp-ll4
SSp-n5
MOp2/3
NDB
SSp-ll1
SSp-m1
SSp-m6a
SSp-n2/3
SSs2/3
Hemisphere:
Left
Right
D
E
F
+2.3mm
+1.2mm
+0.1mm
-1.4mm
-2.4mm
-3.5mm
1mm
G % total input cells – PV
Cortex ipsi
Cortex contra
Thalamus ipsi
Others
4%
6%
13%
77%
H % total input cells – SST
Cortex ipsi
Cortex contra
Thalamus ipsi
Others
8%
4%
23%
65%
I % total input cells – VIP
Cortex ipsi
Cortex contra
Thalamus ipsi
Others
3%
8%
24%
65%

Figure 1. Generating brain-wide input maps to VIP-INs, PV-INs, and SST-INs in MOp. (A) Helper virus was injected unilaterally into right MOp followed by pseudotyped G-deleted rabies virus 3 weeks later. Animals were sacrificed 1 week after the rabies viral injection. **(B)** Example image of a coronal section (left). Registration of the same section to the Allen Brain Atlas with the Wholebrain software (right). Each black dot is a mCherry^{+} input cell that was automatically detected by the software. **(C)** Cell quantification within different regions, subdivisions, and layers for the example section in **(B)**. **(D-F)** Representative example sections showing detected input cells projecting to PV-INs **(D)**, SST-INs **(E)** and VIP-INs **(F)**. **(G-I)** Mean percentage of brain-wide inputs to PV-INs **(G)**, SST-INs **(H)** and VIP-INs **(I)** from broad subdivisions of the brain. Regions with small proportions were grouped together in Others. PV-INs: n = 4 mice, SST-INs: n = 4 mice, VIP-INs: n = 5 mice.

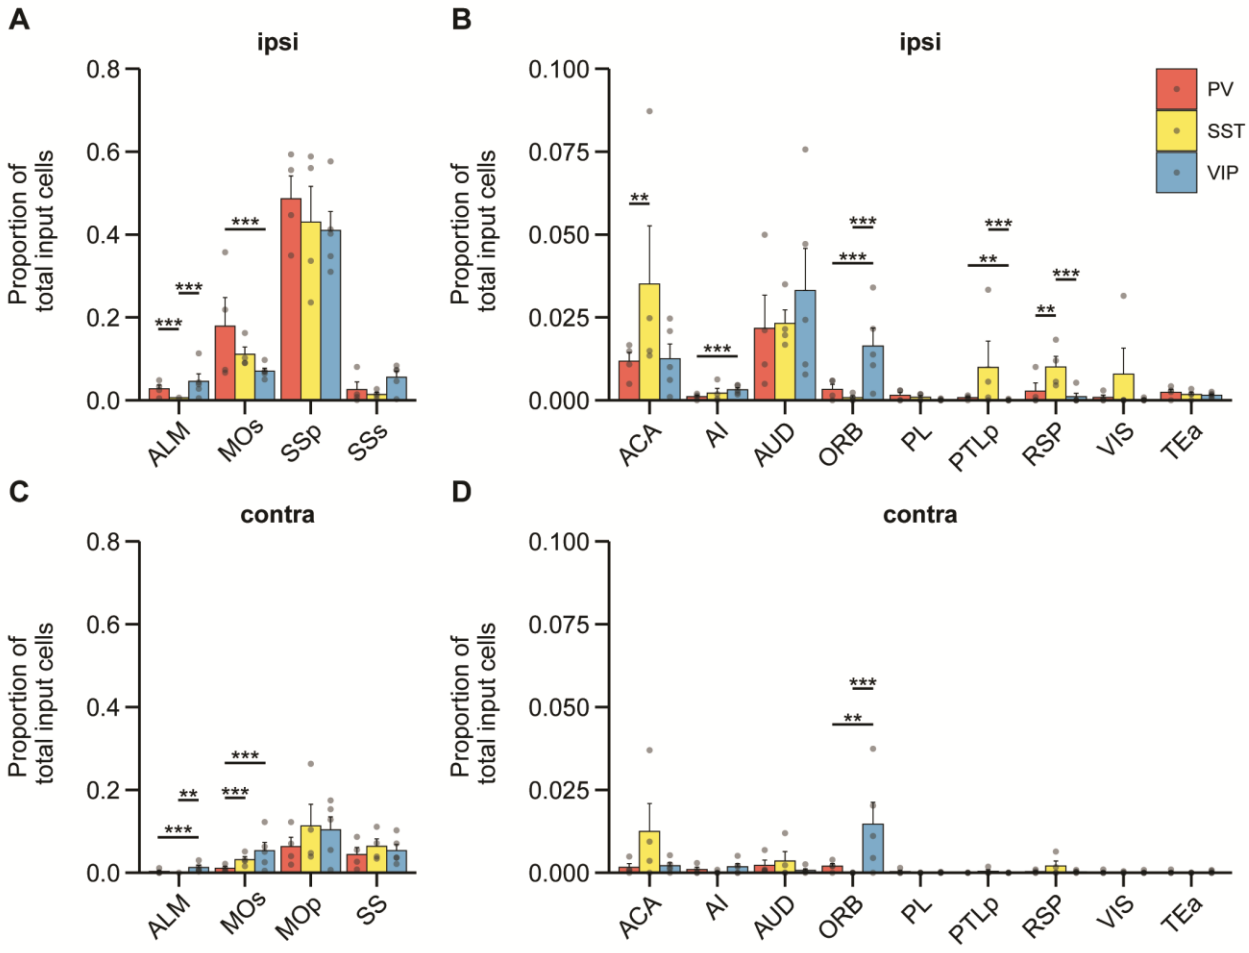

Figure 2. Cortical input to different IN subtypes in MOp. (A) Mean proportion of input cells in ipsilateral somatosensory and motor cortices. **(B)** Mean proportion of input cells in other ipsilateral cortical regions. Regions with negligible cell counts are not shown. **(C)** Mean proportion of input cells in contralateral somatosensory and motor cortices. **(D)** Mean proportion of input cells in other contralateral cortical regions. Regions with negligible cell counts are not

shown. PV-INs: n = 4 mice, SST-INs: n = 4 mice, VIP-INs: n = 5 mice. Bootstrap with Bonferroni correction for multiple comparisons. Each point represents a mouse. Error bars show the SEM. **p < 0.01, ***p < 0.001.

Long-range inputs to different IN subtypes in MOp from the cortex

We next examined which brain regions each IN subtype receives its input from, beginning with broad subdivisions within the brain. We observed that for all IN subtypes, the majority of input originated in the cortex (**Figure 1G - I**); in particular, the greatest source of cortical input came from the sensorimotor regions (**Figure 2A, C**). It is known that ascending tactile sensory information propagates sequentially from the primary somatosensory cortex (SSp) to the secondary somatosensory cortex (SSs). While both regions encode stimulus features, SSp encodes the stimulus more strongly and SSs encodes higher order information such as stimulus-related recall (Condylis et al., 2020) and decision-related activity (Kwon et al., 2016), which is then conveyed to SSp. In our results, we found that on the ipsilateral side, the SSp was the largest source of input for all IN subtypes in MOp and comprised 40.99% ± 2.04, 48.64% ± 2.76%, and 40.40% ± 4.29% of input to VIP-INs, PV-INs, and SST-INs, respectively (**Figure 2A**). In contrast, all IN subtypes received substantially less input from SSs (SSp compared to SSs: $p < 1 \times 10^{-3}$ for all cell types). We also identified major cortical input from the secondary motor cortex (MOs) and ALM (also known as frontal MOs). MOs and ALM are two motor regions that show preparatory activity preceding movement and are thought to be akin to the primate premotor cortex (Guo et al., 2015; Li et al., 2015; Chen et al., 2017). We found that MOs provided inputs to all IN subtypes in MOp, and PV-INs received the most among them all. Interestingly, we observed that MOs provided more inputs to the INs in MOp than ALM. Specifically, on the ipsilateral side, PV-INs and SST-INs but not VIP-INs received greater input from MOs compared to ALM (VIP-INs: p = 0.084; PV-INs and SST-INs: $p < 1 \times 10^{-3}$; **Figure 2A**). In contrast, while input from ALM was relatively low compared to MOs, it projected substantially more to PV-INs and VIP-INs than to SST-INs (VIP-IN vs. SST-INs: $p < 1 \times 10^{-3}$; PV-INs vs. SST-INs: $p < 1 \times 10^{-3}$; **Figure 2A**). Hence, unlike SSp and SSs which had similar proportions of input to all IN subtypes in MOp, MOs and ALM demonstrated some subtype-specific differences. Since MOs and ALM have been shown to have differential roles in movement preparation and execution

(Siniscalchi et al., 2016; Chen et al., 2017), our results suggest that specific IN subtypes in MOp may be involved in processing different movement-related information.

Within ipsilateral cortex, we also found other cortical regions that provide differential input biased toward specific IN subtypes in MOp, although they represented a smaller proportion of brain-wide inputs compared to SSp, SSs, MOs and ALM. The anterior cingulate area (ACA) projected to all IN subtypes but significantly more to SST-INs than to PV-INs, while no significant difference was found between SST-INs and VIP-INs (**Figure 2B**). The auditory cortex (AUD) was a major source of input to all three IN subtypes with no significant differences in the proportion of input cells. Intriguingly, the orbital frontal cortex (ORB) projected significantly more to VIP-INs compared to either PV-INs or SST-INs, and ORB was one of the most pronounced inputs to VIP-INs outside of the SSp (**Figure 2B**). Lastly, SST-INs received significantly more input from RSP than PV-INs and VIP-INs. In addition, SST-INs also received significantly more inputs from the posterior parietal cortex (PTLp) compared to VIP-INs and from the anterior cingulate area (ACA) compared to PV-INs (**Figure 2B**). Overall, among ipsilateral cortex, MOs and SSp account for a substantial proportion of brain-wide input to INs in MOp. In addition, we identified biased input from ORB to VIP-INs and RSP to SST-INs.

In the contralateral cortex, MOp was a major source of input to all three IN subtypes (**Figure 2C**). ALM, MOs and SS (SSp and SSs combined) had substantially less input cells compared to the ipsilateral side. Unlike ipsilateral ALM which projected similarly to both VIP-INs and PV-INs, contralateral ALM projected more to VIP-INs than to the other two subtypes. Intriguingly, contralateral MOs projected more to VIP-INs and SST-INs than to PV-INs; this is in contrast to ipsilateral MOs, which had a greater proportion of input to PV-INs and less input to SST-INs and VIP-INs. The relative proportion of input from MOs and ALM also differed; unlike the ipsilateral side, where PV-INs and SST-INs but not VIP-INs received more input from MOs compared to ALM, all IN subtypes received more input from contralateral MOs compared to ALM (VIP-INs: $p = 0.006$; SST-INs and PV-INs: $p < 1 \times 10^{-3}$). These findings demonstrate that subtype-specificity of long-range input from the same region can vary between hemispheres. Outside of ALM, MOs and SS, the proportion of input from other regions within the contralateral cortex were mostly minimal with some exceptions (**Figure 2D**). Noticeably, the contralateral ACA trended towards providing more input to SST-INs compared to both PV-INs and VIP-INs,

akin to its ipsilateral homologue. In addition, contralateral ORB was another exception, as it had a proportion of input cells comparable to its ipsilateral counterpart and also projected significantly more to VIP-INs compared to PV-INs and SST-INs (**Figure 2D**). Hence, the numerous bilateral inputs from ORB further highlight the importance of this projection to VIP-INs.

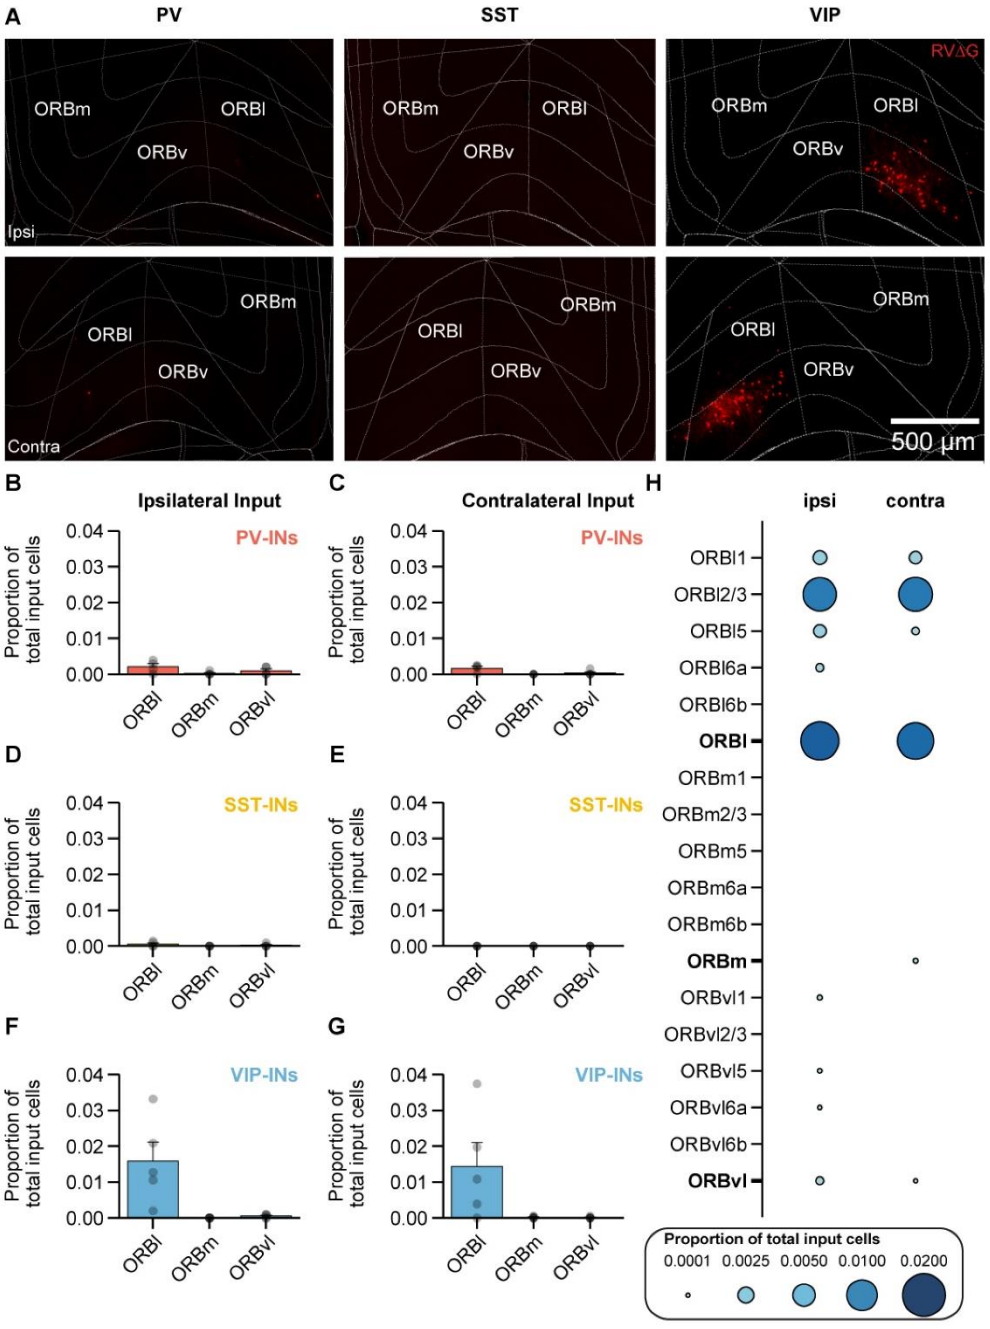

Figure 3. ORBl preferentially projects to VIP-INs. (A) Example images of $mCherry^+$ input cells in ORB that project to PV-INs (left), SST-INs (middle) and VIP-INs (right) in MOp. Top row, ipsilateral ORB. Bottom row, contralateral ORB. **(B-G)** Mean proportion of brain-wide inputs from ipsilateral and contralateral ORB to PV-INs **(B-C)**, SST-INs **(D-E)**, and VIP-INs **(F-G). (H)** Mean proportion of brain-wide inputs to VIP-INs found within different layers of ORBl, ORBm and ORBvl. Bolded labels show the sum across all layers for each subdivision. PV-INs: n = 4 mice, SST-INs: n = 4 mice, VIP-INs: n = 5 mice. Each point represents a mouse. Error bars show the SEM.

Distinct subregion- and layer- specific inputs to different IN subtypes in MOp

ORB has been shown to encode and predict value and reward (Baltz et al., 2018; Namboodiri et al., 2019; Zhou et al., 2019; Wang et al., 2020), and previous work has also demonstrated that VIP-INs in MOp undergo plastic changes in response to reward after associative learning (Lee et al., 2022). To better understand this projection, we sought to further examine the location of the input cells within the ORB (**Figure 3A**). ORB is comprised of three main subdivisions – the lateral, medial and ventrolateral regions (ORBl, ORBm and ORBvl). While the function of individual subregions is still unclear, in monkeys, ORBl has been implicated in reward-guided learning, and ORBm has been implicated in reward-guided decision making (Noonan et al., 2010). In rats, ORBvl has been shown to be involved in goal-directed behavior following contingency switches (Parkes et al., 2018; Zimmermann et al., 2018). In our results, we found that on the ipsilateral side, VIP-INs received almost all of its input from ORBl with no detectable input from ORBm and minimal input from ORBvl. The proportion of input cells from ORBl to VIP-INs was significantly higher than to PV-INs and SST-INs (VIP-IN vs PV-IN: $p < 0.002$; VIP-IN vs SST-INs: $p < 1 \times 10^{-3}$; **Figure 3B, D, F**). Similar observations were made on the contralateral side; VIP-INs also received more input from cells in the contralateral ORBl and barely any input from ORBm and ORBvl. The proportion of input cells from the contralateral ORBl to VIP-INs was also significantly higher than to PV-INs and SST-INs (VIP-IN vs PV-IN: $p = 0.007$; VIP-IN vs SST-INs: $p < 1 \times 10^{-3}$; **Figure 3C, E, G**). These results indicate that VIP-INs in MOp receive a considerable amount of input from both ipsilateral and contralateral ORBl but not ORBm and ORBvl, consistent with the hypothesis that VIP-INs may be involved in reward-guided motor learning (**Figure 3F-G**). Since VIP-INs received the most input from ORB, we next asked which layers of ORB project to VIP-INs. We found that on both the ipsilateral and contralateral sides, most of the VIP-IN projecting cells were located in L2/3 of ORBl (**Figure 3H**). On the ipsilateral side, VIP-INs also received some input from L1, L5 and L6a of ORBl,

whereas on the contralateral side, VIP-INs received similar proportions of input from L1 compared to the ipsilateral side but not as much from L5 and L6a (**Figure 3H**).

In addition to ORB, RSP was another region that was identified to be unique as it provided more input to SST-INs compared to both PV-INs and VIP-INs (on the ipsilateral side, SST-INs vs VIP-INs: $p < 1 \times 10^{-3}$; SST-INs vs PV-INs: $p = 0.015$; **Figure 2B**). RSP is a

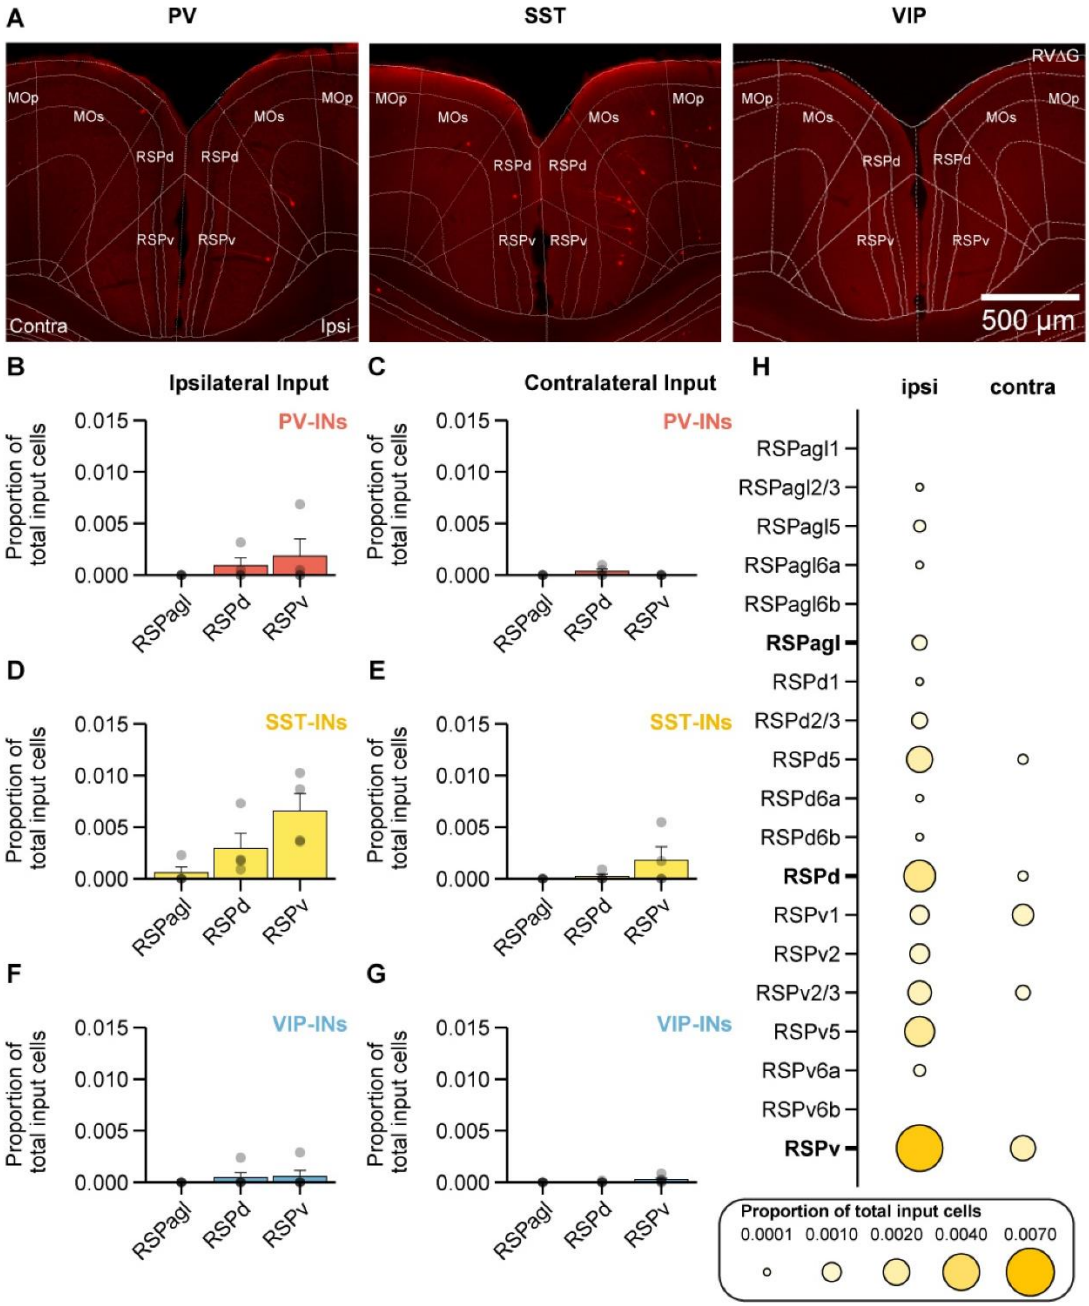

Figure 4. RSPd and RSPv preferentially project to SST-INs in MOp. (A) Example images of $mCherry^+$ input cells in RSP that project to PV-INs (left), SST-INs (middle) and VIP-INs (right) in MOp. **(B-G)** Mean proportion of brain-wide input from ipsilateral and contralateral RSP to PV-INs **(B-C)**, SST-INs **(D-E)**, and VIP-INs **(F-G)**. **(H)** Mean proportion of brain-wide inputs to SST-INs found within different layers of RSPagl, RSPd and RSPv. Bolded labels show the sum across all layers for each subdivision. PV-INs: n = 4 mice, SST-INs: n = 4 mice, VIP-INs: n = 5 mice. Each point represents a mouse. Each point represents a mouse. Error bars show the SEM.

complex brain region that has been implicated in spatial navigation (Vann et al., 2009), associative learning (Lukoyanov and Lukoyanova, 2006; Makino and Komiyama, 2015; Makino et al., 2017; Hattori et al., 2019; Hattori and Komiyama, 2022), and motor learning (Makino et al., 2017). RSP can be further subdivided into the lateral agranular, dorsal, and ventral parts (RSPagl, RSPd, RSPv; **Figure 4A**). Therefore, we also sought to further examine the location of the input cells to SST-INs within the RSP. On the ipsilateral side, SST-INs received the most input from RSPv, followed by RSPd and small amounts of input from RSPagl (**Figure 4D**). In contrast, PV-INs and VIP-INs received no input from RSPagl and minimal inputs from RSPd and RSPv (**Figure 4B, F**). On the contralateral side, SST-INs mainly received input from RSPv and almost no input from RSPd and RSPagl (**Figure 4E)**. Given the biased input to SST-INs from RSP, we then sought to determine the laminar distribution of RSP projection neurons to SST-INs. We observed that in RSPv, input cells were found in all layers on the ipsilateral side except in L6b. Interestingly, on the contralateral side, input neurons were only found in L1 and L2/3.

Long-range inputs to different IN subtypes in MOp from the thalamus

While the majority of brain-wide input to all three IN subtypes originated from cortex, ipsilateral thalamus was the next largest source of input to VIP-INs, PV-INs and SST-INs and provided 8%, 6%, and 4% of the total inputs to VIP-INs, PV-INs, and SST-INs, respectively (**Figure 1G-I**). In contrast, we did not find many input cells in the contralateral thalamus. From all the nuclei in the thalamus, we observed that a greater proportion of inputs arose from the mediodorsal (MD), parafasicular (PF), posterior (PO), ventral anterior lateral (VAL) and ventral posterior (VP) nuclei of thalamus, while relatively small amounts of inputs arose from the central lateral (CL), central medial (CM), lateral dorsal (LD), lateral habenula (LH), paracentral nucleus (PCN), reticular nucleus (RT) and ventromedial (VM) nuclei of thalamus (**Figure 5A**). Intriguingly, we also observed some subtype specificity in the proportion of input. For example, CL projected

more to VIP-INs than to SST-INs, and CM provided modest but significantly more input to VIP-INs than to either PV-INs or SST-INs. MD projected significantly more to VIP-INs compared to SST-INs but not to PV-INs, and PCN also provided modest input that was biased to VIP-INs. PO and VM projected significantly more to VIP-INs and PV-INs compared to SST-INs. Lastly, LD, LH, PF, RT, VAL and VP did not show any subtype specific biases.

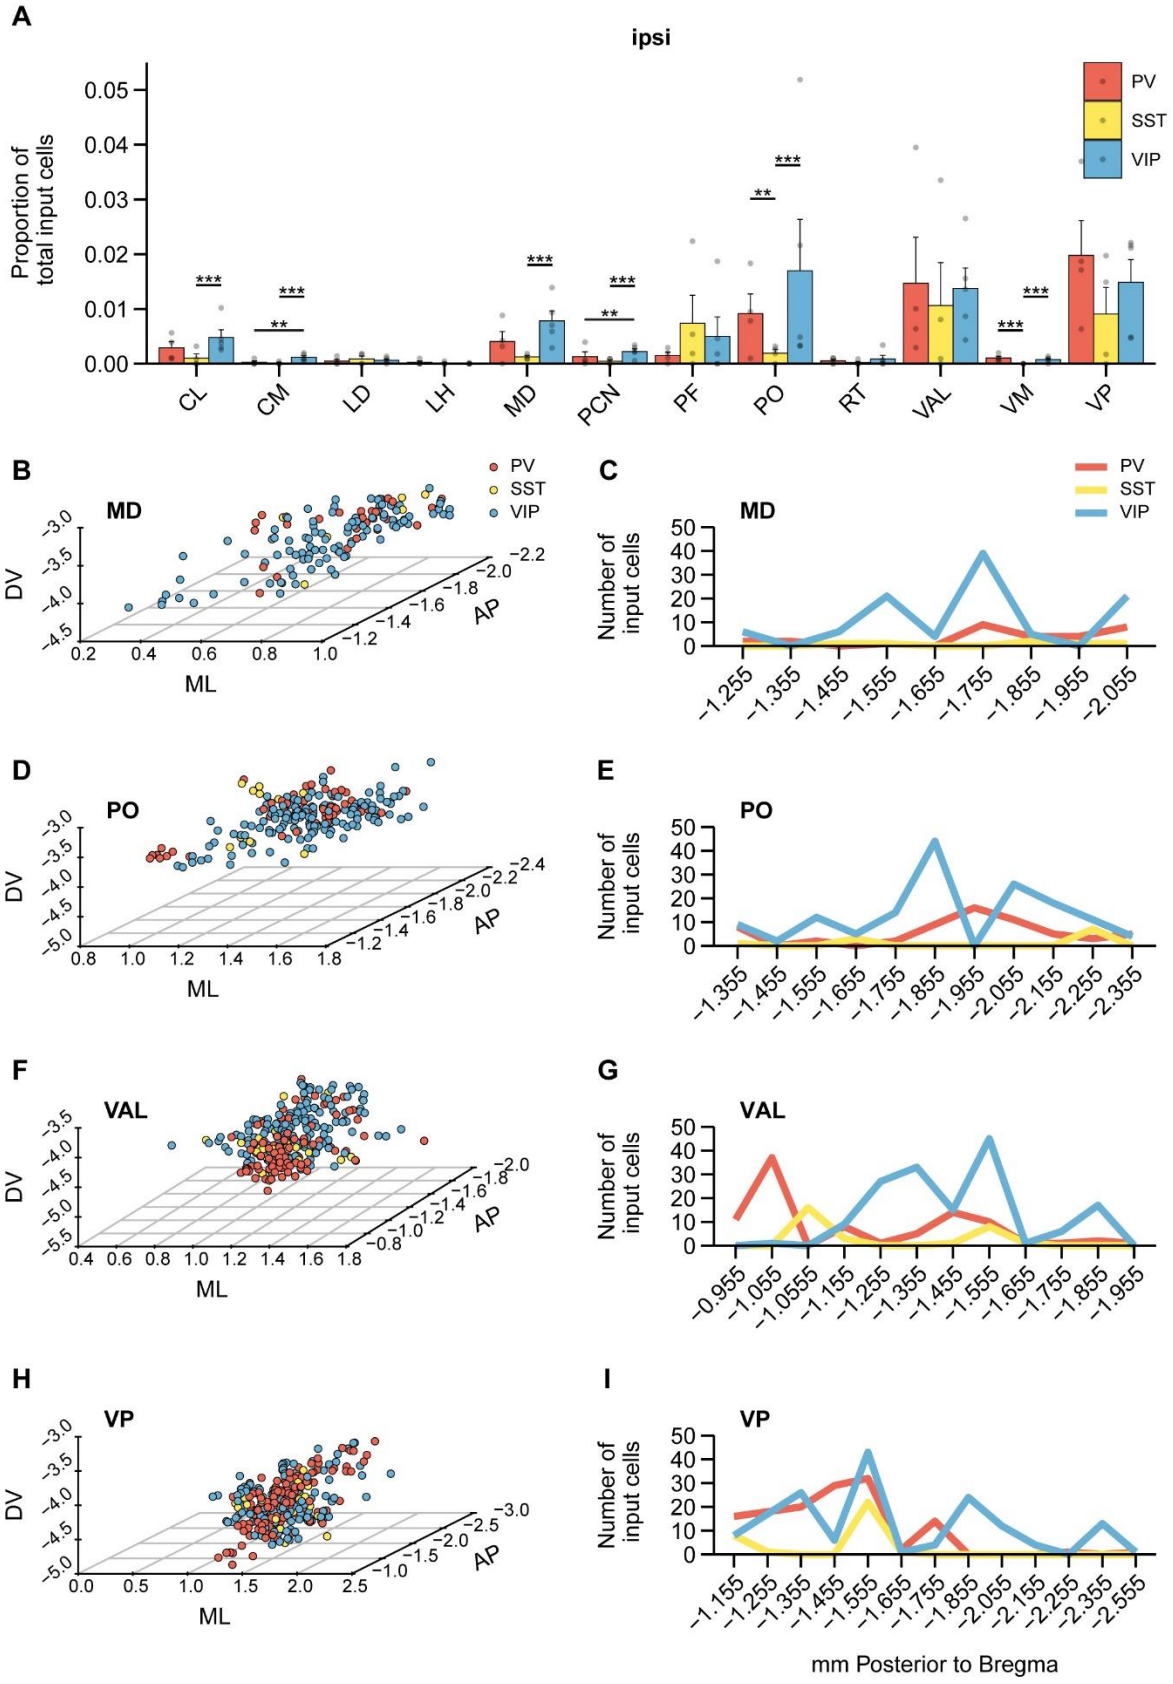
A
ipsi
Proportion of total input cells
PV
SST
VIP
CL
CM
LD
LH
MD
PCN
PF
PO
RT
VAL
VM
VP
B
MD
DV
ML
AP
C
MD
Number of input cells
D
PO
E
PO
F
VAL
G
VAL
H
VP
I
VP
mm Posterior to Bregma

Figure 5. Thalamic input to IN subtypes in MOp. (A) Mean proportion of brain-wide input found in different thalamic nuclei. Nuclei with a negligible number of cells are not shown. Bootstrap with Bonferroni correction for multiple comparisons. Each point represents a mouse**. (B)** 3D spatial distribution of input cells within MD. Axes correspond to the medial-lateral (ML), dorsal-ventral (DV), and anterior-posterior (AP) coordinates relative to bregma. Each point represents a cell. Colour indicates the starter cell type in MOp. **(C)** Number of cells within MD along the anterior-posterior axis for different IN subtypes. **(D)** 3D spatial distribution of input cells within PO. **(E)** Number of cells within PO along the anterior posterior axis. **(F)** 3D spatial distribution of input cells within VAL. **(G)** Number of cells within VAL along the anterior-posterior axis. **(H)** 3D spatial distribution of input cells within VP. **(I)** Number of cells within VP along the anterior-posterior axis. PV-INs: n = 4 mice, SST-INs: n = 4 mice, VIP-INs: n = 5 mice. Error bars show the SEM. **p < 0.01, ***p < 0.001.

It has been shown that various nuclei within thalamus are topographically organized (Angaut et al., 1985; Pierret et al., 2000; Veinante et al., 2000; Lam and Sherman, 2015); hence, we mapped the input neurons within each of the major input nuclei in 3-dimensional space to assess whether there is spatial specificity or clustering among the input neurons. Previous work has shown that MD neurons projecting to PV-INs and VIP-INs in the prefrontal cortex occupy distinct locations along the medial-lateral axis (Mukherjee et al., 2021). Here, we found that MD neurons projecting to PV-INs and VIP-INs in MOp were intermingled along the medial-lateral axis (**Figure 5B**). However, MD neurons projecting to VIP-INs appeared to be more widely distributed along the anterior-posterior axis while neurons projecting to PV-INs were located more towards the posterior end of the nucleus (**Figure 5B, C**). Previous work has shown that first-order POm (the medial division of PO) receives direct input from the brainstem, and is located more anteriorly (centered around ~1.7mm posterior to bregma). Higher-order POm, on the other hand, does not receive input from the brainstem and is located more posteriorly (centered on 2.2mm posterior to bregma; El-Boustani et al., 2020). We found that input neurons to VIP-INs were located within both of these subdivisions, while neurons providing input to PV-INs lay in between these two subdivisions (**Figure 5D, E**). The ventral group of thalamic nuclei including VAL, VM and VP show primarily sensory- and motor-related activity. In addition, within the ventral-lateral subdivision, the anterior regions strongly represent whisking-related activity, and more posterior regions are more limb-related (Tlamsa and Brumberg, 2010). We found that within thalamus, VAL was among the most numerous sources of input to all three IN subtypes, and a cluster of PV-IN projecting neurons was found in the most anterior parts of VAL.

In contrast, VIP-IN projecting neurons were found to be more distributed throughout the middle and most posterior regions of VAL **(Figure 5F, G)**. Noticeably, SST-INs received only a small proportion of input from MD, PO, and VAL (**Figure 5A**). The posterior division of the ventral group (VP) relays sensory-related input to the somatosensory cortex (Jensen and Killackey, 1987; Lee and Sherman, 2008). VP neurons projecting to PV-INs were located in the most anterior parts through to the center of the nucleus along the anterior-posterior axis. VIP-IN-projecting neurons in VP were found more uniformly throughout the anterior-posterior length of VP, and SST-IN-projecting neurons were sparse and located just anterior to the center of the nucleus (**Figure 5H, I**). Interestingly, the center region (~-1.555mm posterior to bregma) appears to have the largest concentration of input neurons, regardless of the IN subtype. Overall, these results demonstrate that INs in M1 receive input from both first and higher order thalamic nuclei, as well as from sensory, motor, and polymodal nuclei. Moreover, neurons in thalamus projecting to different IN cell types in MOp can occupy distinct regions within the same nucleus.

Brain-wide long-range inputs to PNs in MOp

Lastly, we performed a similar experiment and generated a brain-wide map of input to PNs to examine if any regions send exclusive inputs to IN subtypes but not to PNs (**Figure 6A**). We again used the Wholebrain software to unbiasedly count the labelled input cells in all brain regions (**Figure 6B-C**). We found a starter-to-input cell ratio of 1:9, and the total number of input cells for PNs was 3,431 ± 1,671 cells. We found that PNs in MOp received very similar inputs compared to all IN subtypes, which includes a large proportion of input from the ipsilateral SSp and MOs (**Figure 6D**) and a smaller proportion from the ACA and AUD. PNs also received a considerable amount of input from ORB as previously shown (Hooks et al., 2013); however in contrast to VIP-INs, ORB input to PNs was primarily ipsilateral (**Figure 6E, G**), and input cells were located throughout all the layers in ORBl except for L6b (**Figure 6H**). In addition, similar to SST-INs, PNs also received input from RSP but it was proportionally distributed across RSPd and RSPv (**Figure 6I**). Lastly, many of the major thalamic input nuclei to INs in MOp also projected to PNs, including PO, VAL, and VP, which is consistent with the literature (Hooks et al., 2013; Hunnicutt et al., 2014; Tanaka et al., 2018; **Figure 6J**). Overall, PNs received input from the same major input regions as the IN subtypes.

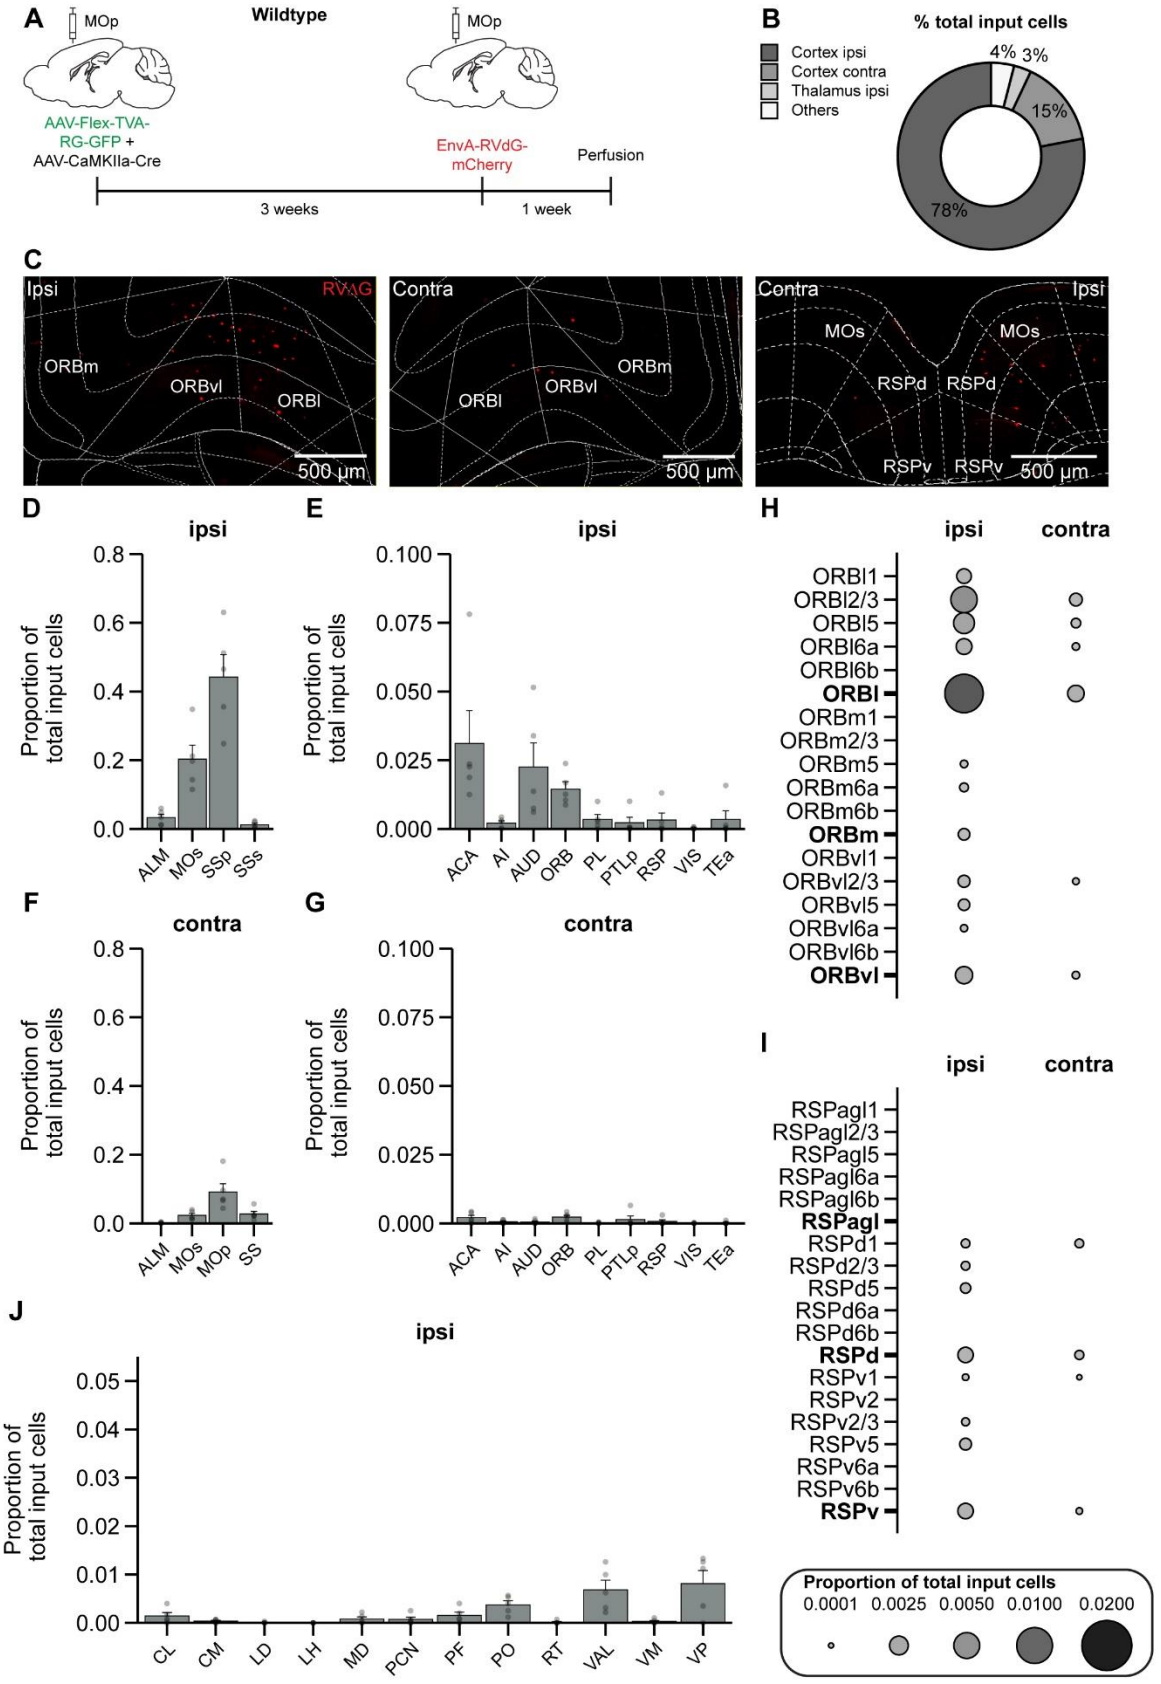
A
MOp
Wildtype
MOp
AAV-Flex-TVA-
RG-GFP +
AAV-CaMKIIa-Cre
EnvA-RVdG-
mCherry
Perfusion
3 weeks
1 week
B
% total input cells
Cortex ipsi
Cortex contra
Thalamus ipsi
Others
4%
3%
15%
78%
C
Ipsi
RVdG
ORBm
ORBvl
ORBl
500 µm
Contra
ORBm
ORBl
ORBvl
500 µm
Contra
Ipsi
MOs
MOs
RSPd
RSPd
RSPv
RSPv
500 µm
D
ipsi
Proportion of
total input cells
0.8
0.6
0.4
0.2
0.0
ALM
MOs
SSp
SSs
E
ipsi
Proportion of
total input cells
0.100
0.075
0.050
0.025
0.000
ACA
AI
AUD
ORB
PL
PTLp
RSP
VIS
TEa
F
contra
Proportion of
total input cells
0.8
0.6
0.4
0.2
0.0
ALM
MOs
MOp
SS
G
contra
Proportion of
total input cells
0.100
0.075
0.050
0.025
0.000
ACA
AI
AUD
ORB
PL
PTLp
RSP
VIS
TEa
H
ipsi
contra
ORBl1
ORBl2/3
ORBl5
ORBl6a
ORBl6b
ORBl
ORBm1
ORBm2/3
ORBm5
ORBm6a
ORBm6b
ORBm
ORBvl1
ORBvl2/3
ORBvl5
ORBvl6a
ORBvl6b
ORBvl
I
ipsi
contra
RSPagl1
RSPagl2/3
RSPagl5
RSPagl6a
RSPagl6b
RSPagl
RSPd1
RSPd2/3
RSPd5
RSPd6a
RSPd6b
RSPd
RSPv1
RSPv2
RSPv2/3
RSPv5
RSPv6a
RSPv6b
RSPv
Proportion of total input cells
0.0001
0.0025
0.0050
0.0100
0.0200
J
ipsi
Proportion of
total input cells
0.05
0.04
0.03
0.02
0.01
0.00
CL
CM
LD
LH
MD
PCN
PF
PO
RT
VAL
VM
VP

Figure 6. Brain-wide input to PNs in MOp. (A) Helper virus was injected unilaterally into right MOp followed by pseudotyped G-deleted rabies virus 3 weeks later. Animals were sacrificed 1 week after the rabies viral injection. **(B)** Mean percentage of brain-wide inputs to PNs from broad subdivisions of the brain. **(C)** Example images of $mCherry^+$ input cells in ipsilateral ORB (left), contralateral ORB (middle) and bilateral RSP that project to PNs (right). **(D)** Mean proportion of brain-wide input cells in ipsilateral somatosensory and motor cortices. **(E)** Mean proportion of input cells in other ipsilateral cortical regions. Regions with negligible cell counts are not shown. **(F)** Mean proportion of input cells in contralateral somatosensory and motor cortices. **(G)** Mean proportion of input cells in other contralateral cortical regions. Regions with negligible cell counts are not included. **(H)** Mean proportion of brain-wide inputs to PNs found within different layers of ORBl, ORBm and ORBvl. Bolded labels show the sum across all layers of each subdivision. **(I)** Mean proportion of brain-wide inputs to PNs found within different layers of RSPagl, RSPd and RSPv. Bolded labels show the sum across all layers of each subdivision. **(J)** Mean proportion of brain-wide input to PNs found in different thalamic nuclei. n = 5 mice. Each point represents a mouse. Error bars show the SEM.

Discussion

It has been well-established that MOp is involved in the execution of volitional movement; however, recent observations that MOp also exhibits reward-related activity that could facilitate reward-based motor learning have only begun to be explored. Previous work using *in vivo* two-photon imaging revealed that VIP-INs in MOp are preferentially responsive to reward compared to the other cell types in MOp, and their responses to reward become more reliable after associative learning (Lee et al., 2022). Hence, in this study, we employed monosynaptic rabies tracing and brain-wide mapping using the Wholebrain software (Fürth et al., 2018) to identify candidate brain regions with preferential projections to VIP-INs in MOp, which might confer reward-related input to VIP-INs. By generating a comprehensive and subtype-specific map of brain-wide inputs to VIP-INs and comparing it to the input maps for PV-INs and SST-INs, we demonstrated that the major inputs to all three IN subtypes originated from sensory, motor, and prefrontal cortex. While many of the identified brain regions provided input to all three IN cell types, some regions were significantly biased toward one particular IN cell type. Among these, we observed dense bilateral input from ORB projecting primarily to VIP-INs. In comparison, we also found biased input from the ipsilateral RSP to SST-INs. With the generation of IN subtype-specific brain-wide input maps to MOp, this study provides a framework for future investigations exploring how different IN subtypes in MOp integrate long-range inputs from various brain

regions and consequently, influence motor output and motor learning. Below, we discuss the conceptual implications of these findings.

One major finding from our work is that ORB provides dense bilateral input preferentially to VIP-INs, and ipsilateral input to PNs. Although previous work identified ORB input to PNs in the vibrissal region of MOp, it remained unclear whether ORB inputs also synapse onto other cell types. Using fluorescent retrograde microbeads, Hooks et al., (2013) found that ipsilateral ORB input neurons projecting to MOp largely originated in deep layers of ORB, while contralateral input was predominantly from superficial layers of ORB. In addition, anterograde tracing of ORBl and ORBvl projections demonstrated the presence of ORB axons throughout all layers of ipsilateral MOp with the greatest presence of axons in deep L5B and L6 (Hooks et al., 2013). In the present study, we found that PNs and VIP-INs are the major recipient cell types of ORB input to MOp, and that both receive input mainly from the lateral part of ORB (ORBl). Interestingly, ORBl neurons projecting to PNs were more numerous on the ipsilateral side and were more evenly distributed throughout L1 to L6a. In contrast, ORB input neurons projecting to VIP-INs were unique in that they resided mainly in the L2/3 of both ipsilateral and contralateral ORBl. These results suggest that the anatomical differences in laminar and hemispheric connectivity from ORB to PNs and VIP-INs in MOp could be related to different roles in processing reward-related signals during learning.

The exact functional role of ORB is still unclear but it is thought to be involved in predicting or updating expected outcomes or values during learning (Baltz et al., 2018; Zhou et al., 2019). Neurons in ORB have been shown to persistently represent both cue and reward throughout learning in an auditory-cued reward associative learning task, and photoinhibition of ORB specifically during the reward-predicting cue period impaired behavioral performance during learning (Namboodiri et al., 2019). Another study found that during an odor-cued reward associative learning task, 29% of ORB neurons demonstrated large amplitude responses to the cue odor after associative learning (Wang et al., 2020). Interestingly, if naïve mice were exposed to the same odor prior to learning and in the absence of reward, only 11% of neurons in ORB showed a response, and these responses were lower in amplitude and trial-to-trial consistency compared to after learning. In line with the notion that ORB represents value, the authors found that ORB's responses to the cue were greater when the mice were thirsty compared to when they

were satiated. Furthermore, optogenetic silencing of ORB during the cue and response period reduced cue-evoked anticipatory licking, suggesting ORB is involved in encoding the relative value of a stimulus and choosing an appropriate response (Wang et al., 2020). Together with our anatomical results, these findings signify that neurons in ORB represent value by integrating internal state with learned associations and might be involved in reward-based motor skill learning.

It has also been hypothesized that ORB may have a role in regulating reinforcement learning through top-down modulation of cortex (Banerjee et al., 2020; Liu et al., 2020). However, the different subdivisions of ORB are infrequently investigated in isolation, and their individual function remains unclear. One study found that in monkeys, ORBl is necessary for reward-based learning while ORBm is necessary for reward-based decision making (Noonan et al., 2010). In S1, ORBl has been shown to project strongly to L2/3 and L5 (Banerjee et al., 2020). Similar to the cue and reward-related responses we observed in MOp during an auditory-cued classical conditioning task (Lee et al., 2022), a subpopulation of neurons in S1 showed outcome or reward-related activity during a texture-based go/no-go task. Furthermore, these outcome-related responses in S1 underwent ORB-dependent remapping during a reversal learning task, indicating that ORBl input to S1 may function as a teaching signal to modulate and remap S1 activity during reward-based reversal learning. In line with their findings, our results suggest that during motor learning, a teaching signal from both ipsi and contralateral ORBl to VIP-INs in MOp could mediate broad disinhibition. Recent findings from Ren et al. (2022) and Szadai et al. (2022) show that VIP-INs are key regulators in gating a cortex-wide response to reward. Therefore, input from ORB to VIP-INs in MOp during cue and reward could be a potential mechanism for driving VIP-INs and gating plastic changes in MOp during reward-based motor learning.

A previous study also used rabies tracing to map presynaptic inputs to VIP-INs, PV-INs and SST-INs in MOp; however, the results are discordant with some of our key findings. In Duan et al. (2020), they did not observe preferential inputs from ORB to VIP-INs in MOp, but instead, they found that ORB projected mostly to PV-INs, followed by VIP-INs, and then SST-INs. The authors also found that ipsilateral input from ORB to VIP-INs was substantially more abundant than contralateral input, while we found ipsilateral and contralateral inputs from ORB to VIP-INs

were comparable in proportion (**Figure 3F-G**). Additionally, Duan et al. found that ORBvl was the primary source of ORB input to MOp, followed by ORBl input, and little to no input from ORBm. Incongruously, we found that ORB input to MOp arose mainly from ORBl with very small amounts of input from ORBvl and ORBm (**Figure 3F-G**). After carefully examining the experimental procedures, we noticed a significant difference in the injection coordinates. In our study, we targeted MOp using the coordinates 0.3mm anterior and 1.5mm lateral to bregma. This coordinate was determined based on previous motor mapping studies (defined by measuring evoked forelimb movement) using either optogenetic (Harrison et al., 2012) or electrical (Tennant et al., 2011; Brown et al., 2022) stimulation. This coordinate evoked forelimb movement most reliably (Harrison et al., 2012) and is near the center of the forelimb motor representations (Tennant et al., 2011; Brown et al., 2022). Furthermore, using *in vivo* two-photon imaging, several studies have observed motor learning-related plasticity during forelimb motor learning tasks using this coordinate (Peters et al., 2014, 2017; Chen et al., 2015b; Yang et al., 2022), and optogenetic inhibition of this coordinate impairs forelimb function (Peters et al., 2014). Lastly, we also previously observed reward responses in VIP-INs and associative learning-related plasticity in different neuronal subtypes at this coordinate (Lee et al., 2022), which supports our observations that input from ORB to VIP-INs in MOp could provide reward-related signals during reward-based motor learning. In contrast to our injection site, Duan et al., used the coordinate of 1.34mm anterior and 1.75mm lateral to bregma. While stimulation at this coordinate can evoke forelimb movement, it is located at the edge of the forelimb representation (Tennant et al., 2011; Harrison et al., 2012; Brown et al., 2022) and is more strongly associated with jaw movement (Tennant et al., 2011). Hence, the differential observations on the cell-type specificity of post-synaptic targets in MOp and bilateral projection patterns from the two studies may further highlight how different IN subtypes are uniquely involved in regulating local circuitry in different brain regions.

In this study, we have identified a multitude of inputs that target different cell types within MOp. The brain-wide maps of inputs reveal long-range connectivity onto different IN subtypes and provide insight on how input from different regions is parsed within the MOp microcircuitry. Importantly, we identify ORB as a putative candidate region that could drive the reward representation among VIP-INs within MOp. These inputs suggest MOp is not only involved in producing motor commands, but also integrates numerous streams of complex input,

including sensory and reinforcement-related information, to modulate motor behavior and motor learning. Many inputs from a single brain region project onto several IN subtypes; therefore, these streams could engage distinct cell types in MOp based on behavioural state or context. Future work will involve detailed investigation of the functional connectivity of these long-range inputs during behavior to examine whether and how they engage distinct modules within MOp during associative and motor learning.

Materials and Methods

Mouse lines

All animal experiments were approved by the University of Ottawa Animal Care Committee and in accordance with the Canadian Council on Animal Care guidelines. Experimental mice were group-housed in plastic cages with food and water ad libitum in a room with a reversed light cycle (12-12h). PV-Cre (008069), SST-Cre (013044), VIP-Cre (010908) and 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 males and females were used.

Surgery

Mice underwent two surgeries. In the first surgery, they were injected with a helper virus. After three weeks, in the second surgery, animals were injected with an engineered rabies virus. The same surgical procedures were used for both surgeries on the same injection site. Mice were deeply anesthetized using 1-2% isoflurane and given a subcutaneous injection of buprenorphine (0.05 mg/kg) for analgesia. An incision was made, and a small craniotomy was performed at the coordinate 1.5mm lateral and 0.3mm anterior to bregma above the forelimb area of MOp. A glass pipette was loaded and lowered to 500µm below the pia and 100nl of the virus was injected at a rate of 10nl/min. The pipette was left in place for 10 minutes to avoid backflow, then the pipette was raised to 300µm below the pia and an additional 100nl of the virus was injected. The pipette was again left in place for 10 minutes. All injections were performed on the right hemisphere only. The incision was then sutured, bupivacaine ointment was applied topically, and mice recovered on a heated pad. Four hours following surgery, an additional subcutaneous injection of buprenorphine (0.05mg/kg) was given. For the helper virus injections in PV-Cre, SST-Cre and

VIP-Cre, either AAV1-EF1a-DIO-TVA950-T2A-CVS11G (plasmid obtained from Friedrich Miescher Institute for Biomedical Research Vector Core) or AAV2/DJ-hSyn-FLEX-TVA-P2A-eGFP-2A-oG (Canadian Neurophotonics Platform Viral Vector Core Facility) was used. For the helper virus injections targeting PNs, a 1:1 mixture of AAV2/DJ-hSyn-FLEX-TVA-P2A-eGFP-2A-oG and AAV9.CamKII-Cre.SV40 (Addgene) was used. For the engineered rabies virus injection, EnvA-G-deleted-Rabies-mCherry (Salk Institute for Biological Studies or University of Berlin Viral Core Facility) was used.

Histology

One week after the rabies injection, mice were deeply anesthetized and transcardial perfusion was performed with 4% paraformaldehyde (PFA). Brains were kept in 4% PFA overnight at 4°C and transferred to a 30% sucrose and 0.1% sodium azide in phosphate buffered saline (PBS) solution at 4°C. The bottom right side of the brain was cut approximately 1-2mm deep to mark the injected hemisphere. Brains were then sectioned with a microtome with the thickness of 40µm and kept in 0.1% sodium azide in PBS solution. Every fourth section was mounted such that the entire brain was screened at 120µm intervals. Sections were then counterstained using either Neurotrace Blue 435/455 Blue Fluorescent Nissl Stain (ThermoFisher Scientific) or Vectashield Hardset Antifade Mounting Medium with DAPI (Vector Laboratories).

Imaging

Images were obtained at 10X with either the Zeiss AxioImager M2, Zeiss AxioScanner Z1 or Zeiss AxioObserver 7 microscopes. Entire brain sections were tiled using motorized stage controls and stitched using Zeiss ZEN Microscope Software.

Data Analysis

Starter cells were identified as cells with colocalized eGFP, mCherry and DAPI/Neurotrace Blue and counted manually using the multi-point tool in Fiji (Schindelin et al., 2012). Input cells were identified as cells outside of right MOp with colocalized mCherry and DAPI/Neurotrace Blue. Input cells were detected, counted, and registered to brain regions using WholeBrain Software Suite (Fürth et al., 2018) in R. All cell quantifications and image registrations were manually inspected and adjusted as needed. Any incorrectly detected cells that did not colocalize with the counterstain (DAPI or Neurotrace Blue) were removed from the dataset. Subsequent analysis and

figures were made using custom-written code in R and Matlab. All analyses were performed on the proportion of total input cells for each region unless otherwise stated. To calculate the proportion, the number of input cells in a specified region was divided by the total number of input cells in the entire brain for each animal.

Statistics

Comparisons between IN subtypes were performed using one-sided bootstrap. Briefly, distributions F and G, were sampled with replacement and compared under the null hypothesis $H_0: F = G$ for 1,000 replications. The achieved significance level was calculated as the proportion of replications supporting the null hypothesis (Efron and Tibshirani, 1994). P-values were corrected for multiple comparisons using the Bonferroni correction. All statistics were performed in Matlab.

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). C.L. was supported by Ontario Graduate Scholarship and Queen Elizabeth II Graduate Scholarship. S.L.C was supported by Fonds de Recherche du Québec Natural Sciences, Mathematics, Sciences and Engineering Doctoral Training Scholarship (FRQNT).

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.

Conflict of Interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest

GENERAL DISCUSSION

Summary of Results

Early investigations of the brain defined distinct regions as having specialized functions. While this assertion remains largely true, a continually growing body of evidence highlights how profusely interconnected the brain is and how single regions can multiplex diverse and flexible information outside of its specialized modality. In this thesis, we hypothesized that in response to reward, VIP-INs in M1 receive long-range input from reward-related brain regions to activate VIP-INs and disinhibit local PNs, thus enabling plastic changes in M1. We propose this mechanism is key to gating reward-based motor learning. The objective of this thesis was to examine firstly, if different neuronal cell-types in M1 represent reward, how these representations change with associative learning, and finally, which brain regions provide these reward-related inputs to M1. The key findings of this thesis are that during a head-fixed classical conditioning task, VIP-INs and PV-INs were both responsive to reward and reward-related cues. After associative learning, PV-INs became more reliably responsive to the cue, while VIP-INs became more reliably responsive to reward. These changes in PV-INs and VIP-INs were not observed when these stimuli were presented in the absence of an associative learning paradigm. In contrast, PNs were unresponsive to the CS and initially responsive to the reward but suppressed by reward after associative learning. SST-INs had relatively small responses to CS and reward. Additionally, we generated a brain-wide map of input to PV-INs, SST-INs and VIP-INs in M1 and demonstrated that ORB, a region implicated in encoding value (Tremblay & Schultz, 1999; Baltz et al., 2018; Zhou et al., 2019; Namboodiri et al., 2019; Wang et al., 2020) preferentially projects to VIP-INs and may potentially confer reward-related activity to VIP-INs and thereby function as a teaching signal to initiate M1 plasticity. Below, I will discuss the conceptual implications.

Reward Responses in Cortex

Reward is a major driver in animal behaviour. Reward can positively reinforce behaviours associated with the reward (Schultz, 2000), and it has been shown to accelerate and enhance motor learning (Wächter et al., 2009; Palminteri et al., 2011; Abe et al., 2011; Nikooyan & Ahmed, 2015). In 2006, Shuler & Bear demonstrated that 28% of recorded neurons in the

primary visual cortex responded to water rewards after rats learned a visually-cued associative learning task (Schuler & Bear, 2006). This result posed a major finding, as this was the first demonstration that sensory regions, as early in the ascending pathway as the primary visual cortex, are modulated by reward. As discussed in the general introduction, investigations in monkeys have also reported reward-related activity in M1 (Marsh et al., 2015; Ramkumar et al., 2016; Ramakrishnan et al., 2017). Reward-related modulation has now been reported throughout the cortex, including the somatosensory cortex, (Lacefield et al., 2019; Benezra et al., 2021; Schoenfeld et al., 2021), the motor cortex and the olfactory cortex (Ottenheimer et al., 2022). Specifically, among VIP-INs, reward responses have been observed in the auditory cortex and the prefrontal cortex (Pi et al., 2013). A recent preprint that employs random access two-photon imaging and fiber photometry remarkably reports that VIP-INs throughout a dozen different regions in cortex respond to reward (Szadai et al., 2022). These findings are consistent with our result that VIP-INs in M1 represent reward even when mice are not performing a motor task (Lee et al., 2022), supporting the notion that reward may be such a salient signal that it is globally broadcasted throughout cortex.

Previous work has demonstrated that the apical dendrites of PNs in M1 are an important site for plasticity, and unlike the perisomatic dendrites, apical dendrites undergo compartment-specific dendritic spine remodelling during motor learning. SST-INs, the main target of VIP-INs, primarily target the apical dendrites of PNs and optogenetic manipulation of SST-INs disrupted spine remodeling in apical dendrites in M1 (Chen et al., 2015b).

These results suggest the apical dendrites of PNs are an important site for learning and their activity is in part, regulated by VIP-INs. During a whisker-based task where mice must detect the presence of a pole and release a lever to receive a reward, reward has been shown to induce long-latency tuft-wide dendritic events in the apical dendrites of L2/3 and L5 neurons in the barrel cortex (Lacefield et al., 2019). This reward-related activity was not seen in L4 neurons which are the recipients of bottom-up thalamic input; therefore, reward-related activity is likely introduced to S1 through top-down input targeting the superficial layers. These findings suggest that within cortex, reward-related regions such as ORB could provide top-down input in response to reward, and activation of VIP-INs disinhibit the dendritic tufts of PNs. Intriguingly, Lacefield et al. observed that reward-tracking dendrites were also more likely to be responsive to the initial

whisker contact with the pole than non-reward tracking dendrites, suggesting reward responses are more likely in dendrites with task-related activity. Moreover, when the mice were presented with random rewards outside of the learned task structure, a distinct subset of dendrites responded to the random rewards, demonstrating that dendrites show task-specific reward-related activity and that these neurons represent both sensory-related information and reward-related information.

Another intriguing question is how these responses change with learning. Lacefield et al. imaged dendritic activity in mice that had already undergone a pretraining phase involving reward, therefore, the mice in this study may expect reward during behavioural training sessions. We found that in the motor cortex, L2/3 PNs only respond significantly to novel rewards but habituate after repeated CS-reward pairings and had activity levels below baseline. A study from Schoenfeld et al. (2021) demonstrated that in naïve mice, reward induced global, tuft-wide calcium events in the apical dendrites of L5 PNs during a whisker discrimination task with approximately one third of dendrites showing reward-related activity – a greater proportion than the 14% observed by Lacefield et al. (2019). As training progressed, the reward-related activity gradually declined (Schoenfeld et al., 2021), consistent with our finding that L2/3 PNs are significantly reward responsive in the naïve state but after associative learning has occurred, they no longer respond to reward. Dendritic reward-related responses are yet to be examined in the apical dendrites of PNs in M1.

In S1, it has been reported that reinforcement induced long-lasting changes in the activity of L5 apical dendrites in PNs (Benezra et al., 2021). The authors used a whisker discrimination task where mice must discriminate between a CS+ (an air puff across the whiskers paired with a subsequent water reward) and a CS- (an air puff across the whiskers in a different direction, not associated with reward). In naïve mice exposed to the whisker stimulation, there was no bias for either stimulus and individual tufts responded to both the CS+ and the CS-. In contrast, after learning, there was still no overall population bias for either stimulus, but individual dendritic tufts became selective for either the CS+ or CS-. This change in tuft selectivity was not seen when mice were repeatedly exposed to the CS+ and CS- in the absence of reward. These findings suggest that in S1, associative learning induces plastic changes in the apical dendrites of PNs, but this plasticity is dependent on reward. These reward-related plastic changes in PN apical

dendrites in S1 are consistent with our results demonstrating that reward recruits VIP-INs to regulate local plasticity in M1 through inhibition of SST-INs. This suggests that a similar mechanism may drive activity and ensuing plastic changes in the dendrites of PNs in M1 in response to reward. Future investigations performing two-photon calcium imaging in the apical dendrites of PNs in M1 during reward will be necessary to confirm that reward induces dendritic responses in M1.

Long-Range Input and Learning

Reward-related input to PN dendrites in cortex may come form long-range inputs from higher order brain regions with value-related computations. From our rabies tracing results (Chapter 3), candidate regions for reward-related input include ORB, which preferentially projects to VIP-INs more than to PV-INs or SST-INs. ORB also projects to PNs and could, therefore, simultaneously activate VIP-INs, which thereby disinhibit PN dendrites while also directly exciting PNs. RSP was identified as projecting preferentially to SST-INs compared to VIP-INs or PV-INs. The potential role of these brain regions are discussed in-depth in the third chapter of this thesis, but briefly, ORB has a well-characterized role in encoding value (Baltz et al., 2018; Zhou et al., 2019; Namboodiri et al., 2019; Wang et al., 2020) and RSP has a known role in aversive learning and has been shown to be active during learning paradigms, including reinforcement learning (Hattori et al., 2019) and motor learning (Makino et al., 2017), but not during passive visual stimuli (Makino et al., 2015). Others have found RSP was supressed during appetitive, whisker-based associative learning (Gilad & Helmchen, 2020), consistent with our findings that RSP is a major input to SST-INs in M1 and that SST-INs are largely unresponsive during our appetitive head-fixed classical conditioning task. It is possible that if our head-fixed classical conditioning task was repeated using an aversive unconditioned stimulus instead of appetitive, RSP input may drive SST-INs in M1.

We found that PV-INs received a multitude of brain-wide input although we did not find any input regions that preferentially project to PV-INs compared to both VIP-INs and SST-INs. Since VIP-INs and PV-INs both showed cue-related activity, the cue-related input may come from a region that projects to both VIP-INs and PV-INs. The thalamic nucleus MD is a possible candidate. Neurons in MD have been shown to track task uncertainty, and MD neurons projecting to PV-INs in the prelimbic cortex have been shown to engage feed-forward inhibition in the

prelimbic cortex (Mukherjee et al., 2021). The authors speculate that this reduces the signal-to-noise ratio in cortex during decision making. It is possible that the cue engages a similar mechanism in M1, where it is beneficial to reduce the signal-to-noise ratio prior to anticipatory licking and reward. Interestingly, distinct populations of MD neurons project to PV-INs and VIP-INs in the prelimbic cortex (Mukherjee et al., 2021). If MD does in fact confer cue-related activity to M1, distinct PV-projecting and VIP-projecting subpopulations could explain why PV-INs and VIP-INs undergo asymmetric cue-related plasticity. AUD is another possible cue-related input to PV-INs. Our rabies tracing experiments revealed that AUD provided substantial input to all cell types in M1. Previous studies have shown that during classical conditioning, the receptive fields of AUD neurons are plastic and undergo a long-lasting shift toward the frequency of the CS+ (Bakin & Weinberg, 1990; Weiberger & Lepan, 1993; Lee & Rothschild, 2021). Positive reinforcement paired with a tone has also been shown to increase the area in AUD that represents the frequency of the paired tone (Bao et al., 2001; Kisley and Gerstein, 2001). Therefore, increased PV-IN responses to the CS could be the result of increased AUD representation of the CS.

One important caveat is that although we identified monosynaptic inputs to different cell types, we did not perform a functional dissection. Future work will involve *ex vivo* electrophysiological recordings to dissect the functional connectivity between the identified monosynaptic input regions with different IN subtypes in M1. In S1, Wall et al. used engineered rabies virus tracing and showed that PV-INs, SST-INs and VIP-INs all receive comparable levels of input from S2, and M1 (Wall et al., 2016). More recently, Naskar et al. used optogenetics to activate projection-specific inputs to S1 and performed whole cell recordings in different cells and found these inputs can differentially engage different INs. For example, M1 input to S1 strongly recruited VIP-INs but weakly recruited SST-INs. In contrast, input from S2 strongly recruited PV-INs but weakly recruited VIP-INs and SST-INs (Naskar et al., 2021). Hence, while most input regions project to multiple IN subtypes in M1, features such as synapse strength, release probability, and short-term plasticity, such as facilitation versus depression of synaptic input could vary. These experiments will be important to assess synaptic function and highlight cell-type specific differences in synaptic physiology.

Furthermore, the activity and function of long-range inputs can change with context, task contingencies and learning (Makino et al., 2015; Banerjee et al., 2020), therefore, it is critical to assess the function of different projections *in vivo* throughout learning. Here, we identified preferential input from ORB to VIP-INs and RSP to SST-INs in M1. To understand the function of these inputs, dual colour two-photon calcium imaging to simultaneously examine the activity of ORB axons and the somatic activity of VIP-INs in M1 in awake and behaving mice can be used. The advantage of simultaneous dual color calcium imaging while mice are performing the head-fixed classical conditioning task or a head-fixed motor learning task is the relationship between ORB input and VIP-INs activity can be examined. It is possible that during both the associative learning task and the motor learning task, ORB axons in M1 will be activated after receiving reward and ORB axon activity will be highly correlated with VIP-IN somatic activity. These experiments will verify first, whether ORB projections to M1 are reward responsive, and second, if ORB input and VIP-IN activity are correlated *in vivo*. Similar experiments can be used to also test the functional relationship between SST-INs and RSP. Based on our findings that SST-IN responses in M1 are relatively low during the head-fixed classical conditioning task, therefore, strong activity in RSP axons or in SST-IN somas during the head-fixed classical conditioning task are not expected. However, these results could contrast the activity observed during a motor learning task. Based on the findings from Makino et al. (2017) showing that RSP is one of the first cortical regions activated during a motor learning lever-press task, RSP axons in M1 could be active during an M1-dependent motor learning task but may not be active during a classical conditioning task. Findings from Yang et al. (2022), show the transcription factor NPAS4 regulates motor learning-related engrams in SST-INs, where some SST-INs become active and others are supressed. It is possible that RSP regulates the recruitment of SST-INs to the learning-related ensemble and therefore, activation of RSP input to M1 may be dependent on a motor learning task.

Both ORB and RSP function remain enigmatic. A likely hypothesis is that ORB encodes value (Tremblay & Schultz, 1999; Baltz et al., 2018; Zhou et al., 2019; Namboodiri et al., 2019; Wang et al., 2020), while RSP input to M1 may encode the behavioural relevance or the context of the task (Makino et al., 2015). Fitting separate generalized linear models to ORB and RSP axon activity, where task variables such as cue, reward, licking activity, execution of the learned

movement and stage of learning (naïve, early, late, expert), may provide insight on whether any of these task variables are heavily weighted in ORB or RSP input to M1 during these two tasks.

Finally, ORB and RSP input to M1 can be inhibited using optogenetics to test the causal role of these projections during classical conditioning task and in the motor learning. There are a multitude of possible outcomes: performance during associative learning and during motor learning could be differentially affected, and/or specific task variables could be differentially affected. A key to these experiments may be the timing of the photoinhibition. Inhibiting ORB axons in M1 during the reward period may impair reward-based motor learning but not associative learning since M1 is unlikely to be the site of associative learning. It is also unlikely to be the site responsible for expression of the conditioned response since lick-movement cells have been shown to be mostly located in ALM (Komiyama et al., 2010; Mayrhofer et al., 2019). Therefore, while silencing of the soma of neurons in ORB inhibits associative learning (Wang et al., 2020), inhibiting ORB axons specifically to M1 are unlikely to impair associative learning. However, during a motor learning task where reward-related input may promote movement-related plasticity, it is possible that photoinhibition of ORB input to M1 specifically during the reward period will prolong or impair the process of motor learning. Similarly, photoinhibition of RSP input to M1 may not affect the classical conditioning task, however, photoinhibition of RSP input to M1 during a motor learning task may impair performance since RSP has been implicated in conveying behavioural relevance (Makino et al., 2015; Hattori et al., 2019).

Circuit Mechanisms for Reward-Based Learning

Our results show that during a head-fixed classical conditioning task, VIP-INs in M1 respond to reward, demonstrating that reward signals in M1 are not restricted to motor learning. In fact, reward signals may be broadcasted to VIP-INs throughout cortex (Szadai et al., 2022). These results are complementary to findings that VIP-INs throughout cortex are active during the early phases of a motor learning task (Ren et al., 2022). If reward provides a global signal that activates VIP-INs and disinhibits PNs throughout cortex during learning, how then is the spatiotemporal specificity of plasticity maintained? There are multiple theories on plasticity rules in the brain, each of which attempt to explain the spatiotemporal dynamics and conditions required for potentiation and plasticity. Moreover, for a reward to reinforce the reward-seeking behaviour, this form of plasticity would be required to act retroactively, on the timescale of seconds. Previous

theories such as Hebbian plasticity and spike timing dependent plasticity operate on millisecond timescales (Markram et al., 1997; Bi & Poo, 1998) and therefore cannot explain how reward might reinforce behaviour. Other theories offer potential solutions.

Behavioural timescale plasticity has been demonstrated in the hippocampus (Bittner et al., 2017), although it remains to be reproduced in cortex. In PNs in CA1, coordinated input from CA3 to the perisomatic region combined with entorhinal input to the apical dendrites induced plateau potentials in the dendrites, burst firing, and place field induction - a sign that plasticity has occurred, and a spatial memory has been formed (Bittner et al., 2015). By recording from CA1 neurons in mice running on a track, the authors found that plateau potentials in CA1 neurons induced slow decaying, ramp-like, subthreshold depolarization, and this depolarization lasts for seconds and induced place field formation on subsequent trials. Additionally, these events were able to induce place fields in locations on the track where the mouse visited up to seconds prior to the depolarization event, suggesting this plasticity can reinforce behaviours that occurred seconds before the depolarization. The authors' use modelling to explain this extraordinary form of plasticity and show that when a subthreshold, slow decaying input that is insufficient to induce spiking is paired with a second, more global input, the combination of inputs induces burst spiking and plateau potentials, resulting in long-term potentiation (Bittner et al., 2017). It is possible that a similar mechanism occurs in M1 during motor learning, whereby reward recruits long-range input, possibly from ORB to PN apical dendrites, and if paired with slow decaying movement-related perisomatic input, the combination of inputs could induce potentiation. However, this mechanism would be contingent on VIP-INs also responding to reward to mediate the release of inhibition from SST-INs onto the apical dendrites. Indeed, in the hippocampus, VIP-INs have been shown to be important in reward-based spatial learning, and optogenetic inactivation of VIP-INs in CA1 disrupted place field formation in CA1 PNs and impaired spatial learning, while VIP-IN activation enhanced spatial learning (Turi et al., 2019). In M1, VIP-IN mediated inhibition of SST-INs has been shown to be important for motor learning (Ren et al., 2022). Therefore, behavioural time scale plasticity gated by VIP-INs may be a mechanism of learning in M1.

This theory may explain why in Chapter 2, increased VIP-IN activity in M1 was accompanied by a paradoxical decrease in PN activity over the course of learning. One might

expect that increased disinhibition of PNs would result in increased activity in PNs, however, since the classical conditioning task involved reward that is not dependent on forelimb movement, local, perisomatic input to the PNs was likely lacking. Therefore, although reward may induce global VIP-IN mediated disinhibition and reward-responses (Szadai et al., 2022; Ottenheimer et al., 2022), task-relevant local activity is likely necessary for PN potentiation. In the absence of such task-relevant local activity paired with reinforcement, PNs may instead undergo a form of habituation (Makino & Komiyama, 2015; Lee et al., 2022).

In vitro experiments utilizing whole-cell recordings in PNs in M1 can be used to test whether optogenetic activation of ORB axons paired with somatic current injection can induce long-term potentiation in PNs in M1. If this paired stimulation can induce LTP, simultaneous use of inhibitory DREADDs can be used to specifically silence VIP-INs and test whether VIP-IN activity is necessary to gate LTP during this stimulation protocol. These experiments could provide a potential circuit mechanism for how long-range input coordinated with local PN and IN activity cooperate to induce plastic changes in M1.

Another plausible learning rule is the three-factor learning rule (Barto, 1985; Gerstner et al., 2018). The three-factor learning rule builds upon Hebbian and spike-timing dependent plasticity, which posit that presynaptic firing (the first factor) followed shortly by postsynaptic firing (the second factor) can result in long-term potentiation of the synapse, however, three-factor learning proposes that a third, gating factor is necessary. In this framework, the first two factors flag the synapse for potentiation but a third factor, which is postulated to be driven by reward or another type of reinforcement, is globally broadcasted and potentiates any flagged synapses. Neuromodulatory systems are a strong candidate for either providing or enhancing a global reward-related signal. Both the three-factor learning rule and behavioural time scale plasticity rule, highlight the importance of coordinated activity throughout the brain and suggest long-range input, neuromodulators, and local inhibition cooperatively act as gatekeepers of plasticity.

Cell-Type Specific Responses to Neuromodulation

Neuromodulatory systems are a strong candidate for a global reward signal since the serotonergic, noradrenergic, dopaminergic and cholinergic systems each project extensively

throughout the brain. One limitation of this thesis is that the role of neuromodulators was not examined. Different INs have distinct receptor expression and distinct responses to neuromodulators (Alitto & Dan 2012; Muñoz et al., 2017; Gasselin et al., 2021; Bugeon et al., 2020). A comprehensive understanding of cell-type specific effects of different neuromodulators and their dynamics *in vivo* is still lacking.

ACh is a likely candidate for a globally broadcasted third factor. Recently, Ren et al. (2022) have shown that the activity of cholinergic neurons in the basal forebrain is highest early in motor learning and gradually decreases as learning progresses. Ablation of basal forebrain neurons impaired motor learning. Nicotinic antagonists also reduced activity in VIP-INs in M1, suggesting cholinergic modulation of VIP-INs through nicotinic receptors are important for motor learning (Ren et al., 2022). Recent work combining transcriptomics and functional characterization in V1 demonstrated a continuum of subtypes within PV-, VIP-, and SST-INs and different subtypes are differentially modulated by nicotinic or muscarinic antagonists (Bugeon et al., 2022). These findings demonstrate that even within these major subtypes, responses to neuromodulators may be heterogenous.

Dopaminergic input to M1 has been shown to be important in motor learning. Lesioning the VTA impairs motor learning in rats and application of levodopa in M1 rescued motor learning (Hosp et al., 2011). These results are intriguing since levodopa application does not offer temporal specificity, yet it was still able to rescue motor learning. This finding provides support for the three-factor rule since temporal specificity may be conferred by factors one and two – presynaptic and postsynaptic activity, and the third factor is a nonspecific global input. Another investigation found that dopaminergic input to M1 is critical for motor learning but not for the execution of previously learned movements. The authors also found that dopamine promotes LTP in M1 (Molina-Luna et al., 2009). Since dopaminergic activity is closely related to reward (Schultz, 2000), dopamine may provide the globally broadcasted third factor to potentiate task-related neurons in M1 in response to reward during motor learning.

The noradrenergic system is another candidate for learning-related neuromodulation. Yin et al., (2021) implicated reduced noradrenergic input to M1 in delayed motor learning in a mouse model of autism spectrum disorder. The authors found the mice had aberrantly increased synchrony in L2/3 PNs in M1 and both the synchrony and delayed motor learning could be

rescued by activating locus coeruleus (Yin et al., 2021). One possible explanation is that noradrenergic input to M1 modulates inhibitory interneurons and without this input, interneurons cannot regulate PN activity resulting in excessive synchrony among PNs.

Serotonin has been heavily implicated in reward prediction and reward responses although its exact role is still heavily contested (Sweighofer et al., 2008; Kranz et al., 2010). VIP-INs belong to a larger subtype of interneurons that express the ionotropic 5HT3a receptor. Serotonin application has been shown to induce burst spiking in 5HT3aR expressing interneurons but not in PV-INs or SST-INs (Lee et al., 2010). Therefore, serotonin input to M1 is likely to specifically modulate VIP-INs but not PV-INs or SST-INs. Moreover, serotonergic neurons in the dorsal raphe have been shown to encode reward-related events and motor-related events (Ranade & Mainen, 2009). While the role of serotonin in motor learning remains unexplored, serotonergic modulation is a possible candidate for activating VIP-INs in response to reward.

Cell type specific responses to different neuromodulators remain unresolved. *In vitro* experiments performing whole cell recordings in PV-INs, SST-INs and VIP-INs in the presence of different antagonists will be important for better understanding the cell-type specific responses to different neuromodulators. If these experiments identify new cell-type specific modulations, the interactions can then be investigated *in vivo*. The advent of new genetically encoded fluorescent sensors (Feng et al., 2019; Jing et al., 2020; Sun et al., 2020; Wan et al., 2021) enable investigations of different neuromodulators, and if paired with calcium imaging using a red calcium indicator, cell-type specific responses to different neuromodulators can be investigated.

Conclusion

In summary, different cell types in M1 represent reward and reward-related cues during associative learning, even in the absence of motor skill task. Notably, PV-INs were preferentially cue-responsive while VIP-INs were preferentially reward-responsive and both of these stimulus-specific responses strengthened with associative learning. Using cell-type specific neuronal tracing, ORB was identified as a potential long-range input to VIP-INs in M1, that may activate VIP-INs in response to reward, and thus disinhibit the dendrites of PNs. These findings provide evidence for the hypothesis that ORB input to VIP-INs may function as a disinhibitory gate that enable PN plasticity in response to reward. Future investigations will be necessary to functionally

dissect ORB input to VIP-INs during associative learning and during motor learning. Our rabies tracing results provide numerous potential candidates that could activate PV-INs in response to conditioned cue. Future work will be necessary to understand the functional role of PV-IN activation in response to cue and which long-range inputs might relay cue-related input to PV-INs. Finally, SST-INs did not show strong task-related activity during associative learning. Together with recent findings demonstrating the role of SST-INs in motor learning (Chen et al., 2015b; Yang et al., 2022), these results indicate that SST-INs may be more strictly involved in motor learning while PV-INs and VIP-INs additionally represent task variables beyond movement. These findings provide insight on how long-range input could leverage local microcircuitry to modulate motor activity and motor learning based on modalities outside of traditional motor cortex function.

REFERENCES

Abe, M., Schambra, H., Wassermann, E.M., Luckenbaugh, D., Schweighofer, N., and Cohen, L.G. (2011). Reward improves long-term retention of a motor memory through induction of offline memory gains. *Curr Biol* 21, 557-562. 10.1016/j.cub.2011.02.030.

Adesnik, H., and Scanziani, M. (2010). Lateral competition for cortical space by layer-specific horizontal circuits. *Nature* 464, 1155-1160. 10.1038/nature08935.

Adler, A., Zhao, R., Shin, M.E., Yasuda, R., and Gan, W.B. (2019). Somatostatin-Expressing Interneurons Enable and Maintain Learning-Dependent Sequential Activation of Pyramidal Neurons. *Neuron* 102, 202-216.e207. 10.1016/j.neuron.2019.01.036.

Airan, R.D., Thompson, K.R., Fenno, L.E., Bernstein, H., and Deisseroth, K. (2009). Temporally precise *in vivo* control of intracellular signalling. *Nature* 458, 1025-1029. 10.1038/nature07926.

Akerboom, J., Chen, T.W., Wardill, T.J., Tian, L., Marvin, J.S., Mutlu, S., Calderón, N.C., Esposti, F., Borghuis, B.G., Sun, X.R., et al. (2012). Optimization of a GCaMP calcium indicator for neural activity imaging. *J Neurosci* 32, 13819-13840. 10.1523/JNEUROSCI.2601-12.2012.

Alaverdashvili, M., and Whishaw, I.Q. (2008). Motor cortex stroke impairs individual digit movement in skilled reaching by the rat. *Eur J Neurosci* 28, 311-322. 10.1111/j.1460-9568.2008.06315.x.

Alexander, G.M., Rogan, S.C., Abbas, A.I., Armbruster, B.N., Pei, Y., Allen, J.A., Nonneman, R.J., Hartmann, J., Moy, S.S., Nicolelis, M.A., et al. (2009). Remote control of neuronal activity in transgenic mice expressing evolved G protein-coupled receptors. *Neuron* 63, 27-39. 10.1016/j.neuron.2009.06.014.

Alitto, H.J., and Dan, Y. (2012). Cell-type-specific modulation of neocortical activity by basal forebrain input. *Front Syst Neurosci* 6, 79. 10.3389/fnsys.2012.00079.

Alstermark, B., Ogawa, J., and Isa, T. (2004). Lack of monosynaptic corticomotoneuronal EPSPs in rats: disynaptic EPSPs mediated via reticulospinal neurons and polysynaptic EPSPs via segmental interneurons. *J Neurophysiol* 91, 1832-1839. 10.1152/jn.00820.2003.

Angaut, P., Cicirata, F., and Serapide, F. (1985). Topographic organization of the cerebellothalamic projections in the rat. An autoradiographic study. *Neuroscience* 15, 389–401. doi: 10.1016/0306-4522(85)90221-0.

Armbruster, B.N., Li, X., Pausch, M.H., Herlitze, S., and Roth, B.L. (2007). Evolving the lock to fit the key to create a family of G protein-coupled receptors potently activated by an inert ligand. *Proc Natl Acad Sci U S A* 104, 5163-5168. 10.1073/pnas.0700293104.

Askew, C.E., Lopez, A.J., Wood, M.A., and Metherate, R. (2019). Nicotine excites VIP interneurons to disinhibit pyramidal neurons in auditory cortex. *Synapse* 73, e22116. 10.1002/syn.22116.

Atallah, B.V., Bruns, W., Carandini, M., and Scanziani, M. (2012). Parvalbumin-expressing interneurons linearly transform cortical responses to visual stimuli. *Neuron* 73, 159-170. 10.1016/j.neuron.2011.12.013.

Ayzenshtat, I., Karnani, M.M., Jackson, J., and Yuste, R. (2016). Cortical Control of Spatial Resolution by VIP+ Interneurons. *J Neurosci* 36, 11498-11509. 10.1523/JNEUROSCI.1920-16.2016.

Babl, S.S., Rummell, B.P., and Sigurdsson, T. (2019). The Spatial Extent of Optogenetic Silencing in Transgenic Mice Expressing Channelrhodopsin in Inhibitory Interneurons. *Cell Rep* 29, 1381-1395.e1384. 10.1016/j.celrep.2019.09.049.

Bakin, J.S., and Weinberger, N.M. (1990). Classical conditioning induces CS-specific receptive field plasticity in the auditory cortex of the guinea pig. *Brain Res* 536, 271-286. 10.1016/0006-8993(90)90035-a.

Baltz, E. T., Yalcinbas, E. A., Renteria, R., and Gremel, C. M. (2018). Orbital frontal cortex updates state-induced value change for decision-making. *eLife* 7, e35988. doi: 10.7554/eLife.35988.

Banerjee, A., Parente, G., Teutsch, J., Lewis, C., Voigt, F. F., and Helmchen, F. (2020). Value-guided remapping of sensory cortex by lateral orbitofrontal cortex. *Nature* 585, 245–250. doi: 10.1038/s41586-020-2704-z.

Bao, S., Chan, V.T., and Merzenich, M.M. (2001). Cortical remodelling induced by activity of ventral tegmental dopamine neurons. *Nature* 412, 79-83. 10.1038/35083586.

Barto, A.G. (1985). Learning by statistical cooperation of self-interested neuron-like computing elements. *Hum Neurobiol* 4, 229-256.

Bègue, A., Papagiakoumou, E., Leshem, B., Conti, R., Enke, L., Oron, D., and Emiliani, V. (2013). Two-photon excitation in scattering media by spatiotemporally shaped beams and their application in optogenetic stimulation. *Biomed Opt Express* 4, 2869-2879. 10.1364/BOE.4.002869.

Benezra, S., Patel, K., Pérez Campos, I., Hillman E.M.C., Bruno R.M. (2021). Learning enhances behaviorally relevant representations in apical dendrites. bioRxiv 2021.11.10.468144; 10.1101/2021.11.10.468144

Berndt, A., Yizhar, O., Gunaydin, L.A., Hegemann, P., and Deisseroth, K. (2009). Bi-stable neural state switches. *Nat Neurosci* 12, 229-234. 10.1038/nn.2247.

Bi, G.Q., and Poo, M.M. (1998). Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J Neurosci* 18, 10464-10472.

Biane, J.S., Takashima, Y., Scanziani, M., Conner, J.M., and Tuszynski, M.H. (2016). Thalamocortical Projections onto Behaviorally Relevant Neurons Exhibit Plasticity during Adult Motor Learning. *Neuron* 89, 1173-1179. 10.1016/j.neuron.2016.02.001.

Bittner, K.C., Grienberger, C., Vaidya, S.P., Milstein, A.D., Macklin, J.J., Suh, J., Tonegawa, S., and Magee, J.C. (2015). Conjunctive input processing drives feature selectivity in hippocampal CA1 neurons. *Nat Neurosci* 18, 1133-1142. 10.1038/nn.4062.

Bittner, K.C., Milstein, A.D., Grienberger, C., Romani, S., and Magee, J.C. (2017). Behavioral time scale synaptic plasticity underlies CA1 place fields. *Science* 357, 1033-1036. 10.1126/science.aan3846.

Bliss, T.V., and Lomo, T. (1973). Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *J Physiol* 232, 331-356. 10.1113/jphysiol.1973.sp010273.

Bortone, D.S., Olsen, S.R., and Scanziani, M. (2014). Translaminar inhibitory cells recruited by layer 6 corticothalamic neurons suppress visual cortex. *Neuron* 82, 474-485. 10.1016/j.neuron.2014.02.021.

Bouchard, M.B., Voleti, V., Mendes, C.S., Lacefield, C., Grueber, W.B., Mann, R.S., Bruno, R.M., and Hillman, E.M. (2015). Swept confocally-aligned planar excitation (SCAPE) microscopy for high speed volumetric imaging of behaving organisms. *Nat Photonics* 9, 113-119. 10.1038/nphoton.2014.323.

Bouton, M.E., and Moody, E.W. (2004). Memory processes in classical conditioning. *Neurosci Biobehav Rev* 28, 663-674. 10.1016/j.neubiorev.2004.09.001.

Boyden, E.S., Zhang, F., Bamberg, E., Nagel, G., and Deisseroth, K. (2005). Millisecond-timescale, genetically targeted optical control of neural activity. *Nat Neurosci* 8, 1263-1268. 10.1038/nn1525.

Brecht, M., Krauss, A., Muhammad, S., Sinai-Esfahani, L., Bellanca, S., and Margrie, T.W. (2004). Organization of rat vibrissa motor cortex and adjacent areas according to cytoarchitectonics, microstimulation, and intracellular stimulation of identified cells. *J Comp Neurol* 479, 360-373. 10.1002/cne.20306.

Brem, A.K., Ran, K., and Pascual-Leone, A. (2013). Learning and memory. *Handb Clin Neurol* 116, 693-737. 10.1016/B978-0-444-53497-2.00055-3.

Brown, A. R., Mitra, S., Teskey, G. C., and Boychuk, J. A. (2022). Complex forelimb movements and cortical topography evoked by intracortical microstimulation in male and female mice. *Cerebral Cortex*, bhac178. doi: 10.1093/cercor/bhac178.

Bugeon, S., Duffield, J., Dipoppa, M., Ritoux, A., Prankerd, I., Nicoloutsopoulos, D., Orme, D., Shinn, M., Peng, H., Forrest, H., et al. (2022). A transcriptomic axis predicts state modulation of cortical interneurons. *Nature* 607, 330-338. 10.1038/s41586-022-04915-7.

Burette, A., Collman, F., Micheva, K.D., Smith, S.J., and Weinberg, R.J. (2015). Knowing a synapse when you see one. *Front Neuroanat* 9, 100. 10.3389/fnana.2015.00100.

Callaway, E. M., and Luo, L. (2015). Monosynaptic Circuit Tracing with Glycoprotein-Deleted Rabies Viruses. *Journal of Neuroscience* 35, 8979–8985. doi: 10.1523/JNEUROSCI.0409-15.2015.

Cao, V. Y., Ye, Y., Mastwal, S., Ren, M., Coon, M., Liu, Q., et al. (2015). Motor Learning Consolidates Arc-Expressing Neuronal Ensembles in Secondary Motor Cortex. *Neuron* 86, 1385–1392. doi: 10.1016/j.neuron.2015.05.022.

Casas-Torremocha, D., Porrero, C., Rodriguez-Moreno, J., García-Amado, M., Lübke, J.H.R., Núñez, Á., and Clascá, F. (2019). Posterior thalamic nucleus axon terminals have different structure and functional impact in the motor and somatosensory vibrissal cortices. *Brain Struct Funct* 224, 1627-1645. 10.1007/s00429-019-01862-4.

Castro, A.J. (1972). The effects of cortical ablations on tongue usage in the rat. *Brain Res* 45, 251-253. 10.1016/0006-8993(72)90233-8.

Censor, N., Sagi, D., and Cohen, L.G. (2012). Common mechanisms of human perceptual and motor learning. *Nat Rev Neurosci* 13, 658-664. 10.1038/nrn3315.

Chabrol, F. P., Blot, A., and Mrsic-Flogel, T. D. (2019). Cerebellar Contribution to Preparatory Activity in Motor Neocortex. *Neuron* 103, 506-519.e4. doi: 10.1016/j.neuron.2019.05.022.

Chen, J.L., Margolis, D.J., Stankov, A., Sumanovski, L.T., Schneider, B.L., and Helmchen, F. (2015a). Pathway-specific reorganization of projection neurons in somatosensory cortex during learning. *Nat Neurosci* 18, 1101-1108. 10.1038/nn.4046.

Chen, S., Weitemier, A.Z., Zeng, X., He, L., Wang, X., Tao, Y., Huang, A.J.Y., Hashimotodani, Y., Kano, M., Iwasaki, H., et al. (2018). Near-infrared deep brain stimulation via upconversion nanoparticle-mediated optogenetics. *Science* 359, 679-684. 10.1126/science.aaq1144.

Chen, T.-W., Li, N., Daie, K., and Svoboda, K. (2017). A Map of Anticipatory Activity in Mouse Motor Cortex. *Neuron* 94, 866-879.e4. doi: 10.1016/j.neuron.2017.05.005.

Chen, T.W., Wardill, T.J., Sun, Y., Pulver, S.R., Renninger, S.L., Baohan, A., Schreiter, E.R., Kerr, R.A., Orger, M.B., Jayaraman, V., et al. (2013). Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature* 499, 295-300. 10.1038/nature12354.

Chiu, C.Q., Martenson, J.S., Yamazaki, M., Natsume, R., Sakimura, K., Tomita, S., Tavalin, S.J., and Higley, M.J. (2018). Input-Specific NMDAR-Dependent Potentiation of Dendritic GABAergic Inhibition. *Neuron* 97, 368-377.e363. 10.1016/j.neuron.2017.12.032.

Chow, B.Y., Han, X., Dobry, A.S., Qian, X., Chuong, A.S., Li, M., Henninger, M.A., Belfort, G.M., Lin, Y., Monahan, P.E., and Boyden, E.S. (2010). High-performance genetically targetable optical neural silencing by light-driven proton pumps. *Nature* 463, 98-102. 10.1038/nature08652.

Chrobak, J.J., and Buzsáki, G. (1998). Gamma oscillations in the entorhinal cortex of the freely behaving rat. *J Neurosci* 18, 388-398.

Chu, M.W., Li, W.L., and Komiyama, T. (2016). Balancing the Robustness and Efficiency of Odor Representations during Learning. *Neuron* 92, 174-186. 10.1016/j.neuron.2016.09.004.

Chuong, A.S., Miri, M.L., Busskamp, V., Matthews, G.A., Acker, L.C., Sørensen, A.T., Young, A., Klapoetke, N.C., Henninger, M.A., Kodandaramaiah, S.B., et al. (2014). Noninvasive optical inhibition with a red-shifted microbial rhodopsin. *Nat Neurosci* 17, 1123-1129. 10.1038/nn.3752.

Churchland, M.M., Santhanam, G., and Shenoy, K.V. (2006). Preparatory activity in premotor and motor cortex reflects the speed of the upcoming reach. *J Neurophysiol* 96, 3130-3146. 10.1152/jn.00307.2006.

Cichon, J., and Gan, W.B. (2015). Branch-specific dendritic Ca(2+) spikes cause persistent synaptic plasticity. *Nature* 520, 180-185. 10.1038/nature14251.

Condylis, C., Lowet, E., Ni, J., Bistrong, K., Ouellette, T., Josephs, N., et al. (2020). Context-Dependent Sensory Processing across Primary and Secondary Somatosensory Cortex. *Neuron* 106, 515-525.e5. doi: 10.1016/j.neuron.2020.02.004.

Costa, R.M., Cohen, D., and Nicolelis, M.A. (2004). Differential corticostriatal plasticity during fast and slow motor skill learning in mice. *Curr Biol* 14, 1124-1134. 10.1016/j.cub.2004.06.053.

Cottam, J.C., Smith, S.L., and Häusser, M. (2013). Target-specific effects of somatostatin-expressing interneurons on neocortical visual processing. *J Neurosci* 33, 19567-19578. 10.1523/JNEUROSCI.2624-13.2013.

Cruikshank, S.J., Urabe, H., Nurmikko, A.V., and Connors, B.W. (2010). Pathway-specific feedforward circuits between thalamus and neocortex revealed by selective optical stimulation of axons. *Neuron* 65, 230-245. 10.1016/j.neuron.2009.12.025.

Dal Maschio, M., Donovan, J.C., Helmbrecht, T.O., and Baier, H. (2017). Linking Neurons to Network Function and Behavior by Two-Photon Holographic Optogenetics and Volumetric Imaging. *Neuron* 94, 774-789.e775. 10.1016/j.neuron.2017.04.034.

Dana, H., Mohar, B., Sun, Y., Narayan, S., Gordus, A., Hasseman, J.P., Tsegaye, G., Holt, G.T., Hu, A., Walpita, D., et al. (2016). Sensitive red protein calcium indicators for imaging neural activity. *Elife* 5. 10.7554/eLife.12727.

Deschênes, M., Veinante, P., and Zhang, Z.W. (1998). The organization of corticothalamic projections: reciprocity versus parity. *Brain Res* 28, 286-308. 10.1016/s0165-0173(98)00017-4.

Dombeck, D.A., Graziano, M.S., and Tank, D.W. (2009). Functional clustering of neurons in motor cortex determined by cellular resolution imaging in awake behaving mice. *J Neurosci* 29, 13751-13760. 10.1523/JNEUROSCI.2985-09.2009.

Donato, F., Rompani, S.B., and Caroni, P. (2013). Parvalbumin-expressing basket-cell network plasticity induced by experience regulates adult learning. *Nature* 504, 272-276. 10.1038/nature12866.

Doron, G., Shin, J.N., Takahashi, N., Drüke, M., Bocklisch, C., Skenderi, S., de Mont, L., Toumazou, M., Ledderose, J., Brecht, M., et al. (2020). Perirhinal input to neocortical layer 1 controls learning. *Science* 370. 10.1126/science.aaz3136.

Driscoll, L.N., Pettit, N.L., Minderer, M., Chettih, S.N., and Harvey, C.D. (2017). Dynamic Reorganization of Neuronal Activity Patterns in Parietal Cortex. *Cell* 170, 986-999.e916. 10.1016/j.cell.2017.07.021.

Duan, Z., Li, A., Gong, H., and Li, X. (2020). A Whole-brain Map of Long-range Inputs to GABAergic Interneurons in the Mouse Caudal Forelimb Area. *Neurosci Bull* 36, 493-505. 10.1007/s12264-019-00458-6.

Economo MN, Viswanathan S, Tasic B, Bas E, Winnubst J, Menon V, Graybuck LT, Nguyen TN, Smith KA, Yao Z, Wang L, Gerfen CR, Chandrashekar J, Zeng H, Looger LL, Svoboda K (2018) Distinct descending motor cortex pathways and their roles in movement. *Nature* 563:79-84.

Efron, B., and Tibshirani, R. J. (1994). *An Introduction to the Bootstrap*. 0 ed. Chapman and Hall/CRC doi: 10.1201/9780429246593.

El-Boustani, S., Sermet, B. S., Foustoukos, G., Oram, T. B., Yizhar, O., and Petersen, C. C. H. (2020). Anatomically and functionally distinct thalamocortical inputs to primary and

secondary mouse whisker somatosensory cortices. *Nat Commun* 11, 3342. doi: 10.1038/s41467-020-17087-7.

El-Shamayleh, Y., Kojima, Y., Soetedjo, R., and Horwitz, G.D. (2017). Selective Optogenetic Control of Purkinje Cells in Monkey Cerebellum. *Neuron* 95, 51-62.e54. 10.1016/j.neuron.2017.06.002.

Feng, J., Zhang, C., Lischinsky, J.E., Jing, M., Zhou, J., Wang, H., Zhang, Y., Dong, A., Wu, Z., Wu, H., et al. (2019). A Genetically Encoded Fluorescent Sensor for Rapid and Specific In Vivo Detection of Norepinephrine. *Neuron* 102, 745-761.e748. 10.1016/j.neuron.2019.02.037.

Ferezou, I., Haiss, F., Gentet, L.J., Aronoff, R., Weber, B., and Petersen, C.C. (2007). Spatiotemporal dynamics of cortical sensorimotor integration in behaving mice. *Neuron* 56, 907-923. 10.1016/j.neuron.2007.10.007.

Ferrier, D. (1874). The localization of function in the brain. *Proc. R. Soc. Lond*. 22228–232. 10.1098/rspl.1873.0032

Fishell, G., and Rudy, B. (2011). Mechanisms of inhibition within the telencephalon: "where the wild things are". *Annu Rev Neurosci* 34, 535-567. 10.1146/annurev-neuro-061010-113717.

Förster, D., Dal Maschio, M., Laurell, E., and Baier, H. (2017). An optogenetic toolbox for unbiased discovery of functionally connected cells in neural circuits. *Nat Commun* 8, 116. 10.1038/s41467-017-00160-z.

Fritsch G, Hitzig E. Über die electrische Erregbarkeit des Grosshirns. *Arch Anat Physiol*. 1870; 37:300–32.

Fu, Y., Tucciarone, J.M., Espinosa, J.S., Sheng, N., Darcy, D.P., Nicoll, R.A., Huang, Z.J., and Stryker, M.P. (2014). A cortical circuit for gain control by behavioral state. *Cell* 156, 1139-1152. 10.1016/j.cell.2014.01.050.

Fürth, D., Vaissière, T., Tzortzi, O., Xuan, Y., Märtin, A., Lazaridis, I., et al. (2018). An interactive framework for whole-brain maps at cellular resolution. *Nat Neurosci* 21, 139–149. doi: 10.1038/s41593-017-0027-7.

Galea, J.M., Mallia, E., Rothwell, J., and Diedrichsen, J. (2015). The dissociable effects of punishment and reward on motor learning. *Nat Neurosci* 18, 597-602. 10.1038/nn.3956.

Galvan, A., Hu, X., Smith, Y., and Wichmann, T. (2016). Effects of Optogenetic Activation of Corticothalamic Terminals in the Motor Thalamus of Awake Monkeys. *J Neurosci* 36, 3519-3530. 10.1523/JNEUROSCI.4363-15.2016.

Garrett, M., Manavi, S., Roll, K., Ollerenshaw, D.R., Groblewski, P.A., Ponvert, N.D., Kiggins, J.T., Casal, L., Mace, K., Williford, A., et al. (2020). Experience shapes activity dynamics and stimulus coding of VIP inhibitory cells. *Elife* 9. 10.7554/eLife.50340.

Gasselin, C., Hohl, B., Vernet, A., Crochet, S., and Petersen, C.C.H. (2021). Cell-type-specific nicotinic input disinhibits mouse barrel cortex during active sensing. *Neuron* 109, 778-787.e773. 10.1016/j.neuron.2020.12.018.

Georgopoulos, A.P., Ashe, J., Smyrnis, N., and Taira, M. (1992). The motor cortex and the coding of force. *Science* 256, 1692-1695. 10.1126/science.256.5064.1692.

Georgopoulos, A.P., Kalaska, J.F., Caminiti, R., and Massey, J.T. (1982). On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *J Neurosci* 2, 1527-1537.

Georgopoulos, A.P., Schwartz, A.B., and Kettner, R.E. (1986). Neuronal population coding of movement direction. *Science* 233, 1416-1419. 10.1126/science.3749885.

Gerstner, W., Lehmann, M., Liakoni, V., Corneil, D., and Brea, J. (2018). Eligibility Traces and Plasticity on Behavioral Time Scales: Experimental Support of NeoHebbian Three-Factor Learning Rules. *Front Neural Circuits* 12, 53. 10.3389/fncir.2018.00053.

Gibson, J.R., Beierlein, M., and Connors, B.W. (1999). Two networks of electrically coupled inhibitory neurons in neocortex. *Nature* 402, 75-79. 10.1038/47035.

Gilad, A., and Helmchen, F. (2020). Spatiotemporal refinement of signal flow through association cortex during learning. *Nat Commun* 11, 1744. 10.1038/s41467-020-15534-z.

Glasgow, S.D., McPhedrain, R., Madranges, J.F., Kennedy, T.E., and Ruthazer, E.S. (2019). Approaches and Limitations in the Investigation of Synaptic Transmission and Plasticity. *Front Synaptic Neurosci* 11, 20. 10.3389/fnsyn.2019.00020.

Gloor, C., Luft, A.R., and Hosp, J.A. (2015). Biphasic plasticity of dendritic fields in layer V motor neurons in response to motor learning. *Neurobiol Learn Mem* 125, 189-194. 10.1016/j.nlm.2015.08.009.

Goard, M., and Dan, Y. (2009). Basal forebrain activation enhances cortical coding of natural scenes. *Nat Neurosci* 12, 1444-1449. 10.1038/nn.2402.

Gradinaru, V., Thompson, K.R., and Deisseroth, K. (2008). eNpHR: a Natronomonas halorhodopsin enhanced for optogenetic applications. *Brain Cell Biol* 36, 129-139. 10.1007/s11068-008-9027-6.

Graziano, M.S., Taylor, C.S., and Moore, T. (2002). Complex movements evoked by microstimulation of precentral cortex. *Neuron* 34, 841-851. 10.1016/s0896-6273(02)00698-0.

Graziano, M.S.A. (2016). Ethological Action Maps: A Paradigm Shift for the Motor Cortex. *Trends Cogn Sci* 20, 121-132. 10.1016/j.tics.2015.10.008.

Greenough, W.T., Larson, J.R., and Withers, G.S. (1985). Effects of unilateral and bilateral training in a reaching task on dendritic branching of neurons in the rat motor-sensory forelimb cortex. *Behav Neural Biol* 44, 301-314. 10.1016/s0163-1047(85)90310-3.

Gu, X.L. (2010). Deciphering the corelease of glutamate from dopaminergic terminals derived from the ventral tegmental area. *J Neurosci* 30, 13549-13551. 10.1523/JNEUROSCI.3802-10.2010.

Gunaydin, L.A., Yizhar, O., Berndt, A., Sohal, V.S., Deisseroth, K., and Hegemann, P. (2010). Ultrafast optogenetic control. *Nat Neurosci* 13, 387-392. 10.1038/nn.2495.

Guo, J.Z., Graves, A.R., Guo, W.W., Zheng, J., Lee, A., Rodríguez-González, J., Li, N., Macklin, J.J., Phillips, J.W., Mensh, B.D., et al. (2015). Cortex commands the performance of skilled movement. *Elife* 4, e10774. 10.7554/eLife.10774.

Guo, Z.V., Hires, S.A., Li, N., O'Connor, D.H., Komiyama, T., Ophir, E., Huber, D., Bonardi, C., Morandell, K., Gutnisky, D., et al. (2014). Procedures for behavioral experiments in head-fixed mice. *PLoS One* 9, e88678. 10.1371/journal.pone.0088678.

Guo, Z.V., Inagaki, H.K., Daie, K., Druckmann, S., Gerfen, C.R., and Svoboda, K. (2017). Maintenance of persistent activity in a frontal thalamocortical loop. *Nature* 545, 181-186. 10.1038/nature22324.

Hájos, N., Pálhalmi, J., Mann, E.O., Németh, B., Paulsen, O., and Freund, T.F. (2004). Spike timing of distinct types of GABAergic interneuron during hippocampal gamma oscillations *in vitro*. *J Neurosci* 24, 9127-9137. 10.1523/JNEUROSCI.2113-04.2004.

Han, X., Qian, X., Bernstein, J.G., Zhou, H.H., Franzesi, G.T., Stern, P., Bronson, R.T., Graybiel, A.M., Desimone, R., and Boyden, E.S. (2009). Millisecond-timescale optical control of neural dynamics in the nonhuman primate brain. *Neuron* 62, 191-198. 10.1016/j.neuron.2009.03.011.

Harrison, T.C., Ayling, O.G., and Murphy, T.H. (2012). Distinct cortical circuit mechanisms for complex forelimb movement and motor map topography. *Neuron* 74, 397-409. 10.1016/j.neuron.2012.02.028.

Harvey, C.D., Coen, P., and Tank, D.W. (2012). Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature* 484, 62-68. 10.1038/nature10918.

Hass, C.A., and Glickfeld, L.L. (2016). High-fidelity optical excitation of cortico-cortical projections at physiological frequencies. *J Neurophysiol* 116, 2056-2066. 10.1152/jn.00456.2016.

Hattori, R., and Komiyama, T. (2022). Context-dependent persistency as a coding mechanism for robust and widely distributed value coding. *Neuron* 110, 502-515.e11. doi: 10.1016/j.neuron.2021.11.001.

Hattori, R., Danskin, B., Babic, Z., Mlynaryk, N., and Komiyama, T. (2019). Area-Specificity and Plasticity of History-Dependent Value Coding During Learning. *Cell* 177, 1858-1872.e15. doi: 10.1016/j.cell.2019.04.027.

Hayashi-Takagi, A., Yagishita, S., Nakamura, M., Shirai, F., Wu, Y.I., Loshbaugh, A.L., Kuhlman, B., Hahn, K.M., and Kasai, H. (2015). Labelling and optical erasure of synaptic memory traces in the motor cortex. *Nature* 525, 333-338. 10.1038/nature15257.

Hnasko, T.S., Chuhma, N., Zhang, H., Goh, G.Y., Sulzer, D., Palmiter, R.D., Rayport, S., and Edwards, R.H. (2010). Vesicular glutamate transport promotes dopamine storage and glutamate corelease *in vivo*. *Neuron* 65, 643-656. 10.1016/j.neuron.2010.02.012.

Hooks, B.M. (2018). Dual-Channel Photostimulation for Independent Excitation of Two Populations. *Curr Protoc Neurosci* 85, e52. 10.1002/cpns.52.

Hooks, B.M., Hires, S.A., Zhang, Y.X., Huber, D., Petreanu, L., Svoboda, K., and Shepherd, G.M. (2011). Laminar analysis of excitatory local circuits in vibrissal motor and sensory cortical areas. *PLoS Biol* 9, e1000572. 10.1371/journal.pbio.1000572.

Hooks, B.M., Lin, J.Y., Guo, C., and Svoboda, K. (2015). Dual-channel circuit mapping reveals sensorimotor convergence in the primary motor cortex. *J Neurosci* 35, 4418-4426. 10.1523/JNEUROSCI.3741-14.2015.

Hooks, B.M., Mao, T., Gutnisky, D.A., Yamawaki, N., Svoboda, K., and Shepherd, G.M. (2013). Organization of cortical and thalamic input to pyramidal neurons in mouse motor cortex. *J Neurosci* 33, 748-760. 10.1523/JNEUROSCI.4338-12.2013.

Hosp, J.A., Pekanovic, A., Rioult-Pedotti, M.S., and Luft, A.R. (2011). Dopaminergic projections from midbrain to primary motor cortex mediate motor skill learning. *J Neurosci* 31, 2481-2487. 10.1523/JNEUROSCI.5411-10.2011.

Hubel, D.H., and Wiesel, T.N. (1959). Receptive fields of single neurones in the cat's striate cortex. *J Physiol* 148, 574-591. 10.1113/jphysiol.1959.sp006308.

Hubel, D.H., and Wiesel, T.N. (1963). Shape and arrangement of columns in cat's striate cortex. *J Physiol* 165, 559-568. 10.1113/jphysiol.1963.sp007079.

Hubel, D.H., and Wiesel, T.N. (1968). Receptive fields and functional architecture of monkey striate cortex. *J Physiol* 195, 215-243. 10.1113/jphysiol.1968.sp008455.

Hunnicutt, B. J., Long, B. R., Kusefoglu, D., Gertz, K. J., Zhong, H., and Mao, T. (2014). A comprehensive thalamocortical projection map at the mesoscopic level. *Nat Neurosci* 17, 1276–1285. doi: 10.1038/nn.3780.

Inoue, K.I., Takada, M., and Matsumoto, M. (2015). Neuronal and behavioural modulations by pathway-selective optogenetic stimulation of the primate oculomotor system. *Nat Commun* 6, 8378. 10.1038/ncomms9378.

Iriki, A., Pavlides, C., Keller, A., and Asanuma, H. (1989). Long-term potentiation in the motor cortex. *Science* 245, 1385-1387. 10.1126/science.2551038.

Isa, T., Ohki, Y., Seki, K., and Alstermark, B. (2006). Properties of propriospinal neurons in the C3-C4 segments mediating disynaptic pyramidal excitation to forelimb motoneurons in the macaque monkey. *J Neurophysiol* 95, 3674-3685. 10.1152/jn.00103.2005.

Jackman, S.L., Beneduce, B.M., Drew, I.R., and Regehr, W.G. (2014). Achieving high-frequency optical control of synaptic transmission. *J Neurosci* 34, 7704-7714. 10.1523/JNEUROSCI.4694-13.2014.

Jackson, J., Ayzenshtat, I., Karnani, M.M., and Yuste, R. (2016). VIP+ interneurons control neocortical activity across brain states. *J Neurophysiol* 115, 3008-3017. 10.1152/jn.01124.2015.

Jazayeri, M., Lindbloom-Brown, Z., and Horwitz, G.D. (2012). Saccadic eye movements evoked by optogenetic activation of primate V1. *Nat Neurosci* 15, 1368-1370. 10.1038/nn.3210.

Jensen, K., and Killackey, H. (1987). Terminal arbors of axons projecting to the somatosensory cortex of the adult rat. II. The altered morphology of thalamocortical afferents following neonatal infraorbital nerve cut. *J. Neurosci.* 7, 3544–3553. doi: 10.1523/JNEUROSCI.07-11-03544.1987.

Jing, M., Li, Y., Zeng, J., Huang, P., Skirzewski, M., Kljakic, O., Peng, W., Qian, T., Tan, K., Zou, J., et al. (2020). An optimized acetylcholine sensor for monitoring *in vivo* cholinergic activity. *Nat Methods* 17, 1139-1146. 10.1038/s41592-020-0953-2.

Jing, M., Zhang, P., Wang, G., Feng, J., Mesik, L., Zeng, J., Jiang, H., Wang, S., Looby, J.C., Guagliardo, N.A., et al. (2018). A genetically encoded fluorescent acetylcholine indicator for *in vitro* and *in vivo* studies. *Nat Biotechnol* 36, 726-737. 10.1038/nbt.4184.

Jones, T.A. (2017). Motor compensation and its effects on neural reorganization after stroke. *Nat Rev Neurosci* 18, 267-280. 10.1038/nrn.2017.26.

Kalaska, J.F., Cohen, D.A., Hyde, M.L., and Prud'homme, M. (1989). A comparison of movement direction-related versus load direction-related activity in primate motor cortex, using a two-dimensional reaching task. *J Neurosci* 9, 2080-2102.

Kapogiannis, D., Campion, P., Grafman, J., and Wassermann, E.M. (2008). Reward-related activity in the human motor cortex. *Eur J Neurosci* 27, 1836-1842. 10.1111/j.1460-9568.2008.06147.x.

Karnani, M.M., Jackson, J., Ayzenshtat, I., Hamzehei Sichani, A., Manoocheri, K., Kim, S., and Yuste, R. (2016). Opening Holes in the Blanket of Inhibition: Localized Lateral Disinhibition by VIP Interneurons. *J Neurosci* 36, 3471-3480. 10.1523/JNEUROSCI.3646-15.2016.

Kato, H.K., Chu, M.W., Isaacson, J.S., and Komiyama, T. (2012). Dynamic sensory representations in the olfactory bulb: modulation by wakefulness and experience. *Neuron* 76, 962-975. 10.1016/j.neuron.2012.09.037.

Kato, H.K., Gillet, S.N., and Isaacson, J.S. (2015). Flexible Sensory Representations in Auditory Cortex Driven by Behavioral Relevance. *Neuron* 88, 1027-1039. 10.1016/j.neuron.2015.10.024.

Kawaguchi, Y., and Kubota, Y. (1996). Physiological and morphological identification of somatostatin- or vasoactive intestinal polypeptide-containing cells among GABAergic cell subtypes in rat frontal cortex. *J Neurosci* 16, 2701-2715.

Kawaguchi, Y., and Kubota, Y. (1997). GABAergic cell subtypes and their synaptic connections in rat frontal cortex. *Cereb Cortex* 7, 476-486. 10.1093/cercor/7.6.476.

Kawai, R., Markman, T., Poddar, R., Ko, R., Fantana, A. L., Dhawale, A. K., et al. (2015). Motor Cortex Is Required for Learning but Not for Executing a Motor Skill. *Neuron* 86, 800–812. doi: 10.1016/j.neuron.2015.03.024.

Kisley, M. A., and Gerstein, G. L. (2001). Daily variation and appetitive conditioning-induced plasticity of auditory cortex receptive fields: Variation and plasticity of receptive fields. *European Journal of Neuroscience* 13, 1993–2003. doi: 10.1046/j.0953-816x.2001.01568.x.

Klapoetke, N.C., Murata, Y., Kim, S.S., Pulver, S.R., Birdsey-Benson, A., Cho, Y.K., Morimoto, T.K., Chuong, A.S., Carpenter, E.J., Tian, Z., et al. (2014). Independent optical excitation of distinct neural populations. *Nat Methods* 11, 338-346. 10.1038/nmeth.2836.

Klauer, S., Sengpiel, F., and Hoffmann, K.P. (1990). Visual response properties and afferents of nucleus of the optic tract in the ferret. Exp *Brain Res* 83, 178-189. 10.1007/bf00232207.

Klein, C., Evrard, H.C., Shapcott, K.A., Haverkamp, S., Logothetis, N.K., and Schmid, M.C. (2016). Cell-Targeted Optogenetics and Electrical Microstimulation Reveal the Primate Koniocellular Projection to Supra-granular Visual Cortex. *Neuron* 90, 143-151. 10.1016/j.neuron.2016.02.036.

Koch, C., Poggio, T., and Torre, V. (1983). Nonlinear interactions in a dendritic tree: localization, timing, and role in information processing. *Proc Natl Acad Sci U S A* 80, 2799-2802. 10.1073/pnas.80.9.2799.

Kole, K., Zhang, Y., Jansen, E.J.R., Brouns, T., Bijlsma, A., Calcini, N., Yan, X., Lantyer, A.D.S., and Celikel, T. (2020). Assessing the utility of Magneto to control neuronal excitability in the somatosensory cortex. *Nat Neurosci* 23, 1044-1046. 10.1038/s41593-019-0474-4.

Komiyama, T., Sato, T.R., O'Connor, D.H., Zhang, Y.X., Huber, D., Hooks, B.M., Gabitto, M., and Svoboda, K. (2010). Learning-related fine-scale specificity imaged in motor cortex circuits of behaving mice. *Nature* 464, 1182-1186. 10.1038/nature08897.

Krabbe, S., Paradiso, E., d'Aquin, S., Bitterman, Y., Courtin, J., Xu, C., Yonehara, K., Markovic, M., Müller, C., Eichlisberger, T., et al. (2019). Adaptive disinhibitory gating by VIP interneurons permits associative learning. *Nat Neurosci* 22, 1834-1843. 10.1038/s41593-019-0508-y.

Kranz, G.S., Kasper, S., and Lanzenberger, R. (2010). Reward and the serotonergic system. *Neuroscience* 166, 1023-1035. 10.1016/j.neuroscience.2010.01.036.

Kwon, S. E., Yang, H., Minamisawa, G., and O'Connor, D. H. (2016). Sensory and decision-related activity propagate in a cortical feedback loop during touch perception. *Nat Neurosci* 19, 1243–1249. doi: 10.1038/nn.4356.

Lacefield, C.O., Pnevmatikakis, E.A., Paninski, L., and Bruno, R.M. (2019). Reinforcement Learning Recruits Somata and Apical Dendrites across Layers of Primary Sensory Cortex. *Cell Rep* 26, 2000-2008.e2002. 10.1016/j.celrep.2019.01.093.

Lam, Y.-W., and Sherman, S. M. (2015). Functional topographic organization of the motor reticulothalamic pathway. *Journal of Neurophysiology* 113, 3090–3097. doi: 10.1152/jn.00847.2014.

Larkum, M.E., Zhu, J.J., and Sakmann, B. (1999). A new cellular mechanism for coupling inputs arriving at different cortical layers. *Nature* 398, 338-341. 10.1038/18686.

Lavin, A., Nogueira, L., Lapish, C.C., Wightman, R.M., Phillips, P.E., and Seamans, J.K. (2005). Mesocortical dopamine neurons operate in distinct temporal domains using multimodal signaling. *J Neurosci* 25, 5013-5023. 10.1523/JNEUROSCI.0557-05.2005.

Lawrence DG, Kuypers HG (1968) The functional organization of the motor system in the monkey. II. The effects of lesions of the descending brain-stem pathways. *Brain* 91:15-36.

Lee, C. C., and Sherman, S. M. (2008). Synaptic Properties of Thalamic and Intracortical Inputs to Layer 4 of the First- and Higher-Order Cortical Areas in the Auditory and

Somatosensory Systems. *Journal of Neurophysiology* 100, 317–326. doi: 10.1152/jn.90391.2008.

Lee, C., Lavoie, A., Liu, J., Chen, S.X., and 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. 10.3389/fncir.2020.00018.

Lee, J., and Rothschild, G. (2021). Encoding of acquired sound-sequence salience by auditory cortical offset responses. *Cell Reports* 37, 109927. doi: 10.1016/j.celrep.2021.109927.

Lee, K., Holley, S.M., Shobe, J.L., Chong, N.C., Cepeda, C., Levine, M.S., and Masmanidis, S.C. (2017). Parvalbumin Interneurons Modulate Striatal Output and Enhance Performance during Associative Learning. *Neuron* 93, 1451-1463.e1454. 10.1016/j.neuron.2017.02.033.

Lee, S., Hjerling-Leffler, J., Zagha, E., Fishell, G., and Rudy, B. (2010). The largest group of superficial neocortical GABAergic interneurons expresses ionotropic serotonin receptors. *J Neurosci* 30, 16796-16808. 10.1523/JNEUROSCI.1869-10.2010.

Lee, S., Kruglikov, I., Huang, Z.J., Fishell, G., and Rudy, B. (2013). A disinhibitory circuit mediates motor integration in the somatosensory cortex. *Nat Neurosci* 16, 1662-1670. 10.1038/nn.3544.

Lee, S.H., Kwan, A.C., Zhang, S., Phoumthipphavong, V., Flannery, J.G., Masmanidis, S.C., Taniguchi, H., Huang, Z.J., Zhang, F., Boyden, E.S., et al. (2012). Activation of specific interneurons improves V1 feature selectivity and visual perception. *Nature* 488, 379-383. 10.1038/nature11312.

Legenstein, R., Chase, S.M., Schwartz, A.B., and Maass, W. (2010). A reward-modulated hebbian learning rule can explain experimentally observed network reorganization in a brain control task. *J Neurosci* 30, 8400-8410. 10.1523/JNEUROSCI.4284-09.2010.

Letzkus, J.J., Wolff, S.B., Meyer, E.M., Tovote, P., Courtin, J., Herry, C., and Lüthi, A. (2011). A disinhibitory microcircuit for associative fear learning in the auditory cortex. *Nature* 480, 331-335. 10.1038/nature10674.

Levy, S., Lavzin, M., Benisty, H., Ghanayim, A., Dubin, U., Achvat, S., Brosh, Z., Aeed, F., Mensh, B.D., Schiller, Y., et al. (2020). Cell-Type-Specific Outcome Representation in the Primary Motor Cortex. *Neuron* 107, 954-971.e959. 10.1016/j.neuron.2020.06.006.

Li, N., Chen, T.-W., Guo, Z. V., Gerfen, C. R., and Svoboda, K. (2015). A motor cortex circuit for motor planning and movement. *Nature* 519, 51–56. doi: 10.1038/nature14178.

Lien, A.D., and Scanziani, M. (2013). Tuned thalamic excitation is amplified by visual cortical circuits. *Nat Neurosci* 16, 1315-1323. 10.1038/nn.3488.

Lima, S.Q., Hromadka, T., Znamenskiy, P., and Zador, A.M. (2009). PINP: a new method of tagging neuronal populations for identification during *in vivo* electrophysiological recording. *PLoS One* 4, e6099. 10.1371/journal.pone.0006099.

Lin, J.Y. (2011). A user's guide to channelrhodopsin variants: features, limitations and future developments. *Exp Physiol* 96, 19-25. 10.1113/expphysiol.2009.051961.

Lin, J.Y., Knutsen, P.M., Muller, A., Kleinfeld, D., and Tsien, R.Y. (2013). ReaChR: a red-shifted variant of channelrhodopsin enables deep transcranial optogenetic excitation. *Nat Neurosci* 16, 1499-1508. 10.1038/nn.3502.

Liu, B.H., Huberman, A.D., and Scanziani, M. (2016). Cortico-fugal output from visual cortex promotes plasticity of innate motor behaviour. *Nature* 538, 383-387. 10.1038/nature19818.

Liu, B.H., Li, Y.T., Ma, W.P., Pan, C.J., Zhang, L.I., and Tao, H.W. (2011). Broad inhibition sharpens orientation selectivity by expanding input dynamic range in mouse simple cells. *Neuron* 71, 542-554. 10.1016/j.neuron.2011.06.017.

Liu, D., Deng, J., Zhang, Z., Zhang, Z.-Y., Sun, Y.-G., Yang, T., et al. (2020). Orbitofrontal control of visual cortex gain promotes visual associative learning. *Nat Commun* 11, 2784. doi: 10.1038/s41467-020-16609-7.

Lübke, J., and Feldmeyer, D. (2007). Excitatory signal flow and connectivity in a cortical column: focus on barrel cortex. *Brain Struct Funct* 212, 3-17. 10.1007/s00429-007-0144-2.

Lucas, E.K., and Clem, R.L. (2018). GABAergic interneurons: The orchestra or the conductor in fear learning and memory? *Brain Res* Bull 141, 13-19. 10.1016/j.brainresbull.2017.11.016.

Luft, A.R., Buitrago, M.M., Ringer, T., Dichgans, J., and Schulz, J.B. (2004). Motor skill learning depends on protein synthesis in motor cortex after training. *J Neurosci* 24, 6515-6520. 10.1523/JNEUROSCI.1034-04.2004.

Lukoyanov, N., and Lukoyanova, E. (2006). Retrosplenial cortex lesions impair acquisition of active avoidance while sparing fear-based emotional memory. *Behavioural Brain Research* 173, 229–236. doi: 10.1016/j.bbr.2006.06.026.

Luo, P., Li, A., Zheng, Y., Han, Y., Tian, J., Xu, Z., Gong, H., and Li, X. (2019). Whole Brain Mapping of Long-Range Direct Input to Glutamatergic and GABAergic Neurons in Motor Cortex. *Front Neuroanat* 13, 44. 10.3389/fnana.2019.00044.

Magnus, C.J., Lee, P.H., Atasoy, D., Su, H.H., Looger, L.L., and Sternson, S.M. (2011). Chemical and genetic engineering of selective ion channel-ligand interactions. *Science* 333, 1292-1296. 10.1126/science.1206606.

Mahler, S.V., Vazey, E.M., Beckley, J.T., Keistler, C.R., McGlinchey, E.M., Kaufling, J., Wilson, S.P., Deisseroth, K., Woodward, J.J., and Aston-Jones, G. (2014). Designer receptors show role for ventral pallidum input to ventral tegmental area in cocaine seeking. *Nat Neurosci* 17, 577-585. 10.1038/nn.3664.

Mahn, M., Prigge, M., Ron, S., Levy, R., and Yizhar, O. (2016). Biophysical constraints of optogenetic inhibition at presynaptic terminals. *Nat Neurosci* 19, 554-556. 10.1038/nn.4266.

Makino, H., and Komiyama, T. (2015). Learning enhances the relative impact of top-down processing in the visual cortex. *Nat Neurosci* 18, 1116-1122. 10.1038/nn.4061.

Makino, H., Ren, C., Liu, H., Kim, A.N., Kondapaneni, N., Liu, X., Kuzum, D., and Komiyama, T. (2017). Transformation of Cortex-wide Emergent Properties during Motor Learning. *Neuron* 94, 880-890.e888. 10.1016/j.neuron.2017.04.015.

Manita, S., Suzuki, T., Homma, C., Matsumoto, T., Odagawa, M., Yamada, K., Ota, K., Matsubara, C., Inutsuka, A., Sato, M., et al. (2015). A Top-Down Cortical Circuit for Accurate Sensory Perception. *Neuron* 86, 1304-1316. 10.1016/j.neuron.2015.05.006.

Manvich, D.F., Webster, K.A., Foster, S.L., Farrell, M.S., Ritchie, J.C., Porter, J.H., and Weinshenker, D. (2018). The DREADD agonist clozapine N-oxide (CNO) is reverse-metabolized to clozapine and produces clozapine-like interoceptive stimulus effects in rats and mice. *Sci Rep* 8, 3840. 10.1038/s41598-018-22116-z.

Mao, T., Kusefoglu, D., Hooks, B.M., Huber, D., Petreanu, L., and Svoboda, K. (2011). Long-range neuronal circuits underlying the interaction between sensory and motor cortex. *Neuron* 72, 111-123. 10.1016/j.neuron.2011.07.029.

Markram, H., Lübke, J., Frotscher, M., and Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* 275, 213-215. 10.1126/science.275.5297.213.

Markram, H., Toledo-Rodriguez, M., Wang, Y., Gupta, A., Silberberg, G., and Wu, C. (2004). Interneurons of the neocortical inhibitory system. *Nat Rev Neurosci* 5, 793-807. 10.1038/nrn1519.

Marsh, B.T., Tarigoppula, V.S., Chen, C., and Francis, J.T. (2015). Toward an autonomous brain machine interface: integrating sensorimotor reward modulation and reinforcement learning. *J Neurosci* 35, 7374-7387. 10.1523/JNEUROSCI.1802-14.2015.

Matyas, F., Sreenivasan, V., Marbach, F., Wacongne, C., Barsy, B., Mateo, C., Aronoff, R., and Petersen, C.C. (2010). Motor control by sensory cortex. *Science* 330, 1240-1243. 10.1126/science.1195797.

Mayrhofer JM, El-Boustani S, Foustoukos G, Auffret M, Tamura K, Petersen CCH (2019) Distinct Contributions of Whisker Sensory Cortex and Tongue-Jaw Motor Cortex in a Goal-Directed Sensorimotor Transformation. *Neuron* 103:1034-1043.e1035.

Meister, M. (2016). Physical limits to magnetogenetics. *Elife* 5. 10.7554/eLife.17210.

Meyer, H.S., Schwarz, D., Wimmer, V.C., Schmitt, A.C., Kerr, J.N., Sakmann, B., and Helmstaedter, M. (2011). Inhibitory interneurons in a cortical column form hot zones of inhibition in layers 2 and 5A. *Proc Natl Acad Sci U S A* 108, 16807-16812. 10.1073/pnas.1113648108.

Mingote, S., Chuhma, N., Kalmbach, A., Thomsen, G.M., Wang, Y., Mihali, A., Sferrazza, C., Zucker-Scharff, I., Siena, A.C., Welch, M.G., et al. (2017). Dopamine neuron dependent behaviors mediated by glutamate cotransmission. *Elife* 6. 10.7554/eLife.27566.

Mingote, S., Chuhma, N., Kusnoor, S.V., Field, B., Deutch, A.Y., and Rayport, S. (2015). Functional Connectome Analysis of Dopamine Neuron Glutamatergic Connections in Forebrain Regions. *J Neurosci* 35, 16259-16271. 10.1523/JNEUROSCI.1674-15.2015.

Mitchell, C.J., De Houwer, J., and Lovibond, P.F. (2009). The propositional nature of human associative learning. *Behav Brain Sci* 32, 183-198; discussion 198-246. 10.1017/S0140525X09000855.

Molina-Luna, K., Pekanovic, A., Röhrich, S., Hertler, B., Schubring-Giese, M., Rioult-Pedotti, M.S., and Luft, A.R. (2009). Dopamine in motor cortex is necessary for skill learning and synaptic plasticity. *PLoS One* 4, e7082. 10.1371/journal.pone.0007082.

Moran, D.W., and Schwartz, A.B. (1999). Motor cortical representation of speed and direction during reaching. *J Neurophysiol* 82, 2676-2692. 10.1152/jn.1999.82.5.2676.

Mukherjee, A., Lam, N. H., Wimmer, R. D., and Halassa, M. M. (2021). Thalamic circuits for independent control of prefrontal signal and noise. *Nature* 600, 100–104. doi: 10.1038/s41586-021-04056-3.

Mukohata, Y., Sugiyama, Y., Ihara, K., and Yoshida, M. (1988). An Australian halobacterium contains a novel proton pump retinal protein: archaerhodopsin. *Biochem Biophys Res Commun* 151, 1339-1345. 10.1016/s0006-291x(88)80509-6.

Muñoz, W., Tremblay, R., Levenstein, D., and Rudy, B. (2017). Layer-specific modulation of neocortical dendritic inhibition during active wakefulness. *Science* 355, 954-959. 10.1126/science.aag2599.

Muñoz-Castañeda, R., Zingg, B., Matho, K.S., Chen, X., Wang, Q., Foster, N.N., Li, A., Narasimhan, A., Hirokawa, K.E., Huo, B., et al. (2021). Cellular anatomy of the mouse primary motor cortex. *Nature* 598, 159-166. 10.1038/s41586-021-03970-w.

Murphy, T.H., and Corbett, D. (2009). Plasticity during stroke recovery: from synapse to behaviour. *Nat Rev Neurosci* 10, 861-872. 10.1038/nrn2735.

Nabavi, S., Fox, R., Proulx, C.D., Lin, J.Y., Tsien, R.Y., and Malinow, R. (2014). Engineering a memory with LTD and LTP. *Nature* 511, 348-352. 10.1038/nature13294.

Nagel, G., Szellas, T., Huhn, W., Kateriya, S., Adeishvili, N., Berthold, P., Ollig, D., Hegemann, P., and Bamberg, E. (2003). Channelrhodopsin-2, a directly light-gated cation-selective membrane channel. *Proc Natl Acad Sci U S A* 100, 13940-13945. 10.1073/pnas.1936192100.

Nakamichi, Y., Okubo, K., Sato, T., Hashimoto, M., and Tanifuji, M. (2019). Optical intrinsic signal imaging with optogenetics reveals functional cortico-cortical connectivity at the columnar level in living macaques. *Sci Rep* 9, 6466. 10.1038/s41598-019-42923-2.

Namboodiri, V. M. K., Otis, J. M., van Heeswijk, K., Voets, E. S., Alghorazi, R. A., Rodriguez-Romaguera, J., et al. (2019). Single-cell activity tracking reveals that orbitofrontal neurons acquire and maintain a long-term memory to guide behavioral adaptation. *Nat Neurosci* 22, 1110–1121. doi: 10.1038/s41593-019-0408-1.

Naskar, S., Qi, J., Pereira, F., Gerfen, C.R., and Lee, S. (2021). Cell-type-specific recruitment of GABAergic interneurons in the primary somatosensory cortex by long-range inputs. *Cell Rep* 34, 108774. 10.1016/j.celrep.2021.108774.

Nikooyan, A.A., and Ahmed, A.A. (2015). Reward feedback accelerates motor learning. *J Neurophysiol* 113, 633-646. 10.1152/jn.00032.2014.

Noonan, M. P., Walton, M. E., Behrens, T. E. J., Sallet, J., Buckley, M. J., and Rushworth, M. F. S. (2010). Separate value comparison and learning mechanisms in macaque medial and lateral orbitofrontal cortex. *Proc. Natl. Acad. Sci. U.S.A.* 107, 20547–20552. doi: 10.1073/pnas.1012246107.

O'Connor, D.H., Hires, S.A., Guo, Z.V., Li, N., Yu, J., Sun, Q.Q., Huber, D., and Svoboda, K. (2013). Neural coding during active somatosensation revealed using illusory touch. *Nat Neurosci* 16, 958-965. 10.1038/nn.3419.

O'Connor, D.H., Peron, S.P., Huber, D., and Svoboda, K. (2010). Neural activity in barrel cortex underlying vibrissa-based object localization in mice. *Neuron* 67, 1048-1061. 10.1016/j.neuron.2010.08.026.

Oh, S.W., Harris, J.A., Ng, L., Winslow, B., Cain, N., Mihalas, S., Wang, Q., Lau, C., Kuan, L., Henry, A.M., et al. (2014). A mesoscale connectome of the mouse brain. *Nature* 508, 207-214. 10.1038/nature13186.

Ölveczky, B.P. (2011). Motoring ahead with rodents. *Curr Opin Neurobiol* 21, 571-578. 10.1016/j.conb.2011.05.002.

Ottenheimer, D.J. Hjort, M.M., Bowen A.J., Steinmetz, N.A., Stuber G.D., A stable, distributed code for cue value in mouse cortex during reward learning. *bioRxiv* 2022.07.13.499930; 10.1101/2022.07.13.499930

Pala, A., and Petersen, C.C.H. (2015). *In vivo* measurement of cell-type-specific synaptic connectivity and synaptic transmission in layer 2/3 mouse barrel cortex. *Neuron* 85, 68-75. 10.1016/j.neuron.2014.11.025.

Palminteri, S., Lebreton, M., Worbe, Y., Hartmann, A., Lehéricy, S., Vidailhet, M., Grabli, D., and Pessiglione, M. (2011). Dopamine-dependent reinforcement of motor skill learning: evidence from Gilles de la Tourette syndrome. *Brain* 134, 2287-2301. 10.1093/brain/awr147.

Papagiakoumou, E., Anselmi, F., Bègue, A., de Sars, V., Glückstad, J., Isacoff, E.Y., and Emiliani, V. (2010). Scanless two-photon excitation of channelrhodopsin-2. *Nat Methods* 7, 848-854. 10.1038/nmeth.1505.

Papale, A.E., and Hooks, B.M. (2018). Circuit changes in motor cortex during motor skill learning. *Neuroscience* 368, 283-297. 10.1016/j.neuroscience.2017.09.010.

Pardi, M.B., Vogenstahl, J., Dalmay, T., Spanò, T., Pu, D.L., Naumann, L.B., Kretschmer, F., Sprekeler, H., and Letzkus, J.J. (2020). A thalamocortical top-down circuit for associative memory. *Science* 370, 844-848. 10.1126/science.abc2399.

Parkes, S. L., Ravassard, P. M., Cerpa, J.-C., Wolff, M., Ferreira, G., and Coutureau, E. (2018). Insular and Ventrolateral Orbitofrontal Cortices Differentially Contribute to Goal-Directed Behavior in Rodents. *Cerebral Cortex* 28, 2313–2325. doi: 10.1093/cercor/bhx132.

Perez-Lopez, J.L., Contreras-Lopez, R., Ramirez-Jarquin, J.O., and Tecuapetla, F. (2018). Direct Glutamatergic Signaling From Midbrain Dopaminergic Neurons Onto Pyramidal Prefrontal Cortex Neurons. *Front Neural Circuits* 12, 70. 10.3389/fncir.2018.00070.

Peters, A.J., Chen, S.X., and Komiyama, T. (2014). Emergence of reproducible spatiotemporal activity during motor learning. *Nature* 510, 263-267. 10.1038/nature13235.

Peters, A.J., Lee, J., Hedrick, N.G., O'Neil, K., and Komiyama, T. (2017). Reorganization of corticospinal output during motor learning. *Nat Neurosci* 20, 1133-1141. 10.1038/nn.4596.

Petreanu, L., Huber, D., Sobczyk, A., and Svoboda, K. (2007). Channelrhodopsin-2-assisted circuit mapping of long-range callosal projections. *Nat Neurosci* 10, 663-668. 10.1038/nn1891.

Petreanu, L., Mao, T., Sternson, S.M., and Svoboda, K. (2009). The subcellular organization of neocortical excitatory connections. *Nature* 457, 1142-1145. 10.1038/nature07709.

Pfeffer, C.K., Xue, M., He, M., Huang, Z.J., and Scanziani, M. (2013). Inhibition of inhibition in visual cortex: the logic of connections between molecularly distinct interneurons. *Nat Neurosci* 16, 1068-1076. 10.1038/nn.3446.

Pi, H.J., Hangya, B., Kvitsiani, D., Sanders, J.I., Huang, Z.J., and Kepecs, A. (2013). Cortical interneurons that specialize in disinhibitory control. *Nature* 503, 521-524. 10.1038/nature12676.

Pierret, T., Lavallée, P., and Deschênes, M. (2000). Parallel Streams for the Relay of Vibrissal Information through Thalamic Barreloids. *J. Neurosci.* 20, 7455–7462. doi: 10.1523/JNEUROSCI.20-19-07455.2000.

Polter, A.M., Barcomb, K., Tsuda, A.C., and Kauer, J.A. (2018). Synaptic function and plasticity in identified inhibitory inputs onto VTA dopamine neurons. *Eur J Neurosci* 47, 1208-1218. 10.1111/ejn.13879.

Poort, J., Wilmes, K.A., Blot, A., Chadwick, A., Sahani, M., Clopath, C., Mrsic-Flogel, T.D., Hofer, S.B., and Khan, A.G. (2021). Learning and attention increase visual response selectivity through distinct mechanisms. *Neuron*. 10.1016/j.neuron.2021.11.016.

Prakash, R., Yizhar, O., Grewe, B., Ramakrishnan, C., Wang, N., Goshen, I., Packer, A.M., Peterka, D.S., Yuste, R., Schnitzer, M.J., and Deisseroth, K. (2012). Two-photon optogenetic toolbox for fast inhibition, excitation and bistable modulation. *Nat Methods* 9, 1171-1179. 10.1038/nmeth.2215.

Prönneke, A., Scheuer, B., Wagener, R.J., Möck, M., Witte, M., and Staiger, J.F. (2015). Characterizing VIP Neurons in the Barrel Cortex of VIPcre/tdTomato Mice Reveals Layer-Specific Differences. *Cereb Cortex* 25, 4854-4868. 10.1093/cercor/bhv202.

Raimondo, J.V., Kay, L., Ellender, T.J., and Akerman, C.J. (2012). Optogenetic silencing strategies differ in their effects on inhibitory synaptic transmission. *Nat Neurosci* 15, 1102-1104. 10.1038/nn.3143.

Ramakrishnan, A., Byun, Y.W., Rand, K., Pedersen, C.E., Lebedev, M.A., and 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, E4841-E4850. 10.1073/pnas.1703668114.

Ramkumar, P., Dekleva, B., Cooler, S., Miller, L., and Kording, K. (2016). Premotor and Motor Cortices Encode Reward. *PLoS One* 11, e0160851. 10.1371/journal.pone.0160851.

Ranade, S.P., and Mainen, Z.F. (2009). Transient firing of dorsal raphe neurons encodes diverse and specific sensory, motor, and reward events. *J Neurophysiol* 102, 3026-3037. 10.1152/jn.00507.2009.

Reinhold, K., Lien, A.D., and Scanziani, M. (2015). Distinct recurrent versus afferent dynamics in cortical visual processing. *Nat Neurosci* 18, 1789-1797. 10.1038/nn.4153.

Ren, C., Peng, K., Yang, R., Liu, W., Liu, C., and Komiyama, T. (2022). Global and subtype-specific modulation of cortical inhibitory neurons regulated by acetylcholine during motor learning. *Neuron* 110, 2334-2350.e2338. 10.1016/j.neuron.2022.04.031.

Richards, B.A., Lillicrap, T.P., Beaudoin, P., Bengio, Y., Bogacz, R., Christensen, A., Clopath, C., Costa, R.P., de Berker, A., Ganguli, S., et al. (2019). A deep learning framework for neuroscience. *Nat Neurosci* 22, 1761-1770. 10.1038/s41593-019-0520-2.

Rickgauer, J.P., and Tank, D.W. (2009). Two-photon excitation of channelrhodopsin-2 at saturation. *Proc Natl Acad Sci U S A* 106, 15025-15030. 10.1073/pnas.0907084106.

Ringach, D.L., Hawken, M.J., and Shapley, R. (2003). Dynamics of orientation tuning in macaque V1: the role of global and tuned suppression. *J Neurophysiol* 90, 342-352. 10.1152/jn.01018.2002.

Rioult-Pedotti, M.S., Friedman, D., Hess, G., and Donoghue, J.P. (1998). Strengthening of horizontal cortical connections following skill learning. *Nat Neurosci* 1, 230-234. 10.1038/678.

Ronzitti, E., Conti, R., Zampini, V., Tanese, D., Foust, A.J., Klapoetke, N., Boyden, E.S., Papagiakoumou, E., and Emiliani, V. (2017). Submillisecond Optogenetic Control of Neuronal Firing with Two-Photon Holographic Photoactivation of Chronos. *J Neurosci* 37, 10679-10689. 10.1523/JNEUROSCI.1246-17.2017.

Sakamoto, T., Porter, L.L., and Asanuma, H. (1987). Long-lasting potentiation of synaptic potentials in the motor cortex produced by stimulation of the sensory cortex in the cat: a basis of motor learning. *Brain Res* 413, 360-364. 10.1016/0006-8993(87)91029-8.

Sato, H., Katsuyama, N., Tamura, H., Hata, Y., and Tsumoto, T. (1996). Mechanisms underlying orientation selectivity of neurons in the primary visual cortex of the macaque. *J Physiol* 494 (Pt 3), 757-771. 10.1113/jphysiol.1996.sp021530.

Schindelin, J., Arganda-Carreras, I., Frise, E., Kaynig, V., Longair, M., Pietzsch, T., et al. (2012). Fiji: an open-source platform for biological-image analysis. *Nat Methods* 9, 676–682. doi: 10.1038/nmeth.2019.

Schoenfeld, G., Kollmorgen, S., Lewis, C., Bethge P., Reuss, A.M, Aguzzi, A., Mante, V., Helmchen F. (2021). Dendritic integration of sensory and reward information facilitates learning. *bioRxiv* 2021.12.28.474360; 10.1101/2021.12.28.474360

Schultz, W. (2000). Multiple reward signals in the brain. *Nat Rev Neurosci* 1, 199-207. 10.1038/35044563.

Schuman, B., Machold, R.P., Hashikawa, Y., Fuzik, J., Fishell, G.J., and Rudy, B. (2019). Four Unique Interneuron Populations Reside in Neocortical Layer 1. *J Neurosci* 39, 125-139. 10.1523/JNEUROSCI.1613-18.2018.

Schwarz, L.A., Miyamichi, K., Gao, X.J., Beier, K.T., Weissbourd, B., DeLoach, K.E., Ren, J., Ibanes, S., Malenka, R.C., Kremer, E.J., and Luo, L. (2015). Viral-genetic tracing of the input-output organization of a central noradrenaline circuit. *Nature* 524, 88-92. 10.1038/nature14600.

Schweighofer, N., Bertin, M., Shishida, K., Okamoto, Y., Tanaka, S.C., Yamawaki, S., and Doya, K. (2008). Low-serotonin levels increase delayed reward discounting in humans. *J Neurosci* 28, 4528-4532. 10.1523/JNEUROSCI.4982-07.2008.

Seeman, S.C., Campagnola, L., Davoudian, P.A., Hoggarth, A., Hage, T.A., Bosma-Moody, A., Baker, C.A., Lee, J.H., Mihalas, S., Teeter, C., et al. (2018). Sparse recurrent excitatory connectivity in the microcircuit of the adult mouse and human cortex. *Elife* 7. 10.7554/eLife.37349.

Senn, V., Wolff, S.B., Herry, C., Grenier, F., Ehrlich, I., Gründemann, J., Fadok, J.P., Müller, C., Letzkus, J.J., and Lüthi, A. (2014). Long-range connectivity defines behavioral specificity of amygdala neurons. *Neuron* 81, 428-437. 10.1016/j.neuron.2013.11.006.

Seybold, B.A., Phillips, E.A.K., Schreiner, C.E., and Hasenstaub, A.R. (2015). Inhibitory Actions Unified by Network Integration. *Neuron* 87, 1181-1192. 10.1016/j.neuron.2015.09.013.

Shemesh, O.A., Tanese, D., Zampini, V., Linghu, C., Piatkevich, K., Ronzitti, E., Papagiakoumou, E., Boyden, E.S., and Emiliani, V. (2017). Temporally precise single-cell-resolution optogenetics. *Nat Neurosci* 20, 1796-1806. 10.1038/s41593-017-0018-8.

Shewcraft, R.A., Dean, H.L., Fabiszak, M.M., Hagan, M.A., Wong, Y.T., and Pesaran, B. (2020). Excitatory/Inhibitory Responses Shape Coherent Neuronal Dynamics Driven by Optogenetic Stimulation in the Primate Brain. *J Neurosci* 40, 2056-2068. 10.1523/JNEUROSCI.1949-19.2020.

Shuler, M.G., and Bear, M.F. (2006). Reward timing in the primary visual cortex. *Science* 311, 1606-1609. 10.1126/science.1123513.

Sieveritz, B., García-Muñoz, M., and Arbuthnott, G.W. (2019). Thalamic afferents to prefrontal cortices from ventral motor nuclei in decision-making. *Eur J Neurosci* 49, 646-657. 10.1111/ejn.14215.

Sillito, A.M., Kemp, J.A., and Patel, H. (1980). Inhibitory interactions contributing to the ocular dominance of monocularly dominated cells in the normal cat striate cortex. Exp *Brain Res* 41, 1-10. 10.1007/bf00236673.

Sinclair, G., Leach, J., Jordan, P., Gibson, G., Yao, E., Laczik, Z., Padgett, M., and Courtial, J. (2004). Interactive application in holographic optical tweezers of a multi-plane Gerchberg-Saxton algorithm for three-dimensional light shaping. *Opt Express* 12, 1665-1670. 10.1364/opex.12.001665.

Siniscalchi, M. J., Phoumthipphavong, V., Ali, F., Lozano, M., and Kwan, A. C. (2016). Fast and slow transitions in frontal ensemble activity during flexible sensorimotor behavior. *Nat Neurosci* 19, 1234–1242. doi: 10.1038/nn.4342.

Soliman, G.S., and Trüper, H.G. (1982). Halobacterium pharaonis sp. nov., a new, extremely haloalkaliphilic archaebacterium with low magnesium requirement. Zentralblatt Für Bakteriologie Mikrobiologie Und Hygiene: I. Abt. Originale C: Allgemeine, Angewandte Und Ökologische Mikrobiologie 3, 318-329.

Stachniak, T.J., Ghosh, A., and Sternson, S.M. (2014). Chemogenetic synaptic silencing of neural circuits localizes a hypothalamus→midbrain pathway for feeding behavior. *Neuron* 82, 797-808. 10.1016/j.neuron.2014.04.008.

Stauffer, W.R., Lak, A., Yang, A., Borel, M., Paulsen, O., Boyden, E.S., and Schultz, W. (2016). Dopamine Neuron-Specific Optogenetic Stimulation in Rhesus Macaques. *Cell* 166, 1564-1571.e1566. 10.1016/j.cell.2016.08.024.

Stuber, G.D., Hnasko, T.S., Britt, J.P., Edwards, R.H., and Bonci, A. (2010). Dopaminergic terminals in the nucleus accumbens but not the dorsal striatum corelease glutamate. *J Neurosci* 30, 8229-8233. 10.1523/JNEUROSCI.1754-10.2010.

Sun, F., Zhou, J., Dai, B., Qian, T., Zeng, J., Li, X., Zhuo, Y., Zhang, Y., Wang, Y., Qian, C., et al. (2020). Next-generation GRAB sensors for monitoring dopaminergic activity *in vivo*. *Nat Methods* 17, 1156-1166. 10.1038/s41592-020-00981-9.

Szadai, Z., Pi H.J., Chevy, Q., Ócsai, K., Albeanu F., Chiovini, B., Szalay, G., Katona, G., Kepecs, A., Rózsa B. (2022). Cortex-wide fast activation of VIP-expressing inhibitory neurons by reward and punishment. *bioRxiv* 2022.04.27.489695; 10.1101/2022.04.27.489695

Takahashi, N., Oertner, T.G., Hegemann, P., and Larkum, M.E. (2016). Active cortical dendrites modulate perception. *Science* 354, 1587-1590. 10.1126/science.aah6066.

Tanaka, Y.H., Tanaka, Y.R., Kondo, M., Terada, S.I., Kawaguchi, Y., and Matsuzaki, M. (2018). Thalamocortical Axonal Activity in Motor Cortex Exhibits Layer-Specific Dynamics during Motor Learning. *Neuron*. 10.1016/j.neuron.2018.08.016.

Tasic, B., Menon, V., Nguyen, T.N., Kim, T.K., Jarsky, T., Yao, Z., Levi, B., Gray, L.T., Sorensen, S.A., Dolbeare, T., et al. (2016). Adult mouse cortical cell taxonomy revealed by single cell transcriptomics. *Nat Neurosci* 19, 335-346. 10.1038/nn.4216.

Tasic, B., Yao, Z., Graybuck, L.T., Smith, K.A., Nguyen, T.N., Bertagnolli, D., Goldy, J., Garren, E., Economo, M.N., Viswanathan, S., et al. (2018). Shared and distinct transcriptomic cell types across neocortical areas. *Nature* 563, 72-78. 10.1038/s41586-018-0654-5.

Tennant, K.A., Adkins, D.L., Donlan, N.A., Asay, A.L., Thomas, N., Kleim, J.A., and Jones, T.A. (2011). The organization of the forelimb representation of the C57BL/6 mouse motor cortex as defined by intracortical microstimulation and cytoarchitecture. *Cereb Cortex* 21, 865-876. 10.1093/cercor/bhq159.

Thabit, M.N., Nakatsuka, M., Koganemaru, S., Fawi, G., Fukuyama, H., and Mima, T. (2011). Momentary reward induce changes in excitability of primary motor cortex. *Clin Neurophysiol* 122, 1764-1770. 10.1016/j.clinph.2011.02.021.

Thévenaz, P., Ruttimann, U.E., and Unser, M. (1998). A pyramid approach to subpixel registration based on intensity. *IEEE Trans Image Process* 7, 27-41. 10.1109/83.650848.

Tlamsa, A. P., and Brumberg, J. C. (2010). Organization and morphology of thalamocortical neurons of mouse ventral lateral thalamus. *Somatosensory & Motor Research* 27, 34–43. doi: 10.3109/08990221003646736.

Tom, S.M., Fox, C.R., Trepel, C., and Poldrack, R.A. (2007). The neural basis of loss aversion in decision-making under risk. *Science* 315, 515-518. 10.1126/science.1134239.

Tremblay, L., and Schultz, W. (1999). Relative reward preference in primate orbitofrontal cortex. *Nature* 398, 704-708. 10.1038/19525.

Tremblay, R., Lee, S., and Rudy, B. (2016). GABAergic Interneurons in the Neocortex: From Cellular Properties to Circuits. *Neuron* 91, 260-292. 10.1016/j.neuron.2016.06.033.

Tritsch, N.X., Ding, J.B., and Sabatini, B.L. (2012). Dopaminergic neurons inhibit striatal output through non-canonical release of GABA. *Nature* 490, 262-266. 10.1038/nature11466.

Turi, G.F., Li, W.K., Chavlis, S., Pandi, I., O'Hare, J., Priestley, J.B., Grosmark, A.D., Liao, Z., Ladow, M., Zhang, J.F., et al. (2019). Vasoactive Intestinal Polypeptide-Expressing Interneurons in the Hippocampus Support Goal-Oriented Spatial Learning. *Neuron* 101, 1150-1165.e1158. 10.1016/j.neuron.2019.01.009.

Tye, K.M., Prakash, R., Kim, S.Y., Fenno, L.E., Grosenick, L., Zarabi, H., Thompson, K.R., Gradinaru, V., Ramakrishnan, C., and Deisseroth, K. (2011). Amygdala circuitry mediating reversible and bidirectional control of anxiety. *Nature* 471, 358-362. 10.1038/nature09820.

Vann, S. D., Aggleton, J. P., and Maguire, E. A. (2009). What does the retrosplenial cortex do? *Nat Rev Neurosci* 10, 792–802. doi: 10.1038/nrn2733.

Vazey, E.M., and Aston-Jones, G. (2014). Designer receptor manipulations reveal a role of the locus coeruleus noradrenergic system in isoflurane general anesthesia. *Proc Natl Acad Sci U S A* 111, 3859-3864. 10.1073/pnas.1310025111.

Veinante, P., Jacquin, M. F., and Deschênes, M. (2000). Thalamic projections from the whisker-sensitive regions of the spinal trigeminal complex in the rat. *J. Comp. Neurol.* 420, 233–243. doi: 10.1002/(SICI)1096-9861(20000501)420:2<233::AID-CNE6>3.0.CO;2-T.

Wächter, T., Lungu, O.V., Liu, T., Willingham, D.T., and Ashe, J. (2009). Differential effect of reward and punishment on procedural learning. *J Neurosci* 29, 436-443. 10.1523/JNEUROSCI.4132-08.2009.

Walker, A.E., Fulton, J.F. (1938). Hemidecortication in chimpanzee, baboon, macaque, potto, cat and coati: a study in encephalization. *The Journal of Nervous and Mental Disease*: 87:6, 677-700

Wall, N. R., De La Parra, M., Sorokin, J. M., Taniguchi, H., Huang, Z. J., and Callaway, E. M. (2016). Brain-Wide Maps of Synaptic Input to Cortical Interneurons. *J. Neurosci.* 36, 4000–4009. doi: 10.1523/JNEUROSCI.3967-15.2016.

Wan, J., Peng, W., Li, X., Qian, T., Song, K., Zeng, J., Deng, F., Hao, S., Feng, J., Zhang, P., et al. (2021). A genetically encoded sensor for measuring serotonin dynamics. *Nat Neurosci* 24, 746-752. 10.1038/s41593-021-00823-7.

Wang, G., Zhang, P., Mendu, S.K., Wang, Y., Zhang, Y., Kang, X., Desai, B.N., and Zhu, J.J. (2020). Revaluation of magnetic properties of Magneto. *Nat Neurosci* 23, 1047-1050. 10.1038/s41593-019-0473-5.

Wang, H., Peca, J., Matsuzaki, M., Matsuzaki, K., Noguchi, J., Qiu, L., Wang, D., Zhang, F., Boyden, E., Deisseroth, K., et al. (2007). High-speed mapping of synaptic connectivity using photostimulation in Channelrhodopsin-2 transgenic mice. *Proc Natl Acad Sci U S A* 104, 8143-8148. 10.1073/pnas.0700384104.

Wang, P. Y., Boboila, C., Chin, M., Higashi-Howard, A., Shamash, P., Wu, Z., et al. (2020). Transient and Persistent Representations of Odor Value in Prefrontal Cortex. *Neuron* 108, 209-224.e6. doi: 10.1016/j.neuron.2020.07.033.

Warden, M.R., Selimbeyoglu, A., Mirzabekov, J.J., Lo, M., Thompson, K.R., Kim, S.Y., Adhikari, A., Tye, K.M., Frank, L.M., and Deisseroth, K. (2012). A prefrontal cortex-brainstem neuronal projection that controls response to behavioural challenge. *Nature* 492, 428-432. 10.1038/nature11617.

Weiler, N., Wood, L., Yu, J., Solla, S.A., and Shepherd, G.M. (2008). Top-down laminar organization of the excitatory network in motor cortex. *Nat Neurosci* 11, 360-366. 10.1038/nn2049.

Weinberger, N.M. (2004). Specific long-term memory traces in primary auditory cortex. *Nat Rev Neurosci* 5, 279-290. 10.1038/nrn1366.

Weinberger, N.M., Javid, R., and Lepan, B. (1993). Long-term retention of learning-induced receptive-field plasticity in the auditory cortex. *Proc Natl Acad Sci U S A* 90, 2394-2398. 10.1073/pnas.90.6.2394.

Wheeler, M.A., Smith, C.J., Ottolini, M., Barker, B.S., Purohit, A.M., Grippo, R.M., Gaykema, R.P., Spano, A.J., Beenhakker, M.P., Kucenas, S., et al. (2016a). Genetically targeted magnetic control of the nervous system. *Nat Neurosci* 19, 756-761. 10.1038/nn.4265.

Wheeler, M.A., Smith, C.J., Ottolini, M., Barker, B.S., Purohit, A.M., Grippo, R.M., Gaykema, R.P., Spano, A.J., Beenhakker, M.P., Kucenas, S., et al. (2016b). Genetically targeted magnetic control of the nervous system. *Nat Neurosci* 19, 756-761. 10.1038/nn.4265.

Whishaw, I.Q. (2000). Loss of the innate cortical engram for action patterns used in skilled reaching and the development of behavioral compensation following motor cortex lesions in the rat. *Neuropharmacology* 39, 788-805. 10.1016/s0028-3908(99)00259-2.

Wickersham, I.R., Lyon, D.C., Barnard, R.J., Mori, T., Finke, S., Conzelmann, K.K., Young, J.A., and Callaway, E.M. (2007). Monosynaptic restriction of transsynaptic tracing from single, genetically targeted neurons. *Neuron* 53, 639-647. 10.1016/j.neuron.2007.01.033.

Wietek, J., Beltramo, R., Scanziani, M., Hegemann, P., Oertner, T.G., and Wiegert, J.S. (2015). An improved chloride-conducting channelrhodopsin for light-induced inhibition of neuronal activity *in vivo*. *Sci Rep* 5, 14807. 10.1038/srep14807.

Wilson, N.R., Runyan, C.A., Wang, F.L., and Sur, M. (2012). Division and subtraction by distinct cortical inhibitory networks *in vivo*. *Nature* 488, 343-348. 10.1038/nature11347.

Witham, C.L., Fisher, K.M., Edgley, S.A., and Baker, S.N. (2016). Corticospinal Inputs to Primate Motoneurons Innervating the Forelimb from Two Divisions of Primary Motor Cortex and Area 3a. *J Neurosci* 36, 2605-2616. 10.1523/JNEUROSCI.4055-15.2016.

Withers, G.S., and Greenough, W.T. (1989). Reach training selectively alters dendritic branching in subpopulations of layer II-III pyramids in rat motor-somatosensory forelimb cortex. *Neuropsychologia* 27, 61-69. 10.1016/0028-3932(89)90090-0.

Witten, I.B., Steinberg, E.E., Lee, S.Y., Davidson, T.J., Zalocusky, K.A., Brodsky, M., Yizhar, O., Cho, S.L., Gong, S., Ramakrishnan, C., et al. (2011). Recombinase-driver rat lines: tools, techniques, and optogenetic application to dopamine-mediated reinforcement. *Neuron* 72, 721-733. 10.1016/j.neuron.2011.10.028.

Wolff, S.B., Gründemann, J., Tovote, P., Krabbe, S., Jacobson, G.A., Müller, C., Herry, C., Ehrlich, I., Friedrich, R.W., Letzkus, J.J., and Lüthi, A. (2014). Amygdala interneuron

subtypes control fear learning through disinhibition. *Nature* 509, 453-458. 10.1038/nature13258.

Wolpert, D.M., and Flanagan, J.R. (2001). Motor prediction. *Curr Biol* 11, R729-732. 10.1016/s0960-9822(01)00432-8.

Wood, K.C., Blackwell, J.M., and Geffen, M.N. (2017). Cortical inhibitory interneurons control sensory processing. *Curr Opin Neurobiol* 46, 200-207. 10.1016/j.conb.2017.08.018.

Wozny, C., Beed, P., Nitzan, N., Possnecker, Y., Rost, B.R., and Schmitz, D. (2018). VGLUT2 Functions as a Differential Marker for Hippocampal Output Neurons. *Front Cell Neurosci* 12, 337. 10.3389/fncel.2018.00337.

Xu, F.X., Zhou, L., Wang, X.T., Jia, F., Ma, K.Y., Wang, N., Lin, L., Xu, F.Q., and Shen, Y. (2020). Magneto is ineffective in controlling electrical properties of cerebellar Purkinje cells. *Nat Neurosci* 23, 1041-1043. 10.1038/s41593-019-0475-3.

Xu, N.L., Harnett, M.T., Williams, S.R., Huber, D., O'Connor, D.H., Svoboda, K., and Magee, J.C. (2012). Nonlinear dendritic integration of sensory and motor input during an active sensing task. *Nature* 492, 247-251. 10.1038/nature11601.

Xu, T., Yu, X., Perlik, A.J., Tobin, W.F., Zweig, J.A., Tennant, K., Jones, T., and Zuo, Y. (2009). Rapid formation and selective stabilization of synapses for enduring motor memories. *Nature* 462, 915-919. 10.1038/nature08389.

Xu, X., and Callaway, E.M. (2009). Laminar specificity of functional input to distinct types of inhibitory cortical neurons. *J Neurosci* 29, 70-85. 10.1523/JNEUROSCI.4104-08.2009.

Xue, M., Atallah, B.V., and Scanziani, M. (2014). Equalizing excitation-inhibition ratios across visual cortical neurons. *Nature* 511, 596-600. 10.1038/nature13321.

Yamawaki, N., Borges, K., Suter, B.A., Harris, K.D., and Shepherd, G.M. (2014). A genuine layer 4 in motor cortex with prototypical synaptic circuit connectivity. *Elife* 3, e05422. 10.7554/eLife.05422.

Yang, G., Pan, F., and Gan, W.B. (2009). Stably maintained dendritic spines are associated with lifelong memories. *Nature* 462, 920-924. 10.1038/nature08577.

Yang, H.W., and Lemon, R.N. (2003). An electron microscopic examination of the corticospinal projection to the cervical spinal cord in the rat: lack of evidence for cortico-motoneuronal synapses. Exp *Brain Res* 149, 458-469. 10.1007/s00221-003-1393-9.

Yang, J., Serrano, P., Yin, X., Sun, X., Lin, Y., and Chen, S. X. (2022). Functionally distinct NPAS4-expressing somatostatin interneuron ensembles critical for motor skill learning. *Neuron* 110, 3339-3355.e8. doi: 10.1016/j.neuron.2022.08.018.

Yildirim, M., Sugihara, H., So, P.T.C., and Sur, M. (2019). Functional imaging of visual cortical layers and subplate in awake mice with optimized three-photon microscopy. *Nat Commun* 10, 177. 10.1038/s41467-018-08179-6.

Yin, X., Jones, N., Yang, J., Asraoui, N., Mathieu, M.E., Cai, L., and Chen, S.X. (2021). Delayed motor learning in a 16p11.2 deletion mouse model of autism is rescued by locus coeruleus activation. *Nat Neurosci* 24, 646-657. 10.1038/s41593-021-00815-7.

Yizhar, O., Fenno, L.E., Prigge, M., Schneider, F., Davidson, T.J., O'Shea, D.J., Sohal, V.S., Goshen, I., Finkelstein, J., Paz, J.T., et al. (2011). Neocortical excitation/inhibition balance in information processing and social dysfunction. *Nature* 477, 171-178. 10.1038/nature10360.

Zeng, H. (2018a). Mesoscale connectomics. *Curr Opin Neurobiol* 50, 154-162. 10.1016/j.conb.2018.03.003.

Zhang, F., Wang, L.P., Brauner, M., Liewald, J.F., Kay, K., Watzke, N., Wood, P.G., Bamberg, E., Nagel, G., Gottschalk, A., and Deisseroth, K. (2007). Multimodal fast optical interrogation of neural circuitry. *Nature* 446, 633-639. 10.1038/nature05744.

Zhang, S., Xu, M., Kamigaki, T., Hoang Do, J.P., Chang, W.C., Jenvay, S., Miyamichi, K., Luo, L., and Dan, Y. (2014). Selective attention. Long-range and local circuits for top-down modulation of visual cortex processing. *Science* 345, 660-665. 10.1126/science.1254126.

Zhou, J., Montesinos-Cartagena, M., Wikenheiser, A. M., Gardner, M. P. H., Niv, Y., and Schoenbaum, G. (2019). Complementary Task Structure Representations in Hippocampus and Orbitofrontal Cortex during an Odor Sequence Task. *Current Biology* 29, 3402-3409.e3. doi: 10.1016/j.cub.2019.08.040.

Zhu, H., Pleil, K.E., Urban, D.J., Moy, S.S., Kash, T.L., and Roth, B.L. (2014). Chemogenetic inactivation of ventral hippocampal glutamatergic neurons disrupts consolidation of contextual fear memory. *Neuropsychopharmacology* 39, 1880-1892. 10.1038/npp.2014.35.

Zhu, Y., Wienecke, C.F., Nachtrab, G., and Chen, X. (2016). A thalamic input to the nucleus accumbens mediates opiate dependence. *Nature* 530, 219-222. 10.1038/nature16954.

Zimmermann, K. S., Li, C., Rainnie, D. G., Ressler, K. J., and Gourley, S. L. (2018). Memory Retention Involves the Ventrolateral Orbitofrontal Cortex: Comparison with the Basolateral Amygdala. *Neuropsychopharmacol.* 43, 373–383. doi: 10.1038/npp.2017.139.