Reward-based learning through layer-specific cortical optogenetic stimulation in mice

A Thesis

submitted to

Indian Institute of Science Education and Research Pune in partial fulfilment of the requirements for the BS-MS Dual Degree Programme

by

Anand Karthik

Indian Institute of Science Education and Research Pune

Dr. Homi Bhabha Road,

Pashan, Pune 411008, INDIA.

Date: April, 2026

Under the guidance of

Supervisor: Prof. Carl Petersen,

École Polytechnique Fédérale de Lausanne (EPFL)

From May 2025 to March 2026

INDIAN INSTITUTE OF SCIENCE EDUCATION AND RESEARCH PUNE

Certificate

This is to certify that this dissertation entitled **"Reward-based learning through layer-specific cortical optogenetic stimulation in mice"** towards the partial fulfilment of the BS-MS dual degree programme at the Indian Institute of Science Education and Research, Pune represents study/work carried out by Anand Karthik at École Polytechnique Fédérale de Lausanne (EPFL) under the supervision of Prof. Carl Petersen, Full Professor, Brain Mind Institute, during the academic year 2025-2026.

Prof. Carl Petersen

Committee:

Prof. Carl Petersen

Dr. Nixon M Abraham

Declaration

I hereby declare that the matter embodied in the report entitled "**Reward-based learning through layer-specific cortical optogenetic stimulation in mice"** are the results of the work carried out by me at the Brain Mind Institute, École Polytechnique Fédérale de Lausanne (EPFL), under the supervision of Prof. Carl Petersen, and the same has not been submitted elsewhere for any other degree. Wherever others contribute, every effort is made to indicate this clearly, with due reference to the literature and acknowledgement of collaborative research and discussions.

Anand Karthik

Reg.no: 20211093

Table of Contents

List of Tables

List of Figures

Abstract

In our everyday lives, we perform multiple sensorimotor transformations, where a particular sensory stimulus triggers an appropriate motor response. A well-studied task for investigating sensorimotor transformations involves training water-restricted mice to lick for a water reward in response to a brief whisker stimulus. Evidence suggests that the whisker primary somatosensory cortex (wS1) contributes significantly to detecting the whisker stimulus in this task. The wS1 is organised into 6 layers, with neurons in Layer 4 (L4) receiving major sensory inputs from the brainstem and the thalamus. The sensory information then rapidly spreads to Layers 2/3 (L2/3), 5 (L5), and 6, which project to many cortical and subcortical structures, leading to a motor command. However, the neuronal population(s) in wS1 that play a critical role in driving goal-directed licking during this detection task remain unknown. We trained mice (expressing an excitatory opsin) from 4 layer-specific (L2/3, L4, L5-IT(Intratelencephalic) and L5-PT(Pyramidal tract)) Cre lines to lick in response to an optogenetic stimulus instead of a whisker stimulus. L4, L5-IT and L5-PT mice were able to respond to the optogenetic stimulus quickly (1st day) while L2/3 mice only responded on the 3rd day. When compared across layers, L4 and L5-IT mice showed a stronger response than L2/3 and L5-PT. L5-IT excitatory neurons may prove to be suitable for further investigation in this transformation due to their known projections to the striatal projection neurons in the dorsolateral striatum (some of which evoke licking). Overall, this project sheds light on the relative strength of the different layers of the wS1 in driving goal-directed licking and thus provides a basis for more focused investigations of critical subpopulations in the wS1 that facilitate sensorimotor transformations during the detection task.

Acknowledgements

I am extremely grateful to my thesis supervisor, Prof. Carl Petersen, for granting me the valuable opportunity to work in the Laboratory of Sensory Processing at EPFL. I would like to thank Dr Sylvain Crochet, Prof. Carl Petersen and Dr Shankar Sachidhanandam for their collective guidance during the project and thesis work. I am thankful to my mentor, Dr Shankar Sachidhanandam, for his teaching and advice throughout the project. I would also like to thank Lana Smith, Myriam Hamon, Morgane Storey, and the other lab members for their valuable feedback and support at various points in my project. I am also deeply grateful to Marianne Nkosi (lab technician) for her assistance in the histology part of my project. Financially, I'm grateful for the KVPY Scholarship and the monthly allowance provided by Prof. Carl Petersen. Finally, I'm thankful to my family and friends for their constant support throughout my stay in Switzerland.

Contributions

Contributor name	Contributor role
Dr Carl Petersen, Dr Sylvain Crochet	Conceptualization Ideas
Dr Sylvain Crochet, Dr Shankar Sachidhanandam	Methodology
Lab GitHub Repository, Anand Karthik	Software
Dr Shankar Sachidhanandam, Anand Karthik	Validation
Anand Karthik, Dr Shankar Sachidhanandam	Formal analysis
Anand Karthik	Investigation
Dr Carl Petersen	Resources
Anand Karthik	Data Curation
Anand Karthik	Writing - original draft preparation
Dr Carl Petersen, Dr Shankar Sachidhanandam, Dr Sylvain Crochet, Anand Karthik	Writing - review and editing
Anand Karthik	Visualization
Marianne Nkosi	Assistance with Histology
Dr Carl Petersen, Dr Shankar Sachidhanandam	Supervision
Dr Sylvain Crochet, Dr Shankar Sachidhanandam	Project administration
Dr Carl Petersen	Funding acquisition

This contributor syntax is based on the Journal of Cell Science CRediT Taxonomy[1].

AI usage attribution

This work was majorly human-generated, with generative artificial intelligence (ChatGPT: Model 5.2) assisting in the generation of Python code for data analysis and plotting. All such code underwent manual verification and testing prior to execution. No Generative AI was used in any other aspect in this thesis. The author maintains full responsibility for the intellectual content, the accuracy of the results described and the final conclusions presented in this work.

[1] https://journals.biologists.com/jcs/pages/author-contributions

Chapter 1 Introduction

In our everyday lives, we perform multiple sensorimotor transformations, where a particular sensory stimulus triggers an appropriate motor response. For example, when driving, we immediately brake when the light turns red at the signal. These sensorimotor transformations rely on learned associations between sensory inputs and motor outputs acquired through experience. Context is another important factor that can lead to different behavioural outcomes for the same sensory stimulus. This flexibility suggests that the neural circuits involved in sensorimotor processing utilise the learned sensory cues with contextual information to generate an appropriate motor response.

These sensorimotor transformations can be studied in mice trained to perform a sensory detection task. Work over the past few decades has identified the primary sensory cortices as playing a key role in sensorimotor transformations during such detection tasks. One such well-studied task involves training water-restricted mice to lick in response to a brief whisker stimulus in order to receive a water reward. Evidence suggests that the whisker primary somatosensory cortex (wS1) contributes to the detection of the whisker stimulus in this task. However, there is also some counter-evidence showing that brain lesioned mice (without wS1) were able to recover their performance in a whisker detection task (Hong *et al.*, 2018). Regardless, the initial drop in performance observed due to lesions and optogenetic inactivation (of excitatory neurons) in wS1 underscores its importance in the whisker detection task. Before diving into more evidence, we must first describe the architecture of wS1 along with the flow of the sensory information from the whisker to wS1 and within wS1.

1.1 Organisation of the Whisker system and wS1

The whisker system provides rodents with spatial and textural cues from their immediate surroundings. Whisking is an active process in which rodents move their whiskers back and forth to sense changes in their surroundings. The deflection of a whisker evokes sequential activity in the neurons of the trigeminal ganglion, the brain stem, and the thalamus before reaching the sensory cortex (wS1) (Fig. **1**b). The incoming sensory information to wS1 is processed in the neuronal microcircuits and relayed to connected downstream cortical and subcortical regions.

Another important aspect of the whisker system is the existence of somatotopic maps at each level. Each whisker on the whisker pad is individually represented in the wS1 by a ‘barrel’. For example, a C2 whisker deflection is reflected as activation in the C2 barrel of the wS1

(Fig. **1**a). This one-to-one mapping for each whisker is also present in the intermediate regions of the whisker pathway (barrelettes in the principal trigeminal brainstem nucleus and barreloids in the Ventral Posterior medial nucleus (VPM) of the thalamus) (Staiger and Petersen, 2021).

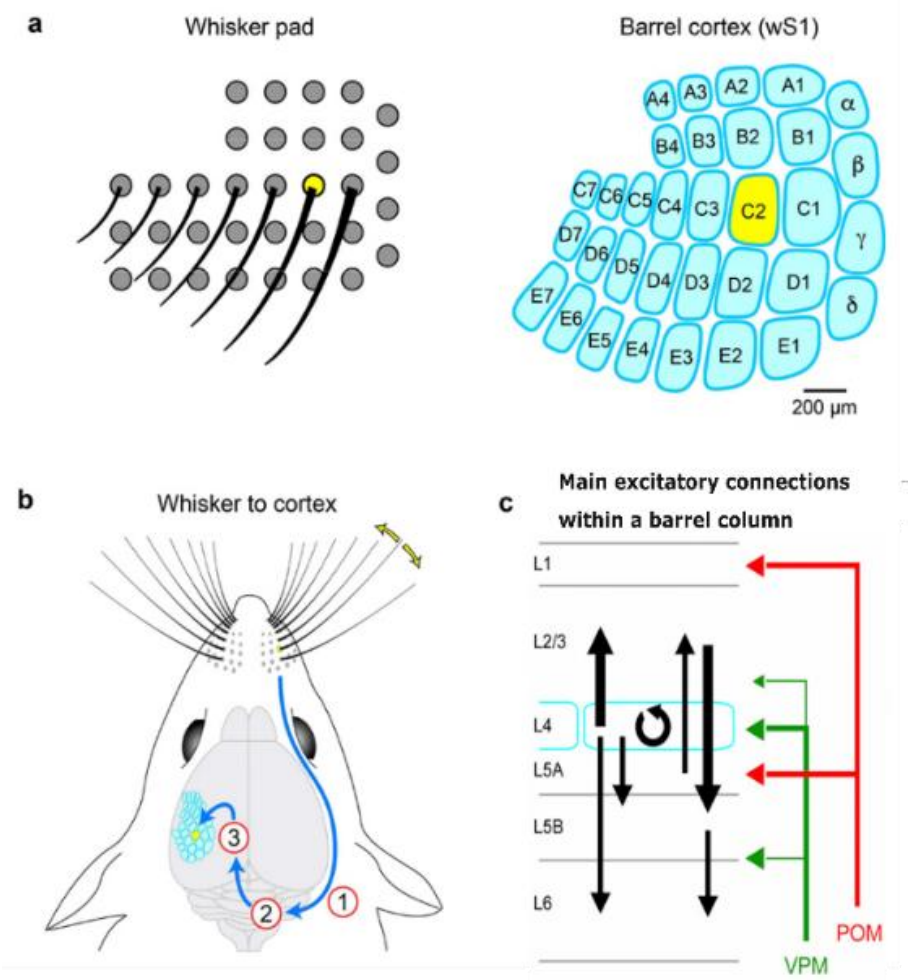

*Figure 1 – Organisation and signalling in the whisker sensory system: **a**) Somatotopic organisation of the whisker system showing the whisker-barrel mapping. Figure from (Petersen, 2019). **b**) Deflection of a whisker evokes activity in: trigeminal ganglion (1); brainstem (2); and thalamus (3), before reaching wS1. Figure from (Petersen, 2019). **c**) Schematic of the excitatory connections between the cortical layers with the main thalamic inputs to a barrel column. Figure from (Petersen and Crochet, 2013).*

Each barrel column in the wS1 is organised into 6 layers, as in most neocortical regions (Fig. **1**c), with specific local and long-range connectivity (Feldmeyer *et al.*, 2013; Staiger and Petersen, 2021; Lefort *et al.*, 2009; Liu *et al*., 2024). The local connectivity of the layers within each barrel and their primary downstream targets are as follows: neurons in layer 4 (L4) receive major sensory inputs from the VPM nucleus of the thalamus (L4 - main input layer). A column of axons from the excitatory L4 neurons innervates the layer 2/3 (L2/3) neurons of the same barrel column. L2 and L3 are mostly composed of excitatory pyramidal projection neurons with dendrites in L1–L3 and local axonal projections in L2, L3 and L5. Layer 2/3 neurons also project long-distance axons to other cortical regions (wS2 and

whisker motor cortex) and the striatum. The excitatory neurons of layer 5 (L5) have apical dendrites that extend up to L1 and receive inputs from all cortical layers and the thalamus. Layer 5 excitatory neurons can be classified into L5-IT (Intratelencephalic) and L5-PT (Pyramidal tract) neurons. L5-IT (similar to L 2/3) project within the wS1 and to other neocortical regions and the striatum. L5-PT neurons mostly show long-distance projections to many subcortical areas, such as the striatum, the posterior medial nucleus of the thalamus (POm), the pons, the superior colliculus and other midbrain nuclei. The higher-order portion of the thalamic POm region sends inputs to L5A (L5-IT neurons) and L1 (Fig. **1**c), acting as a feedback pathway to the wS1. L5B consists of both IT and PT neurons. Lastly, layer 6 (L6) contains a considerable number of neurons that project to the thalamus (corticothalamic neurons, CT) as well as other IT neurons. The L6CT neurons are mostly known to serve a modulatory role rather than driving function (Pauzin and Krieger, 2018; Dimwamwa *et al.*, 2024).

1.2 Evidence for the role of wS1 in sensory processing

Sachidhanandam *et al.*, (2013) trained water-deprived mice to associate whisker deflection with a water reward in a whisker-detection task (Fig. **2**a). They showed that the learning of the whisker stimulus could be substituted by a direct activation of wS1 excitatory neurons through optogenetic stimulation (excitatory opsin ChR2 expressed by all excitatory neurons of the wS1). Similarly, the mice that had learnt to lick in response to the optogenetic stimulus were able to transfer to a whisker stimulus and maintain their performance (Fig. **2**b). Furthermore, pharmacological or optogenetic inhibition of wS1 strongly impaired the mice's performance (Sachidhanandam *et al.*, 2013; Kwon *et al*., 2016; Le Merre *et al*., 2018; Oryshchuk *et al*., 2024). Along with membrane potential recordings, this paper provides the primary evidence for wS1's involvement in encoding a sensory input with reward during a whisker-detection task.

Although the activity of wS1 neurons strongly encodes sensory inputs from the whisker, regardless of the mouse's response, the activity of wS1 neurons also contains information about decision and reward collection when mice lick in response to a whisker stimulus. Oryshchuk *et al*., (2024) performed high-density extracellular electrophysiological recordings from the wS1 (and other areas) from mice during a whisker-detection task. They showed that apart from encoding stimulus properties, there is a group of ‘decision neurons’ in L5 and L6 of wS1 which show increased selectivity of response in Hit (lick in response to stimulus) compared to Miss (no lick) trials just after the whisker stimulus and high selectivity for stimulus-driven vs spontaneous licking right before jaw-opening. In conclusion, a subset of

the wS1 neurons encodes the decision (lick or no lick) for the corresponding whisker stimulus.

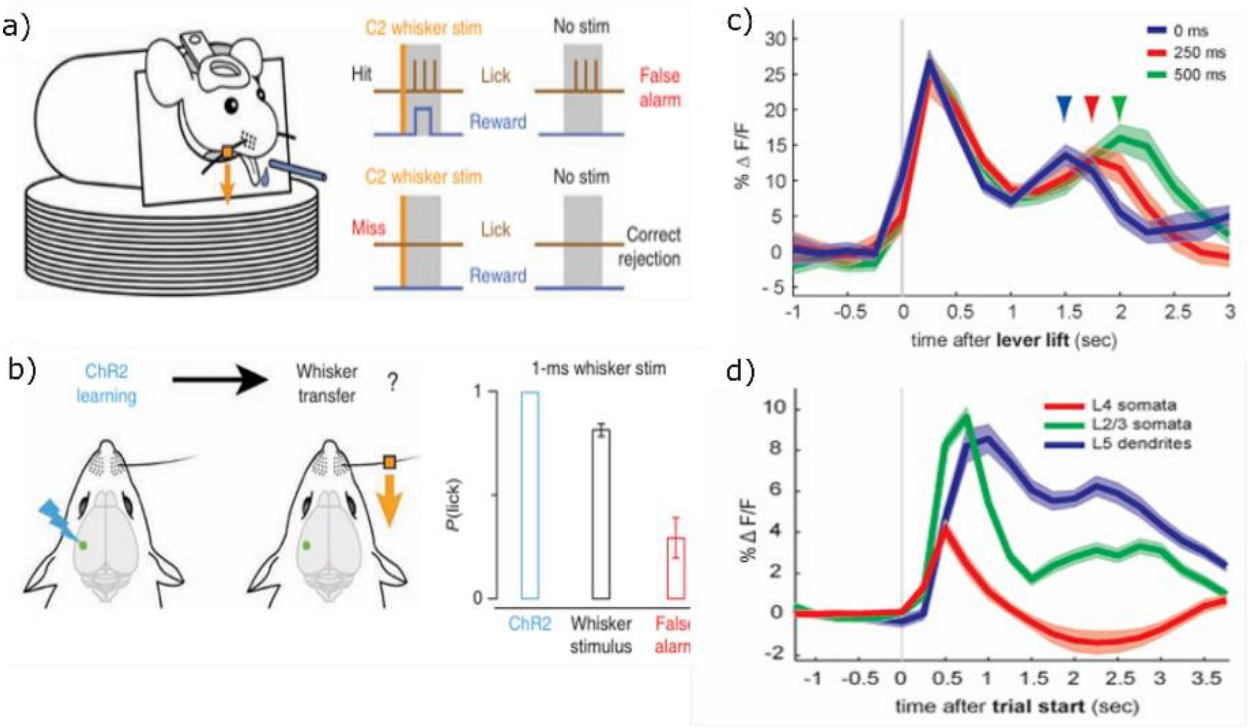

*Figure 2 – Evidence for wS1's role in sensory processing: **a**) Schematic of the mouse during behaviour and the trial outcomes. Figure from (Sachidhanandam et al., 2013). **b**) Left - Mice trained on an optogenetic stimulus transferred to a whisker stimulus, Right – Performance rates for the optogenetic stimulus and the whisker stimulus. Figure from (Sachidhanandam et al., 2013). **c**) Average fluorescence changes in the L5 apical dendrites with varying delay between the lever press and reward delivery. Coloured arrowheads indicate long-latency peaks after response. Figure from (Lacefield et al., 2019). **d**) Comparison of reward-related calcium activity among different cortical layers during task performance. Figure from (Lacefield et al., 2019)*

Lacefield *et al.*, (2019) probed the role of rewards by examining the apical dendrites of Layer 2/3 and Layer 5 neurons in wS1. They recorded calcium signals (GCaMP6f with 2-photon calcium imaging) from the apical dendrites and somata of cortical neurons across different layers of the barrel cortex while mice performed a whisker-detection task. Fig. **2**c shows the calcium signal changes of L5 apical dendrites with different reward timings. Essentially, there was a shift in activity due to the delayed rewards. Fig. **2**d shows the comparison of calcium signals among L2/3 somata, L4 somata and L5 dendrites. The existence of long-latency reward-related fluorescence shown mainly in the apical dendrites hints at a specialised role of apical dendrites in the wS1 (potentially encoding reward collection).

1.3 Downstream flow of sensory information from wS1

1.3.1 Cortico-cortical pathways

One of the main cortical regions that wS1 communicates with is the whisker-related secondary somatosensory cortex (wS2). Kwon *et al*., (2016) recorded calcium activity (GCaMP6 expressed in L2/3 neurons) in wS1 and wS2 of mice performing the whisker

detection task. They observed that S1 activity represented the whisker stimulus more accurately (high response in hit and miss trials), while S2 activity encoded choice more accurately (high response in hit and false alarm trials). To discern the exact neurons involved, they used retrograde tracers to label S1-to-S2 projecting (S2p) and S2-to-S1 projecting (S1p) neurons, and compared their activities with those of unlabelled neurons in both cortical regions. S2p neurons encoded stimulus-related information, while the S2 encoded choice better than the unlabelled neurons of the respective cortices. Overall, this paper shows that the recurrent loop between S1 and S2 could potentially contribute to task performance by amplifying and prolonging stimulus-evoked activity in wS1, thereby initiating the sensorimotor transformation.

Yamashita *et al*., (2016) and Yang *et al*., (2019) are two other studies that found similar results by examining either S2-projection neurons in S1 or top-down modulation from S2 to S1. On the whole, the feedback from wS2 may enhance long-latency depolarisation in the wS1 neurons and refine their activity to encode choice-related information.

1.3.2 Cortico-subcortical pathways

The excitatory neurons in the more superficial layers (L2/3, L4, and L5A) are mostly intratelencephalic-projecting neurons that project to ipsilateral and contralateral cortical regions and the striatum. Meanwhile, the excitatory neurons in deeper layers (L5B and L6) also project to many subcortical regions (thalamus, striatum, midbrain and brainstem) (Liu *et al*., 2024).

Sippy *et al*., (2015) examined striatal projection neurons (SPNs) during goal-directed behaviour. The vast majority of the SPNs express either the D1 or the D2 receptor. D1-SPNs project directly to the substantia nigra pars reticulata (SNr) and form the *direct-pathway (dSPN),* which is traditionally thought to convey a ‘go’ signal for the initiation of motor responses. Meanwhile, D2-SPNs project to the external part of the Globus pallidus (GPe), containing neurons that inhibit the SNr and form the *indirect-pathway (iSPN)*, which is traditionally thought to convey a ‘no-go’ signal that inhibits the motor response. They performed optogenetic manipulations of the dSPNs and iSPNs in the dorsolateral striatum (which receives strong inputs from wS1) of mice trained in a whisker-detection task. Exciting the dSPNs, but not the iSPNs, could substitute for the whisker stimulus and drive licking. Wall *et al*., (2013) utilised a two-virus system (Cre-dependent AAV) and rabies virus to target either D1-SPNs or D2-SPNs, and their monosynaptic inputs, in mice from the respective Cre lines. Due to the specific labelling, they were able to quantify cortical inputs from four different streams (Fig. **3**b). The somatosensory region contributed more to dSPNs

than to iSPNs. Hence, the wS1 probably contributes to the early response of dSPNs. In turn, the dSPN-mediated inhibition of SNr is thought to disinhibit downstream brainstem motor regions to initiate licking. The SNr neurons also likely disinhibit thalamic nuclei, which might influence the long-distance feedback to the apical dendrites in wS1. From both of these papers and others, a hypothesis has been proposed that this cortex-basal ganglia-thalamus pathway might contribute to reward-based learning of whisker-to-lick sensorimotor transformations.

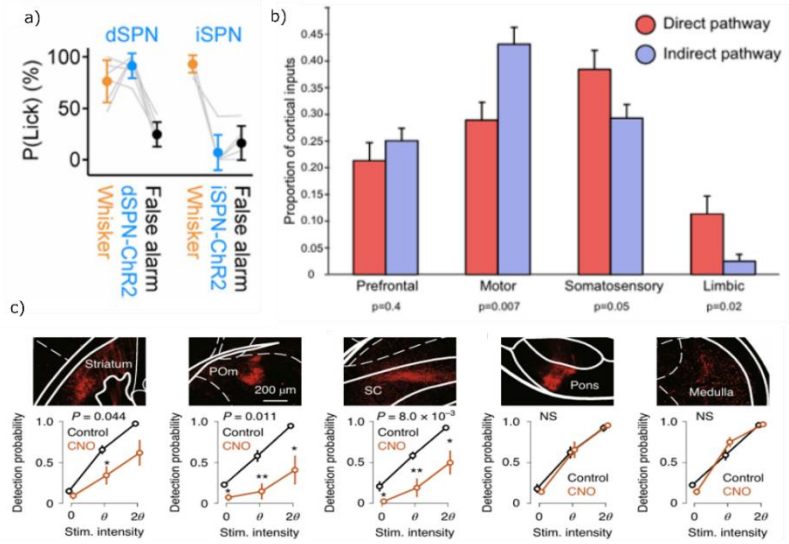

*Figure 3 – Downstream subcortical targets of wS1: **a)** Optogenetic activation of dSPNs causes an increase in P(lick) and leads to a decrease in iSPNs. Figure from (Sippy et al., 2015). **b)** Proportion of cortical inputs to direct and indirect pathway SPNs. The somatosensory cortex contributes slightly higher input to dSPNs (p = 0.05). Figure from (Wall et al., 2013). **c)** Target-specific silencing of axonal outputs of L5-PT neurons expressing mCherry by locally injecting CNO in various cortical regions. The changes in the detection threshold refer to a whisker-detection task with varying stimulus intensities. Figure from (Takahashi et al., 2020).*

Takahashi *et al*., (2020) examined the involvement of subcortical regions directly downstream from wS1 in whisker detection performance. Using calcium activity in apical L5-PT and L5-IT dendrites during a varying intensity whisker detection task, they observed that the detection threshold shifted towards lower values only when the apical L5-PT dendrites were activated. Also, by chemogenetic silencing of the L5-PT outputs individually at various subcortical structures, they confirmed the importance of the deeper wS1 outputs to the striatum, superior colliculus and POm nucleus of the thalamus (Fig. **3**c). On the whole, this evidence of L5-PT apical dendrites receiving long-range projections from subcortical areas implies that L5-PT neurons may guide perceptual detection. At the same time, the various outputs to the subcortical structures also enforce long-range feedback via POm.

1.4 Aim and significance of the project

From the literature review, an important point to note is the differences in input/output connectivity among excitatory neurons across layers in wS1. However, to date, the neuronal

population(s) in wS1 that play a critical role in learning the whisker-detection task remain largely unknown. A possible approach to addressing this question is to train mice to respond to the direct stimulation of a layer-specific neuronal population in wS1, rather than using the sensory whisker stimulus.

Due to multiple differences in connectivity (local and long-range) across layers, a case can be made for any of these layers being the most important in driving goal-directed licking. For instance, L4 excitatory neurons are the main recipients of sensory inputs (barrels) from the thalamus. Hence, direct stimulation of L4 might best mimic the whisker sensory stimulus. L5-PT neurons are shown to have a population of important decision neurons. Hence, the direct stimulation of L5-PT should be efficient at driving licking. L2/3 and L5-IT neurons project to the striatum, which is essential for the 'go' signal for licking. The apical dendrites of L2/3 and L5 play a significant role in relaying long-range feedback from the higher-order thalamus and the striatum. This feedback could increase licking by enhancing activation of the corresponding layers.

Therefore, the aim of this project was to compare the ability of mice to associate a direct optogenetic stimulation with licking for a water reward when stimulating excitatory L2/3, L4, L5-IT or L5-PT neurons in wS1. In the experiments, mice were first trained on an auditory stimulus. After achieving high performance, they were transferred to an optogenetic stimulus to assess whether and to what extent they could transfer learning from the auditory task to the new optogenetic stimulus task.

In this project, we used 4 genetically modified mouse lines expressing Cre recombinase in different populations of excitatory neurons. Liu *et al*., (2024) crossed each of the Cre-driver lines (6) with Cre-dependent tdTomato mice to explore the Cre-recombinase distribution in the transgenic mice. Fig. 4 depicts the layer-specific expression patterns of each Cre line.

Using this paper as the reference, these were the Cre lines chosen for our project: Rasgrf2-2A-dCre (Layer 2/3), Scnn1a-Cre (Layer 4), Tlx3Cre (Layer 5-IT) and Sim1-Cre (Layer 5-PT)

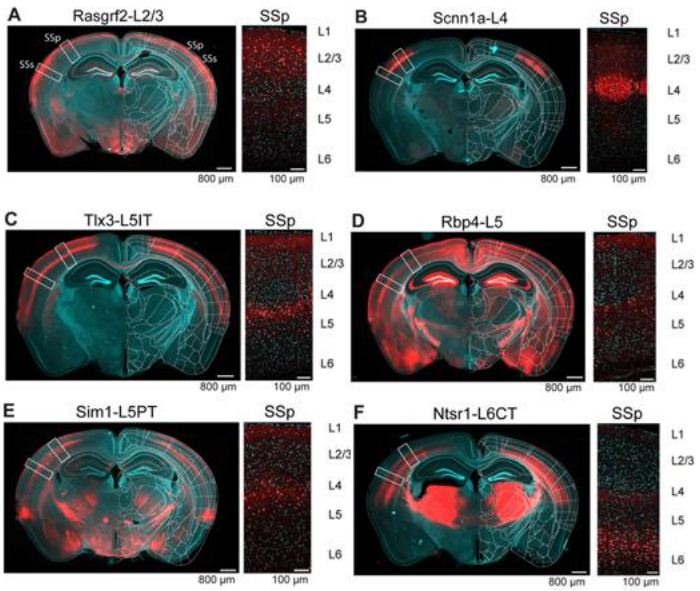

Figure 4: *Cre-dependent tdTomato expression pattern for the layer-specific Cre lines. For each panel, the left image is the full coronal section, and the right image is the magnified (on the Primary somatosensory cortex) image.*

A Cre-dependent AAV vector that encodes an excitatory opsin called ChRmine was injected into a specific layer (dependent on the Cre line) of the C2 barrel region of the wS1. ChRmine was chosen (over standard channel rhodopsin and other variants) due to its long excitation wavelength (625 nm, giving greater light penetration to activate deeper cortical layers) and large photocurrents at comparatively low light intensities (Marshel *et al*., 2019).

With optogenetics, we can attribute the mechanism's causality to a specific population that strongly drives licking during the task. As a result, this project will open new avenues for studying the synaptic mechanisms involved in the learning of such sensorimotor transformations. Mechanisms of synaptic plasticity during task learning in brain areas downstream of the population of neurons stimulated would provide valuable insight into the functioning of the sensorimotor transformation as a whole. Additional factors such as context, internal state of the mouse and stimulus strength-response relationships can be integrated into the paradigm to create a more naturalistic scenario for studying the sensorimotor transformation.

Chapter 2 Materials and Methods

2.1 Materials

Biological materials utilised in the project included mice from different Cre lines and the AAV (Adeno-associated virus) vector that encodes the opsin to be expressed in a Cre-dependent manner.

A total of 25 mice (20 males and 5 females) constituted the main experimental data for the project. All animals were 4-8 weeks old at the start of surgery and were maintained on a 12-hour reversed light-dark cycle in a temperature and humidity-controlled facility.

2.1.1 Mouse Cre lines

Mice from multiple Cre lines were used in this project. Each mouse in these Cre lines expresses a specific promoter-induced Cre recombinase, which allows for localised Cre-dependent expression of the viral opsin.

Rasgrf2-2A-dCre (Rasgrf2-dCre, B6;129S-Rasgrf2tm1(cre/folA)Hze/J) mice express a trimethoprim-induced Cre recombinase primarily in the Layer 2/3 excitatory cells of the cortex.

Scnn1a-Tg3-Cre (Scnn1a-Cre, B6;C3-Tg(Scnn1a-cre)3Aibs/J) mice express Cre recombinase primarily in the cortical layer 4 (which comprises of the barrels in the wS1).

Tlx3-Cre (Tlx3-Cre, B6.FVB(Cg)-Tg(Tlx3-cre)PL56Gsat/Mmucd) mice express Cre recombinase primarily in the Layer 5-IT (Intratelencephalic) neurons of the cortex.

Sim1-Cre (Sim1-Cre, B6.FVB(Cg)-Tg(Sim1-cre)KJ18Gsat/Mmucd) mice express Cre recombinase in the Layer 5-PT (Pyramidal Tract) neurons of the cortex.

Table 1 - Mouse Cre lines summary with global reference IDs and source paper links

Cre recombinase activity	Common name of the Cre line	RRID Number	Source Reference Link
Cortical Layers 2/3	Rasgrf2-2A-dCre	**RRID:IMSR_JAX:022864**	(1)
Cortical Layer 4	Scnn1a-Cre	**RRID:IMSR_JAX:009613**	(2)
Cortical Layer 5-IT	Tlx3-Cre	**RRID:MMRRC_041158-UCD**	(3)
Cortical Layer 5-PT	Sim1-Cre	**RRID:IMSR_JAX:006395**	(4)

2.1.2 Optogenetic AAV Vectors

Double-floxed Inverted Open-reading frame (DIO) AAV vectors encoding the excitatory opsin gene were used as the delivery mechanism to ensure Cre-dependent opsin expression in the C2 region of the whisker somatosensory cortex of the experimental mice. Details of the AAV injections will be further elaborated in section 2.2.2.

Throughout the project, a single AAV encoding ChRmine as the opsin was utilised.

pAAV-Ef1a-DIO-ChRmine-mScarlet-WPRE (Addgene viral prep # 130998-AAV5; http://n2t.net/addgene:130998; RRID: Addgene_130998) is an AAV vector deposited by the Deisseroth Lab. The AAV is Ef1a-driven, Cre-dependent, and expresses ChRmine-mScarlet as a fusion protein. Marshel *et al.,* (2019) first described the ChRmine opsin and its characterisation in vivo.

2.2 Surgical Methods

All experimental procedures were in accordance with the Swiss Federal Veterinary Office (License VD-4107).

2.2.1 Headpost fixation and IOS Imaging

During head-post fixation, mice were anaesthetised under isoflurane. Carprofen was injected subcutaneously based on the mouse's weight to serve as an analgesic (pain reliever). Once the mouse was anaesthetised, the skin was removed to expose the skull. Using the bregma as a reference, the skull was exposed sufficiently to maintain access to the frontal and S1 areas. A headpost was attached to the right half of the exposed skull using cyanoacrylate superglue. Dental cement was applied around the headpost and the general perimeter of the exposed skull.

After a few days of recovery, all of the whiskers on the right whisker pad were trimmed (except the C2 whisker). During intrinsic optical signal (IOS) imaging, changes in reflected red light (630 nm) reflect cortical activity. In our case, the C2 whisker of anaesthetised mice was deflected, leading to increased activity in the C2 barrel in the whisker S1 cortex. Signal changes from multiple instances of whisker deflection and no deflection (control) were used to determine the location of the C2 barrel. This served as the target region for the AAV vector injection.

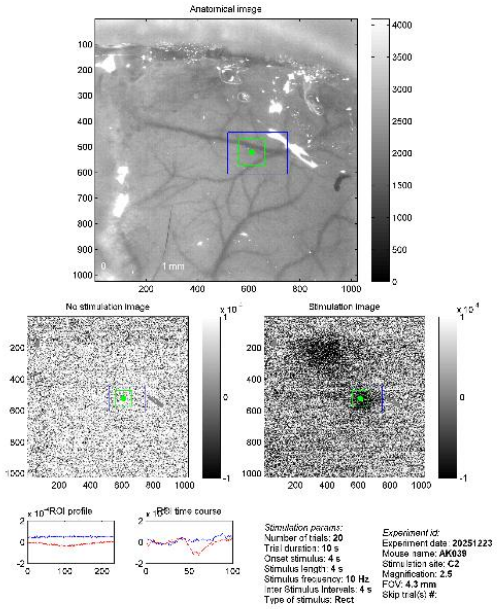

Figure 5: IOS images showing the C2 barrel (marked region) and the auditory cortex (region diagonally above the C2 region) in the 'Stimulated image'. The auditory cortex shows signal changes due to the sound made by the whisker deflector's oscillatory movement. The 'Unstimulated image' is for trials without deflection (control trials). The 'Anatomy image' shows the marked region for C2.

2.2.2 Craniotomy and Viral Injection

The anaesthetised mouse was clamped to a support using the headpost, and the C2 region location was marked on the skull. A craniotomy was performed using a 3 mm biopsy punch to expose the brain's dura. An injection micropipette (internal tip diameter, ~7 µm) was filled with the virus solution and lowered into the C2 barrel column. The virus was diluted to achieve a titre of 4x (relative to the stock concentration of 7×10^{12} vg/mL), and the injection volume was 350 nl (Marshel *et al.*, 2019).

The injection depths below the pia were decided based on the Cre line the mouse belonged to and the depths used in previous literature (such as Liu *et al*., 2024).

- Rasgrf2-2A-dCre (Layer 2/3) – 300 µm
- Scnn1a_Cre (Layer 4) – 500 µm
- Tlx3_Cre and Sim1_Cre (Layer 5) – 800 µm

The micropipette remained in the brain for 5 min before being slowly retracted over 8–10 min to prevent backflow of the virus along the pipette. After the pipette was retracted from the brain, a cranial window (circular coverslip of 3 mm diameter) was placed over the craniotomy, and the edges were sealed with superglue. The condition of the cranial window was monitored in the next **3-5 weeks of expression time**, after which the behaviour was started.

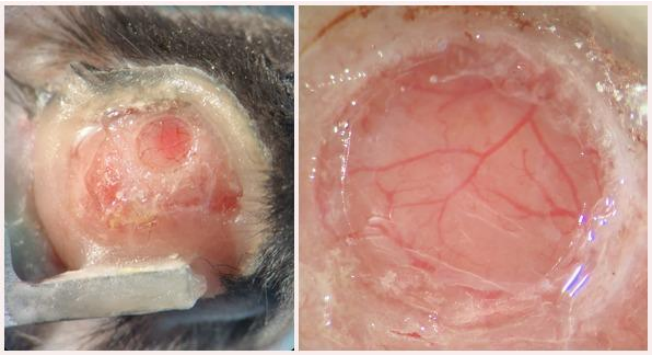

Figure 6: Left - Top view of the mouse head with a cranial window, Right - Magnified onto the window.

2.2.3 Trimethoprim (TMP) injection protocol

The Rasgrf_2A_dCre mice express a TMP-stabilised Cre recombinase. So, a week after the viral injection, these mice were injected with TMP intraperitoneally for 3 consecutive days. The TMP dosage was 0.25 mg/g body weight. A 10% TMP stock solution was prepared every day before injection by dissolving it in DMSO (dimethyl sulfoxide). This was then mixed with 0.9% saline (NaCl) to make the final injection mixture.

For a mouse with a body weight of 'x' grams,

10x µl (Total injected volume) = 2.5x µl of 10% TMP solution + 7.5x µl of 0.9% Saline

2.2.4 Post-Behaviour Histology

After the mice finished the behavioural training, they were perfused, and their brains were extracted and stored in 1x PBS. The lab technician (Marianne Nkosi) prepared 100 µm-thick coronal brain sections. The slices corresponding to the barrel cortex (range determined based on the Allen Brain Atlas) were then mounted onto slides with Vectashield DAPI as the mounting medium (which preserves the fluorescence of ChRmine-mScarlet). The sections were imaged under an epifluorescence microscope with a DsRed filter (545 nm, yellow-green) to examine the opsin expression in the barrel cortex. The images were then analysed to quantify expression levels, which were then correlated with behavioural performance.

2.3 Behaviour

2.3.1 Behavioural Setup

The behavioural training was carried out in parallel in 2 setups with the same configuration. On top of a platform, a cardboard roll was attached to the headpost clamp (to keep the mice still in a specific position). The boxes were padded with soundproof material on the inside and closed during the mouse behaviour to prevent external noise (Fig. **7**).

For the delivery of the water reward, a lick spout (calibrated to deliver 5 µl water per hit trial) attached to a piezo film was positioned in front of the mouse. Licks were detected due to the

spout's movement with the piezo film. More details on data acquisition will be provided in section 2.3.3.

All mice received 3 stimulus types (auditory, whisker, and optogenetic). Two cardboard tubes containing earphone speakers were positioned toward the mouse's ears to deliver the **auditory** stimulus and white noise.

Beneath the platform, a magnetic coil was positioned to deliver a magnetic pulse. On certain days of the training, a metal particle was attached to the C2 whisker (on the right whisker pad) of the mouse. The magnetic pulse deflected the metal particle (and the whisker) and hence provided the **whisker** stimulus.

Light (**Optogenetic**) stimuli were delivered by a fibre optic patch cable (Thorlabs; numerical aperture = 0.5, core diameter = 600 µm) coupled to a light-emitting diode (LED, Thorlabs LEDD1B). An LED of wavelength 625 nm (red) for the excitation of ChRmine was utilised. A masking light of the same wavelength was directed at the mouse's face. The tip of the fibre optic cable was positioned above the cranial window for each mouse. Appropriate support and 3D stages for moving the cable were used to maintain the cable's position and keep it accurately above the cranial window.

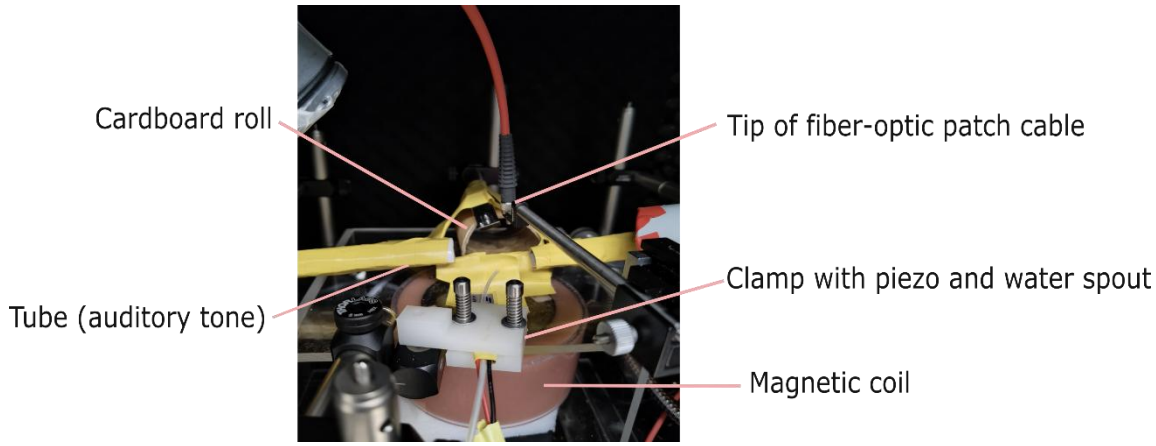

Figure 7: Front view of the behavioural setup – cardboard roll, fibre optic patch cable arching from the top, clamp with the water spout tube and piezo wedged in, two cardboard tubes for delivering the auditory tone, magnetic coil beneath the platform for delivering the whisker stimulus

2.3.2 Calibration of the Stimuli

An audiometer was used to calibrate the auditory stimuli and the white noise at the ends of both tubes (where the mouse's ears would be). White noise was calibrated to 80 dB, and auditory stimuli (10 ms duration, 10 kHz frequency) to 74 dB. This difference was maintained to ensure that the mouse would have to pay attention and distinguish the auditory tone from the louder background white noise.

A Tesla meter (Projekt Electronik) was used to calibrate the strength of the magnetic field pulses. The magnetic pulse (Fig. **8,** bottom) was generated by sending a biphasic sinusoidal voltage pulse (Fig. **8**, top) to the magnetic coil. The amplitude of the voltage pulse was calibrated to generate a magnetic pulse of approximately 25 mT (at peak amplitude, measured at the position of the metal particle attached to the whisker). This was done on both setups to ensure that the whisker is deflected to the same strength, accounting for any unintended differences in coil positioning.

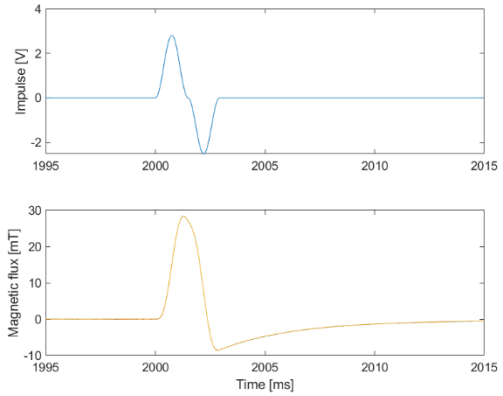

Figure 8 - Example calibration plots: Top- Amplitude of the calibrated impulse used as the input for the magnetic coil during behaviour, Bottom - Magnetic pulse generated by the coil.

The light stimulus was calibrated using a power meter (Thorlabs). During calibration, light intensity was measured as the fibre continuously emitted light. The calibrated LED power used was 10 mW (625 nm for ChRmine). When compared to previous literature (7), the light power we used was higher (to ensure light penetrates deep enough to activate deeper layers such as Layers 4, 5, and 6).

2.3.3 Trial Structure and Training Schedule

After 3-5 weeks of expression time, the mice were water-deprived for 24 hours. Before they started the Go/No-Go detection task, they were habituated and learned to lick the water spout for reward on the **free-licking day**. On that day, the mice were not exposed to any stimulus. The trial parameters were made easier so that the mice learn to lick the water spout for water and get used to the setup in general.

Every mouse had one session (a few hundred trials) per day. No fixed session length was used to terminate an ongoing session, due to variable motivation among the mice for the water rewards across days. The mice were water-restricted to a total of 0.9-1.2 ml of water /day, depending on their initial weight. At the end of each training session, the amount of water collected was calculated (total number of rewards x reward size), and a complement of water was given in an individual cage for 15 min to reach the daily water requirement. The

weight of each mouse was monitored daily, just after the training session, to ensure that the mice did not drop below 80% of their initial weight before water restriction.

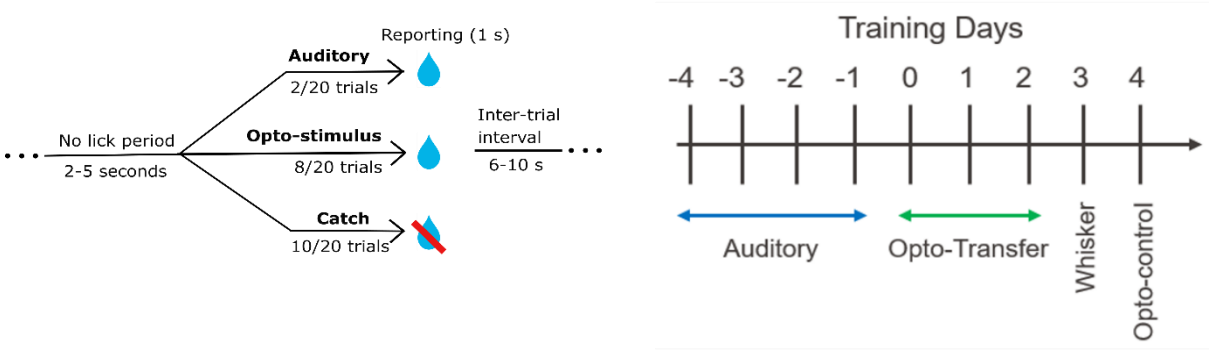

Figure 9: Left – Trial Structure showing the types of trials (on Days 0-2) and distribution within a block (20 trials), Right - Behavioural training schedule

The start of each trial was preceded by a quiet period (2-5 s) during which the mice were not allowed to lick (Fig. **9**). The start of the response window (1 s) was aligned to the stimulus onset for each trial.

Auditory days

A brief auditory tone (10 ms, 10 kHz, 74 dB) was played at the start of Go trials, after which the mouse had to lick the water spout within the response window to receive a small drop (4-5 µl) of water. There was no tone in No-Go trials, and licking in No-Go trials was neither punished nor rewarded. The Go and No-go trials were shuffled pseudo-randomly so that 10 of each would occur every 20 trials (a block). As the training progressed, mice learnt to lick only immediately after the auditory tone, resulting in a high Hit rate (% Correct Go trials) and a low False alarm rate (% Incorrect No-Go trials).

Opto-transfer days

On the transfer day (Day 0), Day 1 and Day 2, opto-stimulus trials were included. For every block (20 trials), there were 8 opto-stimulus trials, 2 auditory trials and 10 catch trials (no stimulus). The opto-stimulus consisted of a train of 5 pulses, each lasting 5 ms, at 50 Hz. The stimulus amplitude was calibrated to deliver 10 mW for activating ChRmine.

During these days, we sought to determine whether there were differences in the ability of mice across different Cre lines to lick in response to the opto-stimulus.

Whisker day

On Day 3, 8 whisker trials were included instead of 8 opto-stimulus trials per block. The proportion of Auditory and catch trials remained the same. In a whisker trial, the C2 whisker was deflected briefly (3 ms, calibrated amplitude).

The Whisker day was used as a secondary transfer to determine if the mice were also able to transfer to a peripheral whisker stimulus (as seen in Sachidhanandam *et al*. (8)).

Opto-control day

On Day 4, the trial structure was the same as on the opto-transfer days (8 opto-stimuli, 2 auditory, and 10 catch trials per block). However, a black piece of tape was introduced in some periods between the tip of the fiber-optic cable and the cranial window to block the light stimulus.

This was done to ensure that none of the mice that performed well (high hit rate) during Days 0-3 used the emitted light as a visual cue. During the time period with this black patch, a mouse which did not use this visual cue would have low performance. After the patch was removed, it would return to high performance. Meanwhile, a mouse that used the visual cue would maintain high performance throughout the session.

2.3.4 Trial Outcomes and Performance Curves

For each stimulus (auditory, optogenetic and whisker), there were 2 possible outcomes – Hit (above noise threshold licking within the response window) and Miss (no licking). The proportion of hits to the total number of stimulus trials was the basis for the hit-rate curves seen in the various plots described in the next section (2.4).

For catch (No-Go) trials, the 2 outcomes were False Alarm and Correct Rejection. The criterion used was the same for both stimulus and catch trials. The proportion of false alarms to the total number of catch trials served as the basis for the False Alarm rate curve shown in the different plots. It acted as a measure of spontaneous or random licking by the mice during the behaviour.

These hit and false alarm rates were calculated using different sampling methods across the plots. The exact nature of the trial sampling is explained in the next section (2.4).

2.4 Behavioural and Statistical Analysis

Behavioural control and data collection were carried out using MATLAB scripts that operated via National Instruments (NI) devices connected to the behavioural PCs. At the end of each session, a set of files containing session-wide and trial-by-trial parameter and outcome information was saved to a folder. These were used to convert into a single NWB (Neurodata Without Borders) file. Finally, the NWB files are analysed on Python (all plots done using Matplotlib).

2.4.1 Single session plots

A standard session table containing the trial parameters and outcomes was used to plot single-session graphs. Firstly, trials with an early lick (licking right before the start of a trial) were filtered out. The performance values were averaged block-wise (20 trials). Essentially, the function picked an average value near the midpoint of a block and used that for the lines on the plot. The performance curves (auditory hit rate, opto-stimulus hit rate) are plotted in this manner against the false alarm rate (spontaneous licking). Only the last auditory day (Day -1) and the first opto-transfer day (Day 0) plots for the example mice are shown in the Results section.

Apart from the performance curves, the outcomes of individual trials were plotted along the x-axis as raster markers (2 contrasting colours for each trial type, denoting a hit or a miss). Lick probability (performance rates) on the Y axis ranges from 0 to 1.

2.4.2 Performance rates (single mouse and layer-averaged)

Every mouse had one session per day. Each data point in the single-mouse plots corresponds to the average performance for a stimulus type (or catch trials) across a session. It shows overall learning of the auditory task and subsequent transfer to the optogenetic and whisker stimuli (with the false-alarm curve serving as a lower bound). The layer-averaged graphs depict the average performance across all mice in that Cre line (layer-specific). The vertical jitter of fainter points around each average indicated the performance of each mouse on that day. Pairwise comparisons were performed between the mean opto-stimulus/whisker hit rate and the mean false alarm rate for Days 0-3 using the Wilcoxon signed-rank test (non-parametric form of paired t-test). Significant (95% significance level) comparisons were indicated on the plot.

2.4.3 d' (single mouse and layer-averaged)

d' (signal detection sensitivity) is a parameter that quantifies the difference between the signal (auditory/opto-stimulus/whisker hit rates) and noise (false alarm rate). The probabilities (performance rates) were converted into Z-scores and subtracted to give a d' value.

$$d' = Z(H) - Z(FA)$$

H - opto-stimulus or whisker hit rate (day-dependent), FA - False alarm rate

Z-score definition

$$\phi(z) = \frac{1}{\sqrt{2\pi}} e^{\frac{-z^2}{2}}$$

where $\phi(z)$ is the area under the standard normal distribution from $-\infty$ to z (z-score)

However, for some sessions in which the mice have learnt to respond to the signal well (H=1) or when the mice do not randomly lick (FA=0), the Z scores become infinite, and thus the d' cannot be computed. When H = 1 or FA = 0 or both

$$\Phi^{-1}(0) = -\infty,\ \Phi^{-1}(1) = +\infty$$

To avoid this issue, an **extreme rate correction** was used. This correction adds or subtracts performance rates by 0.001 (appropriate for the study's session lengths). After that, the rates were averaged across a session, and d' was calculated.

Two d' curves were plotted for the single mouse plots (from each layer) – Auditory d' and Opto d' (Days 0-2) or Whisker d' (Day 3). For the layer-averaged plot, only the Opto d' and Whisker d' were plotted for Days 0-3. Wilcoxon signed-rank tests (two-tailed) were used to compare between days. The p-values obtained were later corrected using the Holm-Bonferroni method. Due to false positives arising from multiple pairwise comparisons (6 in this case), the p-values from the paired t-tests were corrected by a factor (depending on the rank of the p-value) and compared against $\alpha = 0.05$ (significance level used for all statistical tests in results). Additionally, one-sample Wilcoxon signed-rank tests (one-tailed) were used to determine whether the average d' on each day was significantly greater than zero. Significant comparisons for both tests were indicated on the plot.

2.4.4 Reaction time (single mouse and layer-averaged)

In our paradigm, the stimulus and the response window were initiated simultaneously for each stimulus trial. Therefore, for all auditory/opto-stimulus/whisker hit trials, the reaction time was defined as the difference between the first lick above the noise threshold and the stimulus onset. By design, the reaction time ranged from 0 to 1 second. Reaction times from all auditory and opto/whisker hit trials were averaged separately within each session. The 2 reaction time averages were plotted across days for the single mouse and layer-averaged plots. Wilcoxon signed-rank tests were used to compare the Auditory and Opto/Whisker reaction time averages on each day. Significant comparisons were indicated on the plot.

2.4.5 Opto-control plots (Day 4)

As described in 2.3.3, Day 4 was the opto-control day. The sampling method matched the other single-session graphs. The green and red alternating sections on the plot denoted the insertion (red) of the black tape to block the opto-stimuli from the window and the control unblocked condition (green).

2.4.6 Opto v/s blocked (performance rates and d')

The trials (auditory, opto-stimulus and catch) were classified into 2 conditions (opto or blocked) in the session data. Trials (of the same type) from the same condition were grouped and averaged for each mouse. These performance rates were plotted as three separate graphs. For each graph, the thin lines represented the single mouse performance in both conditions. The larger data points represented the mean across all mice in the opto and blocked conditions (error bars represent SEM). Similarly, a fourth graph was plotted for the opto-stimulus d' (single mouse lines and averages with SEM) across both conditions. Two-tailed (auditory and false alarm) and one-tailed (opto-stimulus and opto-d') paired t-tests were used to compare the means from the opto and blocked conditions. Significant comparisons were indicated on the plot.

2.4.7 Grand average d' plots

To compare the transfer to the opto-stimulus and whisker across different layers (Cre lines), different sampling strategies were used.

First hundred trials – Only the first hundred trials of each session were sampled.

Performance window – A performance window was defined from the start to a certain point in the session. The main criterion for terminating the window was the first instance of 2 consecutive auditory trial misses in the second half of the session. The logic was to exclude trials occurring after loss of motivation at the end of a session.

For both methods mentioned above, the d' values for mice from the same Cre line (layer) were averaged per day. The process was applied for days 0-3. The bars represent the mean d', and the error bars denote the standard error of the mean.

Peak performance – A sliding window (20 trials) with a step of 1 trial was used to calculate averaged opto-stimulus/whisker and false-alarm rates, and subsequently a d' value. After all such d' values were found across the session, the highest d' value was chosen as the peak d' for a specific mouse on a specific day. The process of averaging these across the same layer remains the same.

For the statistical analysis, the Mann-Whitney U test with the Holm-Bonferroni correction was used to perform pairwise comparisons among the layers on each day. Any significant p-values found after applying this correction are indicated on the plots with the standard p-value asterisks.

2.4.8 Performance rates averaged across layers

Average performance rates from each Cre line (layer) were plotted across Days 0-3 as three separate plots (Auditory, Opto/Whisker and False alarm). Multiple pairwise comparisons (among the layers on each day) were done using the Mann-Whitney U test, followed by the Holm-Bonferroni correction. Any significant p-values found after applying this correction are indicated on the plots with the standard p-value asterisks.

2.4.9 Opto-stimulus d' v/s Whisker d'

An average d' for a session (day) was calculated for Day 0 (Opto-stimulus d') and Day 3 (Whisker d'). Each point on the plot represents a single mouse. The data points were coloured by the layer (Cre line) to which each mouse belonged. The X and Y axis limits were kept the same to accurately gauge any patterns, if present. A similar figure, using Day 2 (instead of Day 0) for the Opto-stimulus d', was plotted as well. Spearmann's Correlation was the test used to check for any correlation present in both the Day 0 vs Day 3 and Day 2 vs Day 3 plots. This test was chosen for its suitability to our data (outliers, low N, non-normal distribution). The Correlation coefficient (ρ) and the resulting p-value (at the 95% significance level) are shown on the plots.

2.4.10 Opto-stimulus trials to d' = 1

For each mouse, the sessions from Day 0 to 2 were concatenated in chronological order. The opto-stimulus d' was calculated block-wise (20 trials) from the first opto-stimulus trial on Day 0. The first block at which the d' was equal to or higher than 1 was identified. The number of opto-stimulus trials that occurred till the start of this block was calculated. These were averaged across mice within the same layer and plotted as the mean ± SEM. The Mann-Whitney U test with the Holm-Bonferroni correction was applied to compare layers pairwise. Significant comparisons that persist after this correction are indicated on the plot.

2.4.11 Reaction time across lines

The reaction time values for each stimulus were averaged across the session and, subsequently, across the mice from the same Cre line (layer). One figure depicts the average auditory reaction times across Days 0-3. The other figure depicts the opto-stimulus (Days 0-2) and whisker (Day 3) reaction time averages for all Cre lines. The Mann-Whitney U test with Holm-Bonferroni correction was applied for pairwise comparisons across lines on each day. Significant comparisons that persist after this correction are indicated on the plot.

Chapter 3 Results

3.1 Behaviour and the Trial Structure

Headpost-implanted mice were injected with the AAV expressing excitatory opsin (ChRmine) in the C2 region of the barrel cortex at different depths depending on their Cre line (elaborated more in Section **2.2.2**). After 3-5 weeks of expression, the mice were water-deprived for 24h before starting the behaviour.

Head-fixed, water-restricted mice were trained to lick in response to an auditory tone (10 ms, 10 kHz) to obtain a water reward (Fig. **10**a). Catch trials (no reward or punishment) with no stimulus were interleaved with the auditory trials to measure spontaneous licking from the mice. Each trial was preceded by a quiet window (2-5 s) of no licking, followed by a response window (1 s) during which the mouse had to lick to receive the water reward. Consecutive trials were separated by an Inter-trial interval (6-10 s) (Fig. **10**c)

During Days -4 to -1 (labelled as auditory in Fig. **10**b), there were 10 auditory trials and 10 catch trials for each block (20 trials). To prevent association with any trial pattern, they were distributed pseudo-randomly within each block (to ensure 10 trials of each type occurred). By Day -1, all the mice that had learnt to lick in response to the auditory stimulus were transferred to an optogenetic stimulus on Day 0 (Transfer Day). The opto-stimulus trials consisted of a train of 5 light (625 nm) pulses, each 5 ms in duration, delivered at 50 Hz via a fibre-optic patch cable positioned directly on the cranial window. The power at the tip of the fibre-optic patch cable was 10 mW (Marshel *et al.*, 2019). Essentially, we wanted to test which mice (layer-specific) were able to transfer their learning from the auditory task and respond to the opto-stimulus to obtain the same water reward. During Days 0-2 (labelled as opto-transfer in Fig. **10**b), the ratio of Auditory to Opto-stimulus to Catch trials was 2:8:10 for each block (20 trials).

On Day 3 (Whisker in Fig. **10**b), the opto-stimulus trials were substituted by whisker stimulus (3 ms) trials, which consisted of a metal particle on the C2 whisker being deflected by a magnetic stimulus (caused by a magnetic coil). This was done to test the opto-stimulus to whisker transfer experiments as reported in Sachidhanandam *et al*., (2013).

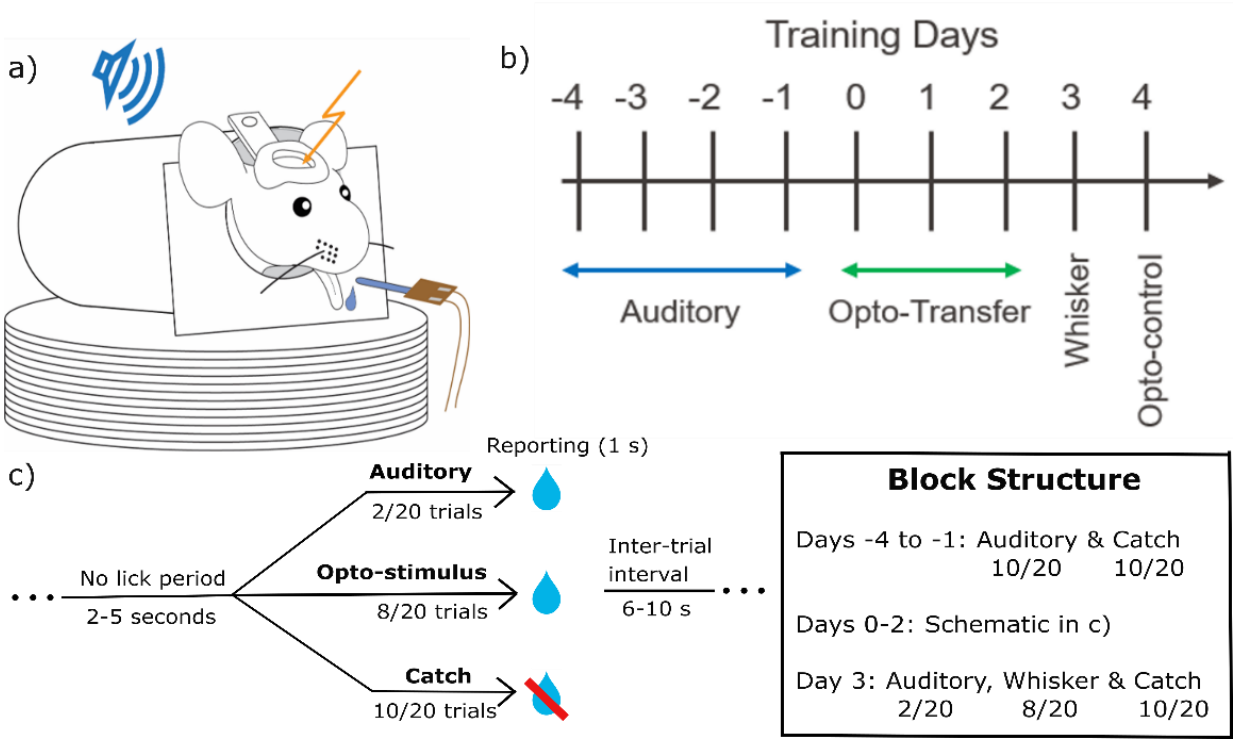

Figure 10 – Behaviour and trial structure: a) Schematic of a head-fixed mouse during behaviour, b) Behavioural training schedule, c) Trial Temporal Structure showing the types of trials (on Days 0-2) and distribution within a block (20 trials). The 'Block Structure' box specifies the distribution of trials within each block for all days.

Lastly, on Day 4, as a control (to identify mice that may be detecting the light flash of the opto-stimulus from peripheral vision), black tape was used to cover the cranial window for a certain number of trials at a time during the session. The trial distribution was identical to the Days 0-2 distribution. Hence, the opto-stimulus was blocked and unblocked in alternating parts of the behavioural session. Mice which licked in response to the light flash using peripheral vision would lick throughout the session and hence fail the control.

3.2 Single mouse and Layer-averaged plots (all 4 Layers)

In this section, a representative example mouse was chosen for each layer. Each figure contains plots of an example single-mouse performance and the layer-specific averaged performance across different behavioural sessions.

3.2.1 Layer 2/3 (Rasgrf2-2A-dCre)

The figures **11**a–**11**f correspond to the histology and performance plots for the layer 2/3 representative mouse (AK010). The remaining ones (g-i) depict the layer-averaged performance.

Fig. **11**a depicts a brain section (of the example mouse) imaged under an epifluorescence microscope with a DsRed filter (545 nm, yellow-green) to examine the opsin expression in the barrel cortex. The opsin is clearly expressed in the L2/3 cells of the barrel cortex (wS1). The darker segmented regions below the expression are the barrels located in L4.

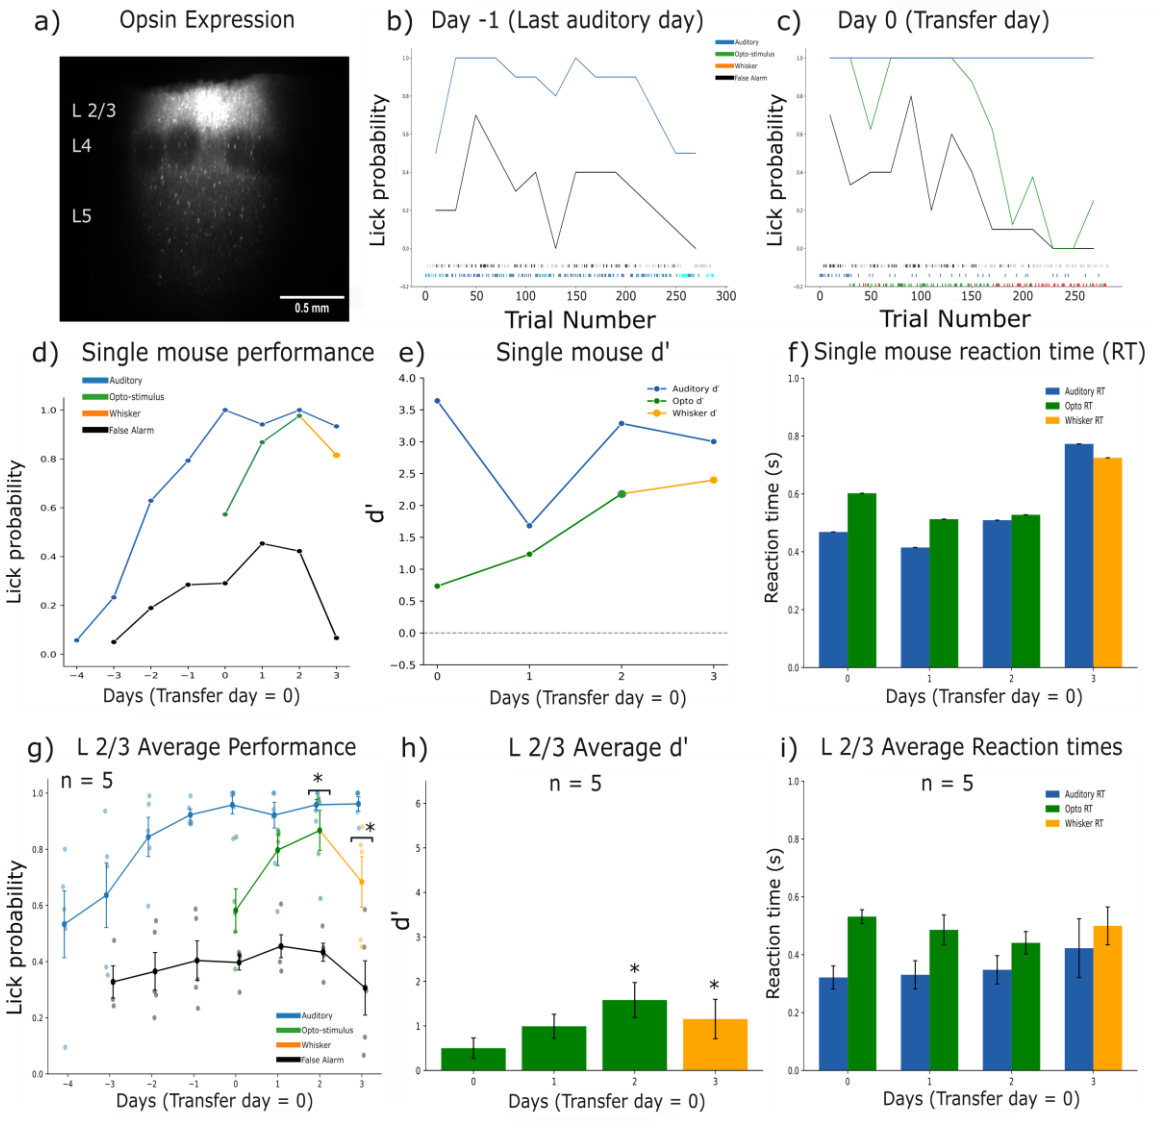

*Figure 11 - Layer 2/3 plots: **a**) Brain section (Scale – 0.5 mm) of the example mouse showing ChRmine expression in Layer 2/3 of the barrel cortex. **b**) Lick probability for Auditory (Blue) and Catch (Black) trials for the example mouse on Day -1. Lick probability indicates performance rates for each trial type (Auditory hit rate and False alarm rate). Raster markers above the trial number axis indicate single-trial outcomes (auditory hit (blue), miss (cyan), false alarm (black) and correct rejection (grey)). **c**) Lick probability across 3 trial types (Auditory, Catch and Opto-stimulus (green)) on Day 0. The colours for auditory and catch trials remain the same as **b**. Raster markers for opto-trial outcomes are shown as green (hit) and red (miss). **d**) Example mouse performance (session averages) across all days (Auditory hit rate – Blue, Opto hit rate – Green, Whisker hit rate – Yellow and False Alarm rate – Black). **e**) d' of the example mouse across days 0-3 for 3 stimulus types. **f**) Reaction time of the example mouse across days (session averages). **g**) Layer-averaged version of **d**. The data points depict the mean, and the error bars depict SEM. The significance asterisks indicate comparisons between the opto-stimulus/whisker hit rates and the false alarm means (one-tailed Wilcoxon signed-rank test) on that day. **h**) Layer-averaged version of **e** (only opto/whisker d' averages plotted, no auditory d'). Height of the bar depicts the mean with the error bar showing SEM. Significance asterisks on top of bars indicate that the d' is significantly greater than zero (one-sample Wilcoxon signed-rank test). The other significance asterisks (with a bracket beneath) denote the pairwise day comparisons (two-tailed Wilcoxon signed-rank tests, followed by the Holm-Bonferroni (HB) correction). The significance level used was 95%. **i**) Layer-averaged version of **f**. Pairwise comparisons on each day using a two-tailed Wilcoxon signed-rank test. Significance asterisks - *(p<0.05), **(p<0.01) and ***(p <0.001).*

On Day -1 (Fig. **11**b), the auditory hit rate (blue) remains high for most of the session, while the false alarm rate (black) fluctuates more but remains within a lower range. This indicates that the mouse (AK010) has learnt to lick specifically in response to the auditory tone. On Day 0 (Fig. **11**c), the opto-stimulus trials (green) are started around the 30th trial. For this mouse, the probability of licking in response to the opto-stimulus is high at the start but

decreases across the session. The false alarm rate also follows a similar trend (hinting at lesser random licking near the session end).

The auditory performance of AK010 shows an increasing trend over the days (Fig. **11**d). The opto-stimulus hit rate also increases across days 0-2, indicating the improvement of the transfer from the auditory to the opto-stimulus. The last point (orange) denotes the whisker hit rate on Day 3 for AK010. A similar picture is seen when the performance rates from all the L2/3 mice are averaged across each day (Fig. **11**g). However, when comparing the opto-stimulus/whisker hit rate and false alarm rates across Days 0-3, significant differences are observed only on Days 2 and 3 (Day 2: Opto-hit rate = 0.87 ± 0.07 (mean ± SEM), False alarm (FA) rate = 0.43 ± 0.03, (Wilcoxon signed-rank test statistic)W = 15, p = 0.03125 and n = 5 mice; Day 3: Whisker rate = 0.68 ± 0.09, FA rate = 0.31 ± 0.09, W = 15, p = 0.03125 and n = 5 mice). So, L2/3 mice did not respond to the opto-stimulus on Days 0 and 1 (Day 0: Opto-hit rate = 0.58 ± 0.08, FA = 0.39 ± 0.03, W = 14, p = 0.0625 and n = 5; Day 1: Opto-hit rate = 0.79 ± 0.05, FA = 0.45 ± 0.04, W = 14, p = 0.0625 and n = 5).

The d' (Auditory and Opto) represents how well the mouse was able to selectively lick for the auditory stimulus and opto-stimulus (Fig. **11**e). The opto-stimulus d' increases gradually across the days (reflecting the same inference from Fig. **11**d). The layer-averaged trend of the opto/whisker d' (Fig. **11**h) also reflects the same observations from Fig. **11**g (d' significantly higher than zero only on Days 2 and 3). Additionally, from the pairwise comparison across days, no significant differences were observed (no significance asterisks with brackets in Fig. **11**h). The closest to a significant comparison was Day 0 vs 2 (Day 0: Opto-d' = 0.5 ± 0.23 (mean ± SEM), Day 2: Opto-d' = 1.58 ± 0.39, HB corrected Wilcoxon signed-rank test, p = 0.375, n = 5 mice). Therefore, while there is a visible trend (across Days 0-2), more data would be needed to validate it statistically. Lastly, the Whisker d' being significantly greater than 0 (Day 3 in Fig. **11**h) means that the L2/3 mice responded to the whisker stimulus.

The reaction times (RT) (auditory, opto-stimulus and whisker) don't show any consistent trend in the single mouse (Fig. **11**f) and layer-averaged (Fig. **11**i) plots. While the RTs on Day 3 for the example mouse (Fig. **11**f) look higher than the other days, the average shows no such difference (Fig. **11**i). Pairwise comparisons between the auditory and opto/whisker RTs didn't show any significant differences (Day 0: Auditory RT = 0.32 ± 0.04 (mean ± SEM), Opto RT = 0.53 ± 0.02, (Wilcoxon signed-rank test statistic) W = 0, p = 0.0625 and n = 5; Day 1: Auditory RT = 0.33 ± 0.05, Opto RT = 0.49 ± 0.05, W = 0, p = 0.0625 and n = 5; Day 2: Auditory RT = 0.35 ± 0.05, Opto RT = 0.44 ± 0.04, W = 0, p = 0.0625 and n = 5; Day 3: Auditory RT = 0.42 ± 0.10, Whisker RT = 0.50 ± 0.06, W = 3, p = 0.3125 and n = 5).

3.2.2 Layer 4 (Scnn1a-Cre)

The figure layout is the same across all layers. Fig. **12**a depicts a brain section (of the example mouse AK040) imaged under an epifluorescence microscope with a DsRed filter (545 nm, yellow-green) to examine the opsin expression in the barrel cortex. The opsin is strongly expressed in the L4 cells of the barrel cortex (wS1). The bright, separated compartments are the different barrels that express the opsin in the wS1 of this mouse.

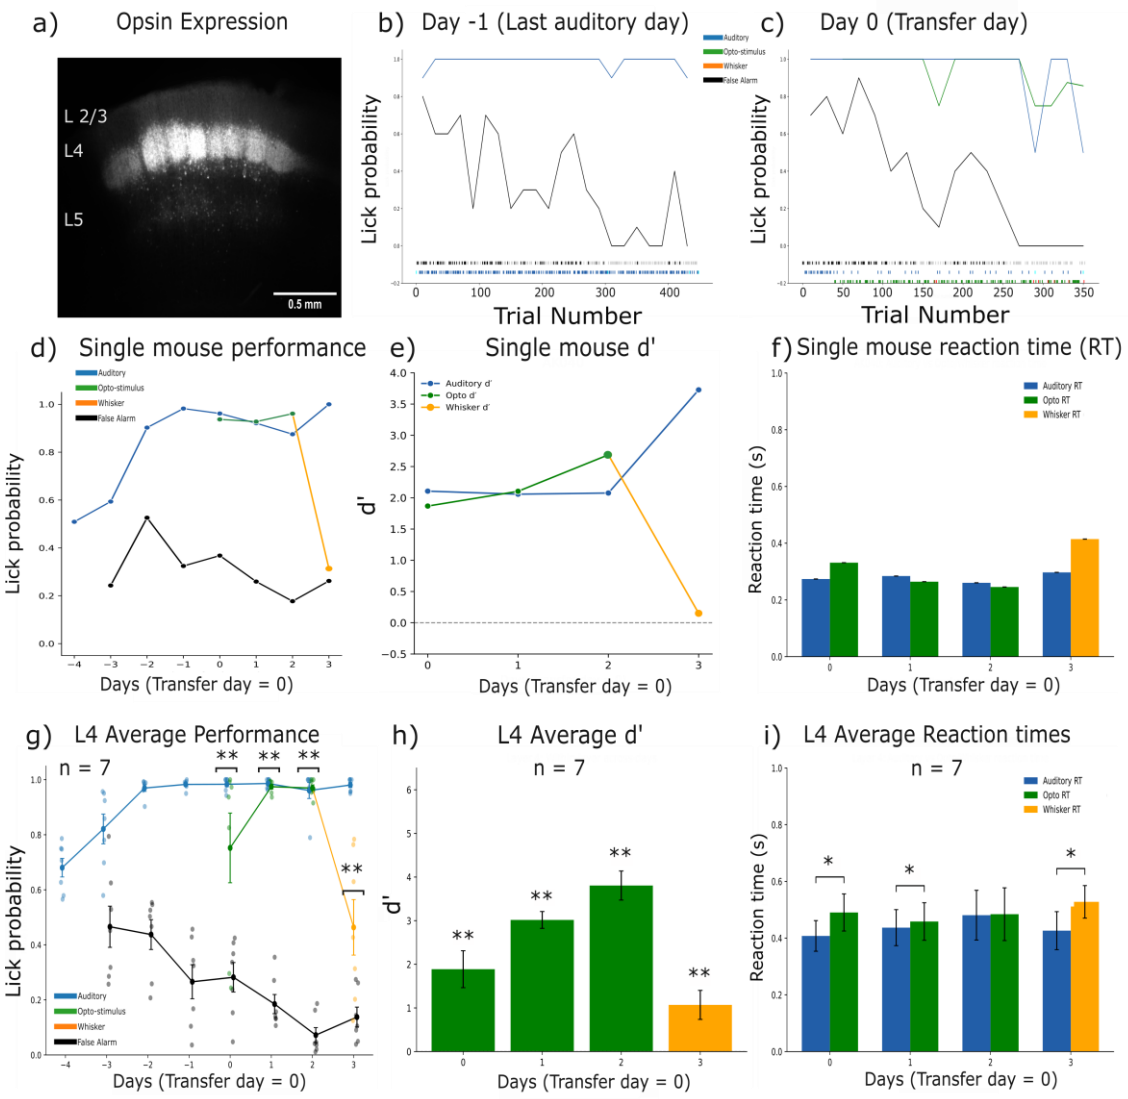

*Figure 12 – Layer 4 plots: **a**) Brain section (Scale – 0.5 mm) of the example mouse showing ChRmine expression in Layer 4 of the barrel cortex. **b**) Lick probability for Auditory (Blue) and Catch (Black) trials for the example mouse on Day -1. Lick probability indicates performance rates for each trial type (Auditory hit rate and False alarm rate). Raster markers above the trial number axis indicate single-trial outcomes (auditory hit (blue), miss (cyan), false alarm (black) and correct rejection (grey)). **c**) Lick probability across 3 trial types (Auditory, Catch and Opto-stimulus (green)) on Day 0. The colours for auditory and catch trials remain the same as **b**. Raster markers for opto-trial outcomes are shown as green (hit) and red (miss). **d**) Example mouse performance (session averages) across all days (Auditory hit rate – Blue, Opto hit rate – Green, Whisker hit rate – Yellow and False Alarm rate – Black). **e**) d' of the example mouse across days 0-3 for 3 stimulus types. **f**) Reaction time of the example mouse across days (session averages). **g**) Layer-averaged version of **d**. The data points depict the mean, and the error bars depict SEM. The significance asterisks indicate comparisons between the opto-stimulus/whisker hit rates and the false alarm means (one-tailed Wilcoxon signed-rank test) on that day. **h**) Layer-averaged version of **e** (only opto/whisker d' averages plotted, no auditory d'). Height of the bar depicts the mean with the error bar showing SEM. Significance asterisks on top of bars indicate that the d' is significantly greater than zero (one-sample Wilcoxon signed-rank test). The other significance asterisks (with a bracket beneath) denote the pairwise day comparisons (two-tailed Wilcoxon*

*signed-rank tests, followed by the Holm-Bonferroni (HB) correction). The significance level used was 95%. **i)** Layer-averaged version of **f**. Pairwise comparisons on each day using a two-tailed Wilcoxon signed-rank test. Significance asterisks - *(p<0.05), **(p<0.01) and ***(p <0.001).*

The Day -1 (Fig. **12**b) session for the mouse AK040 depicts that the mouse has learnt to respond to the auditory stimulus. On Day 0 (Fig. **12**c), AK040 strongly responded to the opto-stimulus (Auditory and opto-stimulus curves at the same level) throughout the session.

AK040's strong transfer to the opto-stimulus is consistent across Days 0-2 (Fig. **12**d). But, it did not respond well to the whisker stimulus on Day 3 (Fig. **12**d). A similar picture is seen when the performance rates from all the L4 mice are averaged across each day (Fig. **12**g). Unlike L2/3 mice, L4 mice strongly responded to the opto-stimulus right from Day 0 to Day 2 (Day 0: Opto-hit rate = 0.75 ± 0.13 (mean ± SEM), False alarm (FA) rate = 0.28 ± 0.05, (Wilcoxon signed-rank test statistic) W = 28, p = 0.007813 and n = 7; Day 1: Opto-hit rate = 0.97 ± 0.01, FA rate = 0.18 ± 0.03, W = 28, p = 0.007813 and n = 7; Day 2: Opto-hit rate = 0.96 ± 0.01, False alarm (FA) rate = 0.07 ± 0.03, W = 28, p = 0.007813 and n = 7). On average, the L4 mice also responded strongly to the whisker stimulus (Day 3 in Fig. **12**g) (Day 3: Whisker hit rate = 0.46 ± 0.1, FA = 0.14 ± 0.04, W = 28, p = 0.007813 and n = 7).

The d' (Auditory and Opto) for AK040 (Fig. **12**e) shows the same trends seen in Fig. **12**d. The layer-averaged trend of the opto/whisker d' (Fig. **12**h) also reflects the same observations from Fig. **12**g (Opto-d' and Whisker d' being significantly higher than zero across Days 0-3). Pairwise comparisons across days did not yield any significant results (no significance asterisks with brackets). The closest to significance comparisons were Day 1 vs 3 and Day 2 vs 3 (Day 1: Opto-d' = 3.02 ± 0.19 (mean ± SEM), Day 2: Opto-d' = 3.8 ± 0.33, Day 3: Whisker d' = 1.07 ± 0.33, HB corrected Wilcoxon signed-rank test, p = 0.09375 (for both Day 1 vs 3 and Day 2 vs 3), n = 7). Overall, L4 mice showed a strong, statistically equivalent response to the optogenetic and whisker stimuli. Furthermore, the response to the opto-stimulus remained constant across Days 0-2 (opto-transfer days).

The reaction times (RT) (auditory, opto-stimulus, and whisker) show a similar trend in AK040 (Fig. **12**f) and L4 mice on average (Fig. **12**i). The paired comparisons between the auditory and opto RT means show a clear trend of decreasing differences across Days 0-2 (significant differences seen on Days 0 and 1 but not on Day 2). This implies that the L4 mice respond to the opto-stimulus at similar speeds to the auditory stimulus across days. The difference returns on Day 3 when they are exposed to the whisker stimulus (Fig. **12**i) (Day 0: Auditory RT = 0.41 ± 0.05 (mean ± SEM), Opto RT = 0.49 ± 0.06, (Wilcoxon signed-rank test statistic) W = 0, p = 0.015625 and n = 7; Day 1: Auditory RT = 0.44 ± 0.06, Opto RT = 0.46 ± 0.07, W = 2, p = 0.046875 and n = 7; Day 2: Auditory RT = 0.48 ± 0.09, Opto RT =

0.48 ± 0.09, W = 12, p = 0.8125 and n = 7; Day 3: Auditory RT = 0.43 ± 0.07, Whisker RT = 0.53 ± 0.06, W = 0, p = 0.015625 and n = 7).

3.2.3 Layer 5-IT (Tlx3-Cre)

Fig. **13**a depicts a brain section (of the example mouse AK044) imaged under an epifluorescence microscope with a DsRed filter (545 nm, yellow-green) to examine the opsin expression in the barrel cortex. The opsin is expressed in the L5 region of the wS1. The thin apical dendrites emerging from the expression layer outline the barrel boundaries and confirm the morphology observed in other studies of layer 5-IT neurons.

The Day -1 (Fig. **13**b) session for the mouse AK044 shows that the mouse has learnt to respond to the auditory stimulus. On Day 0 (Fig. **13**c), AK044 strongly responded to the opto-stimulus (auditory and opto-stimulus curves at the same level) throughout the session.

AK044's strong transfer to the opto-stimulus is consistent across Days 0-2 (Fig. **13**d). But, it did not respond much to the whisker stimulus on Day 3 (Fig. **13**d). A similar trend is seen when the performance rates from all the L5-IT mice are averaged across each day (Fig. **13**g). Like L4 mice, L5-IT mice responded (not as strong as L4) to the opto-stimulus right from Day 0 to Day 2 (Day 0: Opto-hit rate = 0.77 ± 0.06 (mean ± SEM), False alarm (FA) rate = 0.44 ± 0.05, (Wilcoxon signed-rank test statistic) W = 15, p = 0.03125 and n = 5; Day 1: Opto-hit rate = 0.96 ± 0.01, FA rate = 0.28 ± 0.06, W = 15, p = 0.03125 and n = 5; Day 2: Opto-hit rate = 0.98 ± 0.01, False alarm (FA) rate = 0.13 ± 0.05, W = 15, p = 0.03125 and n = 5). On average, the L5-IT mice responded (less strongly than L4) to the whisker stimulus (Day 3 in Fig. **13**g) (Day 3: Whisker hit rate = 0.42 ± 0.13, FA = 0.24 ± 0.05, W = 15, p = 0.03125 and n = 5).

The d' (Auditory and Opto) for AK044 (Fig. **13**e) shows the same trends seen in Fig. **13**d. The layer-averaged trend of the opto/whisker d' (Fig. **13**h) also reflects the same observations from Fig. **13**g (Opto-d' and Whisker d' being significantly higher than zero across Days 0-3). Pairwise comparisons across days did not yield any significant results (no significance asterisks with brackets). Every comparison except Day 0 vs 3 had the same p-value, which was the closest to significance. (Day 0: Opto-d' = 0.98 ± 0.23 (mean ± SEM), Day 1: Opto-d' = 2.59 ± 0.44, Day 2: Opto-d' = 3.59 ± 0.37, Day 3: Whisker d' = 0.56 ± 0.26, HB corrected Wilcoxon signed-rank test, p = 0.375 (for all except Day 0 vs 3) and p = 0.625 (Day 0 vs 3), n = 5). Overall, L5-IT mice showed a statistically equivalent response to the optogenetic and whisker stimuli. The response to the opto-stimulus remained at similar levels across Days 0-

2. L5-IT mice showed a less significant (yet considerable) response to auditory and whisker stimuli compared to L4 mice (Figures **12**h and **13**h).

The reaction times (RT) (auditory, opto-stimulus, and whisker) show a similar trend in AK044 (Fig. **13**f) and L5-IT mice on average (Fig. **13**i). However, unlike L4 mice (Fig. **12**i), none of the paired comparisons between the auditory and opto/whisker RTs is significant for L5 mice (Fig. **13**i). While there are some visual similarities in the L4 (Fig. **12**i) and L5-IT (Fig. **13**i) RT plots, data from more L5-IT mice are needed to statistically evaluate whether the trend in differences in L5-IT matches that in L4 (Day 0: Auditory RT = 0.37 ± 0.07 (mean ± SEM), Opto RT = 0.53 ± 0.07, (Wilcoxon signed-rank test statistic) W = 0, p = 0.0625 and n = 5; Day 1: Auditory RT = 0.37 ± 0.07, Opto RT = 0.37 ± 0.07, W = 7, p = 1 and n =5; Day 2: Auditory RT = 0.32 ± 0.06, Opto RT = 0.33 ± 0.05, W = 4, p = 0.4375 and n = 5; Day 3: Auditory RT = 0.37 ± 0.07, Whisker RT = 0.51 ± 0.07, W = 0, p = 0.0625 and n = 5).

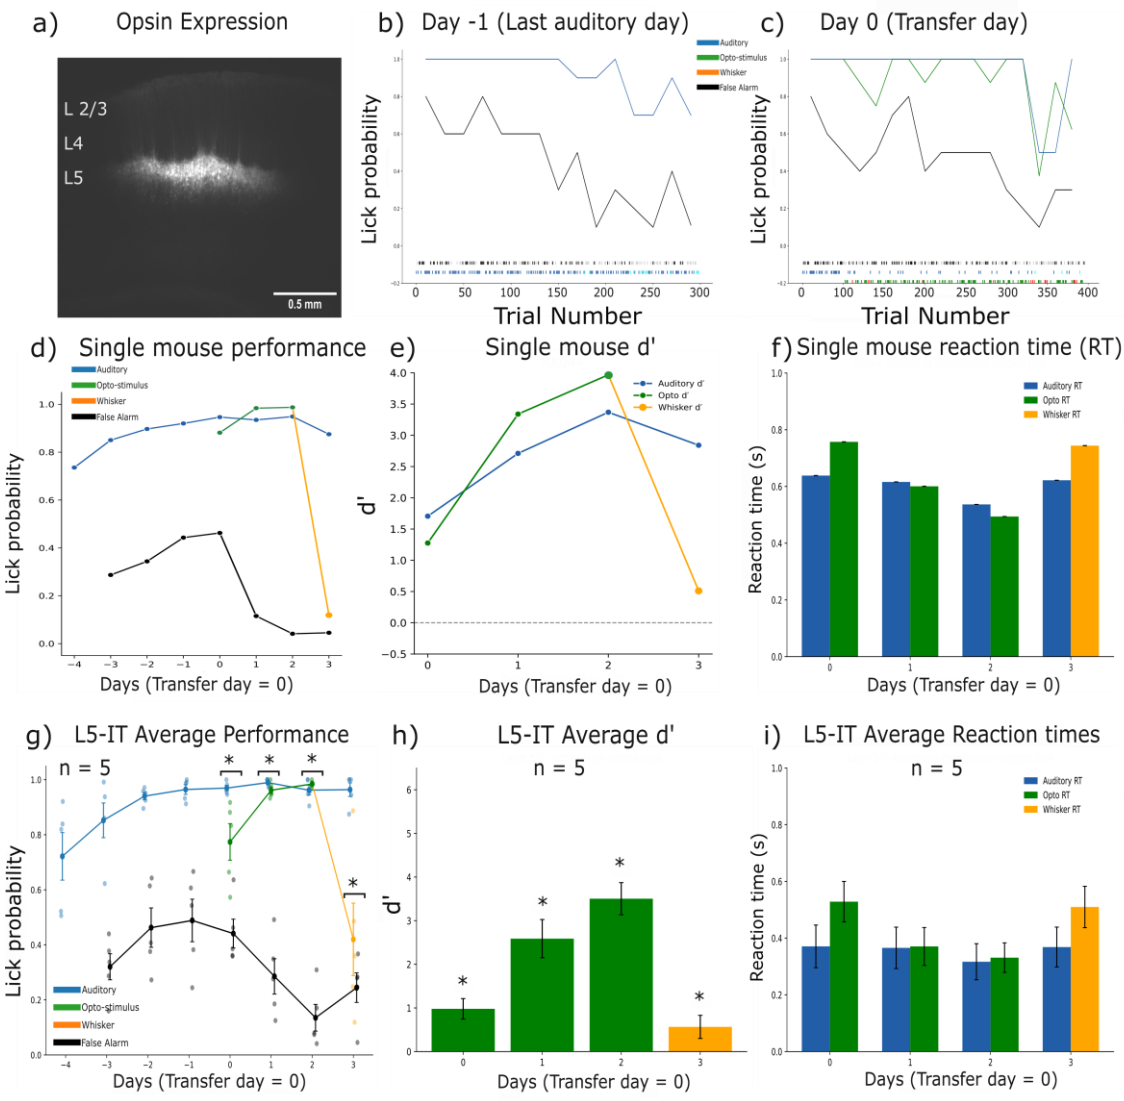

*Figure 13 – Layer 5-IT plots: **a**) Brain section (Scale – 0.5 mm) of the example mouse showing ChRmine expression in Layer 5 of the barrel cortex. **b**) Lick probability for Auditory (Blue) and Catch (Black) trials for the example mouse on Day -1. Lick*

probability indicates performance rates for each trial type (Auditory hit rate and False alarm rate). Raster markers above the trial number axis indicate single-trial outcomes (auditory hit (blue), miss (cyan), false alarm (black) and correct rejection (grey)). ***c****) Lick probability across 3 trial types (Auditory, Catch and Opto-stimulus (green)) on Day 0. The colours for auditory and catch trials remain the same as* ***b****. Raster markers for opto-trial outcomes are shown as green (hit) and red (miss).* ***d****) Example mouse performance (session averages) across all days (Auditory hit rate – Blue, Opto hit rate – Green, Whisker hit rate – Yellow and False Alarm rate – Black).* ***e****) d' of the example mouse across days 0-3 for 3 stimulus types.* ***f****) Reaction time of the example mouse across days (session averages).* ***g****) Layer-averaged version of* ***d****. The data points depict the mean, and the error bars depict SEM. The significance asterisks indicate comparisons between the opto-stimulus/whisker hit rates and the false alarm means (one-tailed Wilcoxon signed-rank test) on that day.* ***h****) Layer-averaged version of* ***e*** *(only opto/whisker d' averages plotted, no auditory d'). Height of the bar depicts the mean with the error bar showing SEM. Significance asterisks on top of bars indicate that the d' is significantly greater than zero (one-sample Wilcoxon signed-rank test). The other significance asterisks (with a bracket beneath) denote the pairwise day comparisons (two-tailed Wilcoxon signed-rank tests, followed by the Holm-Bonferroni (HB) correction). The significance level used was 95%.* ***i****) Layer-averaged version of* ***f****. Pairwise comparisons on each day using a two-tailed Wilcoxon signed-rank test. Significance asterisks - *(p<0.05), **(p<0.01) and ***(p <0.001).*

3.2.4 Layer 5-PT (Sim1-Cre)

Fig. **14**a depicts a brain section (of the example mouse AK030) imaged under an epifluorescence microscope with a DsRed filter (545 nm, yellow-green) to examine the opsin expression in the barrel cortex. The opsin is expressed sparsely in Layer 5 of the wS1 for this mouse. The faint expression present above the L2/3 region might correspond to the apical dendrites of the L5-PT neurons.

The Day -1 (Fig. **14**b) session for the mouse AK030 shows that the mouse has learnt to respond to the auditory stimulus. On Day 0 (Fig. **14**c), AK030 doesn't respond to the opto-stimulus. The opto-hit and false alarm (FA) rates fluctuate within the same range, indicating the same. The auditory hit rate, which remained at 1.0, confirms that the mouse did not lose motivation during the session.

AK030 showed a highly fluctuating transfer to the opto-stimulus across Days 0-2 (Fig. **14**d). It also did not respond much to the whisker stimulus on Day 3 (Fig. **14**d). The average performance rates across all L5-PT mice on each day paint a different picture (Fig. **14**g). L5-PT mice responded to the opto-stimulus right from Day 0 to Day 2 (Day 0: Opto-hit rate = 0.61 ± 0.08 (mean ± SEM), False alarm (FA) rate = 0.38 ± 0.03, (Wilcoxon signed-rank test statistic) W = 33, p = 0.019531 and n = 8; Day 1: Opto-hit rate = 0.65 ± 0.1, FA rate = 0.33 ± 0.04, W = 35, p = 0.007813 and n = 8; Day 2: Opto-hit rate = 0.85 ± 0.06, False alarm (FA) rate = 0.33 ± 0.06, W = 36, p = 0.003906 and n = 8). On average, the L5-PT mice also responded well to the whisker stimulus (Day 3 in Fig. **14**g) (Day 3: Whisker hit rate = 0.58 ± 0.11, FA = 0.37 ± 0.08, W = 36, p = 0.003906 and n = 8).

The d' (Auditory and Opto) for AK030 (Fig. **14**e) shows the same trends seen in Fig. **14**d. The layer-averaged trend of the opto/whisker d' (Fig. **14**h) also reflects the same observations from Fig. **14**g (Opto-d' and Whisker d' being significantly higher than zero across Days 0-3). Pairwise comparisons across days yield significant results for Day 0 vs 2 and Day 1 vs 2

(Day 0: Opto-d' = 0.63 ± 0.23 (mean ± SEM), Day 1: Opto-d' = 1.04 ± 0.39, Day 2: Opto-d' = 1.92 ± 0.49, HB corrected Wilcoxon signed-rank test, p = 0.046875 (for both Day 0 vs 2 and Day 1 vs 2), n = 8). Overall, L5-PT mice showed a statistically equivalent response to the optogenetic and whisker stimuli. Unlike other layers, the response to the opto-stimulus increased across Days 0-2 for L5-PT mice.

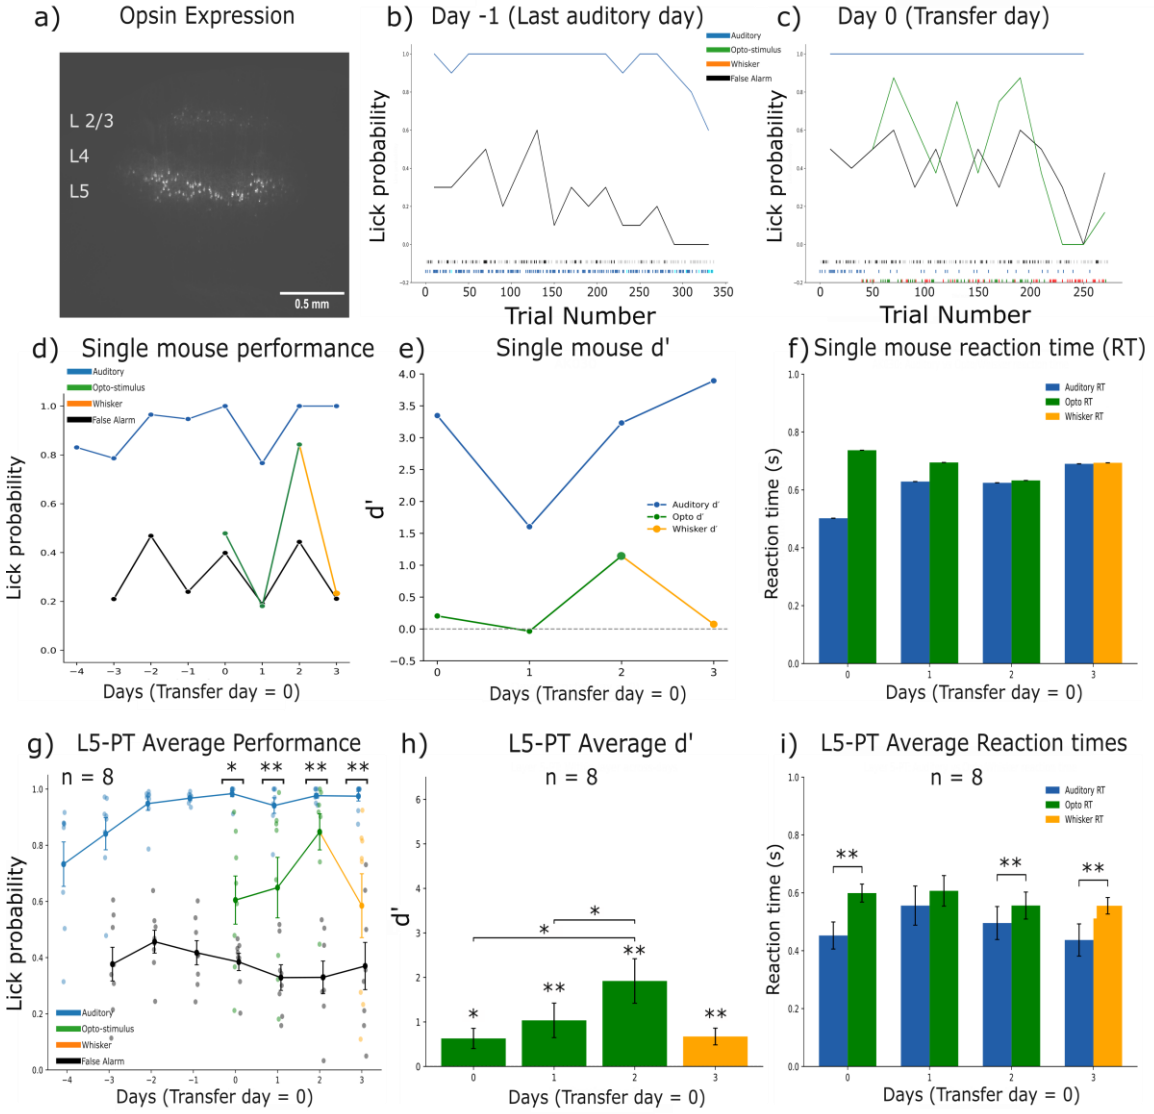

*Figure 14 – Layer 5-PT plots: **a**) Brain section (Scale – 0.5 mm) of the example mouse showing ChRmine expression in Layer 5 of the barrel cortex. **b**) Lick probability for Auditory (Blue) and Catch (Black) trials for the example mouse on Day -1. Lick probability indicates performance rates for each trial type (Auditory hit rate and False alarm rate). Raster markers above the trial number axis indicate single-trial outcomes (auditory hit (blue), miss (cyan), false alarm (black) and correct rejection (grey)). **c**) Lick probability across 3 trial types (Auditory, Catch and Opto-stimulus (green)) on Day 0. The colours for auditory and catch trials remain the same as **b**. Raster markers for opto-trial outcomes are shown as green (hit) and red (miss). **d**) Example mouse performance (session averages) across all days (Auditory hit rate – Blue, Opto hit rate – Green, Whisker hit rate – Yellow and False Alarm rate – Black). **e**) d' of the example mouse across days 0-3 for 3 stimulus types. **f**) Reaction time of the example mouse across days (session averages). **g**) Layer-averaged version of **d**. The data points depict the mean, and the error bars depict SEM. The significance asterisks indicate comparisons between the opto-stimulus/whisker hit rates and the false alarm means (one-tailed Wilcoxon signed-rank test) on that day. **h**) Layer-averaged version of **e** (only opto/whisker d' averages plotted, no auditory d'). Height of the bar depicts the mean with the error bar showing SEM. Significance asterisks on top of bars indicate that the d' is significantly greater than zero (one-sample Wilcoxon signed-rank test). The other significance asterisks (with a bracket beneath) denote the pairwise day comparisons (two-tailed Wilcoxon signed-rank tests, followed by the Holm-Bonferroni (HB) correction). The significance level used was 95%. **i**) Layer-averaged*

*version of **f**. Pairwise comparisons on each day using a two-tailed Wilcoxon signed-rank test. Significance asterisks - *(p<0.05), **(p<0.01) and ***(p <0.001).*

The reaction times (RT) (auditory, opto-stimulus, and whisker) show a slightly different trend in AK044 (Fig. **14**f) and L5-PT mice on average (Fig. **14**i). While the pairwise comparisons between the auditory and opto/whisker RTs do show significant differences on Days 0, 1 and 3, there is no consistent trend of differences across days (Fig. **14**i) (Day 0: Auditory RT = 0.45 ± 0.05 (mean ± SEM), Opto RT = 0.6 ± 0.03, (Wilcoxon signed-rank test statistic) W = 0, p = 0.007813 and n = 8; Day 1: Auditory RT = 0.56 ± 0.07, Opto RT = 0.61 ± 0.05, W = 8, p = 0.195313 and n = 8; Day 2: Auditory RT = 0.5 ± 0.06, Opto RT = 0.56 ± 0.05, W = 0, p = 0.007813 and n = 8; Day 3: Auditory RT = 0.44 ± 0.06, Whisker RT = 0.55 ± 0.03, W = 0, p = 0.007813 and n = 8).

3.3 Opto-control (Day 4) Analysis

On Day 4 (opto-control), we wanted to identify mice (if any) which licked in response to the opto-stimulus light flash using peripheral vision (rather than the activation of wS1 layer neurons). Hence, the session was divided into phases. During alternate phases, black tape was inserted between the cranial window and the end of the fibre-optic cable (which emits the light pulses). This insertion ensured that no light would permeate beneath the cranial window, but the residual light bouncing off the surface would remain the same as in the unblocked condition.

The single-session performance (Figures **15**a and **15**b) of 2 example mice (L4 and L5-IT) shows alternating phases of opto (green) and blocked (red). At transitions, the opto-hit rate (green) either rises (opto to blocked) or falls quickly (blocked to opto), indicating that the mice are unable to detect the opto-stimulus and therefore stop licking in the opto-stimulus trials during the blocked phases.

The auditory hit rates, false alarm rates and opto-stimulus hit rates were calculated and averaged separately for the opto and blocked conditions for each mouse (represented by a line in Figures **15**c, **15**d and **15**e). The mean auditory hit rate doesn't change with the conditions (Fig. **15**c) as expected (auditory hit rate (opto phase) = 0.95 ± 0.17 (mean ± SEM), auditory hit rate (blocked phase) = 0.95 ± 0.16, t-statistic = -0.15, p = 0.881894 and n = 24). The false alarm drops slightly (but significantly) in the blocked phase compared to the opto phase (Fig. **15**d) (FA rate (opto phase) = 0.26 ± 0.03 (mean ± SEM), FA rate (blocked phase) = 0.2 ± 0.03, t-statistic = 3.89, p = 0.000733 and n = 24). The mean opto-stimulus hit rate drops drastically in the blocked phase compared to the opto phase (Fig. **15**e) (opto-hit rate (opto

phase) = 0.9 ± 0.03 (mean ± SEM), opto-hit rate (blocked phase) = 0.21 ± 0.03, t-statistic = 15.96, p = 3.11×10^{-14} and n = 24).

The opto-d' also drops significantly from the opto to blocked condition (Fig. **15**f) (d' (opto phase) = 2.6 ± 0.28 (mean ± SEM), d' (blocked phase) = -0.01 ± 0.07, t-statistic = 3.97, p = 1.22×10^{-9} and n = 24). The clustering of the d' values at d' = 0 indicates that there are no mice with high d' in the Blocked condition (all of them pass the control).

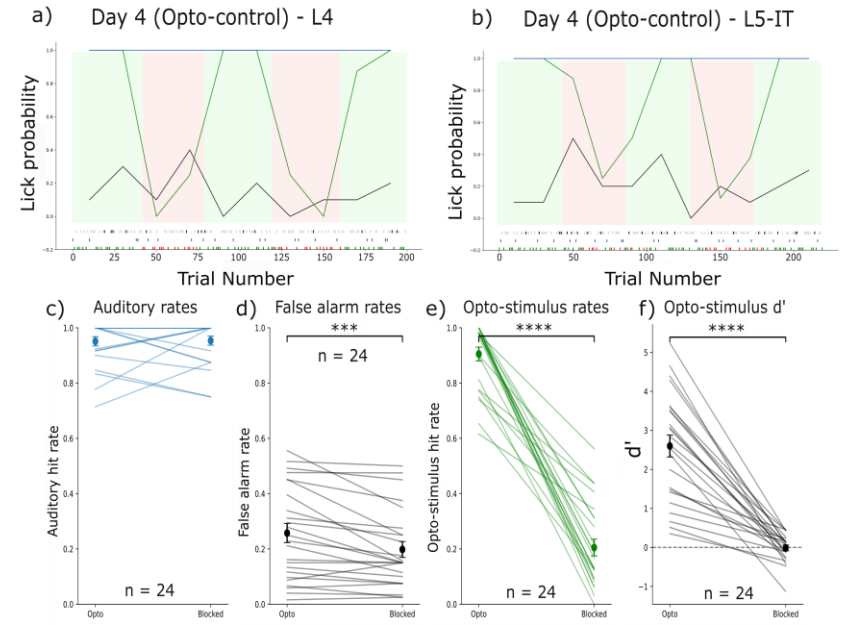

*Figure 15 – Opto-control plots: **a**, **b**) Single session performance rate plots of 2 example mice (Layer 4 and Layer 5-IT). The green shaded regions indicate the regular opto condition (unblocked), and the red regions indicate the blocked condition. Auditory hit rate in blue, Opto-stimulus rate in green and False alarm rate in black. Raster markers follow the same colour code as the single-session plots in Section 3.1. **c**) Opto (unblocked) vs blocked averaged auditory hit rates per mouse. Each line represents a single mouse. The data points represent the means, and the error bars represent SEM. A paired t-test was used to compare the means between the two conditions. **d**) Opto (unblocked) vs blocked averaged FA rates per mouse. Each line represents a single mouse. The data points represent the means, and the error bars represent SEM. A paired t-test was used to compare the means from the two conditions. **e**) Opto (unblocked) vs blocked averaged opto-stimulus rates per mouse. Each line represents a single mouse. The data points represent the means, and the error bars represent SEM. A paired t-test was used to compare the means from the two conditions. **f**) Same as e, but plots average opto d' instead of opto hit rates. A paired t-test was used to compare the means from the two conditions. Significance represents the result of a paired t-test. Significance asterisks - *(p<0.05), **(p<0.01), ***(p<0.001) and ****(p<0.0001).*

On the whole, all the mice passed the control and were therefore included in the grand-average graphs of the next section.

3.3 Grand average plots and additional analysis

Various behavioural parameters were analysed to compare across mice from different layers (Fig. **16**). The average auditory hit rates, opto/whisker hit rates, and false alarm rates for each layer were plotted groupwise across Days 0-3 (Figures **16**b, **16**c, and **16**d). These averages (and SEMs) are the same values observed on Days 0-3 in the individual layer plots (Figures **11**g, **12**g, **13**g and **14**g). The auditory hit rate averages lie around the same level for all the

layers (Fig. **16**b), confirming that they were all able to learn the auditory task. The multiple pairwise comparisons among all layers on each day did not yield any significant differences.

The opto-hit rate (Days 0-2, Fig. **16**c) shows broad variation in response to the opto-stimulus across layers, especially on Day 1. While all layers except L2/3 responded to the opto-stimulus from Day 0, no significant differences were found among them (Day 0, Fig. **16**c). A groupwise comparison (Kruskal-Wallis ANOVA test) to assess whether one of the layers was significantly different from the others on Day 0 also yielded an insignificant p-value ($H = 3.5759$, $p = 0.311$). However, Day 1 showed significant differences in many pairwise comparisons. L4 was significantly higher than L2/3 and L5-PT (L4 opto-hit rate = 0.75 ± 0.13 (mean ± SEM), n = 7; L5-PT opto-hit rate = 0.6 ± 0.09, n = 8; L2/3 opto-hit rate = 0.58 ± 0.08, n = 5; HB corrected Mann-Whitney U test, $p = 0.015$ (L4 vs L2/3) and $p = 0.019$ (L4 vs L5-PT)). L5-IT was higher than L2/3 (L5-IT opto-hit rate = 0.77 ± 0.07, n = 5; HB corrected $p = 0.032$ (L5-IT vs L2/3)), and no significant difference was observed between L4 and L5-IT (HB corrected $p = 1$). The pairwise differences (Mann-Whitney U) and groupwise differences (Kruskal-Wallis) were insignificant on Day 2 (Fig. **16**c) (all layers had a strong response to the opto-stimulus). The Day 3 comparisons for the whisker hit rate (Fig. **16**c) were insignificant. Meanwhile, the false alarm rate averages showed some sporadic differences on Days 1 and 2 (Fig. **16**d) across the layers.

To account for differences in the observed false alarm rates, the d' averages were plotted across Days 0-3 using 3 sampling methods (Figures **16**a, **16**e, and **16**f).). The performance window (Fig. **16**a) criterion was used to trim session duration at the end of a session when the mice lost motivation (used for all averaged data). Similar to the hit rate plots, the d' values (Fig. **16**a) are the same values observed on Days 0-3 in the individual layer plots (Figures **11**h, **12**h, **13**h and **14**h). The significant differences of d' from the pairwise comparisons (Fig. **16**a) match the differences observed for the opto-hit rates (Fig. **16**c) on Days 0 and 1. However, on Day 2, L4 (unlike L5-IT) maintains the difference with L2/3 and L5-PT (L4 opto-hit rate = 3.8 ± 0.33 (mean ± SEM), n = 7; L5-PT opto-hit rate = 1.92 ± 0.49, n = 8; L2/3 opto-hit rate = 1.58 ± 0.39, n = 5; HB corrected Mann-Whitney U test, $p = 0.015$ (L4 vs L2/3) and $p = 0.046$ (L4 vs L5-PT)). No significant differences were observed on Day 3 (Fig. **16**a), indicating that the transfer to the whisker stimulus didn't differ much across the layers.

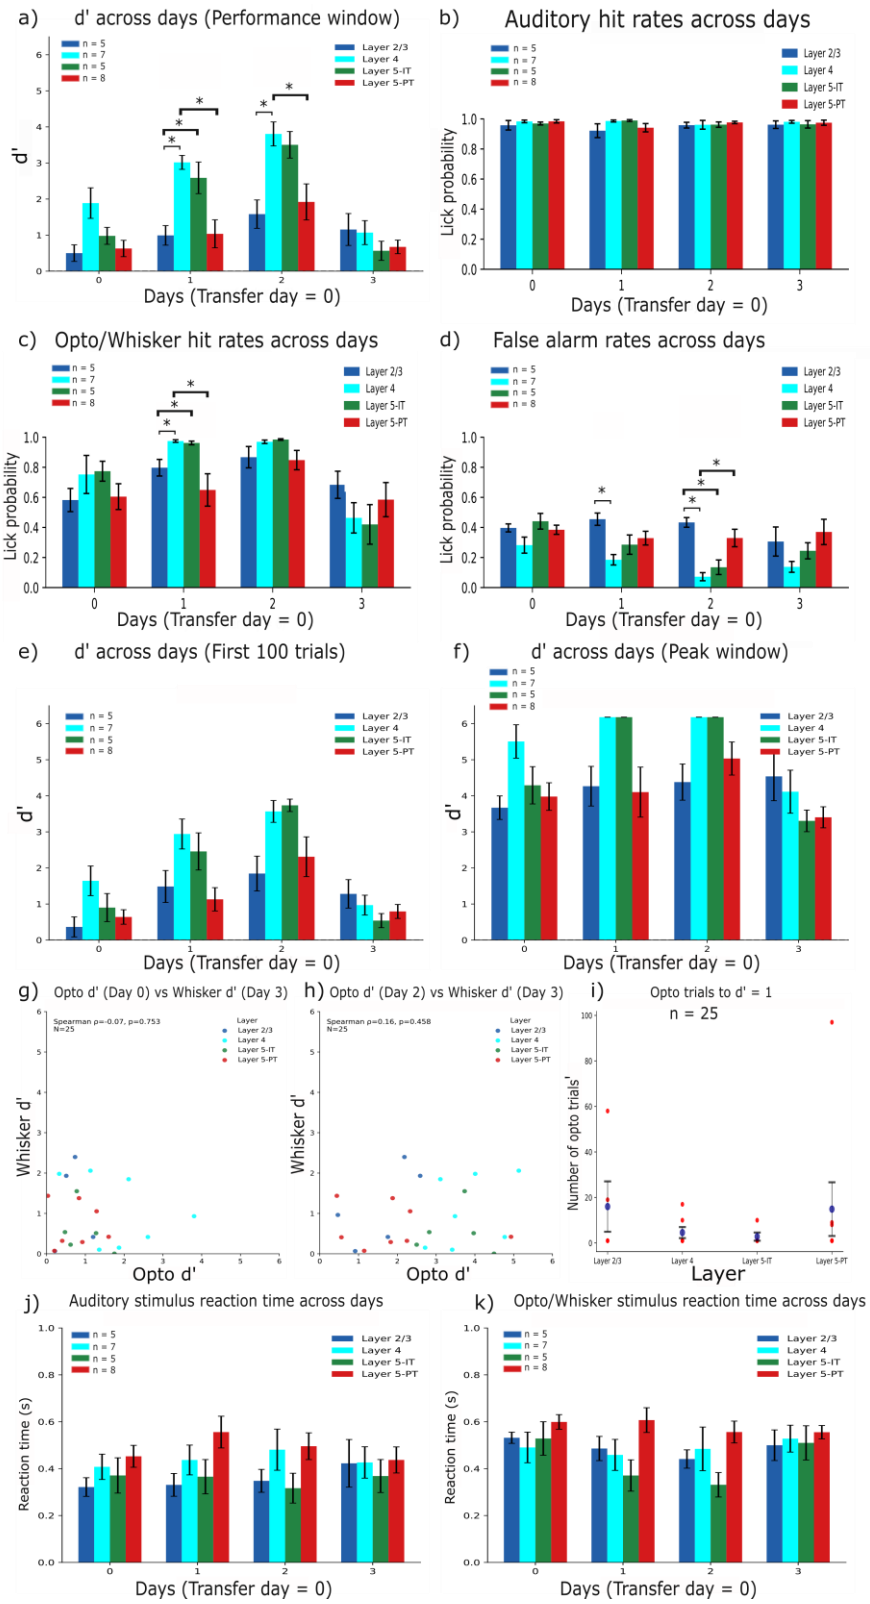

*Figure 16 – Grand average plots and additional analysis: **a**) d' across days of all layers (based on the performance window criterion). Pairwise comparisons among the four layers on each day were performed using the Mann-Whitney U test with the Holm-Bonferroni (HB) correction. **b, c & d)** Performance rates across days of all layers (Auditory (**a**), Opto/Whisker (**b**) and False alarm (**c**)). Same statistical tests as **a**. **e**) d' across days of all layers (first 100 trials of each session). Same stats as **a**. **f**) d' across days of all layers (peak window – highest d' block of each session). Same stats as **a**. **g, h)** Whisker d' (Day 3) vs Opto d' (Day 0 (**g**) and Day 2 (**h**)). Each data point represents the average whisker d' and opto d' (of the respective day) for a single mouse. Spearmann's correlation was used to calculate the correlation coefficient and the p-value (95% significance) across all 25 mice. **i)** Number of opto-stimulus trials to d' = 1. Trials counted from the first opto-stimulus of Day*

*0. Mann-Whitney U test with HB correction for multiple pairwise comparisons. **j, k**) Auditory and Opto/Whisker reaction time averages of all 4 layers. Same statistical tests as **a**. Significance asterisks - *(p<0.05), **(p<0.01) and ***(p <0.001).*

The d' across layers didn't show any significant differences in any comparison when only the first 100 trials of each session were sampled (Fig. **16**e) (the difference in response to the opto-stimulus was due to consistently higher performance by L4 and L5-IT). The peak window (defined as a 20-trial window with the highest d' for a session) sampling method (Fig. **16**f) also did not yield any significant differences. This indicates that even some mice (L2/3 and L5-PT), which didn't transfer well to the opto-stimulus, had brief periods of high performance.

While it is clear that the transfer to the whisker stimulus was moderate and showed no differences across layers, we wanted to check if there was any correlation between the opto d' and whisker d' (Figures **16**g and **16**h). Spearman's correlation was applied to all the data points (n=25 mice). The correlation coefficient (ρ) lay very close to zero (-0.07 (Day 0 Opto d') and 0.16 (Day 2 Opto d')), indicating a very weak monotonic relationship, if present. The p-values (0.753 (Day 0 vs 3) and 0.458 (Day 2 vs 3)) clearly indicate that the weak correlation observed is mostly due to noise. To conclude, there was no correlation detected between the Whisker d' and Opto-stimulus d' (both Days 0 and 2) for all mice

An auxiliary parameter called the number of opto-stimulus trials to d' = 1 was calculated for all layers (Fig. **16**i). Essentially, the first block for which the d' was equal to or higher than 1 (starting from Day 0) was identified, and the start trial index of the block was plotted for each mouse. No significant differences were observed (Mann-Whitney U test with Holm-Bonferroni correction). Similar to the single-layer reaction time (RT) plots (Figures **11**i, **12**i, **13**i and **14**i), auditory and opto/whisker RTs were plotted across layers for each Day (0-3) (Figures **16**j and **16**k). No standout trends were seen across layers or days. No significant differences were found in pairwise comparisons on any day for either plot.

3.4 Histological analysis

To complement the behavioural analysis results, we quantified opsin expression in brain sections from the mice post-behaviour (perfusion and brain extraction). Coronal sections of 100 μm thickness were obtained from the extracted brains, mounted on slides and imaged under an epifluorescence microscope with a DsRed filter (545 nm, yellow-green) to examine the opsin (ChRmine-mScarlet) expression in the barrel cortex. The quantification process was as follows: A rectangular region of interest (ROI) was drawn to contain the respective layer in which the opsin was expressed in that specific mouse (Fig. **17**a – ROI contains the L2/3 region for AK010). Then, the intensity profile was plotted (on ImageJ) to visualise the

intensity values of the pixels contained in the ROI. A Gaussian (with offset to account for background noise) curve was fit to this distribution (Fig. **17**b – Intensity profile (black) and the Gaussian (blue) fit of the ROI depicted in Fig. **17**a).

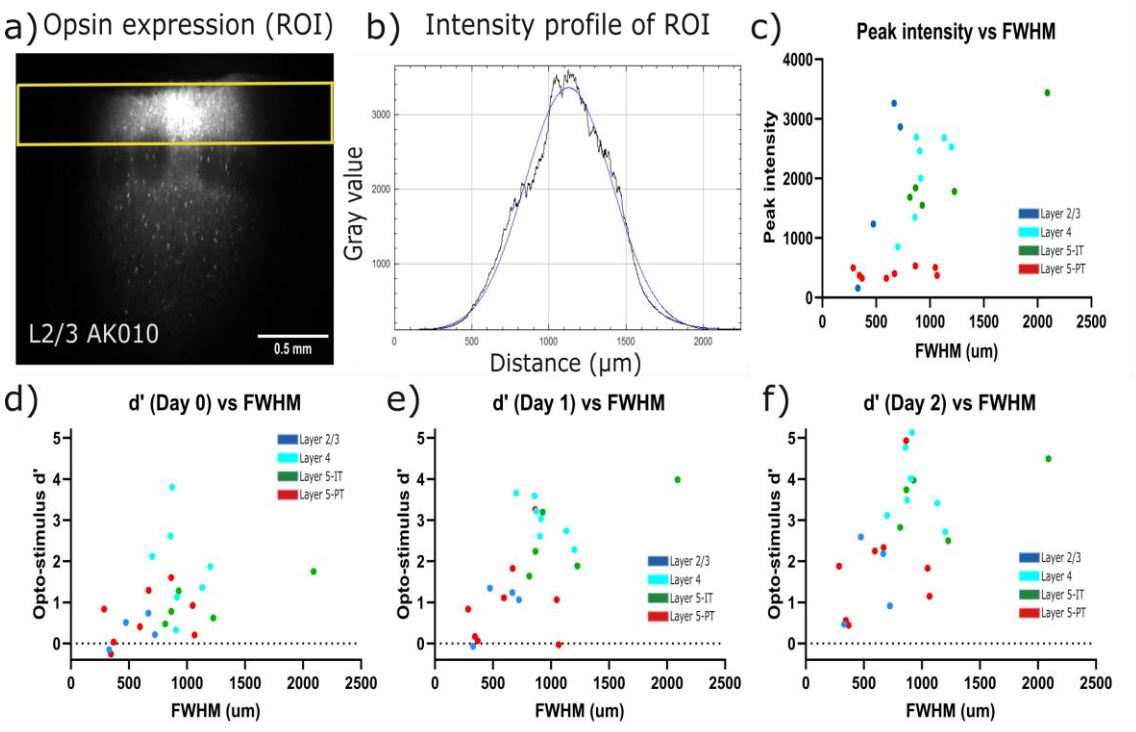

*Figure 17 – Histological analysis: **a**) Region of interest (ROI) depicted as a yellow rectangle containing the L2/3 region of the L2/3 mouse AK010 (same as Fig. **11**a). **b**) Intensity profile (black) and the Gaussian fit (blue) corresponding to the ROI depicted in **a**. Intensity values of the pixels correspond to a 16-bit Grayscale image. **c**) Peak intensity vs Full width at half-maximum (FWHM) for mice from all the layers. **d**) Opto-stimulus d' (Day 0) vs FWHM for mice from all the layers. The y-value of each data point represents the average opto-stimulus d' for a specific mouse on Day 0. **e**) Same as **d**, but uses Day 1 instead of Day 0. **f**) Same as **d**, but uses Day 2 instead of Day 0.*

Using this Gaussian fit, 2 parameters (Peak intensity and FWHM) were calculated for each mouse. The Peak intensity was calculated as the difference between the peak value (the highest intensity value at the centre of the Gaussian) and the offset (value of the Gaussian at the extremes – represents the background noise). The Full-width at half-maximum (FWHM = 2.355*σ, σ – Standard deviation) is the standard parameter that measures the distance between 2 x-values on either side of the centre of the Gaussian (whose y-values are half the peak value). FWHM essentially quantifies the spread of significant opsin expression in the ROI. The peak intensity was plotted against FWHM for all mice (Fig. **17**c) to identify any patterns within each layer. Spearman Correlation tests were performed for each layer separately (since the expression profile changed considerably across the Cre lines (Figures **11**a, **12**a, **13**a and **14**a)), but no significant correlations were observed (L2/3: Correlation coefficient, $\rho = 0.8$, (p-value) $p = 0.3333$ (two-tailed at 95% significance); L4: $\rho = 0.5714$, $p = 0.2$; L5-IT: $\rho = 0.5$, $p = 0.45$; L5-PT: $\rho = 0.1905$, $p = 0.6646$).

We wanted to identify mice (if any) that did not express much opsin (indicated by low FWHM) and yet licked strongly in response to the opto-stimulus (indicated by high opto-d'). So, the Opto-stimulus d' values (Days 0,1 and 2) for each mouse were plotted against their FWHM values (Figures **17**d, **17**e and **17**f). Spearman correlation tests did not detect any significant correlations for most of the layers on all the 3 days (Day 0 - L2/3: Correlation coefficient, $\rho = 0.4$, (p-value) $p = 0.75$ (two-tailed at 95% significance); L4: $\rho = -0.5$, $p = 0.2667$; L5-IT: $\rho = 0.7$, $p = 0.2333$; L5-PT: $\rho = 0.381$, $p = 0.3599$; Day 1 - L2/3: $\rho = 0.2$, $p = 0.9167$; L5-IT: $\rho = 0.7$, $p = 0.2333$; L5-PT: $\rho = 0.0952$, $p = 0.8401$; Day 2 - L2/3: $\rho = 0.2$, $p = 0.9167$; L4: $\rho = -0.2143$, $p = 0.6695$; L5-IT: $\rho = 0.4$, $p = 0.5167$; L5-PT: $\rho = 0.2143$, $p = 0.6191$). L4, on Day 1, showed the only significant correlation ($\rho = -0.8929$, $p = 0.0123$). While this is a strong negative correlation, we observe that all L4 mice have relatively high FWHM values compared to other lines (Fig. 17e).

On average, we observe no significant correlations (except L4 on Day 1) between opto-stimulus performance and FWHM (a measure of the spread of expression). This could indicate limitations of the analysis based on Gaussian fitting (elaborated in the Discussion section).

Chapter 4 Discussion

In this project, we aimed to identify the excitatory neuronal population in wS1 most critical for driving goal-directed licking in a whisker-detection task. However, instead of using a whisker stimulus, we used a different approach of making the trained mice (with an auditory stimulus) lick in response to an optogenetic stimulus. This optogenetic stimulus activated ChRmine-expressing excitatory neurons in different layers of the wS1. The excitatory opsin was expressed in only one of the four layers (L2/3, L4, L5-IT, and L5-PT) per mouse using layer-specific Cre driver lines. Head-fixed water-restricted mice were first trained on an auditory task (4 days) and then transferred to an optogenetic stimulus on Day 0 (Fig. **10**). After 3 days of this, they were finally trained on a whisker stimulus (similar to Sachidhanandam *et al*., 2013).

4.1 Comparisons with Sachidhanandam et al. (2013)

The activation of different layers of wS1 led to varying responses to optogenetic stimuli in our experiments (Days 0-2 in Figures **11**g, **12**g, **13**g and **14**g). Meanwhile, the response to the whisker stimulus (Day 3, Fig. **16**a) was moderate and similar across the layers. This does not exactly replicate the findings of Sachidhanandam *et al*., (2013) (Fig. **2**b). In our experiments, the mice were trained on an opto-stimulus (3 days), much longer than the whisker stimulus (a single day). The opto-stimulus performance of the mice had no correlation with the whisker performance (Figures **16**h and **16**i). So, it might be that the mice needed more sessions (days) with the whisker stimulus to appropriately change their stimulus and reward expectations. However, regardless of the training duration, there are some key differences when the experiments in Sachidhanandam *et al*., (2013) and our experiments are compared. Mainly, they utilised the Emx-Cre line, which allowed for the opsin expression in excitatory neurons from all cortical layers of the wS1 (unlike our layer-specific Cre lines). A broad activation of all layers would likely result in a stronger lick response to the opto-stimulus (potentially affecting its transfer to the whisker stimulus). But at the same time, they used a more conventional excitatory channelrhodopsin (ChR2) rather than ChRmine (in our experiments). ChR2 shows lower photocurrents (lesser activation) than ChRmine for the same LED power. While the weaker opsin does not exactly counterbalance activation across all layers, it indicates that the extent of activation in these layers was lower at the neuronal level. Furthermore, our experiments used 5 light pulses (compared to a single pulse in Sachidhanandam., *et al* 2013) of 5 ms each for each optogenetic stimulus. This further accentuates the difference in activation. Last, but not least, Sachidhanandam *et al*., trained

mice directly on the opto-stimulus, followed by the whisker stimulus. Whereas, we trained them with an auditory stimulus first and eventually transferred them to an opto-stimulus and a whisker stimulus. In our experiments, the auditory task is essentially a general learning task during which the mice adjust to the task structure, including the various trial parameters, and learn to lick in response to a stimulus. The transfer from the auditory to the optogenetic stimulus simplifies the inferences derived from the opto-stimulus performance of these mice. Since they have learnt to respond to a stimulus efficiently, drops in opto-stimulus rates are primarily due to weaker optogenetic stimulus effects. Differences in learning capabilities would not be a factor since all mice would have achieved peak auditory performance by Day -1.

On the whole, our findings broadly agree with the results from Sachidhanandam *et al*. (2013) in terms of the whisker stimulus being able to substitute for the optogenetic stimulus. But the extent to which the mice transferred their learning to a whisker stimulus was lower in our experiments. Due to the multiple differences (experimental materials and training sequence) described above, more evidence is needed to explain the cause of this difference.

4.2 Across-layer differences in opto-stimulus learning

The opto-stimulus hit rate (Panel g of Figures **11**, **12**, **13** & **14**) and d' averages (Panel h of Figures **11**, **12**, **13** & **14**) showed that L4, L5-IT and L5-PT mice responded to the opto-stimulus right from Day 0. Meanwhile, L2/3 showed significant licking only on Day 2. While this is a positive result and directly addresses the aim, the multiple pairwise comparisons (Fig. **16**a) paint a more nuanced picture. On Day 0, L2/3 shows no significant differences with the other three layers. Even within the 3 layers that responded to the opto-stimulus, L4 and L5-IT show a higher d' (and opto-hit rate) on Day 1 compared to L5-PT (indicating a hierarchy). At the same time, L2/3 and L5-PT never show any significant differences. All this evidence suggests that activating L4 and L5-IT excitatory neurons produced a stronger drive for goal-directed licking than activating L2/3 or L5-PT.

However, such an interpretation comes with some limitations. Firstly, the number of mice per layer was low and variable. The largest difference was between L2/3 and L5-IT (n = 5) and between L5-PT (n = 8). With such low numbers, outliers strongly influenced the mean (especially in L2/3 and L5-IT). Accordingly, non-parametric statistical tests were utilised, which compute and compare median-rank-like quantities between two groups (independent or dependent) rather than means. While this ensures that significance is not incorrectly determined, the problem of low representation in some layers relative to others cannot be ignored. Many trends observed in the plots, reported as non-significant, might grow or vanish

with more mice. The viral expression is also an underlying factor. Even though the same volume of viral opsin was injected into each mouse, the amount of opsin expression can vary depending on the Cre line. For example, Sim-cre (L5-PT) showed consistently lower peak expression intensities in the brain sections (Fig. **17**c). It is also important to note that while there weren't any significant correlations of the opto-stimulus performance of the mice with their FWHM (Figures **17**d, **17**e and **17**f), a more nuanced measure of the expression is necessary. Sim1-cre mice showed sparse expression in L5, but the intensity at the neuronal level was similar to that of other Cre lines. Meanwhile, Scnn1a-cre mice showed dense expression in multiple barrels (L4). Hence, our ROI-based analysis (Gaussian fitting) may underestimate opsin expression in Sim1-Cre mice. Therefore, counting the number of neurons infected (expressing the excitatory opsin) would yield more accurate measures of opsin expression across all 4 Cre lines. Subsequently, it can be assessed whether the observed differences in learning are best explained by identity or by the number of infected neurons.

Nevertheless, it is plausible to interpret the results in the light of previous work. L4, the main input layer that receives sensory information from the VPM in the thalamus (Fig. 1c), is an obvious candidate for strongly driving goal-directed licking. Apart from some counter-evidence (Hong *et al*., 2018) regarding the non-essentiality of the wS1 in the whisker-detection task, most of the literature shows that the wS1 is an important contributor to this sensorimotor transformation. The main input layer, being a strong driver of licking, adds more credibility to the role of wS1.

But, in view of the micro-circuits and the excitatory connections in each barrel, L2/3 activation does not drive licking as much as L4, hinting at the importance of the other cortical connections of L4. L4 excitatory neurons have a lot of connections with each other apart from the projections to L5 (Fig. 1c). Along with L4, L5-IT activation also showed a significant (but marginally weaker) drive in licking to the opto-stimulus. Perhaps projections from L4 to L5 are crucial for the efficient activation and processing in L5. At the same time, although L5-PT showed significantly lower performance than L4 and L5-IT, the mice still responded to the opto-stimulus. Therefore, delineating the exact differences in the processing performed by L5-IT and L5-PT via neural activity recordings or optogenetic manipulations (inhibition) during the task might reveal that a neuronal population in L5 is crucial for driving the sensorimotor transformation (apart from L4, of course).

4.3 L5 and dopamine-modulated *dSPNs*

The L5-IT and L5-PT excitatory neurons have been shown to project to the Striatal projection neurons(SPNs) as described in Section 1.3.2. However, dopamine's role as a neuromodulator

is essential for the proper functioning of these projections. Dopamine is released by the midbrain dopaminergic neurons near the synapses between the dSPNs and the cortical or thalamic axonal projections. The pairing of presynaptic and postsynaptic high-frequency firing in the presence of dopamine can induce Long-term potentiation (LTP) at the glutamatergic inputs (originating from cortex or thalamus) onto *dSPNs*(direct pathway). Following this, glutamatergic inputs exhibit increased efficacy, thereby increasing the amplitude of excitatory postsynaptic potentials (EPSPs) in *dSPNs*. This form of synaptic plasticity, which persists longer, is thought to contribute to learning.

The hypothesised schematic circuit (Fig. **18**) depicts the connections among the brain structures involved in a whisker-detection task. (Bech *et al.*, 2023). Importantly, the previously described role of the *dSPNs* in initiating the 'go' signal for licking is shown by the inhibition of SNr by the *dSPNs* (D1R-MSNs), which in turn causes the disinhibition of brainstem motor nuclei that send signals for motor execution (licking).

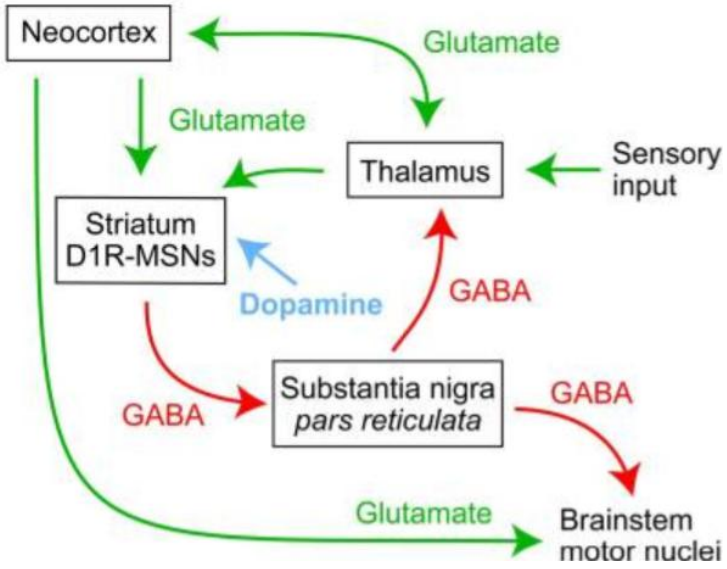

Figure 18 - A hypothesised schematic diagram that accounts for some aspects of learning and goal-directed licking in the whisker detection task. Figure from the review Bech et al., (2023)

To determine the processing performed by L5 neurons that drive licking, future work could focus on L5 projections to the striatum (Neocortex to Striatum D1R-MSNs in Fig. **18**) and on long-range feedback from the higher-order POm to L5 apical dendrites (Thalamus to Neocortex).

4.4 Conclusions and Future Work

In conclusion, this project narrowed the regions to search for a neuronal population in wS1 that would strongly drive goal-directed licking. Using a common behavioural paradigm and optogenetic manipulations, we directly compared the driving capabilities of 4 excitatory

neuronal populations. L4 and L5-IT excitatory neurons showed a greater effect than L2/3 and L5-PT in driving goal-directed licking during the reward-based task. However, one of the main limitations of this project is the lack of control mice that express a Cre-dependent reporter (rather than an excitatory opsin). While we have used another type of control (black tape to cover the window), reporter-expressing mice not being able to respond to the optogenetic stimulus would be the final evidence that the mice responding to the optogenetic stimulus in our paradigm are licking only due to the activation of opsin-expressing neurons in some layer of the wS1.

While this project has identified certain layers to drive licking more strongly than others, the goal would be to identify the exact subpopulation(s) in the wS1 that are most critical to this process. Optogenetic inhibition of such genetically defined populations would provide a second line of evidence to demonstrate the necessary and sufficient conditions that the critical subpopulation(s) of neurons might satisfy. Subsequently, refinement of such a population can be achieved by 2-photon optogenetics (manipulating smaller sub-populations).

A way to distinguish distinct subpopulations in the wS1 that drive goal-directed licking would be to use contextual rewards. The opto-stimulus trials can be divided into two contexts (rewarded and non-rewarded), which alternate in a session. As a result, we may discover that some subpopulations, when excited, drive licking regardless of context, while others drive licking only during the rewarded context. The deeper layers (such as L5-PT) have been implicated in decision-making, which is strongly influenced by context. Multiple techniques, such as optogenetic manipulations, neuronal recordings and calcium imaging, can be utilised to investigate the specific roles of such ‘lick-driving’ subpopulations of neurons in the wS1.

Once these subpopulations are identified and distinguished from one another, the downstream projections to other cortical and subcortical structures can be investigated. Similar to the experiments reported in Takahashi et al. (2020), silencing the downstream projections individually during a whisker-detection task will shed light on which pathways and structures are crucial for the subpopulations to drive licking across various contexts. Uncovering the pathways controlled by neurons in the sensory cortex (such as wS1) that drive motor actions will provide substantial insight into our understanding of sensorimotor transformations.

References

1) Harris, JA, Hirokawa, KE, Sorensen, SA, Gu, H, Mills, M, Ng, LL, Bohn, P, Mortrud, M, Ouellette, B, Kidney, J, *et al.* (2014). Anatomical characterization of Cre driver mice for neural circuit mapping and manipulation. Front Neural Circuits 8, 76.
2) Madisen, L, Zwingman, TA, Sunkin, SM, Oh, SW, Zariwala, HA, Gu, H, Ng, LL, Palmiter, RD, Hawrylycz, MJ, Jones, AR, *et al.* (2010). A robust and high-throughput Cre reporting and characterization system for the whole mouse brain. Nat Neurosci 13, 133–140.
3) Gong, S, Zheng, C, Doughty, ML, Losos, K, Didkovsky, N, Schambra, UB, Nowak, NJ, Joyner, A, Leblanc, G, Hatten, ME, *et al.* (2003). A gene expression atlas of the central nervous system based on bacterial artificial chromosomes. Nature 425, 917–925.
4) Balthasar, N, Dalgaard, LT, Lee, CE, Yu, J, Funahashi, H, Williams, T, Ferreira, M, Tang, V, McGovern, RA, Kenny, CD, *et al.* (2005). Divergence of melanocortin pathways in the control of food intake and energy expenditure. Cell 123, 493–505.
5) Liu, Y, Bech, P, Tamura, K, Délez, LT, Crochet, S, and Petersen, CC (2024). Cell class-specific long-range axonal projections of neurons in mouse whisker-related somatosensory cortices. eLife 13, RP97602.
6) Sridharan, S, Gajowa, MA, Ogando, MB, Jagadisan, UK, Abdeladim, L, Sadahiro, M, Bounds, HA, Hendricks, WD, Turney, TS, Tayler, I, *et al.* (2022). High-performance microbial opsins for spatially and temporally precise perturbations of large neuronal networks. Neuron 110, 1139-1155.e6.
7) Marshel, JH, Kim, YS, Machado, TA, Quirin, S, Benson, B, Kadmon, J, Raja, C, Chibukhchyan, A, Ramakrishnan, C, Inoue, M, *et al.* (2019). Cortical layer-specific critical dynamics triggering perception. Science 365, eaaw5202.
8) Sachidhanandam, S, Sreenivasan, V, Kyriakatos, A, Kremer, Y, and Petersen, CCH (2013). Membrane potential correlates of sensory perception in mouse barrel cortex. Nat Neurosci 16, 1671–1677.
9) Petersen, CCH (2019). Sensorimotor processing in the rodent barrel cortex. Nat Rev Neurosci 20, 533–546.
10) Petersen, CCH, and Crochet, S (2013). Synaptic Computation and Sensory Processing in Neocortical Layer 2/3. Neuron 78, 28–48.

11) Oryshchuk, A, Sourmpis, C, Weverbergh, J, Asri, R, Esmaeili, V, Modirshanechi, A, Gerstner, W, Petersen, CCH, and Crochet, S (2024). Distributed and specific encoding of sensory, motor, and decision information in the mouse neocortex during goal-directed behavior. Cell Reports 43.
12) Lacefield, CO, Pnevmatikakis, EA, Paninski, L, and Bruno, RM (2019). Reinforcement Learning Recruits Somata and Apical Dendrites across Layers of Primary Sensory Cortex. Cell Reports 26, 2000-2008.e2.
13) Kwon, SE, Yang, H, Minamisawa, G, and O'Connor, DH (2016). Sensory and decision-related activity propagate in a cortical feedback loop during touch perception. Nat Neurosci 19, 1243–1249.
14) Yamashita, T, and Petersen, CC (2016). Target-specific membrane potential dynamics of neocortical projection neurons during goal-directed behavior. eLife 5, e15798.
15) Yang, H, Kwon, SE, Severson, KS, and O'Connor, DH (2016). Origins of choice-related activity in mouse somatosensory cortex. Nat Neurosci 19, 127–134.
16) Sippy, T, Lapray, D, Crochet, S, and Petersen, CCH (2015). Cell-Type-Specific Sensorimotor Processing in Striatal Projection Neurons during Goal-Directed Behavior. Neuron 88, 298–305.
17) Wall, NR, De La Parra, M, Callaway, EM, and Kreitzer, AC (2013). Differential Innervation of Direct- and Indirect-Pathway Striatal Projection Neurons. Neuron 79, 347–360.
18) Takahashi, N, Ebner, C, Sigl-Glöckner, J, Moberg, S, Nierwetberg, S, and Larkum, ME (2020). Active dendritic currents gate descending cortical outputs in perception. Nat Neurosci 23, 1277–1285.
19) Merre, PL, Esmaeili, V, Charrière, E, Galan, K, Salin, P-A, Petersen, CCH, and Crochet, S (2018). Reward-Based Learning Drives Rapid Sensory Signals in Medial Prefrontal Cortex and Dorsal Hippocampus Necessary for Goal-Directed Behavior. Neuron 97, 83-91.e5.
20) Hong, YK, Lacefield, CO, Rodgers, CC, and Bruno, RM (2018). Sensation Movement and Learning in the Absence of Barrel Cortex. Nature 561, 542–546.
21) Staiger, JF, and Petersen, CCH (2021). Neuronal Circuits in Barrel Cortex for Whisker Sensory Perception. Physiological Reviews 101, 353–415.
22) Feldmeyer, D, Brecht, M, Helmchen, F, Petersen, CCH, Poulet, JFA, Staiger, JF, Luhmann, HJ, and Schwarz, C (2013). Barrel cortex function. Progress in Neurobiology 103, 3–27.

23) Lefort, S, Tomm, C, Sarria, J-CF, and Petersen, CCH (2009). The Excitatory Neuronal Network of the C2 Barrel Column in Mouse Primary Somatosensory Cortex. Neuron 61, 301–316.
24) Pauzin, FP, and Krieger, P (2018). A Corticothalamic Circuit for Refining Tactile Encoding. Cell Reports 23, 1314–1325.
25) Dimwamwa, ED, Pala, A, Chundru, V, Wright, NC, and Stanley, GB (2024). Dynamic corticothalamic modulation of the somatosensory thalamocortical circuit during wakefulness. Nat Commun 15, 3529.
26) Bech, P, Crochet, S, Dard, R, Ghaderi, P, Liu, Y, Malekzadeh, M, Petersen, CCH, Pulin, M, Renard, A, and Sourmpis, C (2023). Striatal Dopamine Signals and Reward Learning. Function (Oxf) 4, zqad056.