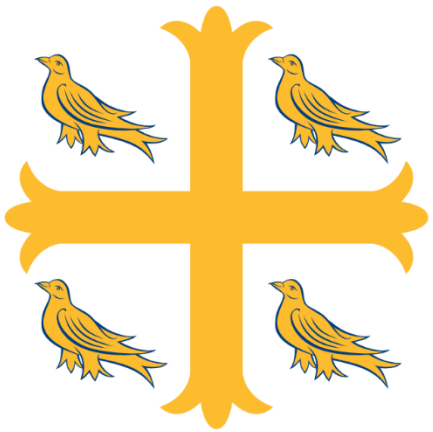
Transformation from cortical sensation to brain-wide perceptual signals





Blake Russell

Department of Physiology, Anatomy and Genetics

University College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Supervisor: Associate Professor Armin Lak

Word count: 29,956

“If you want to work on the cutting edge, be prepared to bleed” – Adam Packer

Acknowledgements

My journey in Oxford started in 2021 when I enrolled on the Masters in neuroscience course. One of the biggest surprises was discovering that my coursemates were some of the kindest, smartest and most fun people I had had the pleasure of knowing. I would not nearly have had the resistance and stubbornness to get through such a rollercoaster of a doctorate without the unwavering support of those separated by nationalities, but united by hate for PhD's: Lara, Lorenzo, Peter, Sasha and Sumedha. You are all absolute legends, some of my closest friends, and it was a pleasure to run the Cortex Club through its generational resurgence by your sides. And to everyone that I met during the masters that I wish could have stayed around longer: Daniel, Parsa, Katie, Svenja and Emma to name but a few, I hope to keep sporadically visiting when we cross paths in the same country/city and I always look forward to hearing about how you're crushing it!

My neuroscience background was entirely molecular when I arrived in Oxford, but I'd had a growing interest in the mind-body problem. So in one of the first weeks during the masters when Peter told me that an Oxford PI had recently put up a pre-print using lasers to study neural representations of perception in cortex with single-cell resolution, I knew I had to talk to them. Even in our first meeting,

I was instantly drawn to Adam's passion for doing cutting-edge neuroscience, with an attitude that any question could be answered with enough technical wizardry. I want to say thank you to Adam for trusting someone who had never worked with animals, microscopes and never heard of optogenetics, to conceptualise and deliver on such a crazy PhD proposal. Your mentorship, guidance and networking help, even after you left, are massively appreciated and have helped me more than you know. I will never forget the hilarity of watching you sprint into the sunset chasing a taxi over Magdalen Bridge on your last night in Oxford.

I want to say a massive thank you to Armin for fully embracing me into Lak Lab when Adam left. Crucially, thank you for trusting and believing in me to deliver on such an ambitious PhD proposal in a new lab that was outside of its niche and would require a lot of infrastructure set up. We have shared a whirlwind of highs and lows over the past few years: from seeing my first spikes, to breaking a record amount of probes, to performing simultaneous all-optical and Neuropixels in behaving mice. You were always at the rig cheering me on in these moments and that's something I will never forget. I would not have been able to get through it without your excellent support, mentorship, strategy, and most importantly your belief that I could achieve anything experimentally that I set my mind to. Your excitement for the project was infectious and never failed to motivate me to push the limits of what we could achieve. It has been a pleasure to watch your lab grow

over the past 4 years and for more and more people to benefit from your leadership.

I want to thank all the members of Packer and Lak labs past and present for their support. Thanks to Gabi, Ivan and Marko for the great banter, lads nights and keeping me grounded and sane in such a high stress environment. Thanks to Sarah for helping me become a mouse whisperer. Thanks to Chiara and Zenia for being two of the kindest people alive, putting up with my constant yapping and for always bringing good vibes. Thanks to Olena, Yuliia and Rathi for always being so helpful, even with last minute emergency perfusion requests on a Friday. Thanks to Jess, Huriye, Sandra and Antara for allowing me to procrastinate with DnD; it was an honour to be a DM to the most fun and chaotic party to ever do it. Thanks to Ethan, Sophie and Patrick for being the best, most hard-working students I could ask for; your motivation and excitement for the project helped it reach new heights. Thanks to Rob and Ash for always helping me at the drop of a hat, even when I'm manically bugging you over text to tell me what a part of the scope does, sending you selfies from inside the laser path, or discussing opsin desensitisation over nights of Spirit Island. And thanks to Michael for swooping into the lab and reviving my love for the project when motivation was at an all-time low. From visualizing my first stimulus responses to designing an entire python package that compiled and plotted my data, your work has been instrumental in keeping the ReportOpto machine running. On top of that, you've

been the best friend and getting through the hardest days of the PhD was made much more manageable with our gym and pub trips. And yes, I would like to see something cool.

Thank you to my parents for our weekly chats; your unwavering support and for helping me mentally overcome some of hardest times that a PhD will put you through. I'll always appreciate you trying to understand what I'm researching and listening to me vent about how messed up the brain is – your perspective on my work, life and everything in between has been invaluable and I couldn't have got through it without you.

Finally and most importantly, thank you to Meg. You have been my rock throughout this whole journey. Deciding to live together for the first time whilst you were working a full-time job and I was doing a PhD did not come without its challenges. Thank you for putting up with the long nights, working 7 days a week for months and crashing as soon as my head hit the sofa. It's not been easy juggling our life together and the demands of this insane project, but without your love, support and sarcastic humour, I probably would've called it quits on the PhD years ago. So thank you for everything that you have done for us over the past few years. I'm glad all the sacrifices along the way were worth it to make us stronger, happier and for us to eventually hold this stupid piece of paper; because this PhD is as much yours as it is mine.

Abstract

Perception is inherently variable; the same external stimulus might be detected or missed in different instances. Previous studies have shown that perceptual variability could be partially explained by whether sensory neuronal signals propagate from primary sensory cortices to higher cortical regions. However, methodological choices in these studies meant that perceptual signals were difficult to isolate due to either pre-cortical fanout produced by external stimuli or activity sensors with dynamics too slow to sufficiently resolve the perceptual period pre-report. Additionally, it remains unknown how perceptual information is communicated between primary sensory cortices and their subcortical outputs. To investigate this, we trained mice to report optogenetic stimulation of primary somatosensory cortex (S1) by turning a wheel to receive a water reward. By requiring the animal to use direct activation of S1 to guide behaviour, the task dictates the directional flow of perceptual information by circumventing fanout at pre-cortical processing stages. We used Neuropixels to simultaneously record electrophysiological neuronal activity from several major S1 cortical and subcortical outputs, including secondary somatosensory cortex (S2), caudoputamen (CP) and the ventral posteromedial thalamic nucleus (VPM) during expert behaviour. We found that stimulation-evoked population-level

neuronal responses reached S2 within 50 ms, shortly followed by CP and later by VPM. At expert performance, mice started to turn the wheel within 100 ms, during which time population activity could not distinguish between hits and misses in S2, CP or VPM. This stimulus-driven activity contained sparse spike sequences in all regions that were reliably evoked within 100 ms post-S1 stimulation. The sequences were predominantly intra-region, occurring first in CP, followed by S2, then VPM. Some sequences were identical in their composition on hits and misses but occurred significantly later on misses. Other sequences occurred only on either hits or misses and also differed in their latency to stimulation. Therefore, these sequences form a rapid and reliable sparse temporal code underlying perception.

To dissect how the structure of optogenetic stimulation dictates the recruitment of these sequences, we used two-photon holographic light shaping. First, we optimised all-optical methodology to deliver reliable two-photon stimulation-evoked activity across trials by reducing cross-talk between the opsin and the calcium reporter. Then we combined this optimised paradigm with retrograde labelling, allowing targeted holographic stimulation of projection-defined S1 neurons in expert mice whilst simultaneously recording subcortical activity with Neuropixels. However, in this final experiment, a small sample size and photoelectric artefacts from the near infra-red laser mean that more data is required before conclusions can be drawn.

The results from these experiments provide evidence that perceptual signals are represented both cortically and subcortically in the top down somatosensory pathway through a rapid and reliable sparse spiking temporal code.

Summary of Contributions

I declare that this thesis is my own and that attribution has been given where information was derived from other sources.

Dr Robert Lees: Implanted cranial windows in 3 of the animals used in chapter 3. **Olena Didenko, Rathi Puliyadi, Yuliia Shevchuk:** Perfusion, slicing and imaging of extracted brains. **Sophie Marchent:** DeepLabCut model creation for paw movement during one-photon behaviour in chapter 2. **Dr Matthias Fritsche:** Helped to write the rigbox behavioural code for the all-optical task in chapter 4. **Dr Adam Packer:** Supervision of opsin desensitisation study and funding. **Associate Professor Armin Lak:** Supervision throughout and funding.

All other surgeries, experimental design, procedures, analysis and visualisation were performed by me.

Contents

1. General introduction 1

1.1 The hard problem 1

1.2 Investigating a percept 6

1.3 Thesis outline 11

2. Reliable, sparse spike sequences underlying perception in the top-down somatosensory pathway 14

2.1 Introduction 15

2.2 Methods 44

2.3 Results 51

2.4 Discussion 68

3. Opsin desensitisation during all-optical interrogation depends on imaging parameters 71

3.1 Introduction 72

3.2 Methods 81

3.3 Results 87

3.4 Discussion 99

4. Work in progress: function- and projection-specific stimulation in a somatosensory CPD task 105

4.1 Introduction 106

4.2 Methods 108

4.3 Results 117

4.4 Discussion 126

5. Conclusion 130

6. Supplementary figures 142

7. References 150

1

General introduction

1.1 The Hard Problem (p. 1)
1.2 Investigating a Percept (p. 6)
1.3 Thesis Outline (p. 11)

1.1 The Hard Problem

Since neuroscience was established as a discipline, there has been remarkable progress in our understanding of the central nervous system including its subcellular molecular pathways[1–4], constituent cell types[5–9], connectomics[10–13], and how both single neurons and neural networks represent information about the external world[14–20] and internal state[21–24]. Technological innovations in recording and manipulating neuronal activity *in vivo*, such as patch-clamp[25], high-density electrophysiology[26], optogenetics[27], two-photon microscopy[28] and functional magnetic resonance imaging[29] have allowed insights into how dysfunction at

these various scales can give rise to disease[30–34]. Yet one question that has persisted in the minds of philosophers and neuroscientists since at least the time of René Descartes[35] (and can be traced back to the musings of Plato[36] and Aristotle[37]), is how the physical substrate of brain matter can give rise to our first-person, subjective perception of reality. The idea that the brain, unlike other organs and amalgams of cells, is uniquely constructed to give rise to the metaphysical "self" with no clear substrate, has historically rendered the **'hard problem of consciousness'**[38] impervious to even the most advanced technological innovations. Consequently, our understanding of disorders of consciousness remains critically sparse[39]. This can lead to a mismatch between prognosis and recovery, rendering the self-fulfilling prophecy of a poor prognosis by withdrawal of life-support an unfortunate reality[40].

In lieu of direct recording and manipulation of a 'consciousness substrate', separate schools of thought have offered formalisms of the relationship between the brain and mind. Dualists, such as René Descartes, believe that the mind and brain are almost entirely independent[35], whereas monists, such as Carl Jung, argue that they are indiscernible[41]. As neuroscientists, we tend to be monists or at least property dualists, where the mind is an emergent property of the brain, if not fully describable by its processes. Popular modern monist theories of consciousness such as Global Neural Workspace[42] and Integrated Information Theory[43] aim to describe consciousness in its entirety, or as Thomas Nagel describes it, the

'something-it-is-like-ness' of existence[44,45]. However, the vague nature of this definition requires a distinctly human purview to be reported and does not lend itself to readily-reversible causal manipulations. As such, evidence for the role of these theories in consciousness is primarily correlational – limited to lesion studies, brain imaging and neurophysiological characteristics[46–48].

As scientific fields mature, they shift from trying to infer causality through observations to directly manipulating aspects of the system itself[49,50]. Throughout the second half of the 20th century, landmark lesion studies such as those detailing Phineas Gage[51] and Henry Molaison[52] were hugely influential on the field of cognitive neuroscience; yet the irreproducibility of lesions alongside their permanent and opportunistic nature render them unconducive to inferences of causality. Even now, methods such as functional magnetic resonance imaging are sometimes seen as 'modern day phrenology' for their lack of ability to distinguish between brain activity causing an output, compensating for an output, or being incidentally related to an output[53]. This is often regarded as the 'causality gap' in neuroscience[54]. In order to close the causality gap on consciousness, we must first agree on what is required for neural dynamics to be 'causal'. Despite causality resting at the crux of many questions in neuroscience[55], its use maps broadly onto two different definitions: causal production and causal dependence[56]. The philosopher Ned Hall defines causal production as an event directly producing another, and causal dependence as counterfactual dependence between two

events[57]. The spirit behind causal production is that it is not possible or practical to map all possible neural necessities for an event and that we should instead focus on the subset that are sufficient. In contrast, causal dependence rightfully notes that interventions to brain areas not directly implicated in an event can modulate its manifestation[58]. This can be described as a causal complex[59]. Although modulation of a causal complex can clearly be useful in therapeutic contexts, perhaps consciousness is too labyrinthine for indirect manipulations to reveal its underpinnings to a satisfactory level. For example, most neuroscientists can agree that the thalamus is part of the causal complex for consciousness[60–64], and is the site of action for the vast majority of anaesthetics[65]. Yet patients regain awareness in up to 2% of surgeries[66], which is independent of the expertise of the anaesthetist[67] and occurs despite the field of anaesthesia being a leader in modelling effector site concentration[68]. This can lead to negative psychological consequences such as post-traumatic stress disorder[69]. In truth, this is because we do not understand why anaesthetics work, aside from their global metabolic depression affects[70] (even though some anaesthetic agents produce the opposite[71]), for which a causal production viewpoint is necessary to remedy.

Nevertheless, finding the neural events 'sufficient' for generating the 'something-it-is-like-ness' of 'self' is a problem still too large in scope for neuroscientists to experimentally dissect. In an attempt to make this a more viable challenge, Francis Crick and Christof Koch reduced their definition of the neural correlates

of consciousness (NCC) to the minimal neuronal mechanisms jointly sufficient for a single conscious percept[72]. This ring-fencing of the metaphysical to an experimentally tractable concept rejuvenated neuroscientific consciousness research and is an approach that has largely been adopted throughout the field[73]. Harkening to the way that René Descartes believed visual signals entered the mind when they reached the pineal gland[74], our goal is to find where and how the neuronal activity for sensation transforms to perception.



René Descartes (1662). Treatise of Man.

1.2 Investigating a Percept

In Crick and Koch's definition of the NCC, the word 'percept' turns a nigh insurmountable challenge into a reasonable ask. By relaxing from trying to understand the neural underpinnings of qualia, to just the perception of a single sensation, pandoras box is metaphorically opened for how to experimentally tackle such an endeavour. The obvious choice is perhaps human studies, where the subject can provide a qualitative perceptual description of complex stimuli of many modalities (ignoring whether we can truly know if others possess phenomenal consciousness[75,76]). For example, a typical approach employed in human studies is functional neuroimaging during observation of stimuli that can flip between percepts, such as Mooney images[77], illusions[78] and binocular rivalry[79]. These types of experiments have revealed percept-specific activation in sensory[78,80], temporal[81,82], frontal regions[83,84] and communication routes between these areas[85,86]. However, lesion studies suggest that frontal areas are not critical for forming a percept[87–90], with some evidence suggesting that frontal activation reflects introspecting and reporting the percept rather than the percept itself[91,92]. This is unlike lesions in temporal and sensory areas which can dramatically change a percept, or even remove the ability to form a specific percept completely, as is the case with prosopagnosia and blindsight[93–100].

The methodological task of separating perception from its products in an experimental setting has been historically challenging[101], with some even

labelling it impossible[102]. Nevertheless, many researchers have tried their hand at solving it. One of the more interesting avenues taken in the field is investigating the neural correlates of dreams – a theoretically 'pure conscious' state. Historically, dreams have been associated with the awake-like brain activity observed during rapid eye movement (REM) sleep[103]. Frontal activation specifically decreases during REM[104], seemingly supporting its role in reporting perception rather than perceiving itself. However, more recent studies suggest that the association between dreams and REM is not so clear cut, with both dreaming and non-dreaming experiences reported in patients awakened during REM and non-REM sleep[105]. By looking at the neural correlates of dreams independent of sleep phase, researchers were able to identify activity in the posterior cortical hot zone as being predictive of mentation[106]; a region that overlaps with sensory areas implicated in content-specific NCCs[107]. Sources for this evidence include studies utilizing a no-report paradigm, where on some trials indirect measures are used to infer perception[91,108].

The posterior hot zone encompasses all sensory cortices, but is primarily localised around higher cortical regions or sensory 'association' cortices[109]. But the relative contributions of primary and higher cortices to percept generation are still under active investigation. For example, human studies have shown that imperceivable changes in visual stimuli are still encoded in primary visual cortex (V1)[110–112] and it has been suggested that perceptual experiences are represented as a

consequence of feedback from higher cortical areas of the visual stream[113,114]. Modality-specific perceptual information is also encoded in primary somatosensory cortex (S1) and primary auditory cortex (A1). However, due to the temporal and spatial resolution limitations of human neuroimaging, and limitations on causal manipulations, it is still unclear whether this depends on feedback from higher cortical areas[115–119]. Animal studies provide the perfect opportunity to gain a more mechanistic causal production understanding of perception within the sensory hierarchy. Despite being unable to receive verbal reports of perception from subjects, animals can be trained to report perception of a variety of stimuli[120–122]. For example, mice can be trained to robustly discriminate stimuli of multiple modalities including visual[123], auditory[124] and somatosensory[125]. Victor de Lafuente and Ranulfo Romo were amongst the first to record *in vivo* neuronal activity from primary sensory and higher cortical regions during a detection task. They trained monkeys to report the presence of a vibrating fingertip stimulus by pressing a button, whilst recording electrophysiological activity from S1 and medial premotor cortex (MPC). They found that S1 activity reflected stimulus strength, but not perceptual report, with the opposite being true for MPC. Additionally, direct electrical microstimulation of MPC reproduced the perceptual report, but MPC activity did not reflect the action of reporting in control experiments[126]. Whilst MPC is not part of the posterior hot zone, this was the first evidence in animals suggesting that as neuronal signals propagate from primary sensory cortices to higher cortical

regions, they better reflect the trial-by-trial percept rather than just the sensation itself. However, given that the neuronal activity they used to perform the analysis often spanned the period of reporting, and in the control experiments the monkeys did not have to perceive the stimuli to inform their decision, it is still unclear whether MPC is primarily contributing to introspection and task-relevant motor execution. In fact, there is evidence that frontal regions may hold a mental representation of the percept rather than direct the act of perceiving itself[127]. In a follow-up experiment, the same researchers found that these 'perceptual' signals did not suddenly appear in MPC but rather gradually evolved throughout the cortical hierachy[128]. Since then, focus has been primarily on somatosensation in mice, where perception first seems to be encoded as early as secondary somatosensory cortex (S2), with feedback loops between S1 and S2 playing a critical role in its emergence[129,130].

While it seems clear that perceptual signals evolve throughout the cortical hierarchy, some sources of discrepancy between various human and animal studies remain; namely whether perceptual signals are present in primary sensory areas and whether frontal areas are actively involved in perception or just task execution. These problems largely boil down to difficulties in differentiating between neural signals associated with stimulus, perception and report. A clear strength of animal studies is the precision with which you can measure brain activity – often down to the level of individual neurons. Recording neuronal

activity during a stimulus detection task may be the most powerful approach for unveiling the causal production behind perception, but it also hinges on a high degree of accuracy and control in generating the percept. After all, the dynamics of individual neurons can be greatly influenced by variables that are difficult to experimentally control, such as eye position[131], head position[132], body position[133], uninstructed movements[134], attention[135] and internal state[136], which will also have an effect on the phenomenological content of the percept itself. This is because a percept is inherently irreducible, where even a small conceptual change to a scene (e.g. colour) will generate an entirely different conscious experience; i.e. the representation of a singular phenomenon feeds back to the representation of all others within a percept[48]. Consequently, if the percept is not tightly controlled across presentations, the neuronal signals underlying it will be unreliable and even more difficult to differentiate from confounding variables.

One step towards keeping a percept identical between presentations is to bypass its external representation, where a lot of the difficult to control variables manifest, and instead instil it directly into a sensory cortex. As well as enabling a high degree of control over the stimulus at a cellular level, direct activation of the brain circumvents fan-out of activity during pre-cortical processing that is present with externally presented stimuli; this additionally affords a high degree of read/write control over the percept, as an animal trained to report such a percept is required to only use features of the evoked cortical activity to guide its behaviour[137]. Thus,

all neuronal signals representing that specific percept can be traced outward from the site of modulation. A behavioural association paradigm that utilises this methodology is known as a cortical perturbation detection (CPD) task. CPD tasks build on the pioneering cortical stimulation-evoked percept work first described in brain surgery patients by Wilder Penfield[138], and have since become the gold standard method for investigating the NCC[139–141].

Herein I describe the use of a CPD task in mice for investigating how the NCC are represented in the cortical hierarchy and primary subcortical outputs of a sensory region. Utilising the combination of an optogenetic CPD task and concurrent high-density electrophysiological recordings, we show that perception of an artificial somatosensory stimulus is encoded in rapid and reliable sparse spike sequences throughout the top-down somatosensory pathway. Furthermore, we begin to investigate how the structure of the optogenetic stimulus determines its propagation, and present methodological optimisations to help control the reliability of such a stimulus.

1.3 Thesis Outline

Previous studies have shown that perception can be partially explained by the propagation of neuronal signals from primary sensory cortices to higher cortical regions. However, in these studies, sensor dynamics or task design limit the extent

to which perceptual signals can be isolated. Additionally, the role of cortical-subcortical communication in perception remains unknown. In **chapter 2**, we investigate this by training mice to report one-photon optogenetic stimulation of S1 whilst recording high density electrophysiological activity in the top-down somatosensory pathway. Mice learn to perform the CPD task successfully. Spike resolution recordings allowed us to analyse neuronal activity within the 100 ms post-stimulation, by which time mice would often start to report. Within this period, population level firing rate activity could not distinguish between hits and misses in S2, CP or VPM. Using novel high sensitivity spike train analysis, we show that rapid and reliable sparse spike sequences distinguish between hits and misses in their identity and onset delay.

Previous studies have shown that cortical signals are temporally unreliable and integrate over time to combat inherent noise. To probe how reliability of stimulation at cellular resolution affects the reliability of the temporal code underlying perception, we aimed to combine two-photon holographic stimulation with two-photon calcium imaging (all-optical) in S1. In **chapter 3**, we demonstrate methodological optimisations that minimise previously understudied effects of cross-talk in all-optical experiments. This allows reliable stimulation-evoked activity to be delivered across trials. In **chapter 4**, we utilise these optimisations to perform targeted stimulation of neurons based on their function or projection, whilst simultaneously recording subcortical activity with

Neuropixels in expert mice. However, a small sample size and photoelectric artefacts in the electrophysiology recordings mean that more data is required to draw conclusions.

2

Reliable, sparse spike sequences underlying perception in the top-down somatosensory pathway



2.1 Introduction (p. 15)
2.1.1 CPD tasks (p. 15)
2.1.2 Temporal coding of perception (p. 33)
2.1.3 Objective (p. 43)
2.2 Methods (p. 44)
2.2.1 Animals, viruses and surgery (p. 44)
2.2.2 Two-photon imaging and whisker mapping (p. 45)
2.2.3 Behavioural training (p. 47)
2.2.4 Analysis, statistics and reproducibility (p. 48)
2.2.5 Timing index (p. 48)
2.2.6 Electrophysiology (p. 48)
2.2.7 Sequence detection (p. 49)
2.2.8 Histology (p. 50)
2.3 Results (p. 51)
2.3.1 Mice can be trained to report optogenetic stimulation of wS1 (p. 51)
2.3.2 Optogenetic stimulation of wS1 produces a similar downstream population response, independent of perception (p. 55)
2.3.3 Perception is represented in the timing and identity of rapid, reliable sparse spike sequences in the top-down somatosensory pathway (p. 59)
2.4 Discussion (p. 68)



2.1 Introduction

2.1.1 CPD tasks

2.1.1.1 Intracortical microstimulation

Wilder Penfield was a neurosurgeon who studied under Charles Sherrington at the University of Oxford, eventually pioneering a method of epileptic tissue removal in the 1930's known as the Montreal procedure[142]. The Montreal procedure was ground-breaking at the time as it combined neurophysiology with neurosurgery, requiring patients to report electrical microstimulation of the cortex during open-brain surgery to map sensory and motor homunculi. This helped guide epileptic tissue removal based on either directly triggering a seizure-like response in the patient or evoking percepts concurrent with the prodrome, such as sensory hallucinations[143,144]. The ability to circumvent pre-cortical processing and evoke sensory percepts through direct cortical stimulation was a boon for researchers trying to identify the NCC, as it theoretically permitted a somewhat high-degree of control over the qualia of a percept thanks to population-level spatial specificity and also necessitated that perceptual signals must propagate from the stimulation point outwards. Crucially, this also provided evidence that activity in sensory cortices is causal for sensory perception as opposed to incidental.

The first work to push the envelope of assessing how realistic artificially evoked percepts are and how finely they can be modulated was by Ranulfo Romo and colleagues. They showed that direct electrical stimulation of S1 was perceptually similar enough to mechanical vibration of the fingertip that monkeys trained to identify the highest frequency stimulation between two mechanical stimuli could also identify the highest frequency between a mechanical and electrical stimulus[145]. Furthermore, psychometric curves for the mechanical and electrical stimuli were virtually identical[146]. In visual cortex, the lowest detection thresholds in an electrical microstimulation CPD task were mostly found in layer 3 and the highest thresholds found in layer 6, although voltage thresholds varied substantially by electrode site, suggesting that this rule cannot necessarily be generalised region-wide, let alone brain-wide[147]. Additionally, CPD tasks seem much more difficult to learn in non-sensory cortices[148]; in fact, human studies have shown that percepts are less-likely to be evoked by non-sensory cortical stimulation, and those that are produced are more complex and multimodal (i.e. emotions, memories)[149]. Contrastingly, in a sensory CPD task, mouse behaviour can be biased towards hits by the juxtacellular stimulation of a single cell[150].

Later, another group showed that percepts can be generalised in the reverse direction in rabbits after fear conditioning: from whisker S1 (wS1) stimulation to physical whisker pad stimulation. The percept evoked by wS1 stimulation did not generalise to other physical somatosensory stimuli such as electrical pulses to the

hind limb[151]. More recently, human studies have reported that electrical microstimulation of S1 with smaller currents can evoke realistic tactile percepts[152–154], unlike the unnatural percepts produced by macro stimulation of the cortical surface seen by Penfield[155]. Together, these results show that even percepts evoked by relatively indiscriminate bulk activation of sensory cortical areas can have enough endogenous validity to be indistinguishable from percepts evoked by external stimulation of the corresponding anatomical area of the homunculus. This is despite electrical stimulation being spatially imprecise, even at low intensities, due to the dominating recruitment of passing axons as opposed to somata[156–159]. Although intra-columnar connectivity is the strongest[160,161], excitatory and inhibitory horizontal fibres can synapse millimetres from their origin[162–164]. Axonal spikes recruited by electrical stimulation can also be conveyed antidromically, resulting in a spatiotemporal "blur" of evoked activity[165]. Fast excitatory responses to non-physiological cortical stimulation are rapidly dampened by long-lasting inhibition[166–171], which is likely a mechanism employed by the brain to prevent runaway propagation of non-physiological activity[172]. This secondary inhibitory response is also behaviour-dependant[173], suggesting that certain contexts can drive the network into states that are more accepting of received information.

Despite the brain making an active effort to quench bulk activity artificially injected into cortical populations, it is evident that some of the evoked signals

must escape local inhibition in order to produce a percept. This may be facilitated by the inhibitory response helping to shape structured patterns that map onto existing circuit architecture. The brain struggles to allow off-manifold activity, even when it is behaviourally beneficial to do so[174], suggesting that the structure of the network constrains activity patterns. Additionally, inhibitory responses to cortical activation can be avoided by selectively recreating naturalistic activity patterns with cellular resolution[175]; which also causally recreates the percept encoded by such patterns, despite stimulating several orders of magnitude less cells than electrical microstimulation[176]. Unfortunately, intracortical microstimulation does not possess the required specificity to selectively reproduce cortical population activity, which means that CPD tasks employing such methods risk not only needless rebound inhibitory activity, but also off-target activation[177], network (de)synchronisation[178], maladaptive plasticity[179] and tissue damage[180].

2.1.1.2 Optogenetics

In order to avoid most of these adverse effects, a new method of direct cortical stimulation was needed that was less invasive with higher specificity than intracortical microstimulation. In 1979, Francis Crick was foraying into the world of neuroscience after his success in the field of genetics; as a newcomer to the

discipline, his opinion was that to further understand the brain, we needed a technique that could manipulate neurons with cell-type specificity[181]. It wasn't until over 25 years later when his dream was realised with the advent of optogenetics – a multi-award winning method that would revolutionise neuroscience forever[182–184].

The origins of optogenetics can be traced back to 1851 when Heinrich Müller first discovered visual pigments in photoreceptor cells in the vertebrate retina[185], now widely referred to as animal or type 2 opsins. Around 25 years later, the role of opsins became more clearly linked to vision, when Franz Boll and Wilhelm Kühne discovered their light-sensitive properties[186,187], noting that the frog retina changed colour under illumination, mirroring the observation that the eye becomes less sensitive under bright conditions (dark adaptation)[188]. George Wald was instrumental in furthering our understanding of the mechanism behind this, discovering that the light-sensitive properties of opsins were dependant on the chromophore retinene[189]– now known as 11-*cis* retinal. But it wasn't until almost 100 years after Boll and Kühne's seminal work that photoreceptor cells were revealed to generate electrical signals[190], elucidating how visual information is transmitted to the brain. This was later discovered to be mediated by the opsin-chromophore complexes embedded in the photoreceptor membrane, known as rhodopsins, which undergo photoisomerization and, for type 2 opsins, facilitate transduction via a G-protein-coupled cascade[191,192]. Thanks to advances in

expressing non-endogenous recombinant proteins through genetic manipulation[193,194], researchers in the lab of Gero Miesenböck were the first to show that neuronal activity could be controlled with light *in vitro* by transfecting rat hippocampal neurons with the genes encoding a *Drosophilia* opsin and its associated intracellular machinery[195]. G-protein-coupled optogenetics is advantageous because stimulation of a single type 2 opsin recruits multiple cyclic guanosine monophosphate-gated ion channels; this means that low intensity stimulation is amplified to produce large responses. The drawback is that amplification through the G-protein-coupled cascade is relatively slow, taking up to seconds from stimulation onset to achieve significant depolarization[196]. Furthermore, the G-protein coupled receptor and its associated machinery are encoded by separate genes, rendering it difficult to express functionally *in vivo*[197].

Meanwhile, in the early 1970's, an opsin analogue was discovered in bacteria[198] (and later fungi[199]). However, this bacteriorhodopsin acted via a different mechanism compared to type 2 opsins, instead using its light-driven conformational change to pump protons out of the cell[200]. This hyperpolarizing action would later be repurposed to inhibit neuronal activity *in vivo*[201]. In 1984, opsins were also found in algae[202]. Although initially assumed to operate similarly to type 2 opsins[202], their fast photocurrent induction[203] and low amplification through the transduction chain[204], combined with high homology with bacteriorhodopsin's retinal binding site and proton-transport network suggested

to researchers in the lab of Peter Hegemann that in this rhodopsin, the photoreceptor and ion channel are one and the same protein[205]. Merely a year later, the same group discovered a second algal "channelrhodopsin" with a faster rise time, slower transient to steady-state decay and a larger pore size that conducts both protons and divalent cations[206]. Unlike bacteriorhodopsin, channelrhodopsins 1 and 2 passively transport cations down their concentration gradient into the cell, depolarizing the membrane. This afforded researchers in the lab of Karl Deisseroth millisecond-timescale optical control of spiking in mammalian neurons *in vitro*[207]– a landmark step towards selectively recreating cortical activity. However, on the flipside to GPCR-mediated optogenetics, channelrhodopsin-2 has single molecule conductance as low as 50 femtosiemens, meaning that 100,000 to 1,000,000 molecules would have to be stimulated in a single neuron to trigger an action potential[208]. And although optogenetic methodology does not cause the same clear physical damage as inserting an electrode, opsin expression can affect circuit morphology[209] and confound results by photostimulation directly heating the tissue[210–212].

Although Miesenböck group was technically the first to optogenetically control animal behaviour *in vivo*, their approach relied on using light to uncage GPCR agonists in flies[213]. The first papers utilising channelrhodopsin-2 to causally manipulate mouse behaviour were published in rapid succession in 2007 by Karl Deisseroth's lab; they used optogenetics to show that activity in vibrissal motor

cortex and lateral hypothalamus was causal to whisker movement and wakefulness respectively[214,215]. Once optogenetics had proved its ability to causally produce reflexive behaviour, researchers in Karel Svoboda's group wanted to test its ability to evoke percepts in a CPD task[216]. With this new method offering greater stimulation selectivity and minimal artefact generation compared to electrical microstimulation[217], they were able to more accurately control and read-out injected activity in sparsely transfected wS1 neurons. Specifically, they used simultaneous optical stimulation and *in vivo* patch clamp of channelrhodopsin-2 expressing neurons to show that, for the majority of cells, a 1 ms light pulse was sufficient to evoke a single action potential. Furthermore, they demonstrated that mice could be trained to report perception (and absence of perception) of a 1 ms light pulse by performing a nose poke at one of two ports. The fact that the stimulation period on both trial types was visually masked with light flashes to the behavioural arena, the number of opsin-expressing cells under the cranial window correlated with task performance, and control mice not expressing the opsin performed significantly worse on the task, suggests that an optogenetically-evoked percept was driving learning. By modelling the spread of light through brain tissue, the authors estimated that an action potential needs to be simultaneously evoked in 297 neurons to generate a percept, whereas only 61 neurons need to be injected with a train of 5 action potentials.

The next optogenetic CPD task experiment was performed by Mark Histed and John Maunsell[218], who trained mice to report 100 ms activation of visual cortex. Building on the findings of Svoboda group, they found that perception was better predicted by total injected spike count in stimulated neurons rather than the temporal arrangement. They found this result surprising, expecting that temporally clustered spikes should be preferentially propagated because cortical neurons nonlinearly integrate their subcortical inputs within the dendritic tree, both to amplify incoming signals past the spike threshold and to trigger bursts of action potentials[219–221]. One may attribute this to direct optogenetic activation bypassing the need for coincident input, however the authors show that behaviourally detectable stimulation did not produce reliable spiking outside of the task, and they instead suggest threshold crossings were determined by background activity. Therefore, for weak stimuli near the behavioural threshold, the neuronal noise present in the cortex may act to linearly temporally integrate small changes in activity of hundreds of thousands of cells over longer timescales[222–226]. Importantly, this linear integration may only occur over ~100 ms, as monkey behavioural performance in a somatosensory optogenetic CPD saturated at ~100 ms pulse duration for a given stimulation power[227].

Although optogenetics provides greater spatial specificity than electrical microstimulation, thanks to cell-type specific promotors[214] and somatic targeting sequences[228], widefield illumination is largely incapable of selectively

reproducing physiological activity (but see[229]), and produces a large ring of inhibitory responses around the stimulation site[230], perhaps due to off-manifold perturbation.

2.1.1.3 Multiphoton optogenetics

Sensory cortices in mice broadly represent external stimuli of their respective modalities through salt-and-pepper spatial responses[230–232]. To selectively reproduce these responses, stimulation with single-cell resolution is required. The first step towards realising this came in 1963 when Warner Peticolos and colleagues debuted two-photon excitation of organic molecules[233].

Going into the 1990's, the main method of recording population-level neuronal activity *in vivo* was electrophysiology, which could simultaneously record tens of units with high temporal resolution but low spatial coverage[234]. Meanwhile, microscopy had been making significant advancements in visualising biological specimens at high spatial resolution. Confocal microscopy could resolve chromatin distribution within neuroblastoma nuclei[235] with less background fluorescence contamination than epifluorescence microscopy thanks to a pinhole restricting transmission of the majority of photons emitted from outside the focal plane[236]. By implementing a scanning galvanometer in the laser path, samples could be imaged at much greater temporal resolutions, moving from ~0.1 Hz with

stage translation setups to 8 Hz with laser scanning confocal microscopy[236]. Although very high axial resolutions could be achieved in theory with a perfectly calibrated confocal microscope with an infinitely small pinhole[237], in reality larger pinholes are needed for sufficient signal to noise within scattering specimens, which results in an axial resolution ~150 um[238]. This does not meet the requirements to achieve single-cell resolution as mouse neuronal soma are within the 5 – 30 um diameter range[239]. On the other hand, two-photon microscopy utilises coincident multiphoton excitation to confine excitation to a focal plane of ~20 um[238,240], resulting in approximately somata-level axial resolution even without a pinhole[237]. Because two-photon excitation of a molecule is demonstrably rare under ambient light (~ once every 10 million years), the chance of two-photons exciting one molecule within ~10^{-16} seconds of each other is amplified by using lasers with pulse durations < 1 ps at ~100 MHz repetition rates[241]. Furthermore, because two-photon excitation necessitates that each photon must contribute half-the required energy to boost the electronic state of the molecule[242], the wavelengths used are approximately double those used for confocal and epifluorescence microscopy. These infrared wavelengths are able to penetrate further into tissue with less scattering[243], allowing imaging at hundreds of microns deeper than one-photon alternatives[244,245]. When combined with scanning galvanometers[28], this technology allows activity to be recorded at spatial scales from single cortical units[246] to thousands of neurons[247] simultaneously *in*

vivo with cellular resolution through a cranial window[248]. By balancing signal-to-noise and power on sample, activity can be visualized even deeper using three-photon microscopy[249]. Neuronal activity is typically visualised using calcium dynamics as a proxy through genetically-encoded calcium indicators (GECI)[250]; GCaMP (circularly permuted green-fluorescent protein GECI) is the most popular calcium indicator, having been iteratively improved upon and engineered to enhance signal-to-noise ratio and accelerate kinetics over the past two decades[251–257]. However, even with the latest, most sensitive iteration of GCaMP, it is difficult to resolve individual spikes due to the relatively slow kinetics (< 5 ms rise time, ~ 80 ms decay)[257]. Reporter dynamics are even more unfavourable in red fluorescent alternatives[258]. Instead, voltage imaging has become more commonplace in multiphoton studies in the last few years[259] now that advancements in microscopy hardware have enabled a high enough framerate to capture indicator dynamics[260,261].

Once cortical population responses to stimuli could be resolved at cellular resolution *in vivo* using multiphoton imaging, the next technical challenge for perception researchers was to use multiphoton optogenetic stimulation to recreate them. John Rickgauer and David Tank were the first to show that it was possible to activate channelrhodopsin-2 molecules by spiralling 920 nm laser light over the soma of transfected neurons *in vitro*[262]. However, to record the responses to stimulation, they had to perform patch-clamp electrophysiology on individual

cells. This is because 920 nm is also the peak two-photon excitation wavelength of GCaMP[257] and thus they would not be able to visualise responses with two-photon microscopy without also activating the opsin. The issue of cross-talk, where the excitation wavelength used for the reporter also excites the opsin, was a major confound preventing the ideal "all-optical" experiment[263]. Luckily, the field of opsin engineering had been developing actuators responsive at higher wavelengths, or "red-shifted opsins", to allow spectrally distinct dual-colour optogenetic control of neuronal populations[264,265]. After their previous success identifying a new opsin in the algae *Volvox carteri*[266], researchers led by Peter Hegemann and Karl Deisseroth hoped that the algae might also express red-shifted opsins. They discovered an opsin, VChR1, with polarity differences to channelrhodopsin-2 in the retinal-binding pocket that caused a redistribution of positive charge along the binding complex, suggesting different colour tuning compared to the *Chlamydomonas* opsins[267]. Indeed, VChR1 was much more red-shifted than channelrhodopsin-2 and even channelrhodopsin-1, with maximum excitation at 535 nm. It also had a slower decay constant than channelrhodopsin-2, meaning that single light pulses could trigger multiple excitatory post-synaptic potentials. This aligns with VChR1 expression being promoted by extended dark periods in its native environment[268] and therefore being beneficial to respond strongly to lower light stimulation. However, VChR1 had weaker expression than channelrhodopsin-2 using the same promotor[267]. Wanting the best of both worlds:

a strongly expressed red-shifted opsin with fast kinetics, Hegemann and Deisseroth once again collaborated, producing a chimaera of channelrhodopsin-1 and VChR1 fused at helix 3. The chimaera was better expressed at the membrane than both channelrhodopsin-2 and VChR1; it also had enhanced peak photocurrents compared to VChR1 whilst retaining its ion selectivity[269]. An added Kir2.1 trafficking signal further increased expression and photocurrents. However, they found that the chimaera's photocurrents decayed >10-fold more slowly than channelrhodopsin-2 and even VChR1. To speed the decay kinetics up and afford greater temporal control over injected activity, they mutated two residues in the chromophore region, naming the resulting molecule C1V1-t/t. C1V1-t/t will here on out be referred to as C1V1.

After proving its ability to produce spikes under two-photon stimulation *in vivo*[270], C1V1 became the original opsin of choice for all-optical experiments[271–274] (but see[275,276]), with greatest photocurrents produced under 1040 nm excitation[270]. John Rickgauer, Karl Deisseroth and David Tank were the first to demonstrate the power of all-optical interrogation *in vivo*, using two-photon stimulation of individual C1V1-expressing hippocampal cells to impose unnatural place fields during navigation of a virtual environment[277]. Importantly, they showed that even though there was still a slight overlap in the two-photon excitation spectra of GCaMP and C1V1, imaging GCaMP at 920 nm at powers on sample < 40 mW produced no significant increase in calcium events in opsin-expressing cells.

Despite scanning mirrors facilitating rapid repositioning of the stimulation laser spot across different targets (1-2 kHz)[277], the time needed to recreate stimulation-evoked population activity would exponentially increase with the number of targets, and would no longer be pseudo-simultaneous at larger group sizes such as the hundreds of S1 neurons activated by whisker stimulation[247]. To circumvent this issue, researchers turned to spatial light modulators (SLM) as a tool to help shape light into arbitrary patterns that could stimulate many cells simultaneously. SLM's are liquid crystal devices that reshape an incoming beam by modulating the phase, producing constructive interference at target sites and destructive interference off-target[278]. Crucially, by redistributing the available power of the beam amongst beamlets (unlike digital micromirror devices which remove off-target light using a physical mask[279]), the total power available is still high enough for multiphoton excitation. Additionally, the liquid crystal layer is electrically programmable, allowing online switching of the projected pattern, unlike previously fixed multiphoton beam splitters[280]. Researchers in the lab of Rafael Yuste were the first to use an SLM to perform truly simultaneous two-photon optogenetic stimulation of multiple neurons at cellular resolution *in vitro*[281] (but see previous work with a fixed beam splitter[280]). Finally, this technology was eventually used *in vivo*, in combination with two-photon imaging of GCaMP, to recreate whisker-evoked activity in a small subset of responsive wS1 neurons[272].

In his seminal paper 'Cause and Effect in Biology'[282], Ernst Mayr lamented that because there were very few techniques that allowed for differentiation between causation and correlation at the time, one could not hope to establish causality in biological phenomena. All-optical selective recreation of cortical population activity has helped leave this view in the past, especially regarding our understanding of perception. For example, neuronal population activity in visual cortex evoked by grating stimuli was found to be causal for perception in the absence of a physical stimulus[175]. This study also helped to differentiate perceptual signals from those associated with report, as artificial cortical recreation of go and no-go stimuli reproduced the corresponding behaviour. Importantly, they also found that the number of stimulated cells needed to elicit a percept was higher in layer 2/3 than layer 5. This may be because stimulation of layer 2/3 neurons recruited neurons in layer 5, whereas layer 5 stimulation did not recruit layer 2/3. The authors suggest that this points to a higher weighted role of deeper layer neurons in perception, potentially due to their subcortical outputs to areas like striatum. Another study showed that cortical networks have an auto-complete function whereby only a small number of 'pattern completion' neurons (minimum 2) need to be stimulated to reliably recall an entire visually evoked ensemble, complete with its associated behaviour[283]. However, it has also been shown in a different cortical region (anterolateral motor cortex) that it is not the stimulus responsivity of the directly activated neurons that is causal for

reproducing the percept associated with the stimulus, but rather the responsivity of cells in the local network recruited by the stimulation[284]. Even with more abstract percepts, like the perception of location within an environment, all-optical interrogation has revealed causal roles of population activity in higher cortical areas. For example, building on the original single-cell all-optical work of John Rickgauer, Karl Deisseroth and David Tank, another lab showed that when hippocampal place cell activity evoked by a reward zone was optogenetically recreated elsewhere, licking activity was produced as if the mice perceived themselves as teleporting into the reward zone[285].

Finally, a few years ago, researchers in the lab of Adam Packer used an all-optical approach to investigate how perceptual signals propagate through the cortical hierarchy in a CPD task[286]. As previously mentioned, direct cortical stimulation circumvents the fan-out of activity during pre-cortical processing that is present with externally presented stimuli. When investigating the neural substrates of perception, this confers two main advantages:

1. There is a high degree of control over the stimulus, and assumedly the percept, as sources of trial-by-trial perceptual variability associated with identical external stimuli are avoided, such as the relative spatial position of sensory receptors and attentional state.

2. Perceptual signals must propagate from the site of stimulation outwards, as an animal trained to report an optogenetic percept is required to only use features of the evoked activity to guide its behaviour.

An all-optical CPD task affords even more control over the stimulus, with single-cell spatial resolution targeting enabling questions to be asked about how the functional identity of stimulated neurons affects propagation and perception. Not only that, but simultaneous calcium imaging allows the impact of ongoing dynamics in the stimulated area on perception to be investigated. The authors focused on the latter, finding that higher pre-stimulus signal-to-noise in S1 decreased the probability of perception, which was encoded by the propagation of activity to S2 (similarly to early physical somatosensory stimulus work[126,128–130]). To achieve higher signal-to-noise, more activity can be optogenetically injected into the network. This explains why the number of neurons stimulated vs the probability of perception has a psychometric function[287] and why trial-by-trial fluctuations in the amount of local photostimulus-evoked activity predict perception[288].

2.1.2 Temporal coding of perception

2.1.2.1 Synchronisation

Previously discussed studies have shown that neuronal firing rates encode perception more strongly than stimulus as activity propagates through the somatosensory cortical hierarchy in both physical[126,128,130] and CPD[286] tasks. Similar findings have also been reported in the visual cortical hierarchy during visual-based decision making tasks[289–291], with thresholds for perception in an electrical microstimulation CPD rising slightly from earlier to later stages[292]. We have also discussed how cortical neurons nonlinearly integrate received signals within the dendritic tree, relying on temporally coincident inputs for reliable propagation[219–221]. The idea that neuronal synchrony may also be a substrate for perception was primarily proposed within the visual system towards the end of the 20th century[293,294]. In terms of ongoing activity in sensory areas, this makes sense as synchronized activity is hypothesised to represent arousal/attention, and would help to boost the signal-to-noise ratio of perceptual signals[286,295,296]. There is also the notion that neuronal synchrony may carry the perceptual information itself. This was spearheaded by the lab of Wolf Singer who found that stimulus-responsive neurons in the cat visual cortex represented stimulus properties through synchronization[293,297,298]. Later studies progressed this line of reasoning to behavioural report studies. For example, Pascal Fries, firstly in the lab of

Robert Desimone, showed that visual stimulus evoked synchronization in the gamma band in V4 was dependant on perception[299]. Then later in his own lab, showed that this was not limited to higher cortical regions, and in fact originated in V1, with V1-V4 synchronization correlating with perception and reaction time in a visual discrimination task[300]. This suggests that presynaptic activation patterns propagate to higher cortical regions in a temporally coordinated effective manner, enabling perception[301]. It seems as though cortical activity evoked by unperceived visual stimuli is lost somewhere along the V1-V4-prefrontal cortex pathway, contrasting to perceptual signals which reached higher cortical regions and then propagated back through recurrent connections[302]. The reliance of perception on feedback propagation from higher to primary cortical regions has also been shown in the somatosensory system[130].

Synchrony has also been investigated in the somatosensory system, especially with relation to the thalamocortical loops. Synchrony between S1 and the primary facial somatosensory nucleus of the thalamus (ventral posteromedial nucleus (VPM)[303]) directly reflects somatosensation[304], with velocity and frequency of whisker deflection being instantiated in the time course and strength of thalamic synchrony[305]. The high density of thalamocortical connections to individual neurons in layer 4 of S1[219,306] provides enough excitatory input to overcome strong local inhibition[307]. Similarly to within the visual hierarchy, synchrony within the somatosensory thalamocortical loop can encode properties of tactile

stimuli; interestingly, decoding of whisker deflections decreased with repeated presentations, whilst decoding properties of the deflections, such as velocity, increased[308]. This suggests that adaptation of thalamocortical responses can help fine tune stimulus representations. In 1984, Francis Crick proposed that the thalamus gates perception of incoming information with an "attentional searchlight", facilitating the transmission of attended stimuli to the cortex through synchronised bursting[309]. Behavioural studies have expanded on this hypothesis, showing that an increase in synchronous oscillations between thalamus and S1 at various power bands correlates with perception of tactile stimuli[310]; specifically bottom up synchronisation between primary thalamic nuclei and S1[311,312]. Recent work has suggested that the temporal precision of these synchronous bursts carries more information than a rate code[313]. It has been proposed that these zero-lag synchronisations between primary thalamic nuclei and S1 could be driven by a common input from higher cortical regions, such as prefrontal cortex[311]. Higher order somatosensory nuclei, such as the posterior complex of the thalamus (PO), are a major output for layers 5 and 6 of somatosensory cortex[314], yet unlike VPM, respond weakly to passive whisker stimulation[315–317]. Instead, PO has been shown to encode goal directed somatosensory behaviour[318,319], with S1-PO connections having a causal role in tactile perception[320]. Recently our understanding of this role has been expanded to multimodal perception, alongside higher order thalamic nuclei of other sensory modalities[321]. This mirrors classical findings in

the cortex that perceptual signals evolve throughout the cortical hierarchy. Indeed, higher order somatosensory and visual thalamic nuclei have a greater modulatory impact on their respective higher order cortical areas compared to their primary cortical counterparts[322]. This may help to synchronise S1 and S2 during goal-directed somatosensation[129].

Dissecting the directional flow of perceptual information in a physical tactile stimulus detection task is difficult between S1, S2, VPM and PO due to their highly interconnected wiring. The other main output of wS1 (outside of the brainstem) is the dorsolateral striatum (DLS) which is unidirectional in its connectivity[323]. This permits more clear investigations as to the role of DLS in propagating perceptual information from somatosensory cortex outwards. DLS is classically implicated in task-specific sensorimotor transformation[324–326], rather than encoding percepts itself. However, these two hypotheses are difficult to tease apart. The firing rates of DLS neurons can differentiate between hits and misses in a whisker stimulation detection task, with optogenetic activation of direct pathway striatonigral neurons (but not indirect pathway striatopallidal neurons) reproducing licking behaviour[327]. But optogenetic activation of S1 corticostriatal terminals reduces behavioural reports when given concurrently with a whisker stimulus, perhaps due to the dense innervation of striatal inhibitory neurons[328]. However, a different study found the opposite, that activation of S1 corticostriatal terminals at the time of whisker stimulation increased reporting whereas

activation of the cell bodies in S1 decreased reporting[329]. But in the absence of a physical stimulus, both cell body and terminal stimulation reproduced the behavioural report. This suggests that optogenetic stimulation of striatal-projecting cells in S1 may disrupt the integration of whisker relevant information at the level of the dendrites. However, the lack of control no-report trials means that it is impossible to differentiate between striatum biasing perception or biasing action selection. There is evidence from studies which investigated corticostriatal synchrony that communication between the regions is more motor-based than it is percept-based. For example, the same type of corticostriatal synchronisation that is observed emerging throughout learning in prefrontal cortex[330,331] is also seen in the motor cortex, but specifically between cortex and striatal cells directly relevant for motor execution in a BMI task[332], suggesting that corticostriatal synchronisation helps map task-appropriate motor programs in the basal ganglia[331]. More research is needed to see whether these findings generalise to sensory corticostriatal networks. Additionally, subcortical-subcortical synchrony may also carry stimulus relevant information; Nicholas Pnvert and Santiago Jaramillo showed that thalamostriatal connections represented sound phase, or timing, through synchronisation whereas corticostriatal neurons represented the same sound's amplitude in their firing rate[333].

2.1.2.2 Sequences

Despite there being a wealth of literature discussing the role of inter-area synchronisation in perceptual processing[301,334,335], comparatively little has been published regarding the role of temporal spike sequences outside of bird songs[336–338] and hippocampal replay[339–341]. In sensory cortices, the identities and spike times of responding neurons have been shown to differ between trials, even when external stimuli are kept constant[222]. This contrasts with *in vitro* experiments where injecting the same current into a neuron produces responses with the same temporal reliability across trials[224]. Responses are likely temporally reliable and precise *in vitro* because there are little to no intrinsic sources of variability, unlike *in vivo* where even minor fluctuations in activity from a small number of the thousands of inputs onto a neuron can affect said neuron's response characteristics to a stimulus[223,224,226,342]. These small perturbations may be amplified in the cortex, facilitating fundamental alterations to spiking activity on a trial-by-trial basis[343] that have been suggested to be a byproduct of the aforementioned cortical synchronisation events[344]. Simon Thorpe and colleagues were some of the first to go against the established rate code dogma and implicate spike timings as a neuronal code of perception. Building on work which showed that complex visual stimulus-evoked activity reached temporal[345] and frontal regions within 100 ms[346], they suggested that propagation through the ten different processing stages prior to these regions combined with the slow

conduction velocities of cortical fibres[347], implied that the remaining 10 ms for computation in each stage was not enough to reliably calculate a rate code[348]. To perform complex tasks like face detection whilst only firing ~1 spike per neuron per processing stage, Thorpe suggested that using inter-stage spike intervals and relative intra-stage firing order as a code was necessary[349]. To computationally validate this prediction, they constructed a four-layer feedforward integrate and fire model that could detect human faces using only one spike per neuron[348].

Meanwhile, Gilles Laurent and colleagues were making very similar discoveries in a vastly different model system – the locust olfactory network. They found that individual olfactory neurons would respond not only with different firing rates for different odours, but that some of the spike times underlying this were consistent across repeated presentations. These consistent spiking events led to odour-specific synchronisation time-scales[350]. Therefore, information about an odour was simultaneously encoded in neuronal identity and precise temporal firing sequence, resulting in a combinatorial spatiotemporal representation[351]. Clearly there is evidence that reliable temporal sequences are evoked in sensory areas by passive stimulation. To test the functional relevance of these sequences, Mark Stopfer and colleagues trained bees to respond to specific odours by extending their probosces. Like locusts, bees also encoded odour identity with reliable spatiotemporal firing. Reliable inter-neuron spike sequences were abolished by treatment with a GABA antagonist, whilst retaining reliable intra-

neuron sequences. This desynchronisation prevented bees from discriminating between similar odours but did not affect dissimilar odorants[352]. This shows that the relative order of neurons in a spiking sequence was fundamental to detailed odour perception, but not broad discrimination. Later, the same group showed that the spike sequences representing different odours diverged throughout their length, becoming more dissimilar later in the sequence without changing their firing rate stimulus tuning[353]. This suggests that ~1 s of time can gradually optimize stimulus representations through spike sequences in sensory networks.

In the mid-2000's, researchers in the group of Kenneth Harris started investigating how temporal sequences are instantiated during spontaneous activity in the rat cortex. They found that the onset of UP states (massive, persistent network activity as opposed to generalised silence in DOWN states) incurred the onset of reliable ~100 ms sequences with up to millisecond precision that decayed over time. Sequences were recruited by passing UP waves, with the same sequences being reliably produced by arriving waves from any direction[354]. These "packets" of sequences occurred sporadically during quiet wakefulness but also at tone onset in auditory cortex. Packets sometimes overlapped spatiotemporally but still reliably differentiated between auditory stimuli, and were theorised to reflect the opening of a "gate" that allowed sensory cortices to broadcast stimulus information to other regions[355]. The functional relevance of packets to perception was later described by Jorrit Montjin and colleagues, who

found that the order of visual cortex packet activation correlated with perception in a go/no-go visual task and the reliability of intra-packet spike times was higher on hits vs misses[356]

These propagating sequences of packets are conceptually similar to synfire chains. Synfire chains were initially proposed by Moshe Abeles in 1982[357] as a method of forming functional assemblies of neurons. In its simplest form, a synfire chain is a feedforward network where each layer has full unidirectional connectivity, allowing synchronous activity to be propagated to successive layers like dominoes[358]. Since being proposed, most of the discussion has centred around the theoretical plausibility of synfire chains as an architecture capable of transmitting a neural code[359–363], due to the analytical difficulty of identifying them *in vivo* (but see[364,365]).

Identifying reliable repeating temporal sequences of spike times in neuronal populations is analytically challenging. The advent of high-density silicon probes has allowed us to record electrophysiological activity with spike resolution from hundreds of neurons simultaneously along a shank that can penetrate up to 10 mm deep in the brain[26]. To cope with this staggering volume of neuronal data, popular sequence finding algorithms rely on searching for sequences of binned activity[366] or compute earth mover distance matrices between spike times evoked during different trial types[367]. Both of these approaches come with their own trade-offs; either spike resolution is lost, which may be crucial to the temporal code[348], or

very little information can be ascertained about the sequences themselves. Additionally, both approaches are insensitive to sub-sequences, i.e. certain spikes evoked by certain neurons that form a reliable temporal pattern[350], which would be drowned out trying to exclusively capture sequences that utilise all spikes. One approach that is able to capture these fine-resolution temporal dynamics is frequent pattern (FP) mining. FP mining was conceptualised by Rakesh Agrawal to identify reoccurring patterns within large databases of customer purchasing transactions[368]. By repurposing this algorithm in a more computationally efficient way utilising FP-growth, one can find neurons which infrequently reliably fire together, even if the majority of neurons and spikes just represent noise[369]. Adding a temporal component in order to find sequences of spikes, as well as synchronisation events, exponentially increases the computational power required for every neuron and time bin added, taking FP-growth ~4.5 hours to process a 15 second recording of 100 neurons[370] (which is still faster orders of magnitude faster than the >60 days it would take to mine without FP-growth[371]). For recordings encompassing more than 100 units across longer timescales, a more computationally efficient FP-growth algorithm is required that parallelises mining operations[372].

2.1.3 Objective

It is clear that CPD tasks are the gold standard behavioural assay for investigating the neural substrates of perception. Although rate codes have primarily been implicated as the most viable neural code, there is evidence that inter- and intra-region spiking synchronisation in tandem with intra-region spiking sequences play a crucial role in perceptual signals evolving throughout the cortical hierarchy. We are the first to combine these lines of enquiry, using Neuropixels to record from S2, VPM, and caudoputamen (CP) during an S1 optogenetic CPD task. And as far as we are aware, we are the first to use spatiotemporal spike pattern detection and evaluation (SPADE) FP-growth mining to investigate spike-resolution inter-area temporal signatures of perception in the top-down somatosensory pathway. Our subsequent aim is to understand how the structure of the optogenetic input influences perception; to do this, we used an all-optical approach to vary the identity of stimulated neurons based on function and projection. One stimulation pattern was identical across trials, whilst all other trials were characterised by a random selection of targets, allowing us to also investigate how the reliability of stimulation at cellular resolution affects the reliability of the temporal code in subcortical regions. Pre-training with widefield optogenetic stimulation ensured that all stimulation patterns were on manifold, and not sequestered into common outputs with local inhibition.

2.2 Methods

2.2.1 Animals, viruses and surgery

Some of the methods were adapted from my previously published work[166]. Animal experimentation was carried out with approval from the UK Home Office and the University of Oxford Animal Welfare and Ethical Review Board. 3 male wild-type mice (C57BL/6J) were used in the pilot study. In the main study, mice were an even mixture of sexes (5 males, 5 females) with transgenic expression of GCaMP6s in excitatory neurons [mouse strains B6; DBATg (tetO-GCaMP6s)2Niell/J crossed with CamK2a-tTa (AI94)] (Jackson Laboratory). At 7 to 14 weeks old, mice were implanted with a cranial window (or an optic fibre for some mice in the pilot study) over the right somatosensory cortex. At the beginning of surgery, mice were anaesthetized with isoflurane (3%, 1L/min) and kept under maintenance anaesthesia for the rest of the procedure (1.5% to 2.5%, 0.3 to 0.5L/min). Preoperatively, buprenorphine (Ceva) (1 unit/g), meloxicam (Boehringer Ingelheim) (1 unit/g), and dexamethasone (MSD) (five units) were administered subcutaneously. The scalp was shaved and cleaned with an antiseptic solution of 2% chlorhexidine gluconate and 70% isopropyl alcohol (Becton Dickinson) before head-fixing the animal via ear bars on a stereotactic rig (Kopf Instruments). Scalp and adipose tissue were excised from the right lateral side of the skull. A titanium headplate weighing ~1g with a 7 mm circular aperture was secured to the skull (−1.4 mm anterior, +4 mm lateral / −1.9 mm

anterior, +3.8 mm lateral relative to bregma) using superglue and dental cement (SuperBond, Prestige Dental, Bengaluru, Karnataka, India). For some mice, a second layer of 20% carbonated cement was used to further adhere the headplate to the skull. A 3 or 4 mm craniotomy was drilled with a dental drill (NSK) at the centre of the headplate aperture, and the dura was removed. A micropipette was either front-loaded with pAAV-CamKIIa-C1V1(t/t)-mScarlet-KV2.1 (Addgene: 124650-AAV9; 2.70×10^{13} vg/mL) diluted 1:11 in sterile PBS, or pAAV-Syn-Chronos-GFP (Addgene: 59170-AAV1; 1.00×10^{13} vg/mL) for the pilot experiments. 500 nL of the virus was injected into the left of the craniotomy for the main experiment, or 150 nL for the pilot mice, 300 μm deep from the cortical surface at a rate of 50 nL/min using a hydraulic micromanipulator (Narashige, Tokyo, Japan). A cranial window comprised of two optically bonded glass coverslips (one the size of the craniotomy and one 1 mm diameter larger) (Warner Instruments, Massachusetts, United States) was used to plug the craniotomy. The cranial window was secured in place using cyanoacrylate (3M) and either carbonated Kwik-cast or dental cement. The animal was allowed to fully recover for at least 3 weeks before any further experiments were conducted.

2.2.2 Two-photon imaging and whisker mapping

The microscope used in this study was a Bruker Ultima 2Pplus equipped with a tuneable wavelength laser for imaging (Chameleon Ultra II, Coherent,

Saxonburg, Pennsylvania 16056, USA). The animal was head-fixed on a tip-tilt stage (TTR001, Thorlabs) under a 16x objective (N16XLWD-PF, Nikon, Tokyo, Japan). Imaging data was captured using Prairie View software (Bruker). Reference two-photon images for GCaMP or Chronos expression (conjugated to GFP) were captured at 920 nm using a 500 to 550 nm emission filter, and C1V1 expression (conjugated to mScarlet) was captured at 765 nm using a 570 to 620 nm emission filter (80 MHz pulse repetition rate, 50 mW average power after the objective). Individual FOV's at 2x optical zoom (580 um x 580 um), 1024 x 1024 resolution and averaged over 124 frames were tiled across the cranial window to form an atlas image. Widefield calcium responses to whisker stimulation were imaged across the entire cranial window using a 4x objective (RMS4X, Thorlabs) and a camera (Orca Flash 4.0, Hamamatsu - 450-490 nm excitation filter, 500-550 nm emission filter). Up to five whiskers were threaded individually into a capillary tube controlled by a piezoactuator (Physik Instrumente) and deflected for 1 second every 10 seconds 10 times. The capillary tube oscillated at 10 Hz across an amplitude of ~300 μm. Responses to whisker stimulation were visualised using ΔF/F baselined to 1 s pre-stimulus. This allowed us to localise the S1/S2 border and confirm that the opsin was expressed in S1.

2.2.3 Behavioural training

All animals were allowed to fully recover for a minimum of 3 weeks post-surgery before being put on water restriction. Mice were kept on a 12:12-h light/dark cycle, with access to food *ad libitum*. Once on water restriction, the weights of the mice were monitored delay and not allowed to fall below 80% of their pre-water restriction bodyweight. For behavioural training, mice were head-fixed with their forepaws resting on a Lego wheel, as previously described[373]. Wheel movement was measured using a rotary encoder (Kubler 0.52400.1122.0360). The precise timing and volume of water delivery were controlled by a solenoid pinch valve. During neuronal recording sessions, two cameras recorded pupil and face/body movement (FLIR, lens: MVL50M23, MVL16M23) in combination with an infrared light source (BW 48 LED Illuminator). A patch cable attached to a fibre-coupled LED (M470F4, Thorlabs – pilot experiment, M530F3, Thorlabs – main experiment) was lowered to just above the cranial window. The behavioural task was programmed and controlled using Rigbox[374]. More details on the hardware and software can be found here[375,376]. Optogenetic stimulation spanned 180 milliseconds, consisting of five 20 millisecond pulses separated by 20 millisecond inter-pulse intervals. To earn a 3 uL water reward, the wheel had to be turned at least 6 mm in total in any direction during the response window.

2.2.4 Analysis, statistics and reproducibility

All analysis was performed using MATLAB, except for SPADE which was performed using python. No statistical method was used to predetermine sample size. No data were excluded from the analyses unless otherwise specified (see supplementary figure 7 for example). The scientists were not blinded to opsin or control allocation during experiments or analysis.

2.2.5 Timing index

For a given list of movement times ***M***, and a given period ***T***, movements were converted to phases (how many rotations of a 2π radians circle):

$$\Phi = 2\pi M/T$$

The timing index ***TI*** represents how reliably spaced the phases are:

$$TI = \frac{1}{N}\sum e^{j\Phi}$$

2.2.6 Electrophysiology

Electrophysiological data was recorded using Neuropixels probes[26]. Before insertion, the shank was rinsed in deionised water, then coated in a red-fluorescent dye (DiI, V22885, Invitrogen). Optimal trajectories for penetrating S2, VPM, and

CP were calculated using Pinpoint[377]. To simultaneously skewer S2 and CP, the probe was inserted 5.2 mm deep at -60 yaw and 32.5 pitch. To simultaneously skewer S2 and VPM, the probe was inserted 4.5 mm deep at -100 yaw and 40 pitch. The probe was lowered at ~5 um/s using a micromanipulator (Sensapex) to 100 µm deeper than target depth and then retracted by 100 µm to reduce tension within the tissue. Electrophysiology recordings were captured using SpikeGLX[378] (30 kHz sampling rate) and preprocessed using CatGT[379] which performs common reference averaging and bandpass filtering. Spike sorting was performed using Kilosort4[380] and curated using Bombcell[381]. Defined Bombcell parameters: minimum signal-to-noise ratio = 5, minimum amplitude = 40, maximum resting period violations = 0.1, hillorLlobetMethod = 0.

2.2.7 Sequence detection

Sequence detection was performed using a computationally optimised version of SPADE (spatiotemporal spike pattern detection and evaluation analysis) [369,370,372,382] with a bin size of 5 ms (other bin sizes were also tested; see supplementary figure 7). SPADE uses frequent pattern mining to identify reliably reoccurring sequences of activity at spike resolution, independent of broader population activity. SPADE was run on spike times within the 100 ms post-stimulus onset, concatenated across all hits within a session, and separately across

all misses. The window length over which the sliding window analysis occurred was 100 ms, setting an upper limit for the duration of identified sequences. Therefore, concatenated spike times were padded between trials with ((100 ms/bin size)*bin size) seconds to avoid sequences being erroneously identified as spanning multiple trials due to the sliding window length. Significant sequences had to occur on at least 50/70/90% of a given trial type and survive comparison with a shuffled distribution to control for firing rate (100 shuffles with 20 ms dithering).

2.2.8 Histology

On the last day of electrophysiology recording, mice were injected with a lethal dose of pentobarbitone and perfused with 4% paraformaldehyde. The brain was extracted and left in 4% PFA for 24 hours. A vibratome was used to slice the brain at a thickness of 70 μm. Slices were mounted and imaged using a red filter to detect DiI and C1V1-mScarlet expression and a green filter to detect GCaMP. Probe trajectory reconstruction and opsin expression localisation and quantification were performed using HERBS[383].

2.3 Results

2.3.1 Mice can be trained to report optogenetic stimulation of wS1

The opsin C1V1 was virally expressed under a cranial window in layer 2/3 wS1 for training in an optogenetic CPD task (figure 1a, supplementary figure 2). During head-fixed behavioural training, mice were positioned with their forepaws on a wheel. Immediately following uncued trial onset was a wait period where any wheel turns over threshold resulted in the trial and wait period being reset. To disincentivise timing strategies, the pre-stimulus wait period was randomly drawn with replacement from a uniform distribution between 4 and 6 seconds. If there was insufficient movement during the wait period to trigger a reset, on two thirds of trials, an optic fibre positioned over the cranial window emitted 5 pulses of 530 nm light over 180 ms (20 ms ON, 20 ms OFF). If the wheel was turned to threshold during the response window starting at stimulation onset, the mouse received 3 uL water reward. On the other hand, if there was insufficient wheel movement during the response period to reach threshold, no reward was delivered. One third of trials were catch trials featuring no optogenetic stimulation. Any movements or lack of movements during the response window on catch trials were unrewarded. A fixed 5 second inter-trial interval followed either timeout of the response window or when wheel threshold was reached (figure 1b). To further make sure that only mice using the stimulation to guide behaviour were classified

as learners, d' was used to assess behavioural performance, with a threshold of 2 for graduation. 2 is a stricter threshold than used in previously published CPD tasks[216,286]. d' is measure of the distance between the hit rate and the false alarm rate. Without a randomised pre-stimulus wait period, catch trials, and using d' as a behavioural metric, mice could readily achieve high hit rates without using the stimulus to guide behaviour (supplementary figure 1). Mice learnt to report 6 mW stimulation in a dark rig above threshold within 4 – 15 days. To further ensure that perception of the optogenetic stimulation was used to guide behaviour rather than visual stimulation, the mice were shaped through 7 subsequent stages that gradually stripped away their ability to rely on visual information, including reducing the stimulation power on sample to 0.1 mW and fitting a light shield between the eye and optic fibre (figure 1c). Indirect retinal activation significant enough to evoke a visual percept is not reported at any of the stimulation powers used under ambient light[384]. However, this may have been unnecessary as none of 4 control mice without opsin expression could meet behavioural criterion even in stage 1 within 20 days (threshold for discontinuation of training) (figure 1d). This suggests that the mice were reaching high behavioural performance by reporting perception of the optogenetic stimulus. Mice that graduated all 8 stages were classified as experts. The number of sessions required to reach expert performance did not correlate with opsin expression, even when only opsin expressed directly under the cranial window was counted (supplementary figure 2).

Despite our high confidence that the mice were using perception of the optogenetic stimulation to guide behaviour, timing wheel movements could still have been a viable strategy to obtain a modest amount of reward. This is because the combination of a fixed inter-trial interval and 4 to 6 second uniformly drawn pre-stimulus wait period meant that there was always a one third chance that making a report between 10 and 11 seconds post-reward delivery would be a hit from stage 4 onwards. To thoroughly investigate the potential use of a timing strategy, we calculated the strongest periodicity with which the mice made reports per session and a timing index which represented how strictly that periodicity was utilised (see methods). We only included sessions in stage 5 or above because by that point the mice could not use visual cues to guide behaviour. As expected, reports delivered with a period of 10 – 12 seconds were the most common. Although periods below 4 seconds were also weakly utilised, they were primarily deployed in the first three stages of learning. As the mice progressed through the stages, their report probability more closely tracked reward probability (given a report). Mice were also significantly more likely to achieve hit after a previous hit compared to after a non-hit, especially in sessions with low behavioural performance, suggesting that the mice may use the previous reward to time their next movement. However, a simulated mouse with no access to stimulation information could not achieve a d' > 0.1 when exclusively moving with an optimal period of 10.5 seconds (supplementary figure 3). Therefore, maybe the higher probability of hits after hits could be attributed to the reward increasing

attention directed towards the optogenetic stimulus. Alongside the fact that control mice with no opsin expression could not achieve a d' > 2 on stage 1 that was easier to 'cheat' with a timing strategy due to the longer response window, taken together, these findings suggest that expert mice were using the optogenetic percept to guide behaviour and not exclusively relying on a timing strategy.

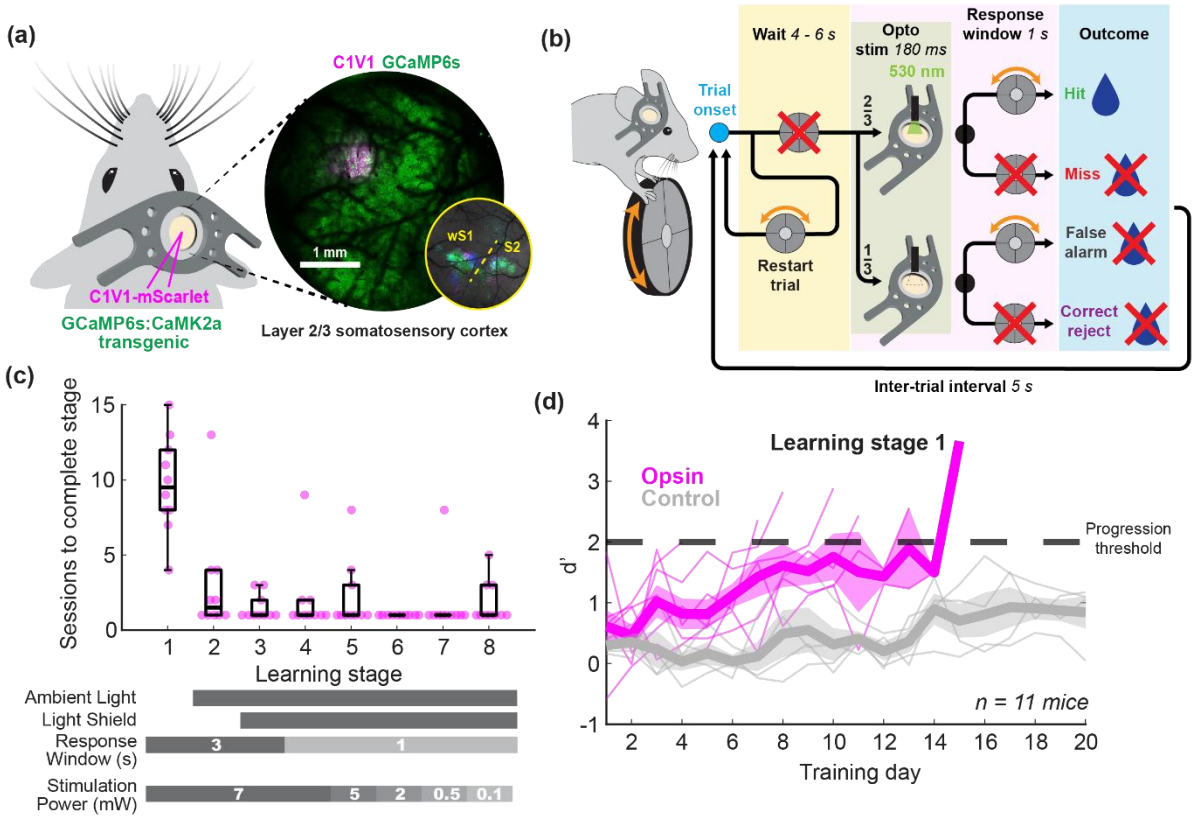


Figure 1. Mice can be trained to report optogenetic stimulation of wS1. (a) The opsin AAV-C1V1-Kv2.1-mScarlet was injected into layer 2/3 of right wS1 of 9 mice transgenic for GCaMP6s in excitatory neurons. **(b)** CPD task structure. **(c)** Number of sessions required to reach a d' ≥ 2 at each stage of training. **(d)** Control mice with no opsin expression failed to reach criterion in stage 1 within the 20 session limit.

2.3.2 Optogenetic stimulation of wS1 produces a similar downstream population response, independent of perception

Initial attempts to record high density electrophysiology data during expert performance of the optogenetic CPD task involved inserting a Neuropixels probe through a small craniotomy adjacent to an implanted optic fibre for stimulation. These initial experiments were fraught with challenges; namely that wS1 and S2 are only ~100 um apart[286], which meant it was difficult to penetrate S2 with the probe whilst a fibre was cemented above S1. This is especially true as the best micromanipulators at the time were drastically unreliable for performing anything other than vertical insertions, which also made it impossible to record from the desired subcortical structures. To overcome these barriers, we consulted with Sensapex to help design and test a new multiprobe rig that was designed for reliable repeatable angled insertions[385] (figure 2a). However, in order to perform Neuropixels recordings, it was necessary to reanaesthetise >15 week old mice after months of training to drill a hole through the glass window above S2. Naturally, this was a high-risk procedure that could potentially be avoided. To circumvent this, we implanted removable cranial windows sealed with a silicon ploymer[386] that maintained cortical health and quality for >100 days. Removing the cranial window for electrophysiology provided a large area of brain over

which to configure multiple simultaneous insertions, allowing concurrent recording of neuronal activity in all of the targeted output regions of S1.

Firing rates increased across all recorded regions at ~50 ms post-stimulation onset on both hits and misses (figure 2b). These responses were much larger than those triggered by movement alone with temporal dynamics that differed between regions. S2 activity matched the time course of the LED pulses whereas CP and VPM primarily responded to the first pulse to different extents. S2 units could be separated into roughly equal proportions of putative excitatory and inhibitory neurons based on their waveform, with putative excitatory neurons responding ~10 ms earlier on hits and misses (supplementary figure 5). CP responded even more strongly than S2 to the initial LED pulse but decreased quickly afterwards. The first LED pulse had a higher weighted influence on the firing rates of CP neurons compared to the subsequent four, implicating stronger adaptation mechanisms than in the other regions. By eye, CP activity from ~250 ms to 1 second post-stimulation could differentiate between hits and misses, similar to what has previously been reported in higher cortical regions in CPD[286] and external stimuli detection tasks[126,130]. However, whether this was due to perception, movement or reward consumption was unknown. To investigate this, we used DeepLabCut on videos of the mice captured during the task to track paw movement[387]. As expected, paw movement ramped shortly after stimulation onset and decreased once reward delivery threshold was reached (supplementary figure

4). We used these movement traces to calculate reaction time, defined as the time between stimulation onset and the beginning of paw movement. On average, it took expert mice 119 ms to significantly move their paws in response to stimulation (supplementary figure 4). To isolate perceptual signals from those representing movement and reward, we conservatively restricted analysis to the 100 ms post-stimulation onset. Stimulation evoked activity reached S2 first, followed by CP, then VPM (figure 2c). S2 and CP responded with sharp increases in firing rate whereas the VPM ramped more slowly. Population-level activity considerably overlapped between hits and misses for all target regions – in amplitude, shape and timing (supplementary figure 5). When aligned to reward, VPM came online first as expected from the ascending somatosensory path[303], shortly followed by S2 then CP (supplementary figure 4).

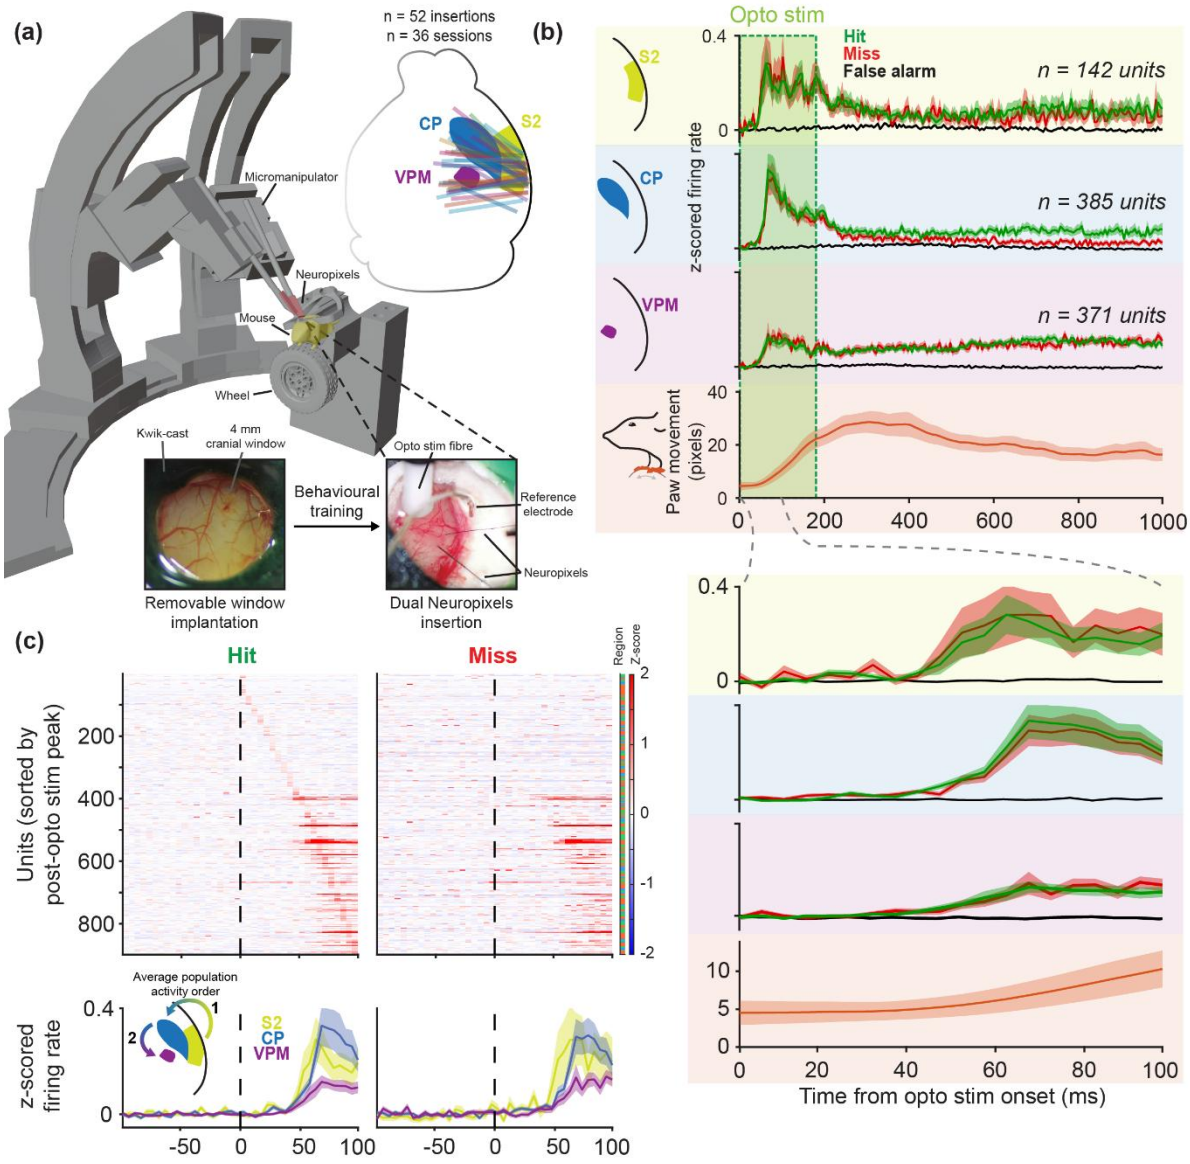


Figure 2. Optogenetic stimulation of wS1 produces a similar downstream population response, independent of perception (a) Sensapex multiprobe *in vivo* rig allowed for angled probe insertions that could simultaneously record from multiple regions within the top-down somatosensory pathway including S2, VPM, and CP. Top: all Neuropixels trajectories. Bottom: Implantation of readily-removable cranial window maintained cortical health and clarity for over 100 days and facilitated easy multiprobe recordings through a large craniotomy. **(b)** Mean z-scored firing rate ± std in the 1 s post-stimulation onset baselined to 300 ms pre-stimulation. Bottom: paw movement calculated using DeepLabCut aligned to stimulation onset. **(c)** Top: heatmap activity from all recorded neurons ordered by post-stimulation onset activity. Bottom: z-scored activity of each region aligned to stimulus onset.

2.3.3 Perception is represented in the timing and identity of rapid, reliable sparse spike sequences in the top-down somatosensory pathway

We next looked to investigate whether perception of the optogenetic stimulus was encoded in spike sequences throughout the top-down somatosensory pathway. One popular algorithm used to detect repeating sequences of neuronal activity is seqNMF[388–391]. SeqNMF uses convolutional nonnegative matrix factorization with a penalty term to uncover repeating sequences that are agnostic to event times[366] (supplementary figure 6). However, it only searches for a specific target sequence length. By requiring a specific time period over which sequences should span, seqNMF naturally emphasizes repeated population-level temporal structures and ignores precise timing details. By minimizing global reconstruction error, the algorithm prioritizes large, consistent temporal structures and small deviations in individual spike times are lost. Thus, seqNMF favours patterns that explain large fractions of population variance and is biased towards sequences that occur many times and involve many neurons. Overall, this means that reliable but sparse sequences which contain a small number of neurons and spiking events can be missed. When given all of the neuronal activity recorded within the 1 second response window concatenated, for the vast majority of sessions seqNMF reliably found one unique sequence that lasted for 1 second,

spanned the entire probe and roughly began at stimulation onset. This sequence occurred on the majority of trials and there was no significant difference in the occurrence rate between hits and misses. SeqNMF struggled to find a unique sequence that spanned 100 milliseconds. On the small number of sessions where it identified a short sequence, it only occurred on ~10% of trials. In conclusion, seqNMF failed to uncover any low-rank convolutional patterns that correlated with perception of the optogenetic stimulus.

Another popular algorithm for assessing whether spike sequences represent task[367,392] and behavioural[393] variables, but is yet to be used for latent variables such as perception, is SpikeShip. In contrast to seqNMF, SpikeShip uses precise spike times locked to task events to compute a dissimilarity measure between trials based on optimal transport theory[367] (supplementary figure 6). SpikeShip computes the optimal transport cost to make all the relative spike-timing relationships across neurons identical between two trials. Unfortunately, this means that once again the algorithm is blind to short, sparse sequences or how differences in sequence onset lag might differ between trial types. Nevertheless, the reliability of spike times across all recorded neurons in the 1 s post-stimulation onset could differentiate between hits and misses. But SpikeShip failed to separate spike sequences based on trial type when access was limited to 100 ms. This shows that, taking into account all recorded neurons, relative spike-timing relationships in the top-down somatosensory pathway are equally dissimilar

within and between perceived and not perceived optogenetic stimuli during the perceptual period.

Both of the previous methods are insensitive to reliable sequences only comprised of certain spikes in certain neurons, which are known to be important for representing stimuli[350]. One method that can detect small, sparse sequences of activity is frequent pattern mining (FPM), originally developed to identify reoccurring patterns within large databases of customer transactions[368]. FPM utilises candidate generation:

For a set of transactions ($\boldsymbol{T}$) featuring items ($\boldsymbol{i}$) within an itemset ($\boldsymbol{I}$)

$$\boldsymbol{I} = \{1,2,3,4\}\ \boldsymbol{T}_1 = \{1,2,3\}\ \boldsymbol{T}_2 = \{1,3,4\}\ \boldsymbol{T}_3 = \{1,2,3,4\}$$

Filter out items that occur below the minimum sequence occurrence threshold ($\boldsymbol{N}$).

$$\boldsymbol{N} = 2$$

$$\{\boldsymbol{i} | \text{freq}(\boldsymbol{i}) \geq \boldsymbol{N}\} = \{1,2,3,4\}$$

Generate candidate sequences (***C***) by permuting the remaining items with length ⩾ minimum sequence length (***S***)

$$S = 2$$

$$C = \{1,2\}\{1,3\}\{1,4\}\{2,1\}\{2,3\}\{2,4\}\{3,1\}\{3,2\}\{3,4\}\{4,1\}\{4,2\}\{4,3\}\{1,2,3\}\{1,3,2\}\{1,2,4\}\{1,4,2\}\{1,3,4\}\{1,4,3\}\ldots$$

Filter out candidate sequences that occur < ***N*** across ***T***'s

$$\{C|freq(C)\geq N\} = \{1,2\}\{1,3\}\{2,3\}\{1,4\}\{3,4\}\{1,2,3\}\{1,3,4\}$$

Extract maximal frequent (no superset) and closed (infrequent superset) sequences

$$\text{Maximal frequent sequences} = \{1,2,3\}\{1,3,4\}$$

$$\text{Closed sequences} = \{1,3\}$$

FPM finds sequences within transactions that do not have to include all items. It also avoids redundancy among extracted sequences by isolating maximal frequent and closed patterns. For example, in the transactions above {1,2} occurs twice which is above threshold for both sequence length and number of

occurrences, but it always exists within the same superset {1,2,3} and is therefore discarded. However, for larger transaction databases, permutation and filtering gets exponentially more computationally expensive. FP-growth reduces the computational cost by limiting permutation with FP-tree mining[369]. This spatio-temporal spike pattern detection and evaluation (SPADE) mining algorithm was initially developed by Sonja Grün's lab[369,370,382] and further computationally optimised for cost efficiency by mining each item in parallel and filtering candidates during FP-growth instead of separately after[372]. The computational optimisations meant that SPADE could be ran on all neurons recorded within a session in the 100 ms post-stimulation onset without overloading memory usage. It found both intra- and inter-region reliable sequences of sparse spikes that were either only present on ≥ 50/70/90% of hit trials, miss trials or were present on both (shared). We defined any sequences that consisted of the same neurons in the same order with the same inter-spike intervals on both hits and misses as shared sequences. These sequences spanned multiple neurons both intra- and inter-region and did not survive comparison to a shuffled distribution, controlling for firing rate (figure 3a). Additionally, there was no correlation between the firing rate of a neuron and the number of sequences it participated in (supplementary figure 7). When aligned to the first spike in the sequence (anchor), follower neuron activity rapidly increased shortly after (figure 3b). However, when aligned to rate-matched spikes from the same anchor neuron in epochs

where the sequence wasn't detected (miss trials for hit-only sequences, hit trials for miss-only sequences), follower neuron activity was normally distributed around zero. Taken together, these results show that although neurons participating in sequences have a broad preference to fire together, even outside of sequences, the sparse, reliable sequences of spikes that SPADE identified cannot be produced by firing rate alone. SPADE is also not biased towards overrepresenting sequences in regions containing the most recorded neurons; as expected, most neurons in the dataset come from S2, CP and VPM whereas sequences were particularly enriched in globus pallidus externa (supplementary figure 7), a finding which has previously been noted in rhesus monkeys[394]. Shared sequences, that is sequences comprised of the same neurons in the same order with the same inter-spike intervals on hits and misses, occurred significantly later on misses (figure 3c). In contrast, hit-only sequences occurred significantly later than miss-only sequences. All sequence types had an average duration of 10-30 ms and primarily only featured two unique neurons (figure 3d).

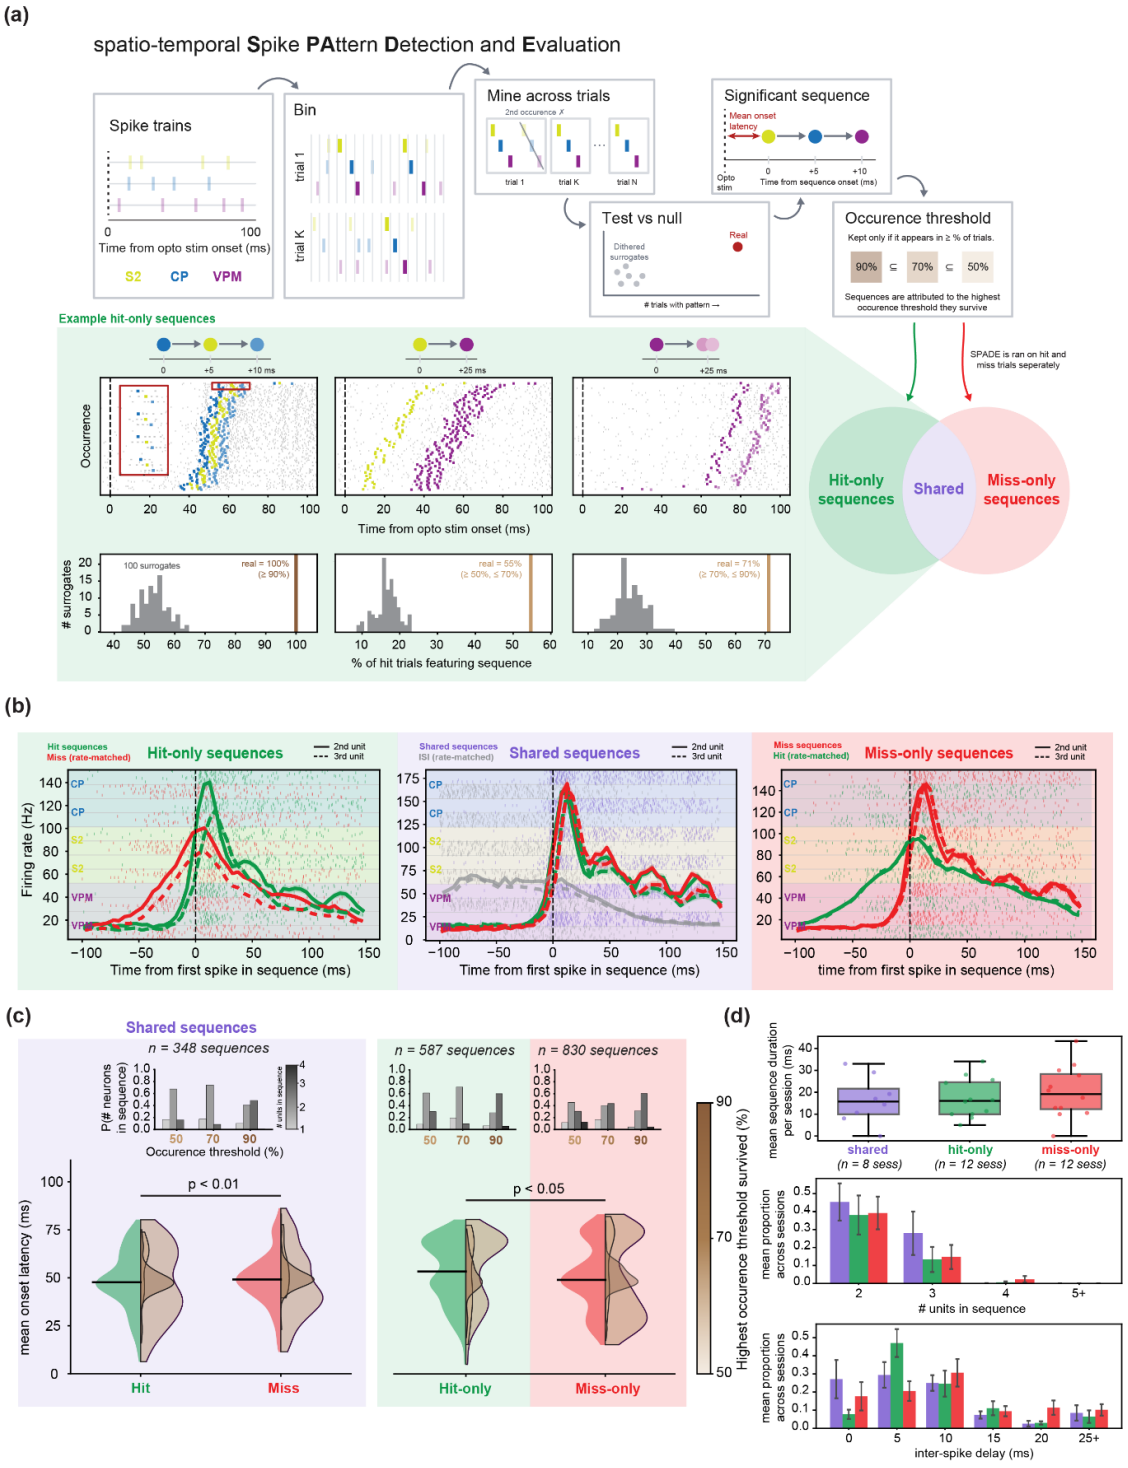


Figure 3. Perception is represented in the timing and identity of rapid, reliable sparse spike sequences in the top-down somatosensory pathway (a) Schematic of sequence detection using frequent pattern growth mining. **(b)** Activity of follower neurons in 3 neuron, 3 spike sequences aligned to the spike of the first neuron (anchor). Opposite-trial traces are aligned to a randomly chosen rate-matched anchor neuron spike. **(c)** Onset latency for shared sequences and independent sequences on hits and misses. $p < 0.01$ = Wilcoxon Sign rank (shared sequences) or Mann Whitney U (hit/miss-only sequences). Top: probability of different length sequences (defined by number of unique units) at different occurrence thresholds. **(d)** Comparison of different sequence types.

To further probe how sparse sequences may represent perception in their onset latency and identity, we split the analysis by target region based on the anchor neuron (figure 4). Unlike population firing rate activity that arrived first in S2, then propagated to CP and VPM, reliable spike sequences occurred in CP first on average, followed by S2 then VPM (figure 4a). The broad finding that shared sequences occurred later on misses held across S2 and CP, but the opposite was found in VPM. Similarly, miss-only sequences occurred significantly later than hit-only sequences in S2 and CP, but there was no different in onset latency in VPM (figure 4b). Interestingly, hit/miss-only sequences were recruited after the shared sequences on average. This may suggest that the onset delay of shared sequences determines whether to recruit hit-specific or miss-specific sequences. Such a hypothesis would support previous findings that show sequences encoding different stimuli diverge over time[353]. However, again the opposite relationship was seen in VPM and hit/miss-only sequences occurred at roughly the same time as shared sequences in S2. Together, this may suggest that shared and hit/miss-only sequences contribute differently to temporal codes of perception in the corticostriatal vs corticothalamic pathway. It has previously been shown that inter-neuronal sequences carry more discriminatory information about odour stimuli than intra-neuronal sequences[352]. Follow up analysis of our dataset should investigate whether inter-neuronal sequences also carry comparatively more weight in a CPD task.

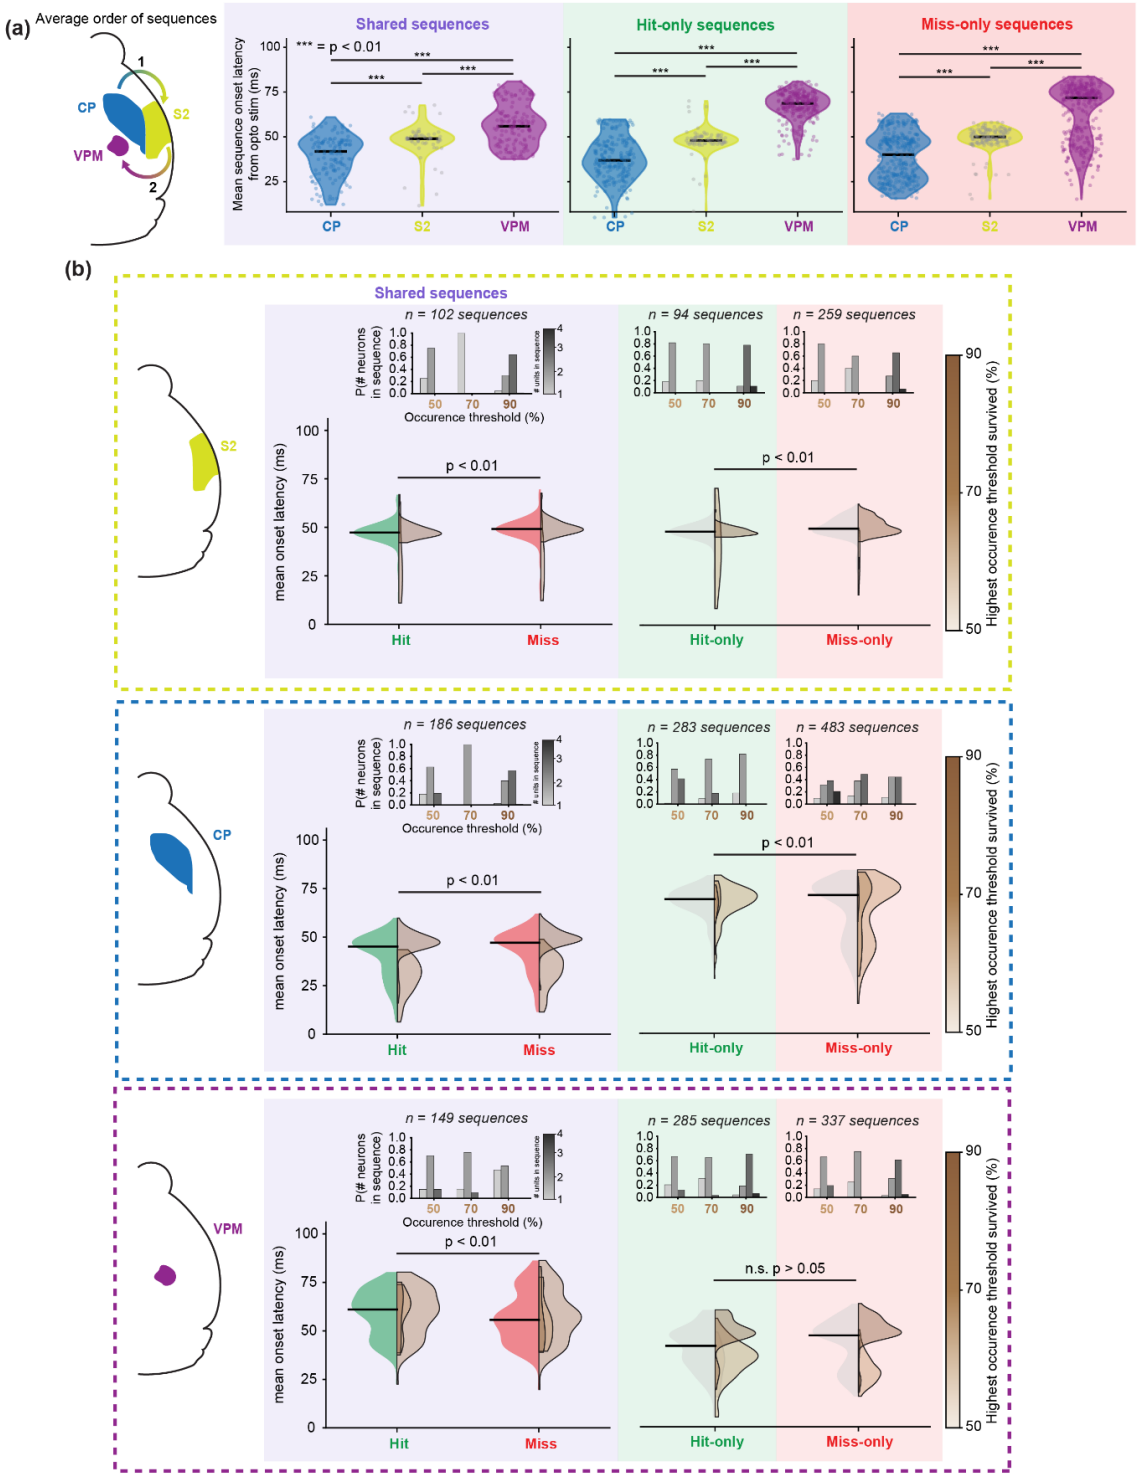


Figure 4. Unique temporal arrangements of sequences in the corticostriatal and corticothalamic pathways (a) Sequences are recruited first in CP, then S2, then VPM. *** = p< 0.01 Kruskal Wallis. **(b)** Onset latency for shared sequences and independent sequences on hits and misses in S2, CP and VPM. p < 0.01 = Wilcoxon Sign rank (shared sequences) or Mann Whitney U (hit/miss-only sequences). Top: probability of different length sequences (defined by number of unique units) at different occurrence thresholds.

2.4 Discussion

In this chapter, we show that rapid, reliable sparse spike sequences in S2, CP and VPM distinguish between perceptual hits and misses within 100 ms of a wS1 optogenetic stimulus. Previous studies have struggled to differentiate between perceptual signals and those involved in reporting perception, due to confounding pre-cortical fanout associated with externally presented stimuli or fluorescent reporters of activity with slow dynamics. Performing Neuropixels recordings during a CPD task allows us to investigate cortical and subcortical communication pre-movement onset with a fixed directional flow of perceptual information from wS1 outwards. However, it is unknown whether our pre-movement activity is confounded with motor-planning signals. It would have been useful to also include no-go trials where perception was reported through immobility to help further disambiguate perceptual and motor signals. Although the decision to not move also requires a motor plan, so implementing this methodology may not be a complete panacea.

One surprising finding from our training paradigm was that learning rate in the CPD task did not correlate with the amount of opsin expression. This was surprising because behavioural performance in an optogenetic CPD task has previously been reported to correlate with the number of opsin-expressing cells under widefield stimulation by Karel Svoboda's group[216] and the number of cells targeted under two-photon holographic stimulation in expert mice pre-trained on

widefield stimulation[286,287]. However, Karel Svoboda's group trained mice using a 1 ms light pulse per trial. Our 100 ms total light ON stimulation likely evoked many more action potentials[269] and thus the activity probably carried much higher signal to noise in the local network[218,227]. It has been shown that weaker stimulation evokes lower firing rates which requires more cells to be simultaneously activated for perceptual signals to successfully overcome inherent cortical noise[216]. This suggests that stronger stimulation utilising multiple pulses of light may be less reliant on high cell numbers to be simultaneously activated for salience.

An important question to consider when digesting these findings is why would the brain have evolved to use sequences of spikes to represent or transmit perceptual information instead of a rate code? As we can see in our data, mice on average react within 100 ms post-stimulation onset, but sometimes as quickly as within the first 20 ms. It is advantageous for an organism to be able to consciously perceive their environment as fast as possible; for example, if an animal is ambushed by a predator, their chances of survival decrease as their reaction time increases. However, when accounting for the time for the signal to undergo sensorimotor transformation, a rate code perceptual gating mechanisms is highly improbable because of the ~10 ms it would take to reliably calculate. If in the fastest cases the signal has undergone sensorimotor transformation within 20 ms, gating perceptual propagation with a temporal code is the only viable solution.

Furthermore, anecdotally the space of perceptual or conscious experience from a first-person subjective point of view is essentially infinitely vast. If we also assume that the perceptual experience itself has to be coded and transformed within this 20 ms, then a temporal code would be a viable mechanism to achieve this due to the almost limitless dimensionality of spike time and sequences that can occur across a handful of neurons even over a such a short duration.

3

Opsin desensitisation during all-optical interrogation depends on imaging parameters



3.1 Introduction (p. 72)

3.1.1 Channelrhodopsin photocycle (p. 72)

3.1.2 Cross-talk (p. 80)

3.1.3 Objective (p. 81)

3.2 Methods (p. 81)

3.2.1 Animals, viruses and surgery (p. 81)

3.2.2 Two-photon imaging and stimulation (p. 83)

3.2.3 GCaMP imaging analysis (p. 85)

3.2.4 Analysis, statistics and reproducibility (p. 86)

3.3 Results (p. 87)

3.3.1 Determining the impact of 920 nm two-photon imaging on 1035 nm photostimulation responses in targeted cells under common all-optical experimental conditions (p. 87)

3.3.2 Two-photon imaging at 920 nm increases 1035 nm photostimulation response desensitization in targeted cells at shorter inter-stimulus intervals (p. 90)

3.3.3 The distribution of non-targeted cells excited or inhibited by 1035 nm stimulation is unaffected by 920 nm imaging dose (p. 93)

3.3.4 Two-photon imaging at 920 nm exacerbates photostimulation response desensitisation of excited non-target cells (p. 96)

3.4 Discussion (p. 99)



3.1 Introduction

3.1.1 Channelrhodopsin photocycle

Parts of this chapter were adapted from my previously published work[166]; I confirm that I wrote the material reproduced. To probe how reliability of stimulation at cellular resolution affects the reliability of the temporal code underlying perception, we wanted to use an all-optical approach in wS1. Directing the stimulation laser through an SLM allows the targeted holographic stimulation of tens of neurons simultaneously[272] and hundreds of neurons pseudo-simultaneously by rapidly switching the hologram[286]. Crucially, concurrent activation of tens of cortical neurons in mice is enough to modulate existing behaviours[283–285,395] and even drive new behaviours[286,287,395,396]. As previously discussed, our aim was to use holographic photostimulation in mice who were experts at the one-photon task to vary the reliability of the percept at a cellular level; to achieve this, one stimulation pattern would be kept identical across trials whilst all other trials would be characterised by a random selection of the same number of targets. However, when trying to generate a reliable and repeatable stimulation pattern across trials, just targeting the same cells may not be sufficient.

C1V1 was the original choice of opsin for all-optical interrogation[271–274] and has remained a popular choice over the last decade[397]. Photostimulation responses of targeted C1V1-expressing neurons have previously been shown to decrease over

trials in an all-optical paradigm *in vivo* with a 10 second inter-stimulus interval (ISI)[395] and require ~30 seconds to return to baseline under one-photon stimulation *in vitro*[398]. Not only does this pose a problem with generating reliable and repeatable activity patterns over trials with reasonable ISI's, it also lessens all-optical's unique methodological advantage. Using an SLM enables the recreation of cortical population activity generated by externally presented stimuli, which positions all-optical experiments with a unique advantage in exploring causal relationships between neuronal activity and behaviour[399]. However, in order to artificially generate endogenous-like firing patterns, the power and scan pattern of the photostimulation laser have to be carefully calibrated using electrophysiology[272,400]. Therefore, any variable that can affect an opsin's stimulus-evoked photocurrent can have a potentially large impact on the validity of all-optical studies. One such variable is opsin desensitisation: the decrease in opsin activation due to repeated exposure.

When channelrhodopsins receive continuous light stimulation, their fast-peaking transient photocurrent rapidly decays to a steady-state level. The steady-state current is stable, even across successive stimulations, whereas the size of the transient current exponentially decreases with each excitation. This is due to opsin desensitisation and was first observed by George Nagel and colleagues in channelrhodopsin-2[205,206]. They found that the time needed to fully recover the transient current between stimulations correlated with membrane potential and

extracellular pH. From this, they surmised that channelrhodopsin-2 desensitisation was mediated by protonation of an extracellular amino acid, with Glu-123 as a likely candidate due to its pivotal role in facilitating both the transient and steady state currents. Based on these findings, they devised a simple three-state photocycle model of channelrhodopsin-2, consisting of a ground state ***C***, an open transient-conducting state ***O*** and a steady state conducting desensitised state ***D***. However, their data showed that the recovery from ***D*** to ***C*** in the dark was slow (τ = 2 seconds), from which the model predicted a much smaller transient current than was experimentally observed. The authors proposed two theories to explain this: the experimental transient current was bigger than predicted by the model due to ***D*** to ***C*** intermediates being catalysed by an additional photon when recovery is not in complete darkness, essentially forming a two-photon photocycle. This was based on solid reasoning as two-photon photocycles were known to exist in halorhodopsin[401] and proteorhodopsin[402]. The alternative explanation was a two-photocycle model with two ground states and two open states, with one of the open states being less conductive. The desensitised ground state would slowly transition back to the original ground state in darkness. A population of channelrhodopsin molecules simultaneously occupying both photocycles with different efficiencies could explain both the high transient to steady state current ratio and slow recovery time.

Peter Hegemann and colleagues were the first to note that the extent of channelrhodopsin-1 desensitization increased as ISI decreased[403]. They also proposed a two-photocycle of model of channelrhodopsin as they yielded that a single photocycle could not capture all major channelrhodopsin-1 photocurrent properties such as fast transient current decay, high transient to steady state current ratio, fast steady state current decay at light off and slow resensitisation of the transient current in the dark. Their two-photocycle model expanded on Nagel et al's, stipulating that recovery from the desensitised ground state required > 10 seconds whereas entry into the desensitised open state was ~30 times faster, resulting in accumulation in the desensitised photocycle with ISI's less than the recovery time. However, their model had the opposite problem to Nagel's; it predicted transient currents that were larger than experimentally observed. Their predicted transient current widths also varied with light intensity which was not seen experimentally. To account for this, they introduced a fast back-reaction between the desensitised and normal open states, rapidly establishing an equilibrium during the transient current. Unfortunately, this updated model did not capture the experimentally observed biexponential relationship between ISI and recovery from the desensitized ground state, where the initial recovery probability increased slower than would be predicted by a single exponential. The slow component was explained by including a third dark state in equilibrium with the desensitized ground.

Later, the extent of desensitization in channelrhodopsin-2 was also found to be dependent on ISI, taking ~30 seconds to fully recover[404]. The authors also found that even though transient current amplitude decreased with successive stimulations, because the steady state current was already depolarising enough to get the membrane close to threshold, 12 Hz 10 ms light pulses reliably produced action potentials *in vitro*. This was supported by Konastantin Nikolic and colleagues who mathematically modelled channelrhodopsin-2 and found that action potentials were reliably generated at stimulation frequencies up to ~15 Hz, with the limiting factor being deactivation rate from open to ground rather than recovery from the desensitised ground[405]. Nikolic concluded that the slow recovery of the transient current in channelrhodopsin-2 could only be mathematically explained by its dependence on light flux, supporting Nagel's initial proposal of a three-state single two-photon photocycle.

Back with Georg Nagel, Christian Bamann and colleagues saw Hegemann's proposal of a two-photocycle model and decided to revisit their own single cycle model of channelrhodopsin using spectroscopy[406]. They identified two spectrally distinct channelrhodopsin-2 intermediates. However, the time-course of absorbance changes indicated the presence of four intermediates, suggesting that three were spectrally overlapping given the low spectral resolution of the flash-photolysis measurements in use. By correlating the spectral time course with channel characteristics, they found that their four intermediates could be

explained using a non-linear one-photocycle model. Similar to before, in their updated model the ground state ***D*** is transferred to the open state ***O*** by photoactivation, only this time firstly undergoing rapid deprotonation and deprotonation through the intermediate ***OX***. ***O*** must first pass through a desensitised ground state ***L*** before slowly returning to ***D***. Interestingly, they found that pulses of green light to ***O*** shortcut the photocycle immediately to ***D***, bypassing the kinetically limiting ***L*** state. This supported Nagel's previous hypothesis regarding a two-photon photocycle. But they conceded that although their photocycle captured the events of channelrhodopsin-2 for a single turnover, it did not describe the existence of the steady-state current or desensitisation of the transient current which can't be observed in a single turnover.

Returning to Peter Hegemann's group, Oliver Ernst and colleagues performed their own spectroscopy analysis of VChR due to its relative ease of expression in standard systems like *E. coli*, yeast or COS cells[266]. Like with the earlier channelrhodopsins, they concluded that a linear three-state photocycle didn't capture all of the observed dynamics which were still consistent with Hegemann's original two-photocycle model. However, they conceded that Nikolic's non-linear two-photon photocycle described the observed photocurrents adequately and was elegant in that it involved less rate constants than a two-photocycle model. However, Nikolic's model did not explain the pH-induced shift of the action spectrum. In response, Ernst and colleagues proposed a hybrid of the two-

photocycle and two-photon photocycle models, featuring two ground states (one desensitised and one not) in an equilibrium that could be shifted depending on pH. Both ground states could be transitioned to a shared open state through photoactivation, but with less efficiency from the desensitized ground.

Following Christian Bamman's initial observation of the green light shortcut, André Berndt and colleagues further investigated how stimulation with different wavelengths of light could affect the channelrhodopsin-2 photocycle. They found that the fastest shortcut rate from open to ground was seen with 530 nm light pulses. Longer wavelengths of light, such as 546 nm, produced slower but more complete inactivation, consistent with reduced off-target absorption by the ground states[407].

However Katja Stehfest and Peter Hegemann argued that if green light converts the open state to normal ground, then it should prevent the desensitisation of transient current amplitude across successive stimulations, which it does not. They also posit that for observed single-channel conductance to be higher during the transient current than the steady state current, there must be two conducting states that are populated with different efficiencies which is only characterised in a two-photocycle model. To reconcile the four intermediates identified through spectroscopy and solve the aforementioned problems, they elaborated on the original two-photocycle model proposed by Hegemann. In this final iteration of the model, each photocycle's open state can be shortcut back to its relative ground

state with green light[408]. This captures why green light stimulation doesn't avoid desensitisation, as it exclusively transitions open to normal ground and desensitised open to desensitised ground, maintaining the ratio of dark adapted to light adapted molecules.

Since 2010, our understanding of the mechanisms behind opsin desensitisation has largely gone unchanged, settling on a two-photocycle model where ground and desensitised photocycles are concurrently populated with different efficiencies by a population of opsin molecules. Because transitioning from the desensitised photocycle to the ground photocycle is ~30 times slower than the reverse scenario[403], ISI's shorter than the time needed for all molecules to recover result in the accumulation of molecules in the desensitised photocycle. This can be circumvented by changing the identity of stimulated neurons between trials, an advantage more easily conferred by holographic stimulation compared with one-photon bulk activation[272]. In an all-optical setup, changing the identity of stimulated neurons by swapping the SLM phase mask would theoretically mitigate opsin desensitisation whilst keeping the imaging field of view constant. However, the contribution of persistent imaging laser excitation to opsin desensitisation remains critically understudied.

3.1.2 Cross-talk

With simultaneous reading and writing of neural activity comes the desire for spectrally distinct opsins and fluorescent reporters. The issue of cross-talk, where the excitation wavelength used for the reporter also excites the opsin, has long been an issue with all-optical studies. In an attempt to alleviate this issue, many developments have been made to try and strike a balance between spectral distinction and optimal sensor/actuator dynamics. As discussed previously, the field of opsin engineering had been developing red-shifted actuators to allow spectrally distinct dual-colour optogenetic control of neuronal populations. First to be discovered was VChR1 with maximal excitation at 535 nm[267], which was then fused with channelrhodopsin-1 to create C1V1 – a chimaera with stronger expression and faster kinetics[269].

Because C1V1 is red-shifted, it has a high absorption cross-section at the near infrared wavelengths utilized for two-photon excitation[409], with the greatest photocurrents occurring at 1040 nm[270]. Although 1040 nm is far from the peak of GCaMP's two-photon excitation spectra (both calcium-bound and unbound), C1V1 still has a significant cross-section at 920 nm – GCaMP's optimal two-photon excitation wavelength[257]. It has previously been shown that two-photon excitation at ~900 nm causes an increase in the photocurrents of red-shifted opsins[175,410], including C1V1[277]; specifically, firing rates of C1V1-expressing

neurons increase with increasing laser dwell time per cell[272] and power on the sample[274]. However, it remains unknown whether imaging parameters that are subthreshold for significantly increasing firing rates of C1V1-expressing neurons are able to affect response dynamics by transitioning opsin molecules into the desensitised photocycle.

3.1.3 Objective

We aimed to investigate whether the rate of opsin desensitisation in photostimulated neurons expressing the opsin C1V1 increases when imaging GCaMP6s at 920 nm with common experimental two-photon imaging parameters that are subthreshold for eliciting firing alone. If so, what are the optimal imaging and stimulation parameters for minimizing desensitization?

3.2 Methods

3.2.1 Animals, viruses and surgery

Animal experimentation was carried out with approval from the UK Home Office and the University of Oxford Animal Welfare and Ethical Review Board. Mice were an even mixture of sexes (eight males and eight females) with transgenic expression of GCaMP6s in excitatory neurons [mouse strains B6;

DBATg (tetO-GCaMP6s)2Niell/J crossed with CamK2a-tTa (AI94)] (Jackson Laboratory). At 7 to 14 weeks old, mice were either implanted with a cranial window over the right somatosensory cortex (n = 12) or right visual cortex (n = 4). At the beginning of surgery, mice were anaesthetized with isoflurane (3%, 1L/min) and kept under maintenance anaesthesia for the rest of the procedure (1.5% to 2.5%, 0.3 to 0.5L/min). Preoperatively, buprenorphine (Ceva) (1 unit/g), meloxicam (Boehringer Ingelheim) (1 unit/g), and dexamethasone (MSD) (five units) were administered subcutaneously. The scalp was shaved and cleaned with an antiseptic solution of 2% chlorhexidine gluconate and 70% isopropyl alcohol (Becton Dickinson) before head-fixing the animal via ear bars on a stereotactic rig (Kopf Instruments). Scalp and adipose tissue were excised from the right lateral side of the skull. A titanium headplate weighing ∼1g with a 7 mm circular aperture was secured to the skull (−1.4 mm anterior, +4 mm lateral / −1.9 mm anterior, +3.8 mm lateral relative to bregma for S1 implantation, +0.3 mm anterior, +3.7 mm lateral relative to lambda for V1 implantation) using superglue and dental cement (SuperBond, Prestige Dental, Bengaluru, Karnataka, India). For some mice, a second layer of 20% carbonated cement was used to further adhere the headplate to the skull. A 3 or 4 mm craniotomy was drilled with a dental drill (NSK) at the centre of the headplate aperture, and the dura was removed. A micropipette was front-loaded with pAAV-CamKIIa-C1V1(t/t)-mScarlet-KV2.1 (Addgene: 124650-AAV9; 2.70×10^{13} vg/mL) diluted 1:11 in sterile PBS. 500 nL of the virus was injected into the left of the craniotomy, 300 μm deep from the

cortical surface at a rate of 50 nL/min using a hydraulic micromanipulator (Narashige, Tokyo, Japan). A cranial window comprised of two optically bonded glass coverslips (one the size of the craniotomy and one 1 mm diameter larger) (Warner Instruments, Massachusetts, United States) was used to plug the craniotomy. The cranial window was secured in place using cyanoacrylate (3M) and dental cement. The animal was allowed to fully recover for at least 3 weeks before any further experiments were conducted. Mice were kept on a 12:12-h light/dark cycle with access to food and water *ad libitum*.

3.2.2 Two-photon imaging and stimulation

The microscope used in this study was a Bruker Ultima 2Pplus equipped with a tuneable wavelength laser for imaging (Chameleon Ultra II, Coherent, Saxonburg, Pennsylvania 16056, USA) and a 1035 nm fixed wavelength laser for photostimulation (Monaco 1035, Coherent). The 1035 nm laser was directed through a SLM (Neuralight 3D Ultra, Bruker, Billerica, Massachusetts, United States). Photostimulation parameters and phase masks were defined and created using custom MATLAB software (NAPARM)[395] before being loaded into external software. Imaging and photostimulation were controlled using Prairie View software (Bruker). Phase masks were loaded onto the SLM using BOSS (Boulder Nonlinear Systems). The animal was head-fixed on a tip-tilt stage (TTR001, Thorlabs) under a 16x objective (N16XLWD-PF, Nikon, Tokyo, Japan).

Stimulation and imaging triggers were coordinated using general purpose input/output software (PACKIO)[411] running on a USB data acquisition card (National Instruments).

A reference two-photon image for opsin expression (conjugated to mScarlet) was captured at 765 nm using a 570 to 620 nm emission filter, averaged over 128 frames at 1024 x 1024 resolution. NAPARM was used to group opsin-expressing neurons into groups of 10. Within each group, 6 mW 1035 nm laser light (2 MHz repetition rate) was simultaneously spiralled over each cell body. The spiral spanned a diameter of 10 μm for a duration of 10 ms, repeating 25 times for a total photostimulation duration of 250 ms. The phase mask changed every 1 s, and each group received 10 sets of spiral stimulations with a 15 s inter-trial interval. Photoresponsive neurons were identified manually from GCaMP6s $\Delta F/F$ stimulus-triggered average images based on a baseline of 1 s pre-stimulus.

Photoresponsive cells were used for experimental testing. During the photoresponse desensitisation experiments, each stimulation event consisted of spiralling 6 mW laser light over each cell body for 10 ms. The phase mask changed every 15 ms to target the other group of cells, and each group received 10 blocks of stimulation in total (over 300 ms).

Responses to the photostimulation were imaged across a 1.4 x 1.4 mm field of view (1024 x 1024 resolution) at 15 Hz using resonant scanning at 920 nm (80

MHz pulse repetition rate, 50 mW average power after the objective), with a 500 to 550 nm emission filter.

3.2.3 GCaMP imaging analysis

Calcium imaging movies were processed using Suite2p[412]; cell bodies were automatically detected and carried forward for analysis. $\Delta F/F$ was computed for each cell body identified by Suite2p using the equation:

$$\Delta F/F = (F - \underline{F})/\underline{F}$$

where F = raw fluorescence, $\underline{F}$ = the mean of F in the 1 s pre-stimulation.

Cells were labelled as "responsive" in the photoresponse decay experiments if their mean $\Delta F/F$ in the 1 s post-stimulation was >= 0.1. The "response" of a population of responsive cells to photostimulation is defined as the area under the curve of the mean $\Delta F/F$ across cells in the 1 s poststimulation.

If multiple cells were wholly or partially in the radius of a beamlet, only the cell with the centroid closest to the centre of the beamlet was labelled as a "target". When analysing "non-target" cells, any non-target cells wholly or partially within the radius of a beamlet were excluded from analysis.

For a cell to be labelled as excited or inhibited, it must be significantly correlated both across all trials (2 s around stimulation) concatenated and across trial one

separately ($p<0.05$ Pearson's correlation). This was to ensure that correlated cells are active on trial 1 so that comparisons can be reliably drawn to the response on trial 10. The correlation coefficient of the cell was taken from the correlation across all trials.

3.2.4 Analysis, statistics and reproducibility

All analysis was performed using MATLAB. No statistical method was used to predetermine sample size. Preliminary data investigating the effect of imaging all 10 trials or just the first and last were removed from analysis due to a small sample size (n-3 mice). The scientists were not blinded to imaging duration or inter-stimulus interval during experiments or analysis.

3.3 Results

3.3.1 Determining the impact of 920 nm two-photon imaging on 1035 nm photostimulation responses in targeted cells under common all-optical experimental conditions

To assess optogenetic stimulus responses, we targeted 16 putative cells in layer 2/3 somatosensory or visual cortex for two-photon stimulation. The cells were split equally between two target groups using separate SLM phase masks (figure 5a). Ten trials were conducted to gather enough data to observe the effects of desensitization on neuronal responses. Cells with a centroid closest to the centre of a 1035 nm beamlet were labelled as "target" cells if they were responsive within the radius of a beamlet. Initially, the stimulations were separated by an interval of 15 s as a common ISI used in all-optical experiments. The responses were collected by imaging using a 920 nm laser (1.4x1.4 mm field of view, 50 mW power on sample, 15 Hz, 1024x1024 pixels, ~3 μs total dwell time per cell; subthreshold dwell time per cell and power on sample for eliciting spikes[272,274]) and defined as the area under the curve of the mean $\Delta F/F$ across all target cells in the 1 s post-stimulation. To investigate the impact of imaging with a 920 nm laser on neuronal responses to targeted photostimulation, we varied the duration of imaging from 2 to 12 s centred on the 1035 nm photostimulation event (figure

5b). On average, the response of target cells decayed across trials, initially dropping by 20% to 30% in the first few trials and plateauing around trial five with responses remaining consistently at ~70% of the initial response (figure 5c). Response desensitisation over trials followed a similar trajectory across all imaging durations until trial five, when responses became comparatively less reduced for the 2 s imaging dose (figure 5d). This upward deflection for trial 10 of the 2 s imaging dose was unexpected but strangely consistent across repeats. It was not seen in later experiments (figure 6d) suggesting that it may have been artefactual in nature, although the root cause is unknown. We confirmed that there were no significant differences in the first trial response between conditions (figure 5e). Responses significantly decreased from trials one to ten, regardless of the imaging condition (figure 5f), with no significant differences in the magnitude of response reduction between imaging conditions (figure 5g). These results suggest that imaging at 920 nm has a negligible impact on 1035 nm photostimulation response desensitisation in targeted cells using a typical ISI.

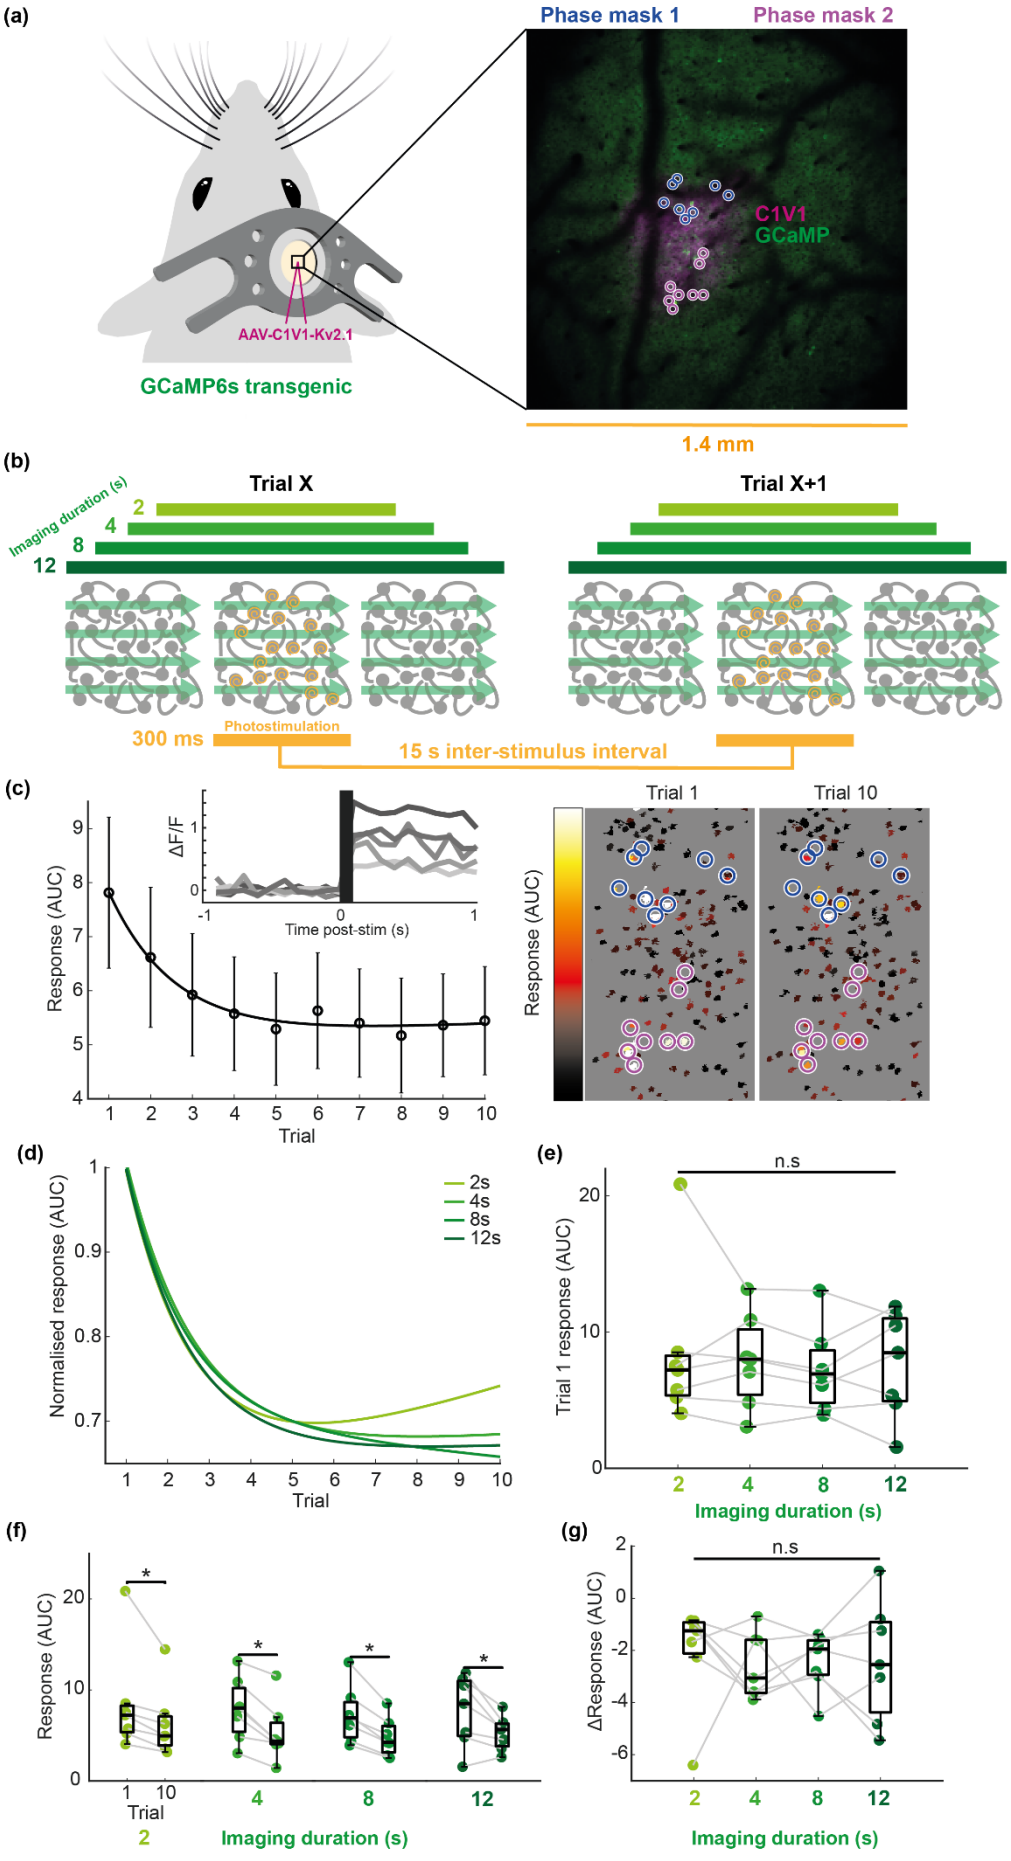

(a)
Phase mask 1
Phase mask 2
C1V1
GCaMP
AAV-C1V1-Kv2.1
GCaMP6s transgenic
1.4 mm
(b)
Trial X
Trial X+1
Imaging duration (s)
2
4
8
12
300 ms
Photostimulation
15 s inter-stimulus interval
(c)
Response (AUC)
ΔF/F
Time post-stim (s)
Trial
Trial 1
Trial 10
(d)
Normalised response (AUC)
2s
4s
8s
12s
Trial
(e)
n.s
Trial 1 response (AUC)
Imaging duration (s)
(f)
Response (AUC)
Trial
Imaging duration (s)
(g)
n.s
ΔResponse (AUC)
Imaging duration (s)

Figure 5. Determining the impact of 920 nm two-photon imaging on 1035 nm photostimulation responses in targeted cells under common all-optical experimental conditions (a) The opsin AAV-C1V1-Kv2.1 was injected into layer 2/3 of either the right somatosensory cortex (n=3) or right visual cortex (n=4) of mice transgenic for GCaMP6s in excitatory neurons. Calcium activity was imaged with a 920 nm laser over a 1.4 mm x 1.4 mm field of view around the opsin-expressing area (somatosensory cortex shown). Sixteen photostimulation targets were equally divided between two SLM phase masks. **(b)** Schematic showing the different raster-scanning imaging durations centred around 1035 nm photostimulation with a 15 s ISI. Each group of cells targeted by one SLM phase mask received ten blocks of 10 ms spiralling stimulation at 6 mW with a 20 ms interval between spirals**. (c)** Left: grand average response (over imaging conditions and then animals, n=7 mice) ± SEM fit with a two-term exponential model. Inset: example targeted cell responses to stimulation. Right: cells from an example session color-coded by their individual responses to stimulation with the same holograms as in **a**. **(d)** Normalized response as a function of trial number, fit with a two-term exponential model. **(e)** Trial 1 response as a function of imaging dose (n.s. = $p > 0.05$, Kruskal-Wallis test). **(f)** Responses on trials 1 and 10 as a function of imaging dose (* = $p < 0.05$, Wilcoxon signed-rank test, n=7 mice). **(g)** ΔResponse (trial 10 − trial 1 response) as a function of imaging dose (n.s = $p > 0.05$, Kruskal–Wallis test, n=7 mice).

3.3.2 Two-photon imaging at 920 nm increases 1035 nm photostimulation response desensitization in targeted cells at shorter inter-stimulus intervals

To further probe the effects of imaging at 920 nm on photostimulation response desensitisation in targeted cells under different experimental paradigms, we tested the two most extreme imaging doses: constant imaging and 2 s imaging, at ISIs from 10 to 40 s (figure 6a). On average across all ISIs, both imaging conditions showed a significant decrease in responses from trials one to ten (figure 6b), with the constant imaging condition causing significantly greater response desensitisation (figure 6c). The response decay curves for the two imaging

conditions diverged as early as trial two (figure 6d). On average across both imaging conditions, ISIs of 10, 20, and 30 s showed significantly decreased neuronal responses to targeted stimulation from trials one to ten, but there was no significant response desensitisation for the 40 s ISI (figure 6e). Constant imaging exacerbated response desensitisation to a greater extent at shorter ISIs (figure 6f). These results demonstrate that desensitisation of photostimulation responses in targeted cells increases as a function of 920 nm imaging dose, but this relationship is only seen with ISIs less than 40 s.

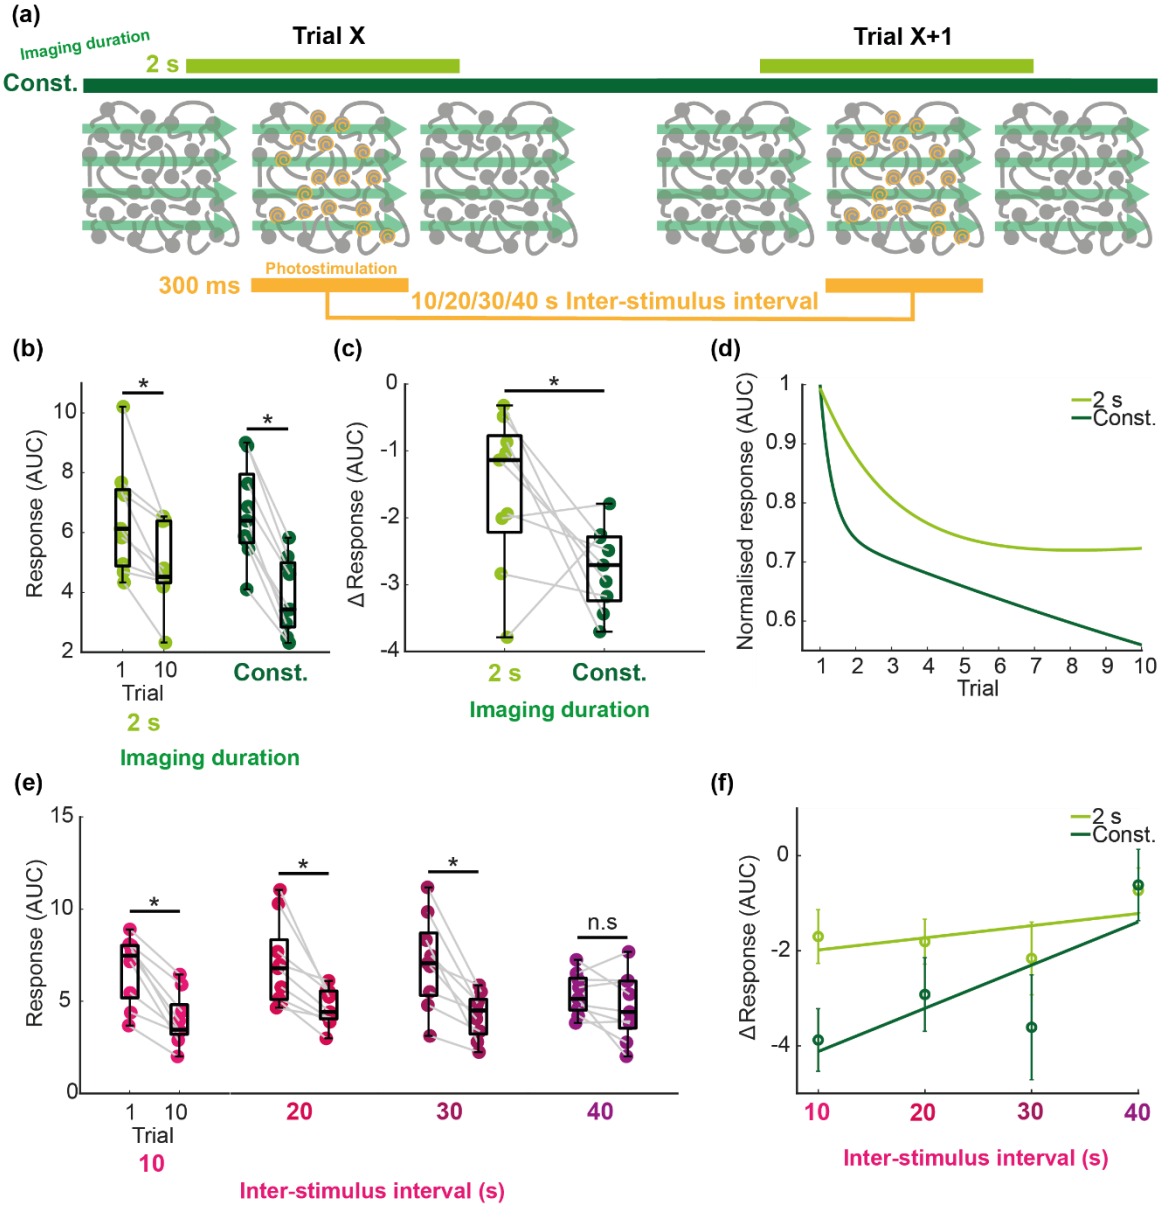


Figure 6. Two-photon imaging at 920 nm increases 1035 nm photostimulation response desensitisation in targeted cells at shorter ISIs. (a) Schematic showing the two 920 nm imaging conditions (either imaging constantly or in a 2 s window centred around photostimulation) and four possible ISI's (10/20/30/40 s). Each group of cells targeted by one SLM phase mask received ten blocks of 10 ms spiraling 1035 nm stimulation at 6 mW with a 20 ms interval between spirals. **(b)** Responses on trials 1 and 10 as a function of imaging dose, averaged across ISI's (* = $p < 0.05$, Wilcoxon signed-rank test, n=9 mice). **(c)** ΔResponses (trial 10 − trial 1 response) as a function of imaging dose, averaged across ISIs (* = $p < 0.05$, Wilcoxon signed-rank test, n=9 mice). **(d)** Normalised response (mean trial response averaged across ISI's then animals/mean trial 1 response averaged across ISI's then animals) as a function of trial number, fit with a two-term exponential model (n = 9 mice). **(e)** Responses on trials 1 and 10 as a function of ISI, averaged across imaging conditions (* = $p < 0.05$, Wilcoxon signed-rank test, n=9 mice). **(f)** Average ΔResponse ± SEM fit with a first-degree polynomial (n=9 mice).

3.3.3 The distribution of non-targeted cells excited or inhibited by 1035 nm stimulation is unaffected by 920 nm imaging dose

Non-target cells were classified as either excited, inhibited, or non-responders (figure 7a). For a cell to be labelled as excited or inhibited, its activity must be significantly positively or negatively correlated with at least one of the target cells (see methods). On average, ~6% of cells in a FOV were excitatory and ~2% were inhibitory. There was no significant correlation between the number of photoresponsive cells stimulated and the proportion of excitatory or inhibitory followers (supplementary figure 8). The probability distribution of excited cells was significantly higher than inhibited cells within 50 μm of the nearest target. The reverse was seen between 50 and 100 μm. This roughly matches what has been previously observed with two-photon[413] and one-photon stimulation[230]. Both excited and inhibited cells exhibited a bimodal distribution which reflects the interaction between proximity to stimulation, area effect and vignetting at the edges of the field of view (figure 7b). Probability distributions were unaffected by imaging at 920 nm (figure 7c). The degree of correlation with the target cells decreased as a function of distance, with excited cells consistently exhibiting greater probability densities at higher correlation coefficients (figure 7d). However, this may be due to GCaMP being optimised to capture positive activity dynamics relative to negative. Naturally, the distribution of opsin-expressing cells

also decreased with distance from the stimulation sites (figure 7b). The nearby peak in excited cells was not purely driven by the distribution of opsin expression, as excited cells consistently showed a different spatial profile to opsin expressing cells within the field of view (figure 7e). Furthermore, the proportion of excited and inhibited cells were not significantly correlated with the proportion of opsin-expressing cells (figure 7h), although it is close to significant with the potential of a larger sample size consolidating this finding. Larger imaging doses flattened and negatively-shifted the ΔResponse distributions of both excited and inhibited cells (figure 7f). Non-target cells also differed in their most correlated targets between imaging conditions (figure 7g). Imaging dose had no effect on the number of excited cells across the field of view (figure 7i) or specifically within the opsin-expressing area (figure 7k), but did increase the number of inhibited cells (figure 7j). Although this is driven by non-opsin expressing cells (figure 7i), suggesting that larger imaging doses preferentially increase the proportion of inhibition in distal proximity to the targets. This is perhaps due to the stimulus response being less due to opsin desensitization and thus a rapid proximal activity quench is less vital to maintain EI balance. And a distal soft clamp on activity would naturally recruit more cells due to area effect.

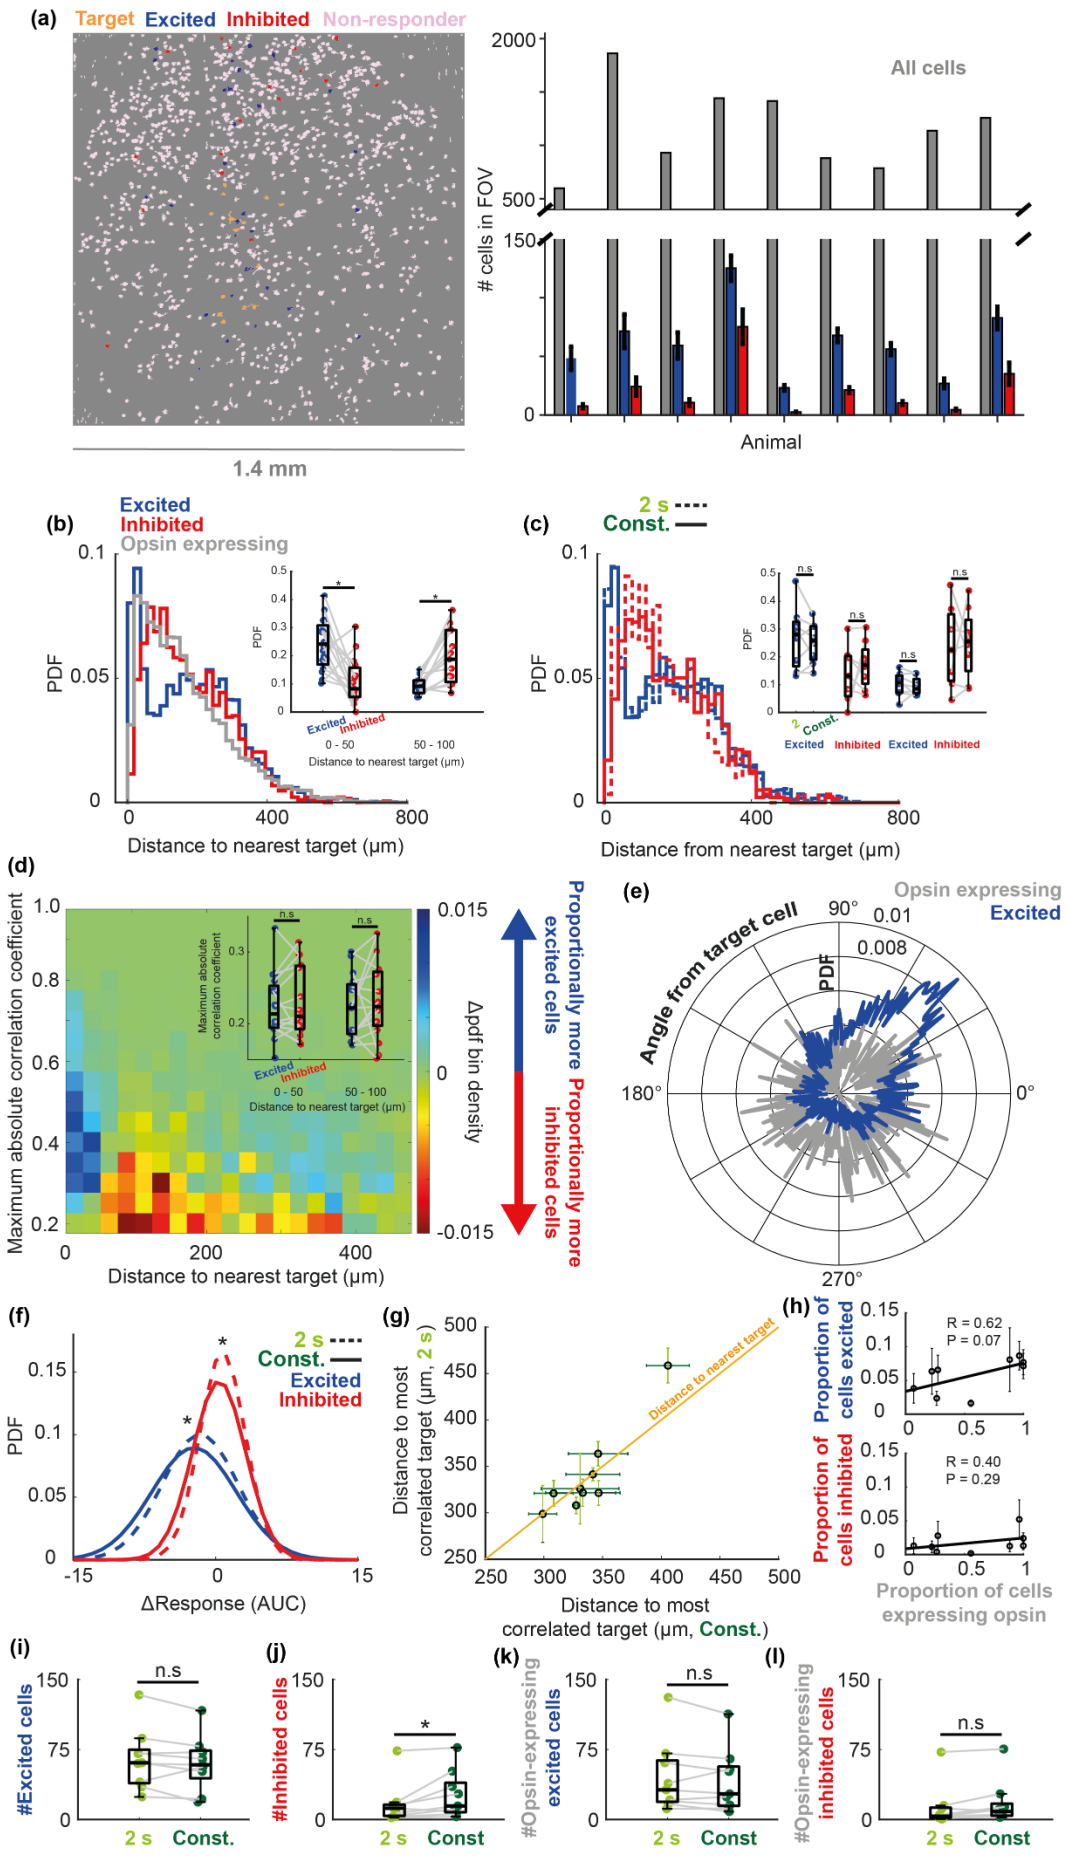

(a)
Target
Excited
Inhibited
Non-responder
1.4 mm
All cells
cells in FOV
Animal
(b)
Excited
Inhibited
Opsin expressing
PDF
Distance to nearest target (μm)
(c)
2 s
Const.
Distance from nearest target (μm)
(d)
Maximum absolute correlation coefficient
Distance to nearest target (μm)
Δpdf bin density
Proportionally more excited cells
Proportionally more inhibited cells
(e)
Opsin expressing
Excited
Angle from target cell
PDF
(f)
ΔResponse (AUC)
(g)
Distance to most correlated target (μm, 2 s)
Distance to most correlated target (μm, Const.)
Distance to nearest target
(h)
Proportion of cells excited
Proportion of cells inhibited
R = 0.62
P = 0.07
R = 0.40
P = 0.29
Proportion of cells expressing opsin
(i)
#Excited cells
(j)
#Inhibited cells
(k)
#Opsin-expressing excited cells
(l)
#Opsin-expressing inhibited cells
n.s
2 s
Const.

Figure 7. The distribution of non-targeted cells excited or inhibited by 1035 nm stimulation is unaffected by 920 nm imaging dose. **(a)** Left: Distribution of target, excited, inhibited, and non-responder cells across the field of view from the same example session as in figures 8a and 8c. Right: Quantification of the number of cells in each animal's FOV and the mean ±SEM excitatory and inhibitory cells across conditions. **(b)** Probability distribution of excited, inhibited and opsin-expressing non-target cells as a function of distance from the nearest target cell. Inset: Each point is data from an individual animal, averaged across all cells across all conditions (* = $p < 0.05$, Wilcoxon signed-rank test, n = 16 mice). **(c)** Same as in **a** but separated by imaging condition. Inset: (* = $p < 0.05$, Wilcoxon signed-rank test, n = 9 mice). **(d)** Maximum correlation strength of an individual non-target cell with any target as a function of distance from the nearest target cell, color-coded by excited bin count - inhibited bin count. Inset: Each point is data from an individual animal, averaged across all cells for a given type across all conditions (* = $p < 0.05$, Wilcoxon signed-rank test, n = 16 mice). **(e)** Polar plot of an example mouse with opsin expression confined to a part of the field of view. The polar plot shows the probability distribution for excited and opsin-expressing cells at different angles from a target cell, averaged across targets across all conditions. The probability distributions for excited and opsin-expressing cells are non-overlapping. **(f)** Probability distribution of excited and inhibited cells under different imaging conditions (* = $p < 0.05$, two-sample Kolmogorov-Smirnov test, pooled across n = 9 mice). **(g)** Mean ±SEM distance of non-target cells from their most correlated target cell under different imaging conditions. Orange line is a representation of what a 1:1 relationship would look like, for example if distance to nearest target was plotted on each axis instead, as this would not vary between imaging conditions. **(h)** Correlation between the proportion of cells in the field of view that are excited/inhibited and the proportion of cells expressing opsin (R = Pearson correlation coefficient). **(i-l)** The number of different classifications of non-target cells under different imaging conditions (* = $p < 0.05$, Wilcoxon signed-rank test, n = 9 mice).

3.3.4 Two-photon imaging at 920 nm exacerbates photostimulation response desensitisation of excited non-target cells

It is possible that as the targeted neurons' response to photostimulation is reduced, their ability to influence wider-scale network properties could change—a major confound for all-optical experiments. To investigate whether imaging at 920 nm had any effect on neuronal responses beyond the cells targeted for 1035

nm photostimulation, we assessed the activity of non-target cells across the field of view. Like target cells, the responses of excited cells significantly decreased from trials one to ten, with the opposite being true for inhibited cells (figure 9a). The degree of response desensitisation in excited cells did not differ between those that expressed opsin and those that did not (figure 9b,c), suggesting that reductions in response amplitude over trials cannot be solely attributed to a ceiling effect produced by sub-threshold depolarisation from the imaging laser. Excited cells did not respond as strongly to photostimulation as target cells and decayed in their responses much faster, plateauing by trial two only slightly above the non-responders (figure 9d). Larger imaging doses caused greater response desensitisation across all ISIs in excited cells (figure 9e,f). Taken together, these results suggest that the amplitude of optogenetically evoked responses in targeted cells non-linearly translates downstream, comparatively exacerbating the rate of response reduction in the local network. This may be due to non-linearities in synaptic transmission, such as short-term plasticity, or dendritic non-linearities. Thus, nonsignificant changes in the response of target cells, such as with a 40 s ISI, can cause dramatic changes in downstream responses. Therefore, during all-optical experiments, a combination of a short imaging dose and large ISI is needed to ameliorate response desensitization across the entire field of view (figure 9g), as a large ISI alone is not sufficient for non-target cells (figure 9h).

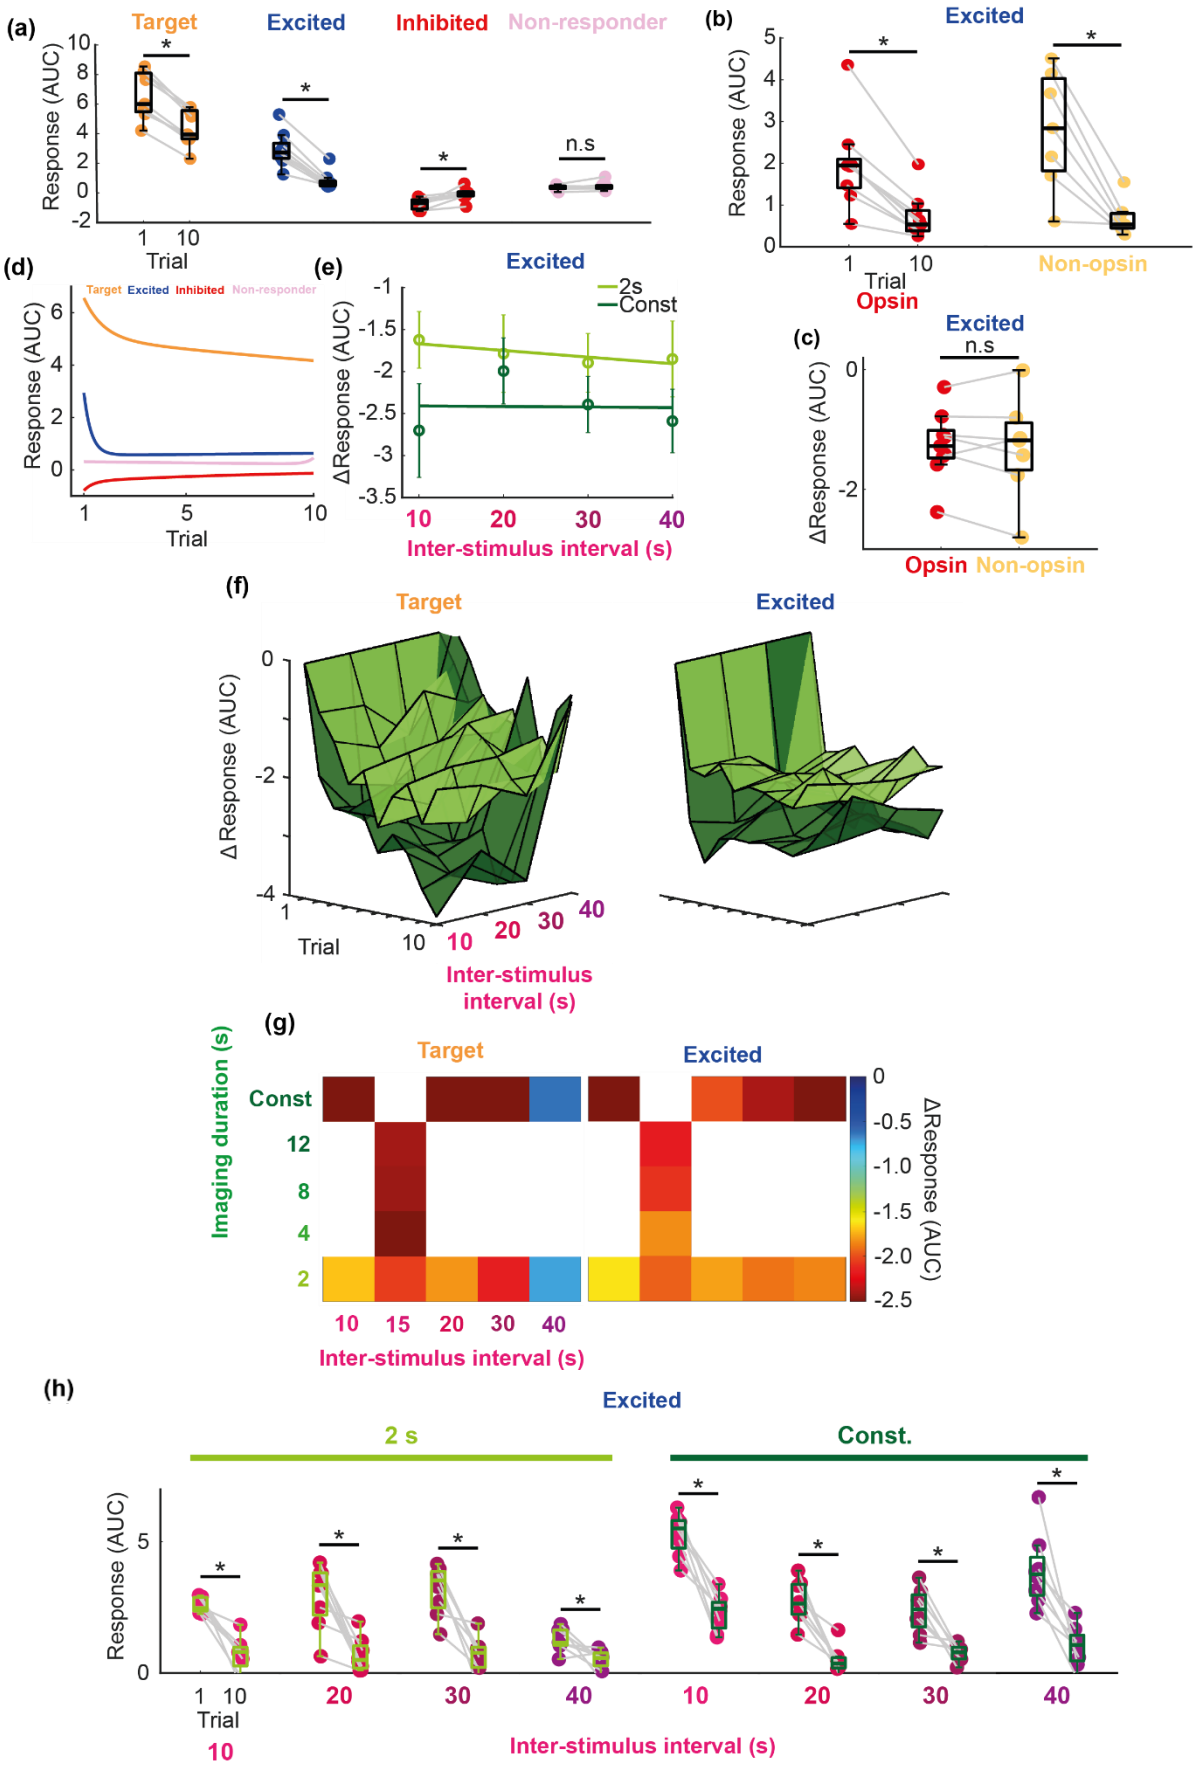

(a)
Target
Excited
Inhibited
Non-responder
Response (AUC)
Trial
n.s
(b)
Excited
Opsin
Non-opsin
(c)
Excited
n.s
ΔResponse (AUC)
Opsin Non-opsin
(d)
Target Excited Inhibited Non-responder
(e)
Excited
2s
Const
Inter-stimulus interval (s)
(f)
Target
Excited
Δ Response (AUC)
Inter-stimulus interval (s)
(g)
Target
Excited
Imaging duration (s)
Const
ΔResponse (AUC)
Inter-stimulus interval (s)
(h)
Excited
2 s
Const.
Response (AUC)
Trial
Inter-stimulus interval (s)

Figure 9. Two-photon imaging at 920 nm exacerbates photostimulation response desensitisation of excited non-target cells. (a) Responses on trials 1 and 10 for the different cell classifications (* = $p < 0.05$, Wilcoxon signed-rank test, n=9 mice). **(b)** Responses on trials 1 and 10 in opsin-expressing and non-opsin-expressing excited cells, averaged across imaging conditions and inter-stimulus intervals (* = $p < 0.05$, Wilcoxon signed-rank test, n = 9 mice for opsin-expressing excited cells; n = 7 mice for non-opsin-expressing excited cells). **(c)** ΔResponses (trial 10 - trial 1 response) in opsin-expressing and non-opsin-expressing excited cells, averaged across imaging conditions and inter-stimulus intervals (* = $p < 0.05$, Wilcoxon signed-rank test). **(d)** Grand average response (over conditions then animals, n=9 mice) fit with a two-term exponential model. **(e)** Average ΔResponse (trial 10 − trial 1 response) ± SEM (n=9 mice) of excited cells, fit with a first-degree polynomial. **(f)** Target cell (left) and excited cell (right) ΔResponses (trial response - trial 1 response) across trials and inter-stimulus intervals under different imaging conditions. **(g)** ΔResponse in target and excited cells as a function of imaging dose and ISI. The same cells were stimulated across imaging conditions for the 10, 20, 30, and 40 s ISIs (n=9 mice). For a different population of cells, the same cells were stimulated across imaging conditions for the 15 s ISI (n=7 mice). **(h)** Excited cell responses on trials one and ten for all ISI's under different imaging conditions (* = $p < 0.05$, Wilcoxon signed-rank test, n=9 mice with somatosensory opsin expression).

3.4 Discussion

The main advantage of all-optical interrogation over other, less precise optogenetic stimulation methods is dissecting causal relationships through the artificial recreation of endogenous-like cortical activity. Causality in systems neuroscience has been historically hard to define, with suggestions to instead broadly agree on methods to disambiguate causality, such as randomisation, rather than dwell on definitions[56]. Whilst randomisation is a conceptually sound gold standard to establish causality[55], it has been difficult to achieve in practice[53]. In this regard, all-optical interrogation is uniquely positioned to determine whether specific neural activity or behaviour is causally produced by other neural activity or is just incidentally linked via a common input. However, recreation of endogenous-like activity requires careful calibration of photostimulation

parameters such as laser power and scanning pattern[272]. As shown here, two-photon imaging at 920 nm can bias neural activity in C1V1-expressing neurons targeted for photostimulation, throwing off the calibration. We also show that this bias can change how injected activity propagates, so it should be taken into account when investigating causal relationships. Our results indicate that all-optical interrogation's unique advantage is lost two-fold if opsin desensitisation is not accounted for: first, if the ISI is not long enough, photostimulation-evoked activity in the targeted cells will decrease across trials, adding the influence of a confounding variable to any suggestions of causal relationships. Second, regardless of ISI, photostimulation-evoked activity in excited cells in the surrounding network will decrease across trials, an effect that is exacerbated by larger imaging doses. Given the diverse response trajectories across target, excited, inhibited, and non-responder cells, these network-level changes in photostimulation response amplitude over trials are unlikely to be solely explained by adverse photoexposure effects such as GCaMP bleaching or phototoxicity. Additionally, adaptation mechanisms that generate a decreasing response of cortical sensory neurons to repetitive external stimulation are primarily driven by changes in subcortical input instead of intrinsic properties and can be overcome with stimulation[414]. The decreases in response caused by opsin desensitisation are more severe than adaptation[415], suggesting that reduced responses to successive photostimulation of targeted cells are likely dominated

by opsin desensitisation. However, without directly measuring photocurrents, alternative contributors to across-trial reductions in photostimulation response cannot be fully excluded. Future studies should use whole-cell voltage-clamp recordings to confirm whether 920 nm two-photon imaging accelerates photocurrent decline in a manner consistent with increased occupancy of the desensitised photocycle.

Since the first algal channelrhodopsin was discovered, there has been a concerted effort to discover and engineer opsins with optimal kinetics for optogenetic modulation of neural circuits[269,398]. Such desirable traits include slower decay constants and larger photocurrents, which provide the high sensitivity preferred in two-photon stimulation experiments. Unfortunately, these characteristics also make opsins more vulnerable to cross-talk, which can complicate the interpretation of findings in all-optical studies. Typical approaches to reduce the impact of cross-talk between opsins and fluorescent reporters of neuronal activity in all-optical experiments include reducing the power of the imaging laser (which also reduces signal-to-noise) and scanning over large fields of view to reduce dwell time[399]. However, here we show that, over a large (1.4 mm x 1.4 mm) field of view, even low levels of C1V1 activation caused by imaging at 920 nm reduce photostimulation responses over trials, likely due to opsin desensitisation. This is likely also true for other red-shifted opsins, as even those with peak excitation at higher wavelengths, such as ChrimsonR and ChRmine, still have spectral overlap

with wavelengths used for GCaMP imaging[175,416–418]. Although red-shifted opsins with faster recovery speeds from the desensitised photocycle than C1V1, such as ChRmine[175], are likely to be less impacted by 920 nm imaging-induced desensitisation. With the aim of minimizing cross-talk, some studies have opted to image GCaMP at 1020 nm and stimulate channelrhodopsin-2 at 920 nm[419]. However, although channelrhodopsin-2 is not efficiently excited at 1020 nm, small photocurrents are still produced[270], which will likely cause similar problems to those seen in the current study. Several all-optical studies use a combination of C1V1, GCaMP, constant two-photon imaging and ISI's of $\leqslant$10 s[283,286,395]. Under these conditions, if the identity of the targeted cells remained constant over successive stimulations, there would likely be a significant decrease in the photostimulation-evoked responses of both the targets and downstream excited cells across trials. Minimising response reduction over successive stimulations across the entire field of view requires combining a long ISI ($\geqslant$40 s) with intermittent short imaging doses.

Here we focus on how reductions in photostimulation-evoked responses can be mitigated in traditional all-optical experiments using red-shifted opsins and green fluorescent reporters of neuronal activity. However, changing the opsin/sensor combination can also be used to mitigate opsin desensitisation via cross-talk. For example, one might use a spectrally distinct combination of blue-shifted opsins with red calcium sensors[275,420]. The downside here is that red calcium sensors are

generally less sensitive with slower kinetics[258]. Another option is using red-shifted step-function opsins (SFO) in combination with green calcium sensors. SFOs can be opened or closed with different wavelengths of light. Therefore, to prevent subthreshold activation of the opsin during imaging outside of the stimulation, one can simultaneously scan with the wavelength that negatively regulates opsin activity. For example, a previous study identified a channelrhodopsin-2 variant (sdChR) that is activated by blue light, but only in the absence of orange light[416]. The downside here is that incorporation of an SFO into a two-photon all-optical experiment would either require the integration of a third wavelength laser into the light path, the use of a fast-switching programmable tuneable wavelength laser for positive/negative stimulation, or an SFO that is negatively modulated by a one-photon source with a wavelength that is outside of the emission filter bounds of the photomultiplier tube.

In conclusion, we demonstrate that a combination of targeted optogenetic ISI's ≥ 40 s and a 2 s imaging dose around each targeted stimulation significantly mitigates response desensitisation of targeted C1V1-expressing cells and ameliorates response desensitisation of excited cells across the field of view during all-optical interrogation. This will reduce an external source of trial-by-trial variability in future all-optical experiments. The extent to which imaging will impact opsin desensitisation in other all-optical experiments will depend on laser wavelengths, power on sample, dwell time per cell (combination of duty cycle,

zoom, resolution, and soma size), spectral overlap of the opsin and reporter, opsin identity, and opsin expression per cell. Therefore, the experiments described herein will have to be repeated for unique all-optical paradigms to find the best trade-off between data quantity and response desensitisation. In the absence of such data, one's best option will be to minimize the imaging dose and increase the ISI as much as is practical.

Although stimulation evoked activity is also dependant on ongoing network activity[286,413], these optimisations will partly help generate reliable cortical stimulus responses across trials. To further increase the reliability of stimulus-evoked activity, future studies should investigate the use of closed loop stimulation based on ongoing neuronal activity[421], behavioural readouts (e.g. pupil dilation) as a proxy for ongoing neuronal activity[296], or enforcing a baseline network state prior to stimulation with task design[285] or dual-colour optogenetics[265].

4

Work in progress: function- and projection-specific stimulation in a somatosensory CPD task

4.1 Introduction (p. 106)
4.1.1 Perceptual importance of functionally characterised sensory neurons (p. 106)
4.1.2 Objective (p. 107)
4.2 Methods (p. 108)
4.2.1 Animals, viruses and surgery (p. 108)
4.2.2 Two-photon imaging and stimulation (p. 110)
4.2.3 Whisker and reward mapping (p. 112)
4.2.4 GCaMP imaging analysis (p. 112)
4.2.5 Analysis, statistics and reproducibility (p. 113)
4.2.6 Behavioural training (p. 114)
4.2.7 Electrophysiology (p. 115)
4.2.8 Histology (p. 116)
4.3 Results (p. 117)
4.3.1 Cortical and subcortical-projecting neurons in wS1 are largely non-overlapping with somatosensory-responsive neurons (p. 117)
4.3.2 Training on a one-photon optogenetic CPD task reduces surround inhibition evoked by two-photon stimulation (p. 120)
4.3.3 Training on a one-photon optogenetic CPD task reduces surround inhibition evoked by two-photon stimulation (p. 123)
4.3.4 Near-infrared wavelengths used for two-photon imaging and stimulation produce large photoelectric artefacts in Neuropixels recordings (p. 126)
4.4 Discussion (p. 126)

4.1 Introduction

4.1.1 Perceptual importance of functionally characterised sensory neurons

All-optical experiments have demonstrated that the behaviour of mice trained to respond to sensory stimuli can be biased by targeted photostimulation of neurons based on their functional identity. For example, targeted stimulation of co-tuned ensembles in visual cortex boosted detection of low contrast visual gratings[413], targeted stimulation of perception-encoding neurons in S1 (but not tactile-responsive neurons) boosted detection of a whisker stimulus[422], and targeted stimulation of non-coding silent S1 neurons boosted performance in a whisker discrimination task[423]. Similar effects have been reported for projection-specific neurons: James Marshel and colleagues showed that the psychometric relationship between number of cells stimulated and perception (see[286,287]) was steeper in layer 5 of visual cortex compared to layer 2/3[175], theorising a weighted contribution of striatal-projecting cells to perception. Likewise, selective stimulation of striatal-projecting cells in primary auditory cortex biased decisions based on their tuning in a frequency discrimination task[424]. Furthermore, inhibition of auditory corticostriatal neurons impaired performance on a frequency discrimination task[425]. Also, optogenetic activation of auditory corticothalamic neurons modulated sound detection and discrimination[426].

Although there is less causal evidence for the role of cortico-cortico projections, passive recording studies have shown that whisker-responsive S1 neurons propagate perceptual information to S2[130]. This suggests that functionally relevant and projection-specific neurons may be inherently higher weighted in propagating perceptual signals. However, this has not been shown in a CPD task where an animal is trained to report widefield activation of thousands of neurons, including but not limited to those with functional or projection-specific identities. Perhaps, despite all of these neurons being stimulated, certain subgroups have a higher perceptual importance and carry the majority of the behaviourally relevant information.

4.1.2 Objective

We aimed to understand how the structure of early cortical stimulus activity governs trial-by-trial perception and subcortical neuronal signals in a somatosensory CPD task. To achieve this, we used an all-optical approach in expert mice previously trained on a one-photon detection task with indiscriminate S1 opsin expression. Mice were either injected with a retrograde tracer in S2, VPM or DLS so that the projection targets of stimulated S1 neurons could be identified. Additionally, the responsiveness of S1 neurons to whisker stimulation and free reward was mapped prior to task onset, allowing the functional identity of S1 neurons to also be identified. In total, this allowed exploration of the effect

of targeted stimulation, based on functional or projection-specific characteristics, on perception of a previously engrained optogenetic stimulus.

Another previously discussed aim was to understand how the reliability of cortical stimulation affected the reliability of the temporal codes underlying perception. To investigate this, we split stimulation in the all-optical task into two trial types: odd trials always received the same stimulation pattern whereas even trials received a random stimulation pattern featuring the same number of targets. During the all-optical task, we simultaneously recorded high density electrophysiological activity in ipsilateral VPM, PO and CP using Neuropixels probes inserted through the contralateral hemisphere.

4.2 Methods

4.2.1 Animals, viruses and surgery

Animal experimentation was carried out with approval from the UK Home Office and the University of Oxford Animal Welfare and Ethical Review Board. A mixture of male (n = 5) and female (n = 2) mice with transgenic expression of GCaMP6s in excitatory neurons [mouse strains B6; DBATg (tetO-GCaMP6s)2Niell/J crossed with CamK2a-tTa (AI94)] (Jackson Laboratory) were implanted with a cranial window over the right somatosensory cortex at 7 to 14 weeks old. At the beginning of surgery, mice were anaesthetized with isoflurane

(3%, 1L/min) and kept under maintenance anaesthesia for the rest of the procedure (1.5% to 2.5%, 0.3 to 0.5L/min). Preoperatively, buprenorphine (Ceva) (1 unit/g), meloxicam (Boehringer Ingelheim) (1 unit/g), and dexamethasone (MSD) (five units) were administered subcutaneously. The scalp was shaved and cleaned with an antiseptic solution of 2% chlorhexidine gluconate and 70% isopropyl alcohol (Becton Dickinson) before head-fixing the animal via ear bars on a stereotactic rig (Kopf Instruments). Scalp and adipose tissue were excised from the right lateral side of the skull. A titanium headplate weighing ∼1g with a 7 mm circular aperture was secured to the skull (−1.4 mm anterior, +4 mm lateral relative to bregma) using superglue and dental cement (SuperBond, Prestige Dental, Bengaluru, Karnataka, India). For some mice, a second layer of 20% carbonated cement was used to further adhere the headplate to the skull. A 4 mm craniotomy was drilled with a dental drill (NSK) at the centre of the headplate aperture, and the dura was removed. A micropipette was front-loaded with pAAV-CamKIIa-C1V1(t/t)-mScarlet-KV2.1 (Addgene: 124650-AAV9; 2.70×10^{13} vg/mL) diluted 1:11 in sterile PBS. 500 nL of the virus was injected into the left of the craniotomy, 300 μm deep from the cortical surface at a rate of 50 nL/min using a hydraulic micromanipulator (Narashige, Tokyo, Japan). Additionally, 150 nL Cholera Toxin Subunit B (Recombinant) Alexa Fluor 647 Conjugate (ThermoFisher: C34778; 100 ug) diluted 1:2 in sterile PBS was injected either into S2 (right side of the craniotomy), dorsolateral striatum (−0.9 mm anterior, +3 mm lateral, +3 mm ventral relative to bregma) or VPM (−1.92 mm anterior,

+3 mm lateral, 75 nL at +3.85 mm ventral and 75 nL at +3.5 mm ventral relative to bregma) at a rate of 50 nL/min. A cranial window comprised of two optically bonded glass coverslips (4 mm diameter and 5 mm diameter) (Warner Instruments, Massachusetts, United States) was used to plug the craniotomy. The cranial window was secured in place using cyanoacrylate (3M) and dental cement. The animal was allowed to fully recover for at least 3 weeks before any further experiments were conducted. Mice were kept on a 12:12-h light/dark cycle with access to food and water *ad libitum*.

After reaching expert performance on the one-photon task, mice were re-anaesthetised and a small craniotomy was drilled over the left hemisphere S1/S2 border (−1.4 mm anterior, −4 mm lateral relative to bregma) and a half headplate (same headplate as before cut in half from front to back) implanted around it for extra secure head-fixation. This second surgery was minor and took less than an hour. Mice were allowed to recover for a minimum of 1 day before continuing experiments.

4.2.2 Two-photon imaging and stimulation

The microscope used in this study was a Bruker Ultima 2Pplus equipped with a tuneable wavelength laser for imaging (Chameleon Ultra II, Coherent, Saxonburg, Pennsylvania 16056, USA) and a 1035 nm fixed wavelength laser for

photostimulation (Monaco 1035, Coherent). The 1035 nm laser was directed through a SLM (Neuralight 3D Ultra, Bruker, Billerica, Massachusetts, United States). Photostimulation parameters and phase masks were defined and created using a custom-modified version of NAPARM[395] before being loaded into external software. Imaging and photostimulation were controlled using Prairie View software (Bruker). Phase masks were loaded onto the SLM using BOSS (Boulder Nonlinear Systems). The animal was head-fixed on a tip-tilt stage (TTR001, Thorlabs) under a 16x objective (N16XLWD-PF, Nikon, Tokyo, Japan). Stimulation and imaging triggers were coordinated using using Rigbox[374]. More details on the hardware and software can be found here[375,376].

A reference two-photon image for opsin expression (conjugated to mScarlet) was captured at 765 nm using a 570 to 620 nm emission filter, averaged over 128 frames at 1024 x 1024 resolution. NAPARM was used to group opsin-expressing neurons into groups of 10. Within each group, 6 mW 1035 nm laser light (2 MHz repetition rate) was simultaneously spiralled over each cell body. The spiral spanned a diameter of 10 μm for a duration of 10 ms, repeating 25 times for a total photostimulation duration of 250 ms. The phase mask changed every 1 s, and each group received 10 sets of spiral stimulations with a 15 s inter-trial interval. Photoresponsive neurons were identified manually from GCaMP6s $\Delta F/F$ stimulus-triggered average images based on a baseline of 1 s pre-stimulus. Photoresponsive cells were used for experimental testing.

Responses to the photostimulation during photoresponse testing and behavioural experiments were imaged across a 1.4 x 1.4 mm field of view (1024 x 1024 resolution) at 15 Hz using resonant scanning at 920 nm (80 MHz pulse repetition rate, 50 mW average power after the objective), with a 500 to 550 nm emission filter. During the behavioural task, imaging was triggered to only occur during the 1 second pre- and post-stimulation to minimise opsin desensitisation.

4.2.3 Whisker and reward mapping

Calcium responses to whisker stimulation and free reward were recorded separately using two-photon imaging of GCaMP prior to the start of the task. All contralateral whiskers were simultaneously deflected for 1 second every 10 seconds 10 times using a paintbrush head controlled by a piezoactuator (Physik Instrumente). The brush was driven by an oscillating signal at 20 Hz across a physical distance of ~300 μm. Separately, 3 uL of free reward was given every 10 seconds in 10 blocks.

4.2.4 GCaMP imaging analysis

Calcium imaging movies were processed using Suite2p[412]; cell bodies were automatically detected and carried forward for analysis. $\Delta F/F$ was computed for each cell body identified by Suite2p using the equation:

$$\Delta F/F = (F - \underline{F})/\underline{F}$$

where F = raw fluorescence, $\underline{F}$ = the mean of F in the 1 s pre-stimulation.

Cells were labelled as "responsive" to optogenetic stimulation, whisker stimulation or free reward if $\Delta F/F$ in the 1 s post-stimulation was >= 0.1. The "response" of a population of responsive cells to photostimulation is defined as the area under the curve of the mean $\Delta F/F$ across cells in the 1 s poststimulation.

If multiple cells were wholly or partially in the radius of a beamlet, only the cell with the centroid closest to the centre of the beamlet was labelled as a "target". When analysing "non-target" cells, any non-target cells wholly or partially within the radius of a beamlet were excluded from analysis.

For a cell to be labelled as excited or inhibited, it must be significantly correlated both across all trials (2 s around stimulation) concatenated and across trial one separately ($p<0.05$ Pearson's correlation). The correlation coefficient of the cell was taken from the correlation across all trials.

4.2.5 Analysis, statistics and reproducibility

All analysis was performed using MATLAB. No statistical method was used to predetermine sample size. No data were excluded from the analysis. The scientists were not blinded to retrograde tracer region expression during experiments or analysis.

4.2.6 Behavioural training

All animals were allowed to fully recover for a minimum of 3 weeks post-surgery before being put on water restriction. Mice were kept on a 12:12-h light/dark cycle, with access to food *ad libitum*. Once on water restriction, the weights of the mice were monitored delay and not allowed to fall below 80% of their pre-water restriction bodyweight. For behavioural training, mice were head-fixed with their forepaws resting on a Lego wheel, as previously described[373]. Wheel movement was measured using a rotary encoder (Kubler 0.52400.1122.0360). The precise timing and volume of water delivery were controlled by a solenoid pinch valve. During neuronal recording sessions, two cameras recorded pupil and face/body movement (FLIR, lens: MVL50M23, MVL16M23) in combination with an infrared light source (BW 48 LED Illuminator). During the one-photon task, a patch cable attached to a fibre-coupled LED (M530F3, Thorlabs) was lowered to just above the cranial window for optogenetic stimulation. During the two-photon task, widefield stimulation was delivered through the objective after passing through a filter cube to produce green light. The behavioural task was programmed and controlled using Rigbox[374]. More details on the hardware and software can be found here[375,376]. Optogenetic stimulation in the one-photon task spanned 180 milliseconds, consisting of five 20 millisecond pulses separated by 20 millisecond inter-pulse intervals. To earn a 3 uL water reward, the wheel had to be turned at least 6 mm

in total in any direction during the response window. Once mice reached a d' of 2 at 0.1 mW stimulation power on sample in the one-photon task, they were graduated to the two-photon task. In the two-photon task, 15% of stimulation events consisted widefield stimulation; 70% consisted of spiralling 6 mW laser light over each cell body for 10 ms. Each cell body received 10 blocks of stimulation in total (over 150 ms). On catch trials (15%) everything was kept the same as two-photon stimulation trials (including galvanometer movement) but laser power was set to 0. This was to prevent the mouse for learning to associate mechanical sounds of the microscope during laser stimulation with reward.

4.2.7 Electrophysiology

Electrophysiological data was recorded using Neuropixels probes[26]. Before insertion, the shank was rinsed in deionised water, then coated in a red-fluorescent dye (DiI, V22885, Invitrogen). Optimal trajectories for penetrating VPM, and CP were calculated using Pinpoint[377]. Neuropixels were inserted through the craniotomy contralateral to the cranial window. To record from CP, the probe was inserted 7.0 mm deep at 70 yaw and 32.5 pitch. To record from VPM and PO, the probe was inserted 5.8 mm deep at 100 yaw and 32.5 pitch. The probe was lowered at ~5 um/s using a micromanipulator (Sensapex) to 100 μm deeper than target depth and then retracted by 100 μm to reduce tension within the tissue.

Electrophysiology recordings were captured using SpikeGLX[378] (30 kHz sampling rate) and preprocessed using CatGT[379] which performs common reference averaging and bandpass filtering. Spike sorting was performed using Kilosort4[380] and curated using Bombcell[381]. Defined Bombcell parameters: minimum signal-to-noise ratio = 2, minimum amplitude = 15, maximum resting period violations = 0.5, hillorLlobetMethod = 0.

4.2.8 Histology

On the last day of electrophysiology recording, mice were injected with a lethal dose of pentobarbitone and perfused with 4% paraformaldehyde. The brain was extracted and left in 4% PFA for 24 hours. A vibratome was used to slice the brain at a thickness of 70 μm. Slices were mounted and imaged using a red filter to detect DiI and C1V1-mScarlet expression and a green filter to detect GCaMP. Slices were also imaged using a far-red filter to detect Alexa Fluor 647. Probe trajectory reconstruction and opsin expression localisation and quantification were performed using HERBS[383].

4.3 Results

4.3.1 Cortical and subcortical-projecting neurons in wS1 are largely non-overlapping with somatosensory-responsive neurons

Spectrally differentiating more than two fluorophores during two-photon imaging is difficult thanks to the broad two-photon excitation spectra of molecular probes over the wavelength range accessible to typical Ti:Sapphire lasers[427] and the limited width of the usable spectral emission window[428]. However, using a combination of different excitation wavelengths paired with a fixed emission filter, we were able to unmix the contribution of different channels and differentiate cells positive for opsin expression (tagged with mScarlet) and Cre-dependant Alex fluor 647, despite global GCaMP expression (figure 10a). On average, a greater number of wS1 cells responded stronger and more sharply to 3 ul free reward than 1 second of 20 Hz whisker stimulation (figure 10b). In fact, reward-responsive cells comprised the highest proportion of functionally-identified cells at ~20%. This is likely due to pre-training on the one-photon optogenetic CPD task, leading to remapping of stimulus-selective neurons to outcome-selective neurons throughout learning[429] (figure 10c). In future cohorts, the same assessments should be repeated to investigate the distribution of reward-responsive cells in naïve mice. The next largest proportion of cells were whisker-

responsive, then S2-projecting, then mixed selective for reward and whisker stimulation, then VPM-projecting, then S2-projecting reward responsive, then DLS-projecting. This contrasts to previous reports which describe more extensive wS1 efferents targeting striatum than S2 and VPM[430], although the present study has a higher sample size. Reward-responsive neurons were proportionally enriched in S2-projectors compared to other projectors and non-projectors. Average population activity across S2-projectors has previously been shown to significantly increase with reward delivery, on top of the increase evoked by whisker deflection alone[431]. However, the segregation of these responses among individual neurons has yet to be reported. We found that neurons with mixed selectivity for both whisker-stimulation and reward were almost entirely restricted to non-projectors.

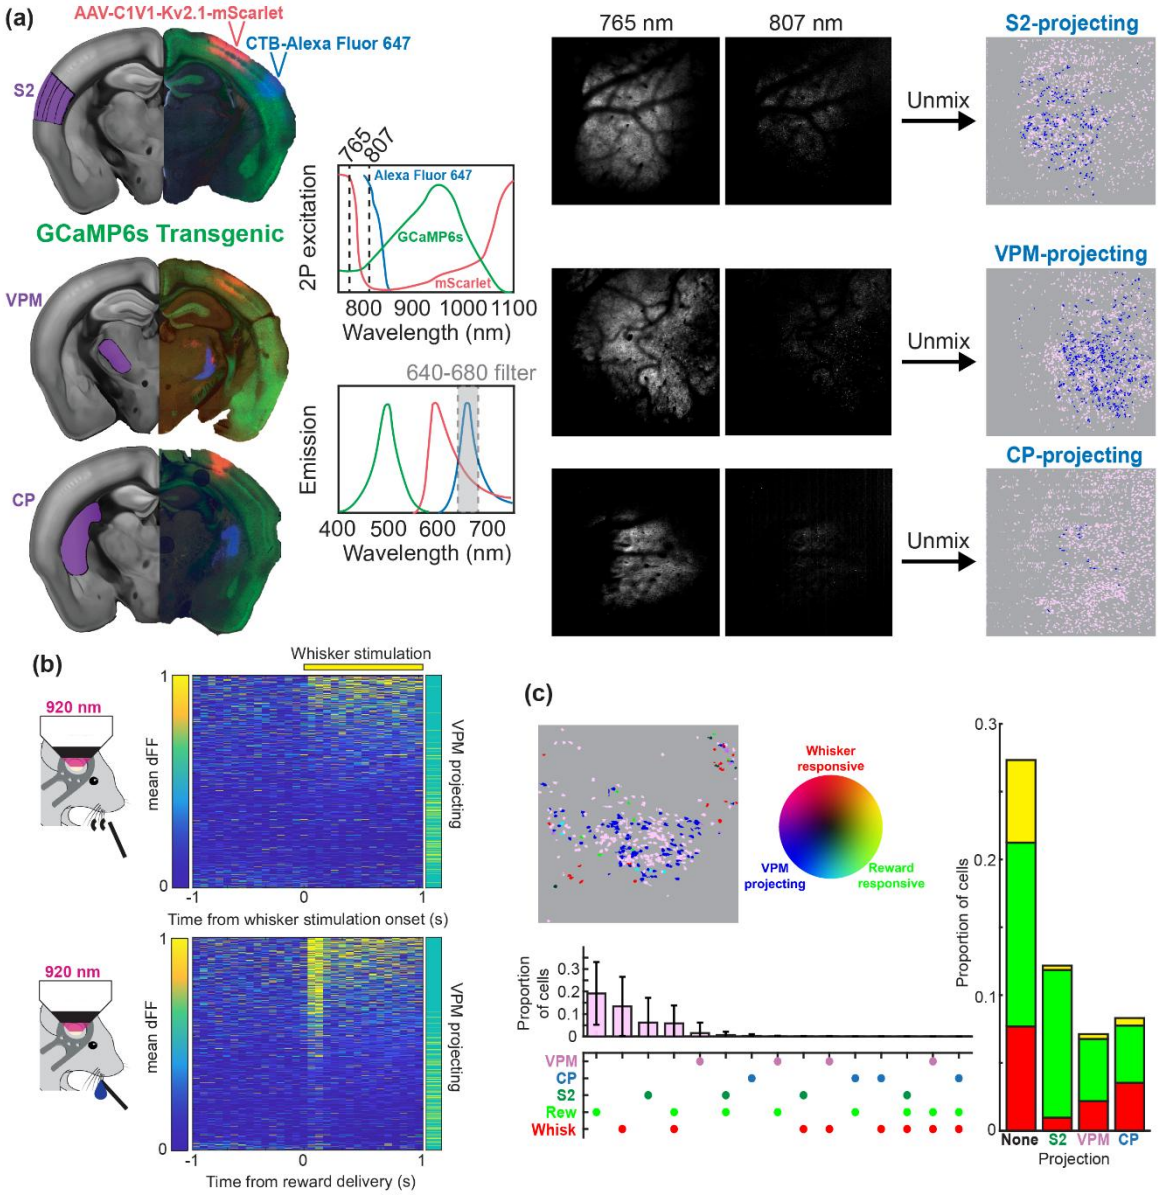


Figure 10. Cortical and subcortical-projecting neurons in wS1 are largely non-overlapping with somatosensory-responsive neurons. (a) Left: CTB-Alexa Fluor 647 was either injected into S2 (n=2), VPM (n=3) or CP (n=2). The opsin AAV-C1V1-Kv2.1-mScarlet was injected into S1. Right: Images were taken at multiple wavelengths with a 640-680 nm emission filter to unmix the fluorescent signals from Alexa Fluor 647 and mScarlet and identify outwardly projecting cells. **(b)** Responses to whisker stimulation and free reward from an example session. **(c)** Left: the relative proportion of somatosensory responsive and projecting cells across the field of view. Right: the relative proportion of somatosensory responsive cells within non-projecting and projecting populations.

4.3.2 Training on a one-photon optogenetic CPD task reduces surround inhibition evoked by two-photon stimulation

After being trained to expert performance on the one-photon somatosensory optogenetic CPD task as described previously, mice were graduated to a version of the task under a two-photon microscope. The chance of one-photon stimulation and catch trials were dropped to 15%; the remaining 70% of trials were allocated to two-photon holographic stimulation (figure 11a). Two-photon imaging was programmatically controlled to only occur from 1 second pre-imaging to 1 second post-imaging with the aim of reducing opsin desensitisation as a result of spectral overlap. The galvanometer mirrors were moved identically on holographic stimulation and catch trials to prevent the animals from associating any mechanical auditory cues with reward delivery[286]. Like one-photon trials, holographic stimulation trials also had a response window of 1 second and were rewarded if the wheel was turned to threshold. The probability of successful detection and report of holographic stimulation increased with number of cells activated and total power on sample of the 1035 nm laser (figure 11b), although with a much weaker relationship than previously reported[286,287]. In total, we acquired data from 26 two-photon training sessions across 7 mice. Because targets were pseudo-randomly selected from all available opsin-expressing

neurons on a trial-by-trial basis, more data is required to build a high enough sample size across cell groups to reliably assess how stimulating different neurons based on their response profile/projection can shift the psychometric curve. Interestingly, unlike in the naïve mice, holographic stimulation did not produce a peak in the distribution of anticorrelated cells 50-100 um away from the targets; however, the proximal relative bias towards excited cells remained (figure 11c). Furthermore, the degree of correlation across excited and inhibited cells did not clearly decrease as a function of distance from the targets like it did in the naïve mice (figure 11d). This suggests that the state of the network may have adapted throughout learning to efficiently receive and propagate the rich optogenetic signal outwards. Potentially, the lack of surround inhibition could be a sign that the injected activity has become on-manifold in expert mice and does not have to be sequestered onto a narrow neural trajectory[173,175,432,433].

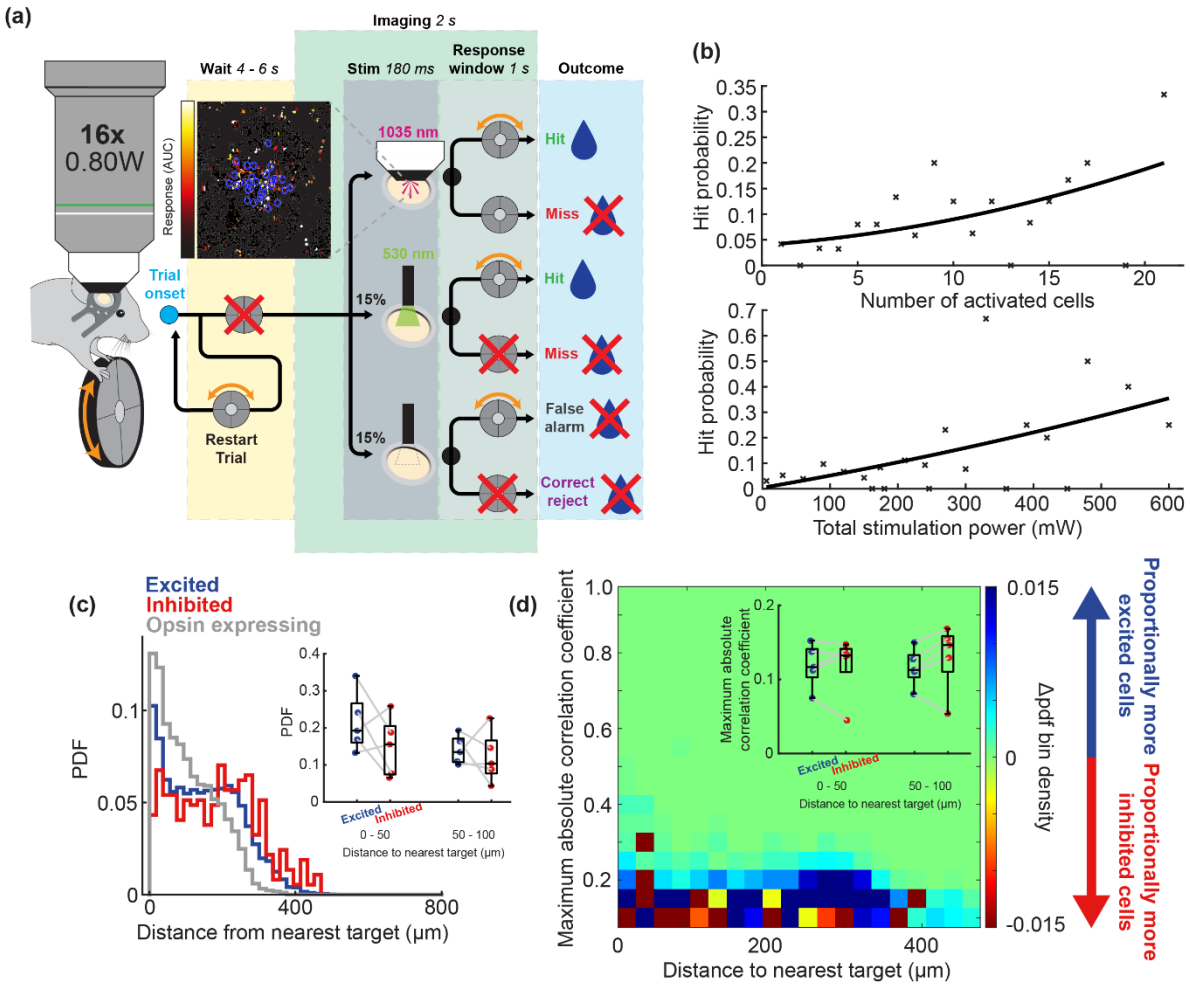


Figure 11. Training on a one-photon optogenetic CPD task reduces surround inhibition evoked by two-photon stimulation. (a) After reaching expert performance on the one-photon CPD task, mice were graduated to a similar task under the two-photon microscope that contained holographic stimulation trials. Inset: responses to holographic stimulation. Beamlet locations are shown as blue rings. **(b)** Hit probability increased on holographic stimulation with the number of targeted cells. **(c)** Probability distribution of excited, inhibited and opsin-expressing non-target cells as a function of distance from the nearest target cell. Inset: Each point is data from an individual animal, averaged across all cells (* = $p < 0.05$, Wilcoxon signed-rank test, n = 5 mice). (d) Maximum correlation strength of an individual non-target cell with any target as a function of distance from the nearest target cell, color-coded by excited bin count - inhibited bin count. Inset: Each point is data from an individual animal, averaged across all cells for a given type (* = $p < 0.05$, Wilcoxon signed-rank test, n = 5 mice).

4.3.3 Near-infrared wavelengths used for two-photon imaging and stimulation produce large photoelectric artefacts in Neuropixels recordings

We wanted to investigate how the structure of two-photon optogenetic stimulation impacted downstream neuronal signatures of perception. Sparse reliable sequences of neuronal spiking activity differentiate between hits and misses throughout the top-down somatosensory pathway (figure 3, figure 4), but is the reliability of these sequences determined by the reliability of the stimulus or just binary perception? To investigate this, we aimed to combine two-photon holographic stimulation of wS1 in expert mice with subcortical high density electrophysiology recordings. Inserting the Neuropixels probe through the same hemisphere as the cranial window would be less invasive thanks to reusing the same craniotomy and would allow for recording from ipsilateral S2. However, navigating around the bulky objective used for two-photon imaging proved difficult (figure 12a). Even air objectives with higher working distances were still too large to permit insertions at the angles required to penetrate subcortical structures in the somatosensory pathway (figure 12b). To circumvent these problems, once mice reached expert performance on the one-photon task, a second craniotomy was drilled over contralateral somatosensory cortex. This facilitated deep Neuropixels insertions that spanned both hemispheres to record from CP and VPM (figure 12c). Unfortunately, this positioned the silicon

electrodes on the shank perpendicular to the objective, producing large photoelectric artefacts during two-photon imaging and stimulation (figure 12d). Photoelectric artefacts are caused by a combination of the deep tissue penetration of near infra-red wavelengths and high power on sample needed to resolve fluorescent probes[434]. Artefactual spikes can be differentiated from real spikes based on their amplitude (figure 12e), because there should be negligible difference between spontaneous and optogenetically evoked waveforms[435,436]. However, filtering out artefactual spikes left no discernible stimulus-evoked change in firing rate on holographic stimulation trials. We theorise that this is a result of preprocessing steps in the electrophysiology pipeline that attenuate spikes due to common reference averaging of the artefactual signal[379]. High-pass filtering interacts badly with large artefacts, spreading the artefact energy in time. Additionally, global common reference averaging means that even epochs without artefacts are referenced against a contaminated signal. This is demonstrated by there also being no discernible stimulus-evoked change in firing rate on one-photon trials (figure 12e) which directly contradicts prior data. Additionally, large artefacts can dominate the dynamic range of the signal, resulting in scaling compression effects during spike sorting. We find that this decreases the amplitude and signal-to-noise ratio of recorded spikes on the two-photon task relative to sessions with only one-photon stimulation (supplementary

figure 9). Future rounds of data collection should focus on preventing photoelectric artefacts during recording[434].

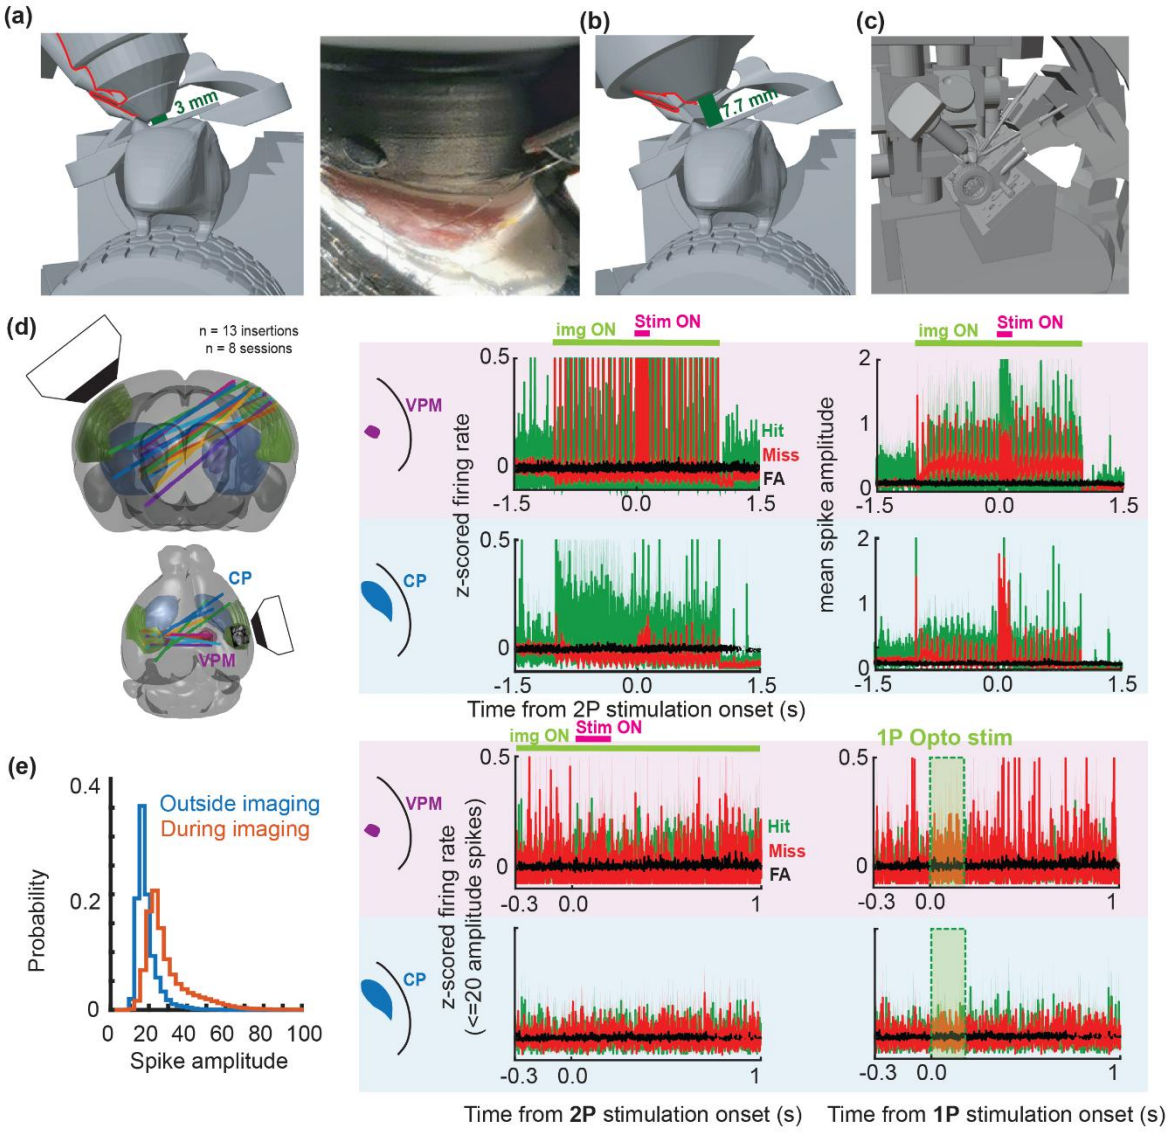


Figure 12. Near-infrared wavelengths used for two-photon imaging and stimulation produce large photoelectric artefacts in Neuropixels recordings (a-b) Insertion angles necessary to record from ipsilateral structures in the somatosensory pathway result in overlapping geometry between the Neuropixels probe and both the Nikon N16XLWD-PF immersion objective and the ThorLabs TL10X-2P air objective. **(c)** Insertion through the contralateral hemisphere avoids overlapping geometry between the Neuropixels probe and the objective. **(d)** Left: reconstructed probe tracts. Right: mean ± sd z-scored firing rates and spike amplitudes around stimulation. Imaging at 920 nm commenced 1 s prior to stimulation onset and lasted until 1s post-stimulation onset. Holographic spiral stimulation at 1035 nm lasted 200 ms. **(e)** Left: spike amplitudes increased during near infra-red imaging/stimulation. Right: mean ± sd z-scored firing rates around stimulation on two-photon and one-photon trials with a ≤ 20 spike amplitude filter.

4.4 Discussion

In order to understand how the structure and reliability of cortical stimulation at a cellular level affects spike sequence temporal codes of perception in a CPD task, we describe a novel methodological approach that combines all-optical interrogation with high density electrophysiological recordings. This culminates in an experimental paradigm with high read and write resolution *in vivo*, permitting investigations with unparalleled precision into how single cortical units collaborate to propagate information downstream. However, further optimisation is needed to reduce the photoelectric artefact produced by near infra-red wavelengths on the silicon electrodes. First observed by Alexandre Edmond Becquerel in 1839[437] and further described by Albert Einstein in 1905[438] (winning him a Nobel Prize), the photoelectric artefact describes the phenomena where photons dislodge electrons on a conductive material causing an artefactual increase in recorded voltage. There are several strategies to avoid multiphoton-induced artefactual spikes described by Takashi Kozai and Alberto Vazquez from the University of Pittsburgh; the main candidates are analytical, relying on evenly illuminating all electrode sites and removing the artefacts post-hoc with common reference averaging or blanking temporally consistent artefacts[434]. However, both these approaches have significant downsides; evenly illuminating all electrode sites simultaneously is impossible with conventional raster scan imaging and would instead rely on widefield illumination. Furthermore, both approaches

would remove or dramatically decrease the amplitude of spikes during the stimulation period. Although this may cause acceptable information loss for larger bin size analyses such as computing firing rate responses, the high temporal resolution acquisition advantage of Neuropixels is significantly attenuated. This is especially problematic for analyses that rely on individual spike times such as SPADE. Theoretically, one could limit the stimulation duration to be shorter than the expected synaptic latency between the stimulation and recording sites, but generating sufficient opsin activation with such an ultra-short pulse through hundreds of microns of tissue is likely to breach laser power safety constraints. Materials-based approaches to mitigating the photoelectric artefact have also been suggested; however, they would primarily be achieved by using optic block coating of the electrode, which has limited biocompatibility and would limit electrical access to surrounding tissue, or by fabricating a high-density probe that utilises ionic conduction electrode sites as opposed to solid-state electrical conductors[434].

One novel use that we employed for cortical holographic stimulation was the somatic targeting of retrogradely labelled projection neurons. Perhaps this has not been attempted before because subcortical projectors are predominantly restricted to deeper cortical layers. Specifically in wS1, projections to striatum are primarily enriched in layer 5, with smaller populations also found in layer 6 and layer 2/3[439], where DLS-projecting neurons can also form collateral

connections with S2[440]. In contrast, canonically wS1 projections to VPM are restricted to layer 6[441,442]. With relatively homogenous distribution of fluorescence throughout the tissue, two-photon imaging in the mouse cortex is limited to a depth of ~450 um[443] which roughly corresponds to layer 4[444]. However, layer 6 corticothalamic neurons can be imaged with two-photon microscopy when selectively labelled[445,446]. This is partly because out of focus signal from the superficial layers is massively reduced, increasing signal-to-background ratio and effective imaging depth. Furthermore, lateral cortical regions such as sensory areas are increasingly compressed at larger window sizes due to curvature of the skull and brain surface, allowing optical access to deeper structures[447]. Therefore, two-photon access to subcortical projecting neurons facilitates additional future experiments investigating the neural and behavioural consequences of precise anatomically-defined single-cell stimulation.

One distinct difference between the findings of the current work and all-optical CPD literature is the relationship between number of cells targeted and the probability of eliciting a percept. These variables classically follow a psychometric function, with two separate studies even both reporting that the simultaneous stimulation of ~22 S1 cells is the threshold for eliciting behavioural reports above chance in an optogenetic CPD task[286,287]. Even though the mice in the present study more-often-than-not achieved dPrime's > 2 on the one-photon trials, they required more cells to be holographically stimulated to produce a

salient enough percept for reliable behavioural report. This is likely because the aforementioned studies also pre-trained mice on holographic stimulation prior to data acquisition, whereby after graduating from the one-photon task, they were trained to reliably report the stimulation of 100-200 neurons before titrating down the number of targets. Pre-training them on holographic stimulation before testing smaller hologram sizes will lower their threshold for perception[287]. The absence of holographic pre-training is a limitation of the current work as it makes the tail of the psychometric curve cover the majority of feasible hologram sizes without providing much additional perceptual information. Future optimisations of the work in this chapter should adopt a holographic pre-training phase. This would also allow essentially unlimited sessions for testing whether selective targeting of projection- or function-defined neurons produce a shift in the psychometric curve before brain health is compromised by repeated probe insertion.

5

Conclusion

In summary, the results described herein provide evidence that perceptual signals are represented both cortically and subcortically in the top-down somatosensory pathway. We are the first to employ FP-growth SPADE analysis on high density *in vivo* electrophysiology recordings across multiple brain regions, elucidating the previously unexplored role that temporally precise, reliable and sparse sequences of spikes play in differentiating between perceived or missed stimuli. However, this evidence is purely correlational. To causally establish the potential role of these sequences as NCC's, one could manipulate them with an all-optical approach that combines subcortical holographic optogenetics[433] with voltage imaging[260]. Unlike calcium imaging, voltage imaging has a high enough temporal resolution to detect individual spikes, which would be required to perform SPADE. As discussed earlier, analyses that detect 'sequences' by minimising global reconstruction error, such as seqNMF, can be run on calcium imaging data

but their bias towards large, consistent temporal structures means that reliable but sparse sequences which contain a small number of neurons and spiking events will be missed. Although the membrane-bound nature of genetically encoded voltage indicators (GEVI) alongside their demand for over 100-fold higher photon flux than GCaMP has previously limited their continuous use to only a few minutes at the required ~100 mW laser power on sample, newer GEVI's with enhanced stability can be recorded for up to an hour[448]. This makes voltage imaging compatible with an optogenetic CPD task for perceptual sequence detection and holographic manipulation. An easier approach that affords less spatially precise stimulation would be to use Neuropixels Opto[449] – a Neuropixels-like shank with 14 evenly-spaced light emitters on the opposite side to the electrodes. The relative ease of this method compared to voltage imaging comes with the trade-off of less than single-cell resolution using the LED's. The true test of the shared spike sequences' perceptual causality would be to optogenetically move them forwards or backwards in time and see if that caused the animal to hit or miss. However, without single-cell resolution stimulation, the reason behind a negative result would be hard to disambiguate; pulling or pushing the 'sequence' may not have had an effect because the temporal code was muddled by the inadvertent recruitment of neurons not typically involved. The all-optical voltage imaging approach provides the resolution to solve this problem, but typically only in a single plane, limiting the available space of viable sequences that can be modulated. We find that only tens of neurons are involved

in sequences despite recording from hundreds of cells; from this we can extrapolate that the likelihood of finding viable sequences that can be replayed in a single voltage imaging FOV is low. Furthermore, identification of such sequences is too computationally intensive for online evaluation and would thus require either finding the same cells across days for voltage imaging or chronically implanting a probe for the Neuropixels Opto approach. This would allow sequences of activity to be analytically identified between recordings and manipulated in future sessions with access to the same units. On the positive side, both approaches would provide valuable insight into how reliable sequences underlying perception and the disappearance of surround inhibition develop throughout learning. An alternative approach for gaining a pseudo-causal understanding of the role that spike sequences play in perception is to fit a recurrent neural network model to the data that captures optogenetically evoked activity[450]. The sequence identification analysis could then be run on the model output to find both shared and independent sequences. Finally, one could temporally shuffle the shared sequences and remove independent sequences whilst conserving firing rate, then assess the model's ability to decode perception.

We find that optogenetically-evoked surround inhibition reduces over learning, clearly demonstrating that the training procedure is changing the network state in wS1. If this is the case, then how do we know that sparse, reliable sequences would differentiate between perceptual hits and misses in a naïve

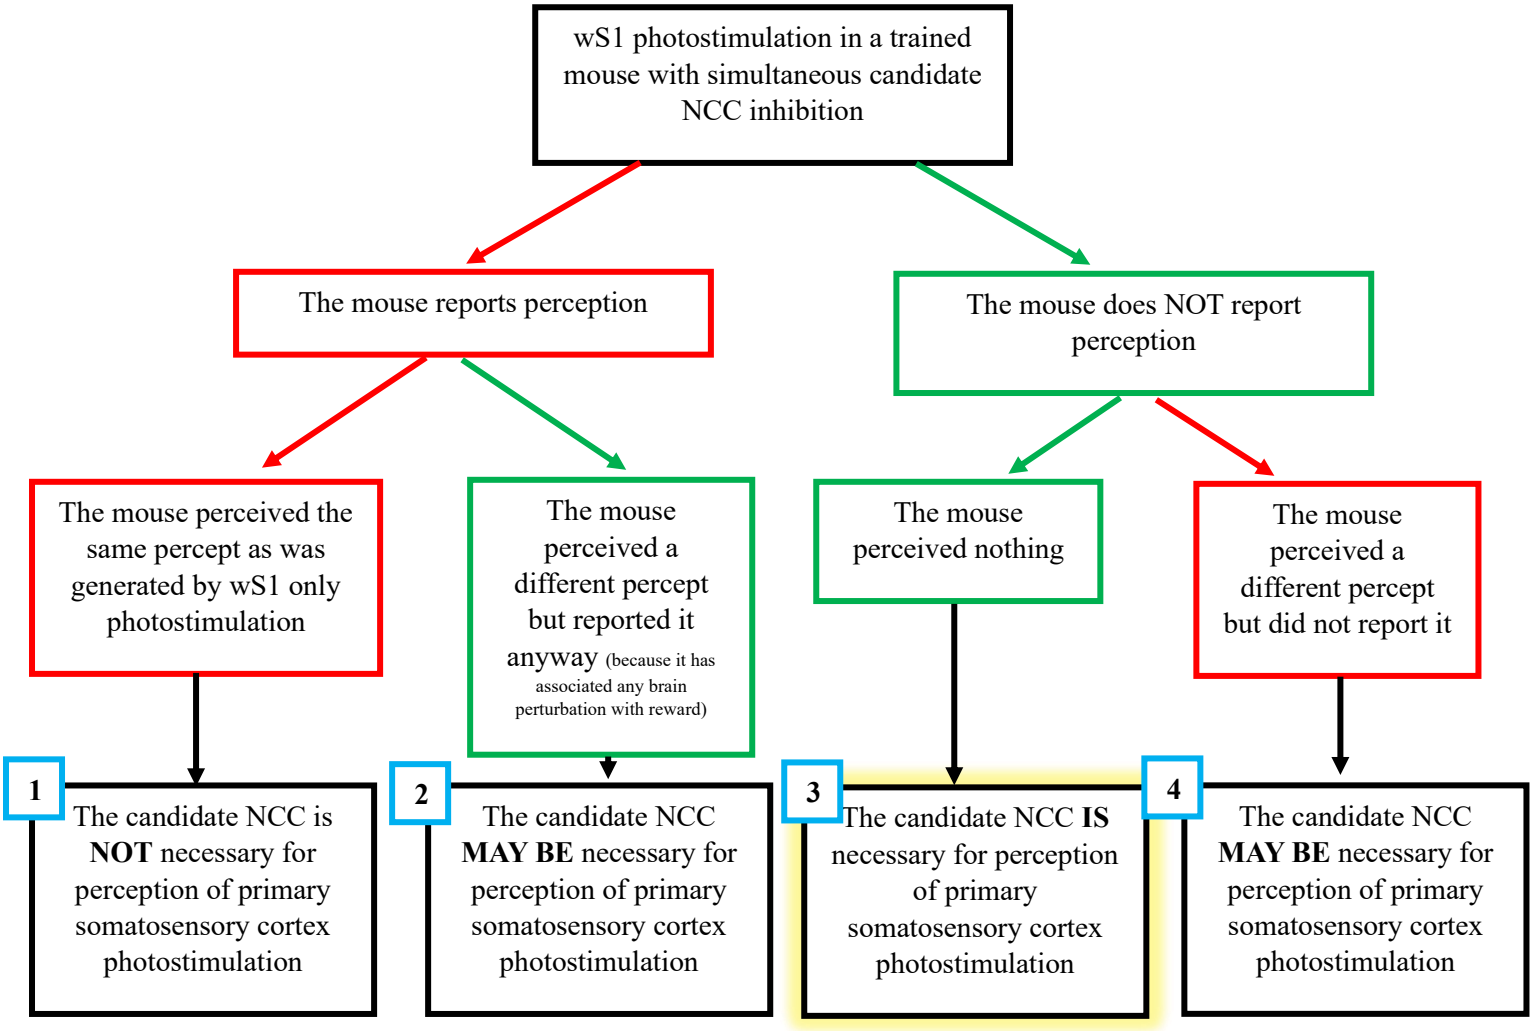
system? This is a genuine limitation of any studies that use artificial perturbations to understand the brain; by not respecting the neural manifold, the produced activity is supernatural and may not generalise to a 'real-world' context[451]. However, the current work at least reveals that the brain has the *capacity* to encode perception in sequences of spiking activity, even if they do not occur under normal circumstances. Perhaps this is still useful information; after all, if the brain is capable of encoding information in this way, then that latent ability has been conserved through evolution for a reason. By understanding this functional capacity of the brain, we can begin to investigate the gating conditions that prevent certain routes being used under physiological conditions, and why a manifold might expand to allow behavioural relevant supernatural dynamics, as is this case with the current findings. Future optogenetic CPD studies should test the surround inhibition response to holographic stimulation throughout learning. Perhaps associating the stimulus with reward through operant conditioning expands the neural manifold to allow artificially injected activity to more readily propagate through the brain, and also encourages all aspects of the stimulus to be encoded rather than just what is sequestered onto the original manifold. This would benefit learning by increasing the richness of the stimulus representation and may correlate with session-by-session learning of the task. The rationale behind the current study using an optogenetic CPD task instead of physical stimulus detection was to more clearly isolate perceptual signals from those involved in pre-cortical processing and reporting of the percept. It also afforded

a high degree of read/write control over the percept, as an animal trained to report such a percept is required to only use features of the evoked cortical activity to guide its behaviour[137]. Thus, all neuronal signals representing that specific percept can be traced outward from the site of modulation. Going forwards, now that it is clear that the brain *can* use rapid, reliable, sparse spike sequences to encode perception, to investigate whether this is employed endogenously, these sequences could be investigated in the somatosensory system during a whisker deflection task. If the sequences follow the same general rules, we can assume that they are perceptual in nature rather than a result of the processing or reporting confounds.

This thesis opened with our current understanding of the neural correlates of consciousness (NCC) and how best to investigate them. But has the current work described herein furthered our understanding in this regard? Rather than elucidating the minimal neuronal mechanisms jointly sufficient for a single conscious percept, we have instead identified candidate NCCs in the form of sparse reliable sequences of spikes in the top-down somatosensory system. To test their causal influence, one could precisely replay them or displace them in time and observe the impact on behavioural reports as a proxy for perception, as previously discussed. However, although this may tell you whether the candidate NCC is *sufficient* for evoking the percept, it does not tell you whether it is *necessary*. This maps back onto Ned Hall's definitions of causal production and

causal dependence respectively. Ideally, one would find a middle ground that reasonably satisfies both definitions. To determine a candidate NCCs causal dependence, one could perform inhibitory optogenetic experiments and observe whether preventing a sequence of spikes reduces the probability of perceptual reports, especially in the case of hit-only sequences that may encode perception in their presence or absence. This could be achieved by expressing an inhibitory opsin in the region housing the sequence, a strategy which has been shown to effectively inhibit subcortical structures and their cortical projections[452]. Again, this would have to be administered through a methodology with simultaneous electrophysiological read capabilities and optogenetic write capabilities like the Neuropixels Opto or all-optical voltage imaging. However, the period of neuronal silence induced by the activation of inhibitory opsins is typically followed by unwanted 'rebound' activity[453]. Alternatively, candidate NCC inhibition may be most efficiently achieved by transfecting local interneurons with an excitatory opsin[454]. This could be facilitated by injecting an excitatory opsin under the control of an interneuron specific promotor such as Dix5/6 into the target structure. However, Dix5/6 is not specific to interneurons outside of the forebrain[455] and inhibitory interneurons are not found in all brain structures, such as most nuclei of the thalamus (except visual thalamus)[456]. Therefore, inhibition by way of expressing an inhibitory opsin in excitatory neurones may be the most appropriate way forward for some candidate NCCs. Regardless, a problem with trying to establish causal dependence through optogenetic inhibition is that if the

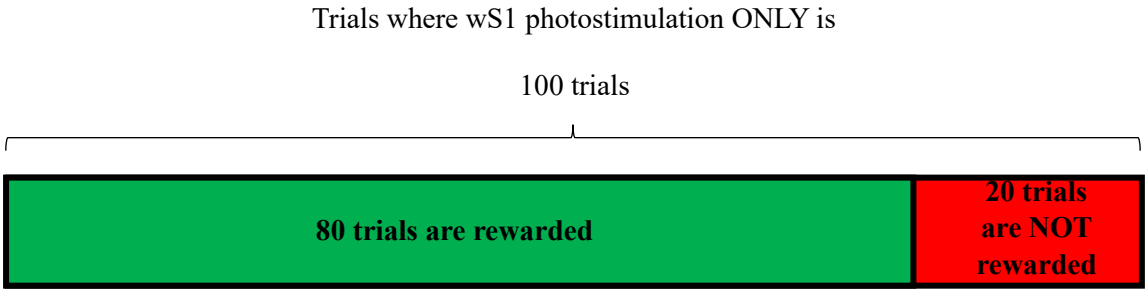
mice still report wS1 stimulation even when the candidate NCC is inhibited, this does not necessarily mean that the candidate NCC is **not** necessary for perception of the photostimulus. This can be illustrated with a flow chart:



One way to disambiguate possibilities 1 and 2 if the animal did report perception with candidate NCC inhibition would be to reduce the mouse's confidence that wS1 stimulation is associated with a reward. For example, one could only reward 80% of wS1 only stimulation trials. Then, when the mouse has become an expert at that paradigm and their report rate determined in a test phase, an extra trial type

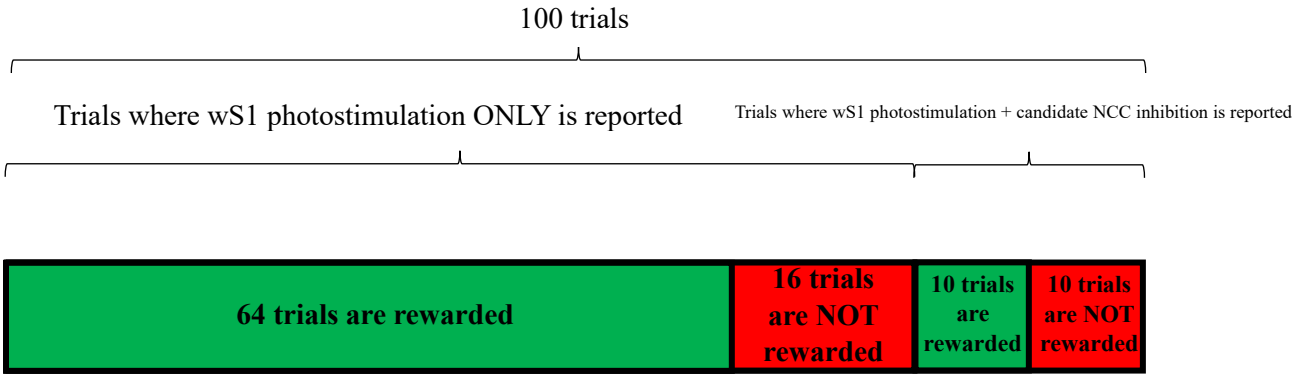
could be introduced where the candidate NCC is inhibited simultaneously with wS1 photostimulation: 50% of trials for this new trial type could be rewarded.

1st phase:



80% of trials producing the wS1 ONLY photostimulation percept are rewarded.

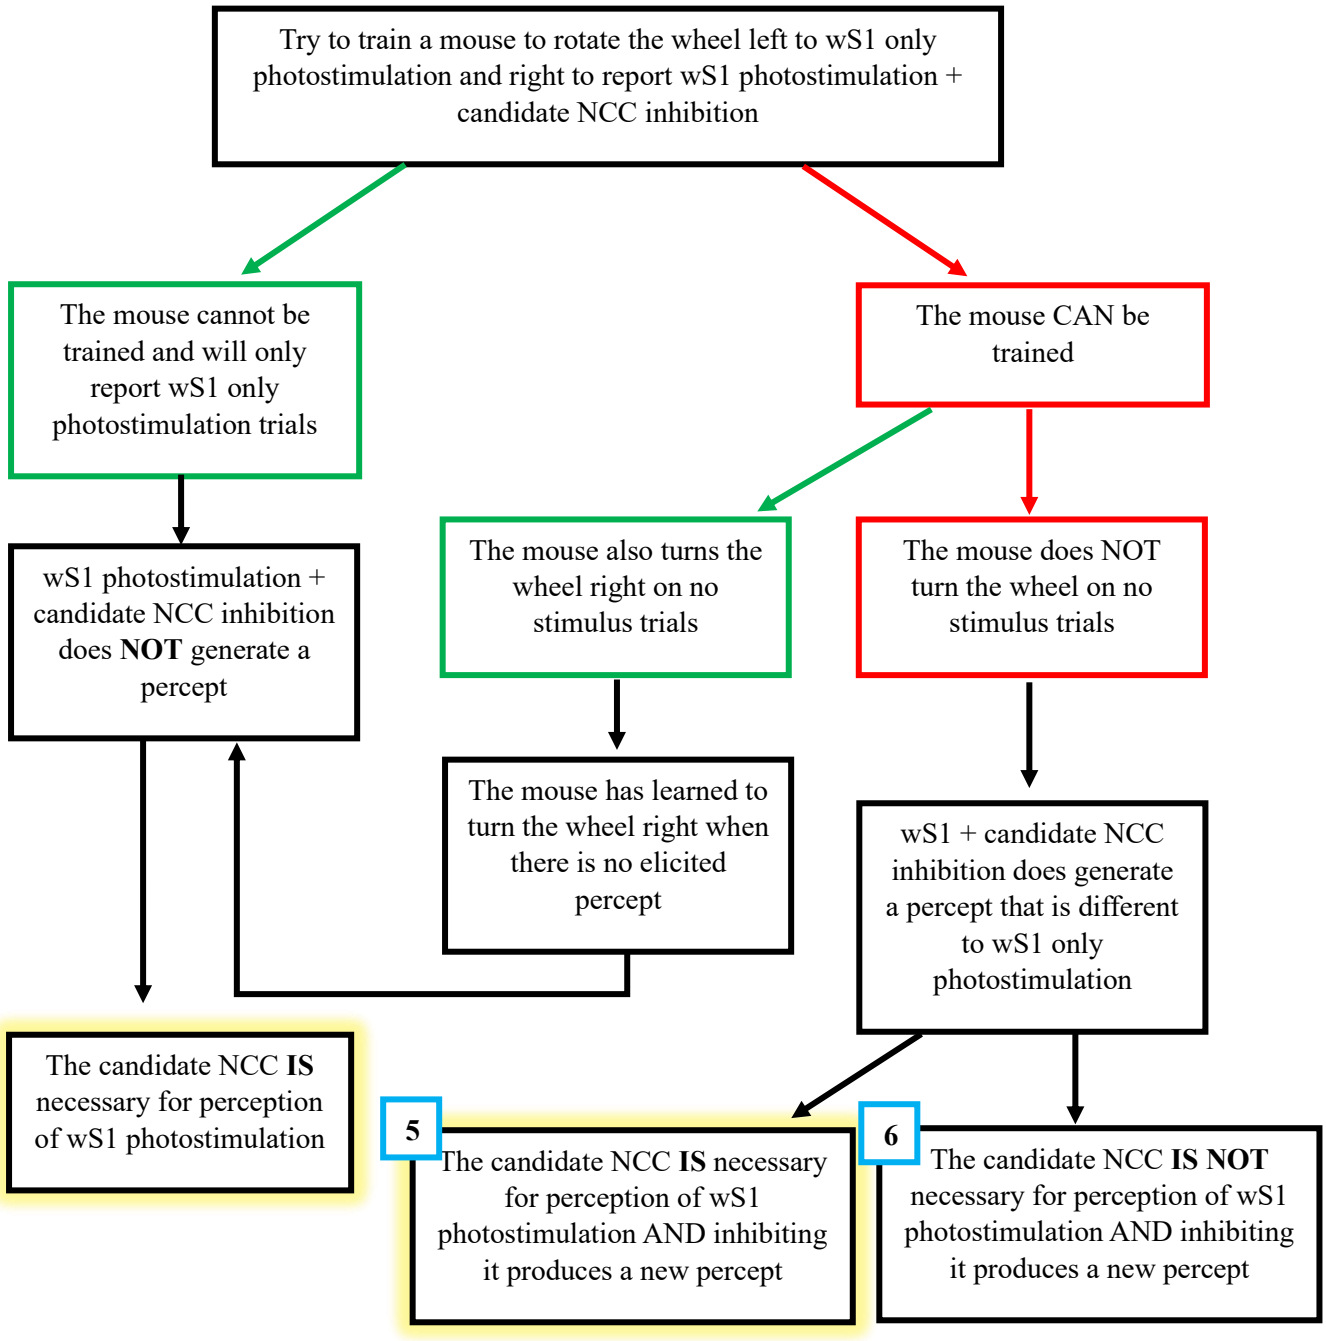
2nd phase:



If the wS1 only photostimulation subjective percept is identical to the wS1 photostimulation + candidate NCC inhibition percept, 74% of trials with that percept are rewarded. If the wS1 only photostimulation percept is different to the wS1 photostimulation + candidate NCC inhibition percept (or lack of percept), 80% of trials with the wS1 only photostimulation percept are rewarded and 50% of trials with the wS1 photostimulation + candidate NCC inhibition percept (or

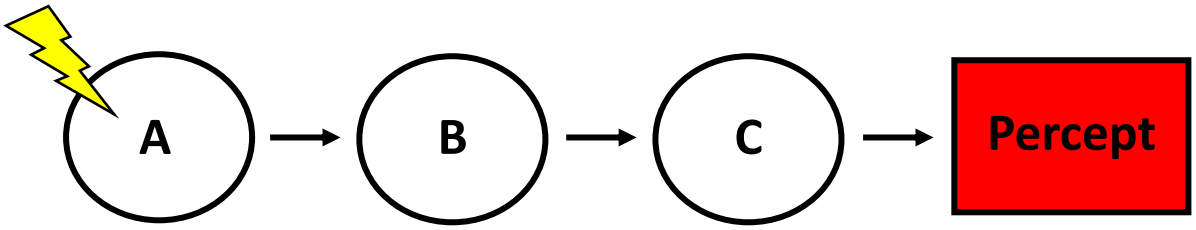
lack of percept) are rewarded. Therefore, if the mouse's report rate over the wS1 only photostimulation trials decreases from the 1st phase to the 2nd phase, it can be assumed that the wS1 only photostimulation percept and the wS1 photostimulation + candidate NCC inhibition percept are identical (as the reward rate for that percept drops from 80% to 74%). Whereas, if the mouse's report rate over the wS1 only trials remains constant between the two phases, it can be assumed that the wS1 only photostimulation percept and the wS1 photostimulation + candidate NCC inhibition percept are different. Even if the second outcome occurs, whether wS1 photostimulation + candidate NCC inhibition even produces a percept or not is a different matter, and requires disambiguating possibilities 3 and 4.

One way to disambiguate possibilities 3 and 4 if the animal did **not** report perception with candidate NCC inhibition would be to try and train the mouse to rotate the wheel left to report wS1 only photostimulation and right to report wS1 photostimulation + candidate NCC inhibition:


Try to train a mouse to rotate the wheel left to wS1 only photostimulation and right to report wS1 photostimulation + candidate NCC inhibition
The mouse cannot be trained and will only report wS1 only photostimulation trials
The mouse CAN be trained
wS1 photostimulation + candidate NCC inhibition does **NOT** generate a percept
The mouse also turns the wheel right on no stimulus trials
The mouse does NOT turn the wheel on no stimulus trials
The mouse has learned to turn the wheel right when there is no elicited percept
wS1 + candidate NCC inhibition does generate a percept that is different to wS1 only photostimulation
The candidate NCC **IS** necessary for perception of wS1 photostimulation
5
The candidate NCC **IS** necessary for perception of wS1 photostimulation AND inhibiting it produces a new percept
6
The candidate NCC **IS NOT** necessary for perception of wS1 photostimulation AND inhibiting it produces a new percept

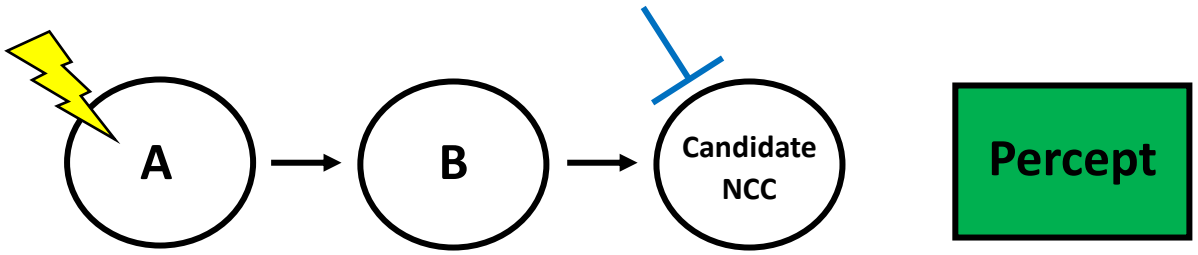
However, if it came down to it, disambiguating between possibilities 5 and 6 would be difficult, as demonstrated here using a simple schematic:

Schema 1:

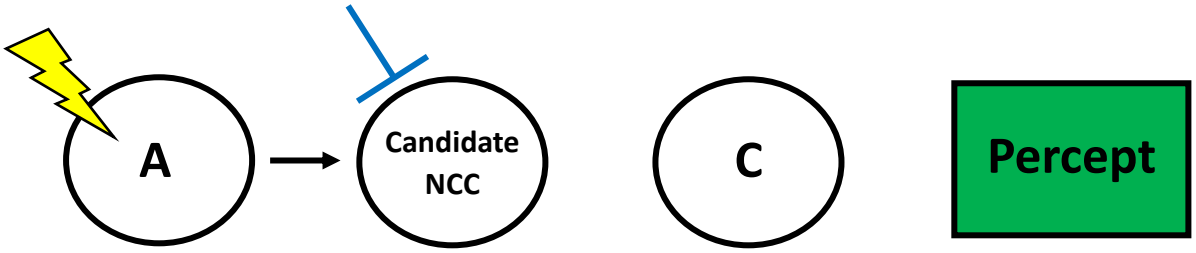


Stimulating brain region A causes propagation of activity through brain regions B and C. This activity reaching brain region C produces percept 'Red'.

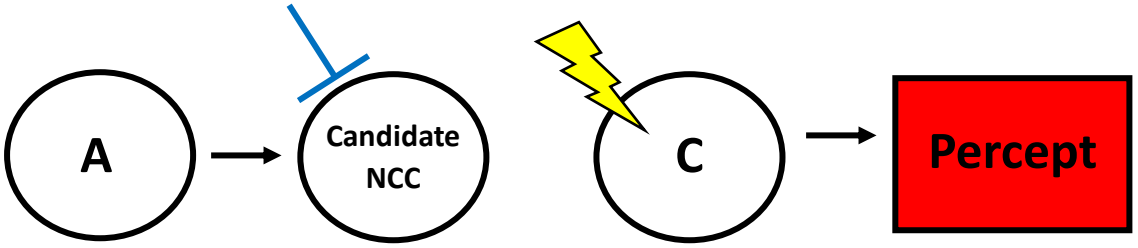
Schema 2 - an example of a situation that could produce possibility 5:



Schema 3 - an example of a situation that could produce possibility 6:



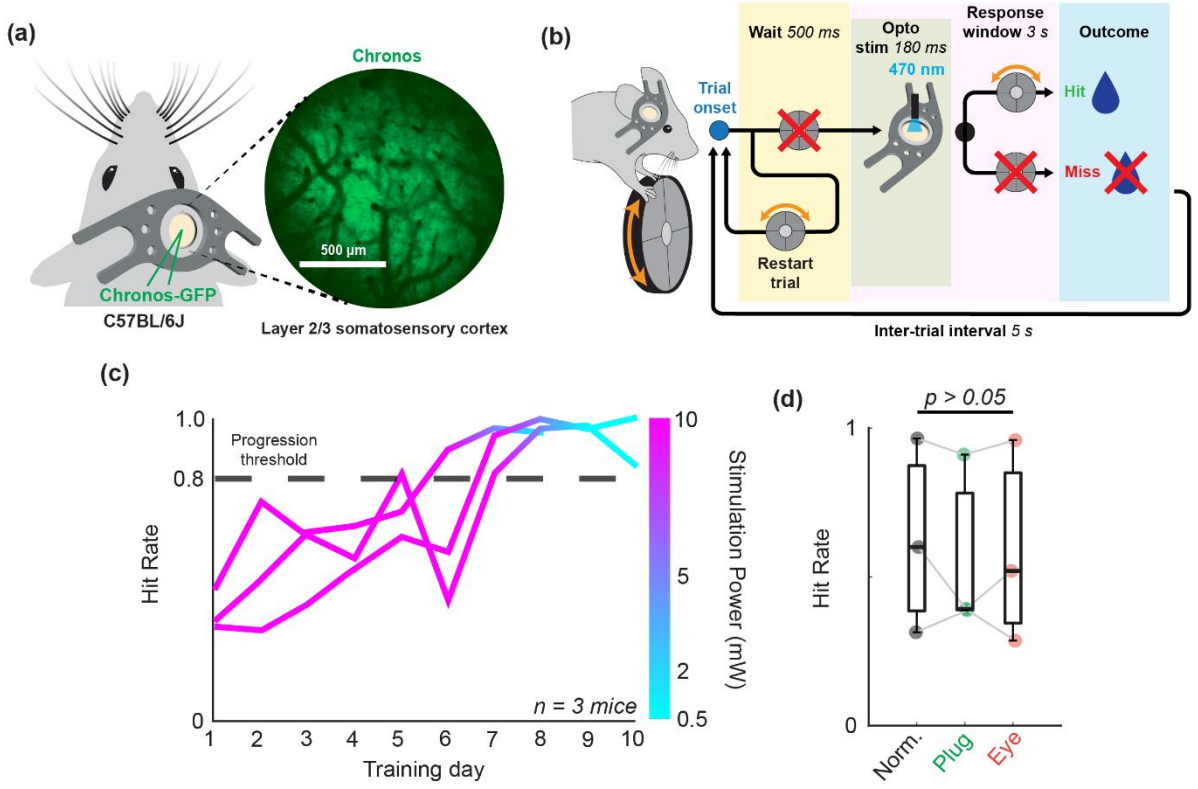
The reason schema 3 could produce possibility 6 is because recreating the activity in brain region C produced during schema 1 would still generate percept 'Red':



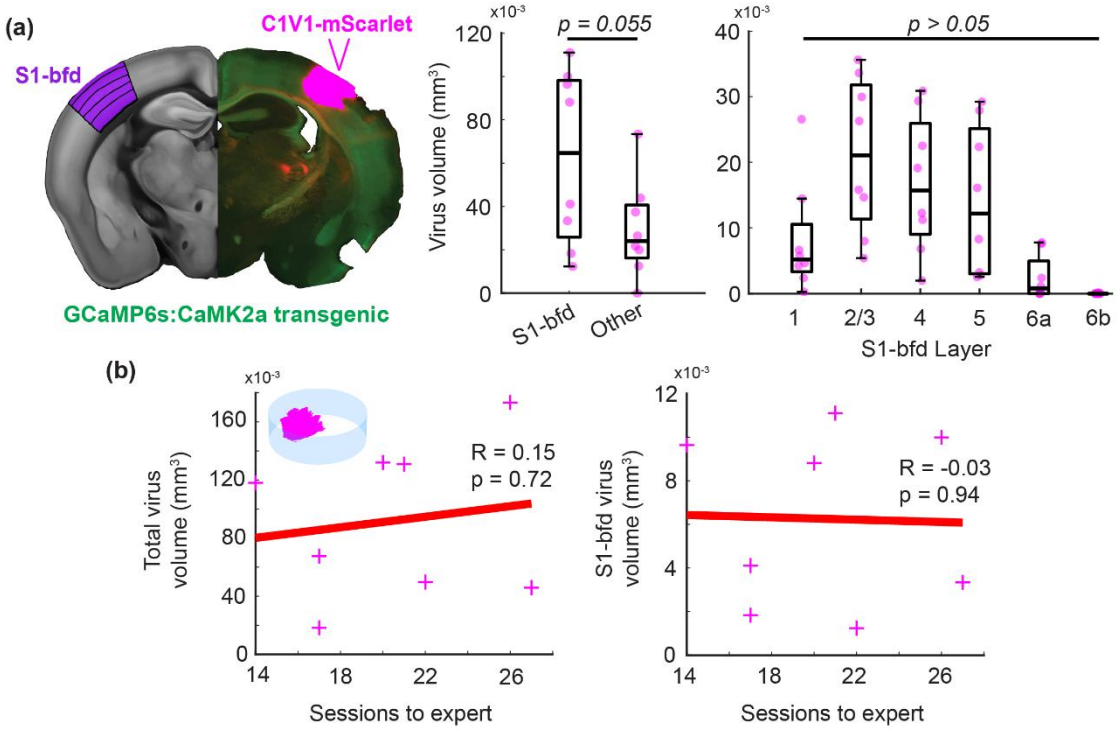
Therefore, differentiating between possibilities 5 and 6 could theoretically be achieved by tracing evoked activity through the network and using holographic optogenetics to recreate it at individual steps whilst selectively inhibiting the candidate NCC sequence(s). All-optical interrogation is uniquely positioned in this regard to determine whether a candidate NCC is both sufficient and necessary for producing a wS1 optogenetic percept, reconciling the dogmatic differences between Ned Hall's approaches to establishing causation in biological phenomena. For the field of perception, an understanding of the neuronal mechanisms underpinning both causal production and causal dependence would constitute a vast leap forward.

6

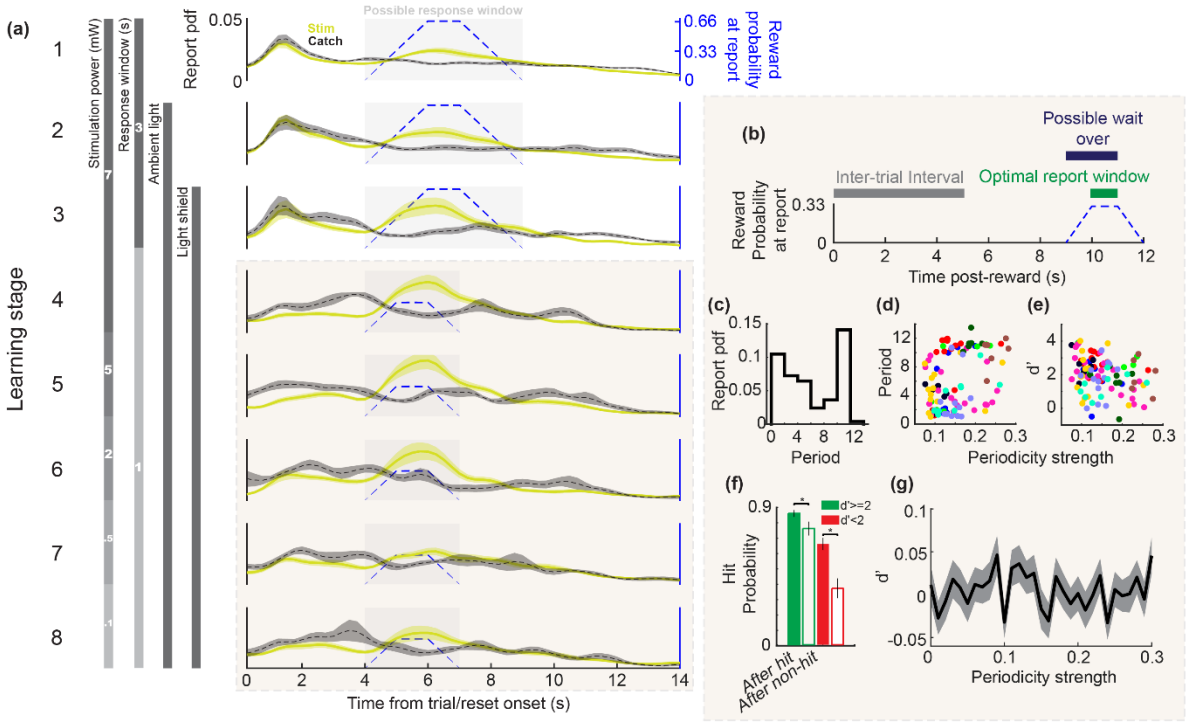
Supplementary figures



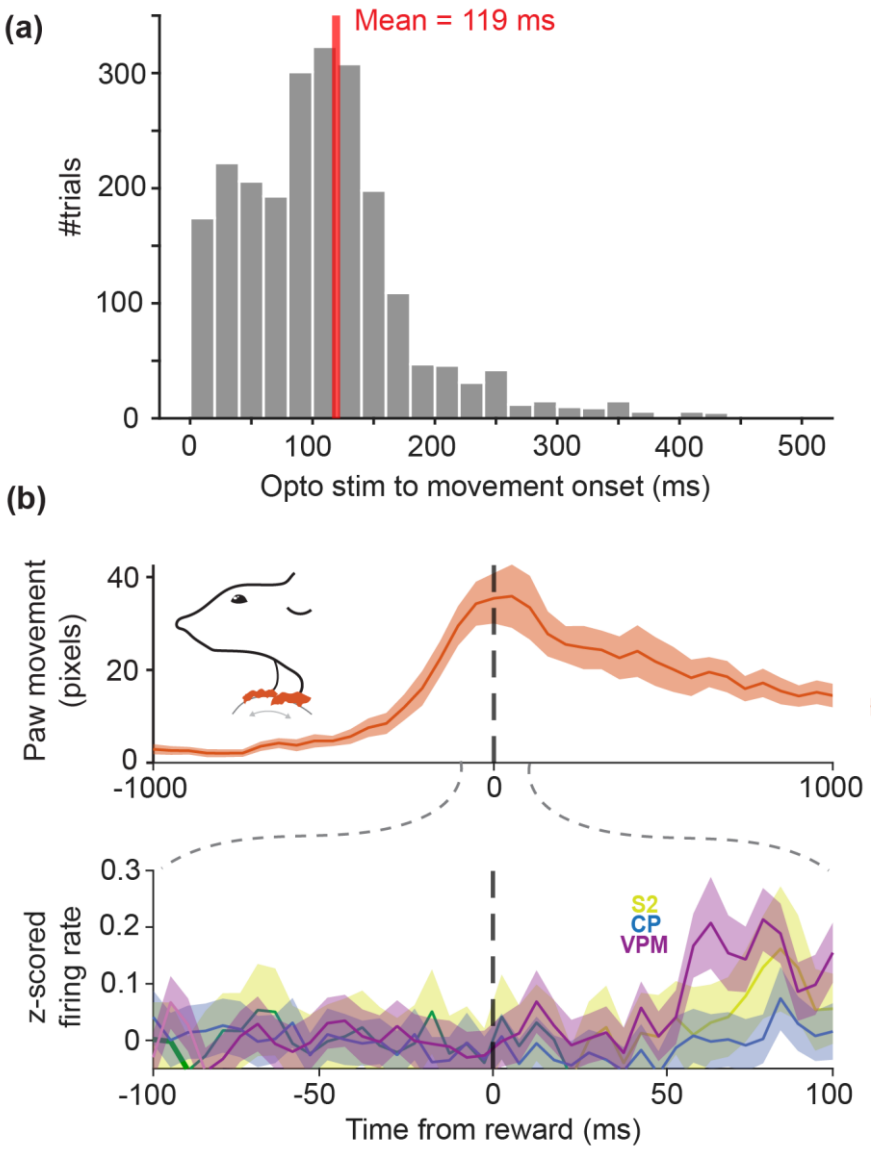
Supplementary figure 1. With a fixed pre-stim wait period and no catch trials, mice can achieve high rates on a CPD task independent of perception. (a) The opsin Chronos was injected into layer 2/3 of right wS1 of 3 wild type mice. **(b)** CPD task structure. **(c)** Once mice reached the behavioural training progression threshold at 10 mW power on sample, their performance did not deteriorate when stimulation power was decreased. **(d)** Task performance in expert mice did not significantly decrease when the optic fibre was moved in front of the eye, or the optogenetic stimulation light was prevented from reaching the brain using a kwik-cast plug, suggesting that the optogenetic stimulation was not used to guide behaviour.



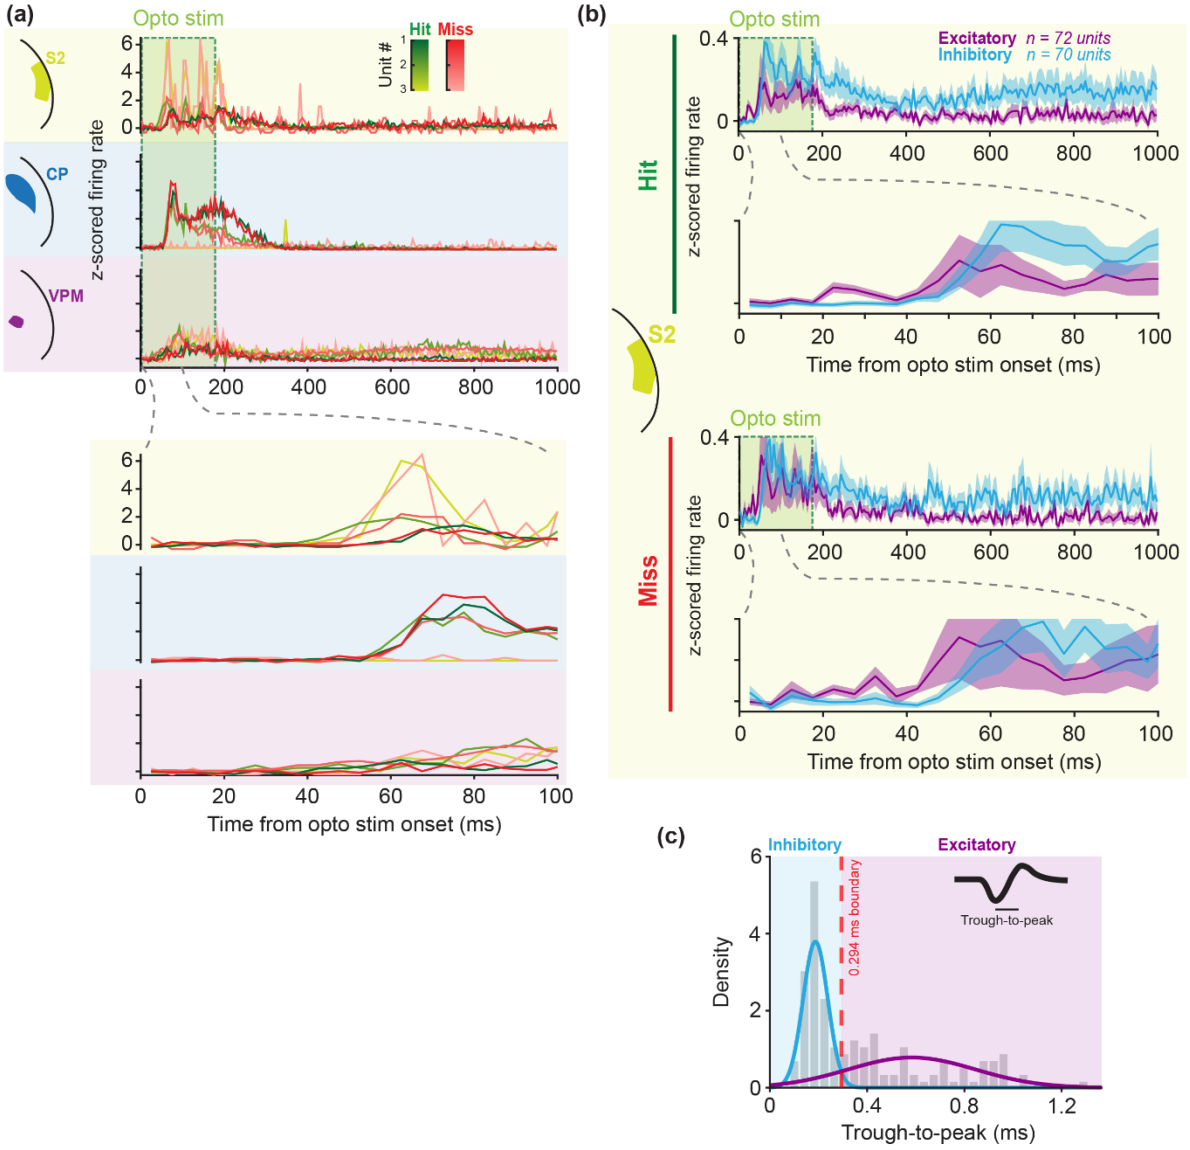
Supplementary figure 2. Number of sessions needed to reach expert performance on the optogenetic CPD task did not correlate with opsin expression. (a) The opsin C1V1 was primarily expressed in layer 2/3 of right wS1 of mice transgenic for GCaMP6s expression in excitatory neurons. However, there was not statistically significantly higher expression in layer 2/3 than other layers or S1-bfd and other regions (n = 9 mice). **(b)** The volume of opsin expression was reconstructed in 3D from histology using HERBS. There was no significant correlation between opsin expression and the number of sessions required to reach expert performance on the optogenetic CPD task, even when only virus expression directly under the cranial window was included (inset).



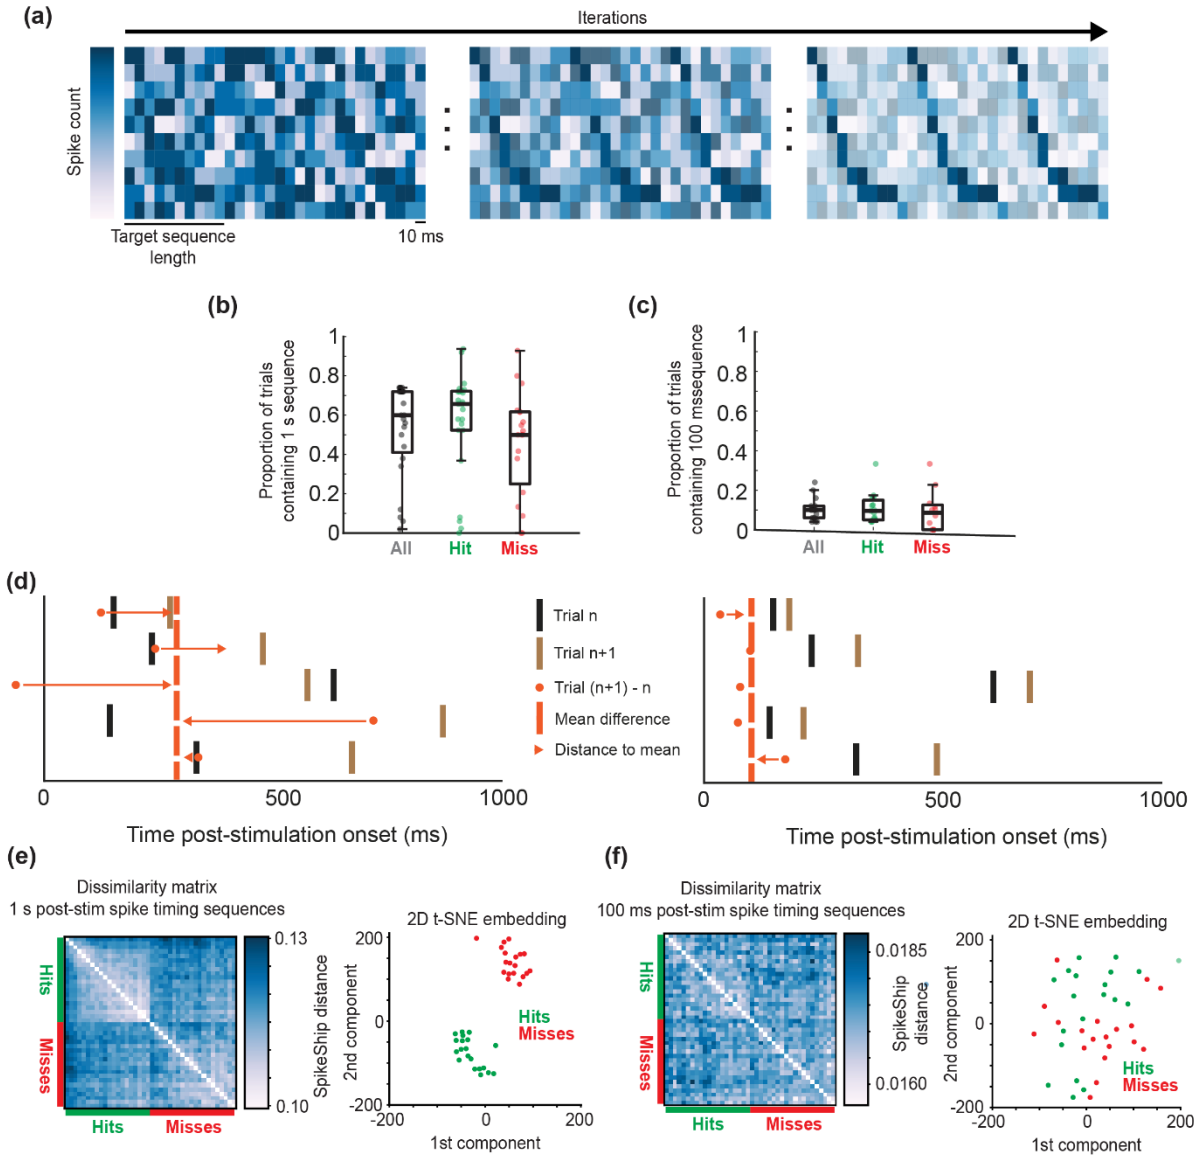
Supplementary figure 3. Expert mice did not exclusively use timing strategies to achieve high behavioural performance. (a) As mice progressed through the learning stages of the optogenetic CPD task, their strategy evolved from rapid wheel turns every ~1 to 2 seconds on both stim and catch trials to following the reward probability time course only on stim trials. This shows that by expert performance, the mice understand the structure of the task. Shaded area = ±SEM. **(b)** Due to the fixed inter-trial interval and uniformly drawn 4 to 6 second pre-stim wait period, reports within 10 to 11 seconds post-reward had a disproportionately high reward probability from learning stages 4 onwards. **(c-d)** The most common period of movements was between 10 to 11 seconds. **(e)** Periodic movements were more common in sessions with worse behavioural performance. **(f)** Hits were more likely to occur after hits vs non hits ($p < 0.05$ Wilcoxon signed-rank test), especially on low behavioural performance sessions, suggesting that the mice could have been using a timing strategy. **(g)** A simulated agent with no access to stimulus information could not achieve above baseline behavioural performance when exclusively utilising a timing strategy.



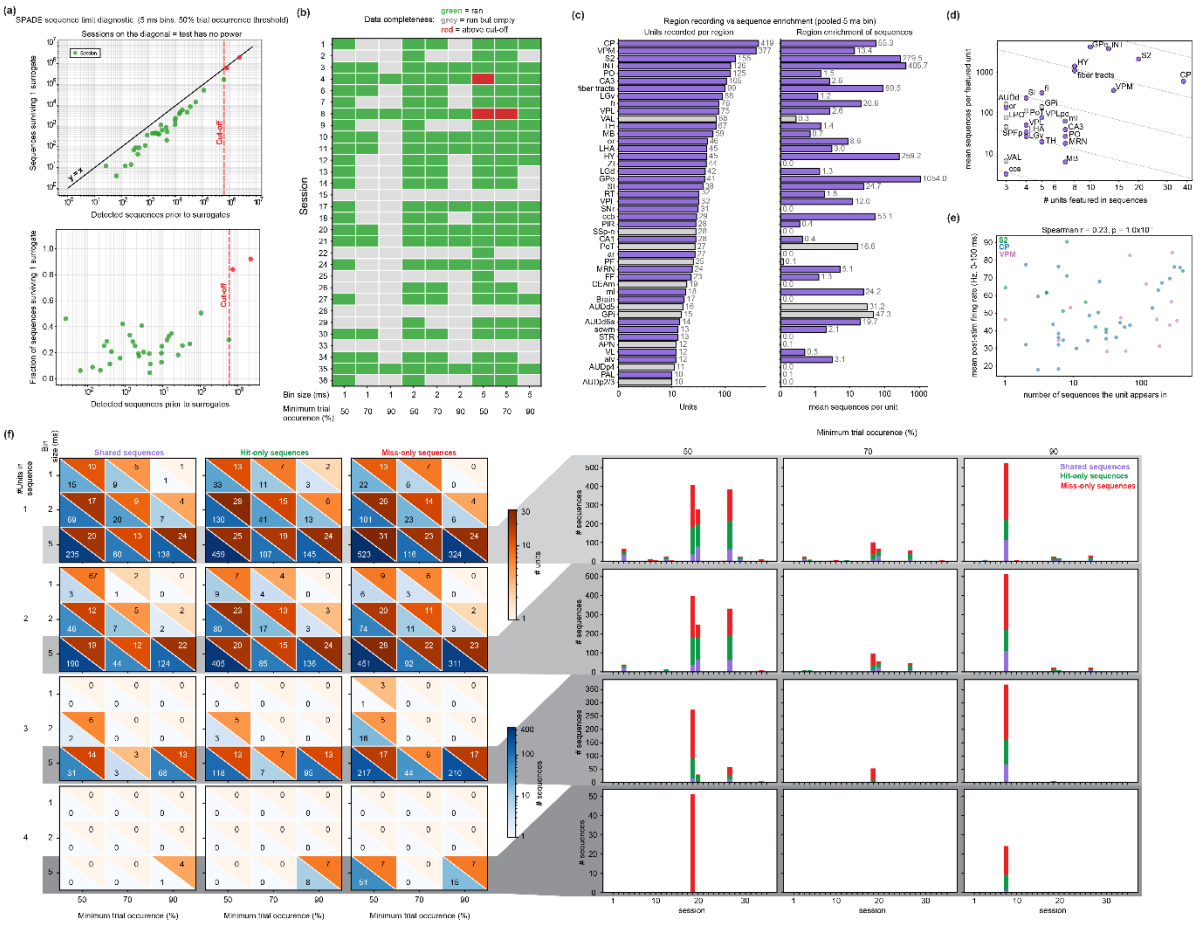
Supplementary figure 4. Expert mice react on average around 100 ms post-stimulation onset. **(a)** Histogram of reaction times computed as the time between optogenetic stimulation onset and paw movement reaching 2 * standard deviation of a pre-stimulation baseline period. **(b)** Upper: paw movement extracted using DeepLabCut and aligned to reward delivery. Lower: z-scored firing rate of S2, CP and VPM baselined to pre-reward delivery.



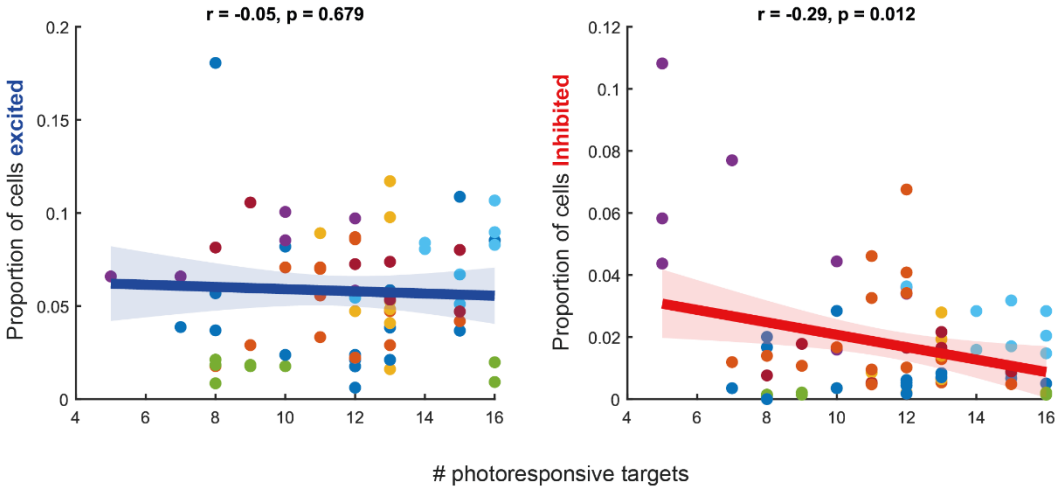
Supplementary figure 5. Putative cortical excitatory neurons in S2 respond faster to optogenetic stimulation of S1 than putative inhibitory units (a) Example single unit z-scored activity on hits and misses. **(b)** Putative excitatory neurons in S2 responded on average ~10 milliseconds faster than putative inhibitory neurons. **(c)** Units in S2 were divided into putative excitatory and inhibitory neurons based on the crossover point in an estimated bimodal distribution of trough-to-peak ratio.



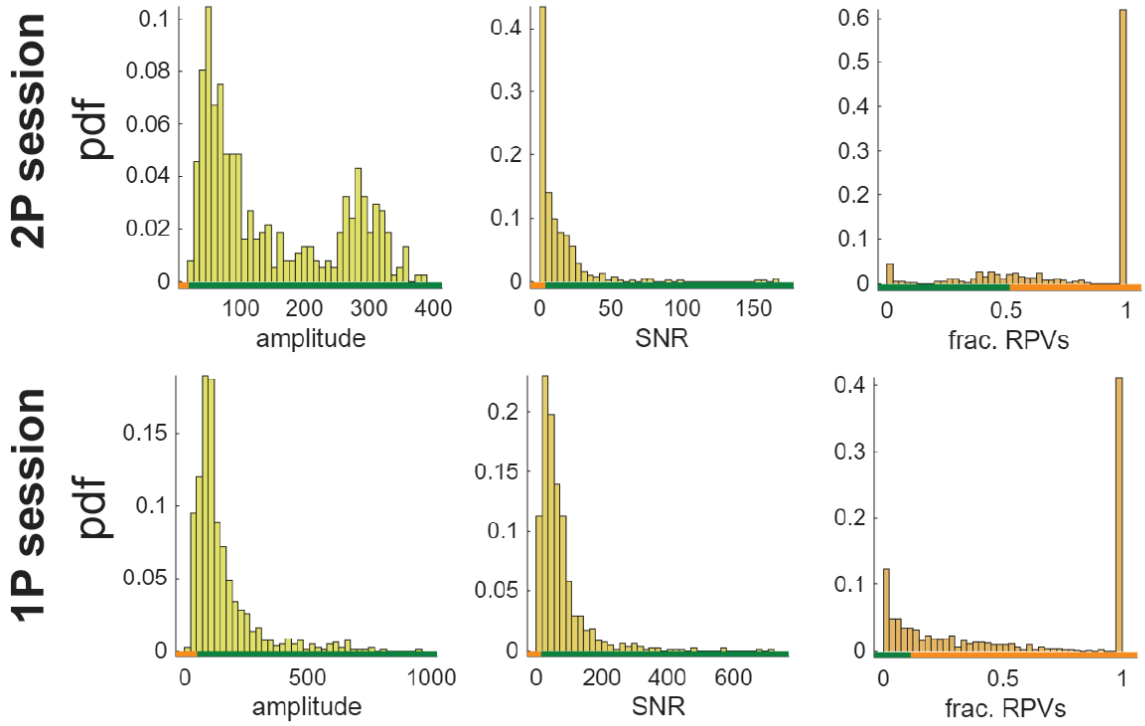
Supplementary figure 6. Popular neuronal sequence finding algorithms are unable to find reliable spike sequences within the 100 ms post-stimulation onset. (a) Schematic of SeqNMF's sequence detection algorithm using non-negative matrix factorisation. **(b)** Proportion of trials for each session where SeqNMF identified a 1 second activation sequence spanning all neurons post-stimulation. There was no significant difference in the proportion of hit trials and miss trials where SeqNMF identified a sequence. **(c)** Proportion of trials for each session where SeqNMF identified a 100 millisecond sequence spanning all neurons post-stimulation. SeqNMF only identified reliable sequences occurring on ~10% of trials. **(d)** Schematic of SpikeShip's sequence detection algorithm using earth-mover distance. **(e)** Left: Dissimilarity matrix computed on spiking activity in the 1 second post-stimulation onset for the first 20 hits and misses across sessions. Right: 2D t-SNE embedding using the dissimilarity matrix. SpikeShip can differentiate between hits and misses based on their global earth mover distance in the 1 second post-stimulation. **(f)** Same as **e** but for 100 milliseconds post-stimulation onset. SpikeShip cannot differentiate between hits and misses based on their global earth mover distance in the 100 ms post-stimulation.



Supplementary figure 7. SPADE parameter-space exploration. (a) When SPADE detects a large number of sequences, they are not effectively filtered when compared to a shuffled surrogate. Any sessions with initially identified sequences where >50% survived the first shuffle were excluded from further analysis. **(b)** Sessions were more likely to contain sequences as identified by SPADE when the bin size was increased and the minimum trial occurrence threshold was lowered. **(c)** Left: all regions recorded from ordered by the number units. Right: same ordering but displaying mean number of sequences that a single unit appears in per region. **(d)** Regions with more units taking part in sequences also contain units that appear in more sequences on average. **(e)** There is no correlation between the number of sequences that a unit appears in and its average firing rate. **(f)** Left: exploration of SPADE parameter-space. More sequences were identified with larger bin sizes and less stringent constraints on the minimum occurrence of trials and minimum number of featured neurons. Right: session by session contribution to hit-only, miss-only and shared sequences.



Supplementary figure 8. There is no correlation between the number of successfully stimulated photoresponsive targets and the proportion of excited or inhibited cells across the field of view.



Supplementary figure 9. Bombcell outputs for a representative two-photon behavioural session (top) and a one-photon behavioural session (bottom). The mean amplitude and signal-to-noise ratio (SNR) of detected units was lower in two-photon sessions, likely because of a combination of the photoelectric artefact and common-reference averaging of a noisy signal.

7

References

1. Mayer, M. L., Westbrook, G. L. & Guthrie, P. B. Voltage-dependent block by Mg2+ of NMDA responses in spinal cord neurones. *Nature* **309**, 261–263 (1984).
2. Dash, P. K., Hochner, B. & Kandel, E. R. Injection of the cAMP-responsive element into the nucleus of Aplysia sensory neurons blocks long-term facilitation. *Nature* **345**, 718–721 (1990).
3. Miller, C. A. & Sweatt, J. D. Covalent Modification of DNA Regulates Memory Formation. *Neuron* **53**, 857–869 (2007).
4. Hodgkin, A. L. & Huxley, A. F. A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of Physiology* **117**, 500–544 (1952).
5. Cajal, S. R. y. *Estructura de los centros nerviosos de las aves (1888)*. (Jiménez y Molina, 1924).

6. Virchow, R. *Gesammelte Abhandlungen zur wissenschaftlichen Medicin*. (Grote, 1862).
7. Deiters, O. & Schultze, M. J. S. *Untersuchungen über Gehirn und Rückenmark des Menschen und der Säugethiere*. (F. Veiweg, 1865).
8. GOLGI, C. Sulla sostanza grigia del cervello. *Gazetta Medica Italiana* **33**, 244–246 (1873).
9. Yao, Z. *et al.* A high-resolution transcriptomic and spatial atlas of cell types in the whole mouse brain. *Nature* **624**, 317–332 (2023).
10. Schlegel, P. *et al.* Whole-brain annotation and multi-connectome cell typing of Drosophila. *Nature* **634**, 139–152 (2024).
11. Hagmann, P. *et al.* Mapping the Structural Core of Human Cerebral Cortex. *PLOS Biology* **6**, e159 (2008).
12. Oh, S. W. *et al.* A mesoscale connectome of the mouse brain. *Nature* **508**, 207–214 (2014).
13. Bae, J. A. *et al.* Functional connectomics spanning multiple areas of mouse visual cortex. *Nature* **640**, 435–447 (2025).
14. Mountcastle, V. B. Modality and topographic properties of single neurons of cat's somatic sensory cortex. *Journal of Neurophysiology* **20**, 408–434 (1957).
15. Hubel, D. H. & Wiesel, T. N. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J Physiol* **160**, 106-154.2 (1962).

16. Simons, D. J. Response properties of vibrissa units in rat SI somatosensory neocortex. *Journal of Neurophysiology* **41**, 798–820 (1978).

17. PENFIELD, W. & BOLDREY, E. SOMATIC MOTOR AND SENSORY REPRESENTATION IN THE CEREBRAL CORTEX OF MAN AS STUDIED BY ELECTRICAL STIMULATION1. *Brain* **60**, 389–443 (1937).

18. Merzenich, M. M., Knight, P. L. & Roth, G. L. Representation of cochlea within primary auditory cortex in the cat. *Journal of Neurophysiology* **38**, 231–249 (1975).

19. O'Keefe, J. & Dostrovsky, J. The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Research* **34**, 171–175 (1971).

20. Hafting, T., Fyhn, M., Molden, S., Moser, M.-B. & Moser, E. I. Microstructure of a spatial map in the entorhinal cortex. *Nature* **436**, 801–806 (2005).

21. Anand, B. K. & Brobeck, J. R. Hypothalamic Control of Food Intake in Rats and Cats. *Yale J Biol Med* **24**, 123–140 (1951).

22. McGinley, M. J. *et al.* Waking State: Rapid Variations Modulate Neural and Behavioral Responses. *Neuron* **87**, 1143–1161 (2015).

23. Schultz, W., Dayan, P. & Montague, P. R. A Neural Substrate of Prediction and Reward. *Science* **275**, 1593–1599 (1997).

24. Kepecs, A., Uchida, N., Zariwala, H. A. & Mainen, Z. F. Neural correlates, computation and behavioural impact of decision confidence. *Nature* **455**, 227–231 (2008).

25. Neher, E. & Sakmann, B. Single-channel currents recorded from membrane of denervated frog muscle fibres. *Nature* **260**, 799–802 (1976).

26. Jun, J. J. *et al.* Fully integrated silicon probes for high-density recording of neural activity. *Nature* **551**, 232–236 (2017).

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

28. Denk, W., Strickler, J. H. & Webb, W. W. Two-Photon Laser Scanning Fluorescence Microscopy. *Science* **248**, 73–76 (1990).

29. Ogawa, S., Lee, T. M., Kay, A. R. & Tank, D. W. Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proceedings of the National Academy of Sciences* **87**, 9868–9872 (1990).

30. Yizhar, O. *et al.* Neocortical excitation/inhibition balance in information processing and social dysfunction. *Nature* **477**, 171–178 (2011).

31. Homayoun, H. & Moghaddam, B. NMDA Receptor Hypofunction Produces Opposite Effects on Prefrontal Cortex Interneurons and Pyramidal Neurons. *J. Neurosci.* **27**, 11496–11500 (2007).

32. Coughlin, B. *et al.* Modified Neuropixels probes for recording human neurophysiology in the operating room. *Nat Protoc* **18**, 2927–2953 (2023).

33. Busche, M. A. *et al.* Clusters of Hyperactive Neurons Near Amyloid Plaques in a Mouse Model of Alzheimer's Disease. *Science* **321**, 1686–1689 (2008).

34. Greicius, M. D., Srivastava, G., Reiss, A. L. & Menon, V. Default-mode network activity distinguishes Alzheimer's disease from healthy aging: Evidence from functional MRI. *Proceedings of the National Academy of Sciences* **101**, 4637–4642 (2004).

35. Descartes, R. *Principia philosophiae*. (J. Jansonius, 1656).

36. Plato. *Phaedrus*. (OUP Oxford, 2002).

37. Aristotle. *De Anima (On the Soul)*. (Penguin UK, 1986).

38. Chalmers, D. J. Facing up to the problem of consciousness. *Journal of Consciousness Studies* **2**, 200–219 (1995).

39. Edlow, B. L., Claassen, J., Schiff, N. D. & Greer, D. M. Recovery from disorders of consciousness: mechanisms, prognosis and emerging therapies. *Nat Rev Neurol* **17**, 135–156 (2021).

40. Elmer, J. *et al.* Association of early withdrawal of life-sustaining therapy for perceived neurological prognosis with mortality after cardiac arrest. *Resuscitation* **102**, 127–135 (2016).

41. Jung, C. G. *On the Nature of the Psyche*. (Routledge, 2014).

42. Baars, B. J. *A Cognitive Theory of Consciousness*. (Cambridge University Press, 1993).

43. Tononi, G. An information integration theory of consciousness. *BMC Neurosci* **5**, 42 (2004).

44. Nagel, T. What Is It Like to Be a Bat? *The Philosophical Review* **83**, 435–450 (1974).

45. Solms, M. & Friston, K. How and why consciousness arises: Some considerations from physics and physiology. *Journal of Consciousness Studies* **25**, 202–238 (2018).

46. Dehaene, S. & Naccache, L. Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition* **79**, 1–37 (2001).

47. Dehaene, S. & Changeux, J.-P. Experimental and Theoretical Approaches to Conscious Processing. *Neuron* **70**, 200–227 (2011).

48. Tononi, G., Boly, M., Massimini, M. & Koch, C. Integrated information theory: from consciousness to its physical substrate. *Nat Rev Neurosci* **17**, 450–461 (2016).

49. Miesenböck, G. & Kevrekidis, I. G. OPTICAL IMAGING AND CONTROL OF GENETICALLY DESIGNATED NEURONS IN FUNCTIONING CIRCUITS. *Annual Review of Neuroscience* **28**, 533–563 (2005).

50. Bellman, R. *Some Vistas of Modern Mathematics: Dynamic Programming, Invariant Imbedding, and the Mathematical Biosciences*. (University Press of Kentucky, 1968).

51. Damasio, H., Grabowski, T., Frank, R., Galaburda, A. M. & Damasio, A. R. The Return of Phineas Gage: Clues About the Brain from the Skull of a Famous Patient. *Science* **264**, 1102–1105 (1994).

52. Scoville, W. B. & Milner, B. Loss of Recent Memory After Bilateral Hippocampal Lesions. *Journal of Neurology, Neurosurgery & Psychiatry* **20**, 11–21 (1957).

53. Siddiqi, S. H., Kording, K. P., Parvizi, J. & Fox, M. D. Causal mapping of human brain function. *Nat Rev Neurosci* **23**, 361–375 (2022).

54. Etkin, A. Addressing the Causality Gap in Human Psychiatric Neuroscience. *JAMA Psychiatry* **75**, 3–4 (2018).

55. Marinescu, I. E., Lawlor, P. N. & Kording, K. P. Quasi-experimental causality in neuroscience and behavioural research. *Nat Hum Behav* **2**, 891–898 (2018).

56. Barack, D. L. *et al.* A call for more clarity around causality in neuroscience. *Trends in Neurosciences* **45**, 654–655 (2022).

57. Kim, J., Korman, D. Z. & Sosa, E. *Metaphysics: An Anthology*. (John Wiley & Sons, 2011).

58. Spiegel, D. Modulation of Nociception in Multiple Brain Systems—The Strain in Pain. *JAMA Neurol* **75**, 1309–1310 (2018).

59. Rizzi, D. A. & Pedersen, S. A. Causality in medicine: Towards a theory and terminology. *Theor Med Bioeth* **13**, 233–254 (1992).

60. Bernat, J. L. Chronic disorders of consciousness. *The Lancet* **367**, 1181–1192 (2006).

61. Parvizi, J. & Damasio, A. Consciousness and the brainstem. *Cognition* **79**, 135–160 (2001).

62. Schiff, N. D. *et al.* Behavioural improvements with thalamic stimulation after severe traumatic brain injury. *Nature* **448**, 600–603 (2007).

63. Redinbaugh, M. J. *et al.* Thalamus Modulates Consciousness via Layer-Specific Control of Cortex. *Neuron* **106**, 66-75.e12 (2020).

64. Alkire, M. T., Hudetz, A. G. & Tononi, G. Consciousness and Anesthesia. *Science* **322**, 876–880 (2008).

65. General anesthesia and the neural correlates of consciousness. in *Progress in Brain Research* vol. 150 229–597 (Elsevier, 2005).

66. Porzionato, A., Manani, G. & Montisci, M. Medical litigation as a result of awareness during general anaesthesia. *Med Sci Law* **47**, 265–268 (2007).

67. Takechi, A., Inoue, S. & Kawaguchi, M. Anaesthesia management by residents does not alter the incidence of self-reported anaesthesia awareness: A teaching hospital-based propensity score analysis. *Anaesthesia Critical Care & Pain Medicine* **38**, 63–68 (2019).

68. Gambús, P. L. & Trocóniz, I. F. Pharmacokinetic–pharmacodynamic modelling in anaesthesia. *British Journal of Clinical Pharmacology* **79**, 72–84 (2015).

69. Tasbihgou, S. R., Vogels, M. F. & Absalom, A. R. Accidental awareness during general anaesthesia – a narrative review. *Anaesthesia* **73**, 112–122 (2018).

70. General anesthesia and the neural correlates of consciousness. in *Progress in Brain Research* vol. 150 229–597 (Elsevier, 2005).

71. Långsjö, J. W. *et al.* Effects of subanesthetic ketamine on regional cerebral glucose metabolism in humans. *Anesthesiology* **100**, 1065–1071 (2004).

72. Crick, F. & Koch, C. A framework for consciousness. *Nat Neurosci* **6**, 119–126 (2003).

73. Seth, A. K. Consciousness: The last 50 years (and the next). *Brain and Neuroscience Advances* **2**, 2398212818816019 (2018).

74. Verbeek, T. & Bos, E.-J. Treatise on Man. in *The Cambridge Descartes Lexicon* (ed. Nolan, L.) 725–727 (Cambridge University Press, Cambridge, 2015). doi:10.1017/CBO9780511894695.247.

75. Campbell, K. *Body and Mind*. (Macmillan Education UK, London, 1970). doi:10.1007/978-1-349-00678-6.

76. KIRK, R. Sentience and Behaviour. *Mind* **LXXXIII**, 43–60 (1974).

77. Mooney, C. M. Age in the development of closure ability in children. *Canadian Journal of Psychology / Revue canadienne de psychologie* **11**, 219–226 (1957).

78. Hasson, U., Hendler, T., Bashat, D. B. & Malach, R. Vase or Face? A Neural Correlate of Shape-Selective Grouping Processes in the Human Brain. *J Cogn Neurosci* **13**, 744–753 (2001).

79. Wheatstone, C. XVIII. Contributions to the physiology of vision. —Part the first. On some remarkable, and hitherto unobserved, phenomena of binocular vision. *Phil. Trans. R. Soc.* 371–394 (1838) doi:10.1098/rstl.1838.0019.

80. Tong, F. & Engel, S. A. Interocular rivalry revealed in the human cortical blind-spot representation. *Nature* **411**, 195–199 (2001).

81. Tong, F., Nakayama, K., Vaughan, J. T. & Kanwisher, N. Binocular Rivalry and Visual Awareness in Human Extrastriate Cortex. *Neuron* **21**, 753–759 (1998).

82. Wang, X. *et al.* Category Selectivity of Human Visual Cortex in Perception of Rubin Face–Vase Illusion. *Front. Psychol.* **8**, (2017).

83. Lumer, E. D., Friston, K. J. & Rees, G. Neural Correlates of Perceptual Rivalry in the Human Brain. *Science* **280**, 1930–1934 (1998).

84. Sterzer, P. & Kleinschmidt, A. A neural basis for inference in perceptual ambiguity. *Proceedings of the National Academy of Sciences* **104**, 323–328 (2007).

85. Imamoglu, F., Kahnt, T., Koch, C. & Haynes, J.-D. Changes in functional connectivity support conscious object recognition. *NeuroImage* **63**, 1909–1917 (2012).

86. Boly, M. *et al.* Preserved Feedforward But Impaired Top-Down Processes in the Vegetative State. *Science* **332**, 858–862 (2011).

87. BRICKNER, R. M. BRAIN OF PATIENT A. AFTER BILATERAL FRONTAL LOBECTOMY; STATUS OF FRONTAL-LOBE PROBLEM. *AMA Arch NeurPsych* **68**, 293–313 (1952).

88. Hebb, D. O. & Penfield, W. HUMAN BEHAVIOR AFTER EXTENSIVE BILATERAL REMOVAL FROM THE FRONTAL LOBES. *Arch NeurPsych* **44**, 421–438 (1940).

89. *Selective Partial Ablation of the Frontal Cortex: A Correlative Study of Its Effects on Human Psychotic Subjects*. xiv, 517 (Paul B. Hoeber, Inc., Oxford, England, 1949).

90. Markowitsch, H. J. & Kessler, J. Massive impairment in executive functions with partial preservation of other cognitive functions: the case of a young patient with severe degeneration of the prefrontal cortex. *Exp Brain Res* **133**, 94–102 (2000).

91. Frässle, S., Sommer, J., Jansen, A., Naber, M. & Einhäuser, W. Binocular Rivalry: Frontal Activity Relates to Introspection and Action But Not to Perception. *J. Neurosci.* **34**, 1738–1747 (2014).

92. Riggall, A. C. & Postle, B. R. The Relationship between Working Memory Storage and Elevated Activity as Measured with Functional Magnetic Resonance Imaging. *J. Neurosci.* **32**, 12990–12998 (2012).

93. Bodamer, J. Die Prosop-Agnosie. *Arch. f. Psychiatr. u. Z. Neur.* **179**, 6–53 (1947).

94. Damasio, A. R., Tranel, D. & Damasio, H. Face Agnosia and the Neural Substrates of Memory. *Annual Review of Neuroscience* **13**, 89–109 (1990).

95. Farah, M. J. *Visual Agnosia*. (MIT Press, Cambridge, MA, USA, 2004).

96. WEISKRANTZ, L., WARRINGTON, E. K., SANDERS, M. D. & MARSHALL, J. VISUAL CAPACITY IN THE HEMIANOPIC FIELD FOLLOWING A RESTRICTED OCCIPITAL ABLATION. *Brain* **97**, 709–728 (1974).

97. Stoerig, P. & Cowey, A. Blindsight in man and monkey. *Brain* **120**, 535–559 (1997).

98. ZEKI, S. A CENTURY OF CEREBRAL ACHROMATOPSIA. *Brain* **113**, 1721–1777 (1990).

99. Critchley, M. Clinical Investigation of Disease of the Parietal Lobes of the Brain. *Medical Clinics of North America* **46**, 837–857 (1962).

100. Fujii, T. *et al.* Auditory Sound Agnosia without Aphasia Following a Right Temporal Lobe Lesion. *Cortex* **26**, 263–268 (1990).

101. Adolphs, R. The unsolved problems of neuroscience. *Trends in Cognitive Sciences* **19**, 173–175 (2015).

102. Cohen, M. A. & Dennett, D. C. Consciousness cannot be separated from function. *Trends in Cognitive Sciences* **15**, 358–364 (2011).

103. Aserinsky, E. & Kleitman, N. Regularly Occurring Periods of Eye Motility, and Concomitant Phenomena, During Sleep. *Science* **118**, 273–274 (1953).

104. Nir, Y. & Tononi, G. Dreaming and the brain: from phenomenology to neurophysiology. *Trends in Cognitive Sciences* **14**, 88–100 (2010).

105. Stickgold, R., Malia, A., Fosse, R., Propper, R. & Hobson, J. A. Brain-Mind States: I. Longitudinal Field Study of Sleep/Wake Factors Influencing Mentation Report Length. *Sleep* **24**, 171–179 (2001).

106. Siclari, F. *et al.* The neural correlates of dreaming. *Nat Neurosci* **20**, 872–878 (2017).

107. Koch, C., Massimini, M., Boly, M. & Tononi, G. Neural correlates of consciousness: progress and problems. *Nat Rev Neurosci* **17**, 307–321 (2016).

108. Noy, N. *et al.* Ignition's glow: Ultra-fast spread of global cortical activity accompanying local "ignitions" in visual cortex during conscious visual perception. *Consciousness and Cognition* **35**, 206–224 (2015).

109. Koch, C. What Is Consciousness? *Nature* **557**, S8–S12 (2018).

110. Jiang, Y., Zhou, K. & He, S. Human visual cortex responds to invisible chromatic flicker. *Nat Neurosci* **10**, 657–662 (2007).

111. He, S. & MacLeod, D. I. A. Orientation-selective adaptation and tilt after-effect from invisible patterns. *Nature* **411**, 473–476 (2001).

112. Haynes, J.-D. & Rees, G. Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nat Neurosci* **8**, 686–691 (2005).

113. Tang, H. *et al.* Spatiotemporal Dynamics Underlying Object Completion in Human Ventral Visual Cortex. *Neuron* **83**, 736–748 (2014).

114. Lamme, V. A. F. & Roelfsema, P. R. The distinct modes of vision offered by feedforward and recurrent processing. *Trends in Neurosciences* **23**, 571–579 (2000).

115. Blankenburg, F., Ruff, C. C., Deichmann, R., Rees, G. & Driver, J. The Cutaneous Rabbit Illusion Affects Human Primary Sensory Cortex Somatotopically. *PLOS Biology* **4**, e69 (2006).

116. Schaefer, M., Noennig, N., Heinze, H.-J. & Rotte, M. Fooling your feelings: Artificially induced referred sensations are linked to a modulation of the primary somatosensory cortex. *NeuroImage* **29**, 67–73 (2006).

117. Meyer, K., Kaplan, J. T., Essex, R., Damasio, H. & Damasio, A. Seeing Touch Is Correlated with Content-Specific Activity in Primary Somatosensory Cortex. *Cereb Cortex* **21**, 2113–2121 (2011).

118. Yoo, S.-S., Freeman, D. K., McCarthy, J. J. I. & Jolesz, F. A. Neural substrates of tactile imagery: a functional MRI study. *NeuroReport* **14**, 581 (2003).

119. Wiegand, K. & Gutschalk, A. Correlates of perceptual awareness in human primary auditory cortex revealed by an informational masking experiment. *NeuroImage* **61**, 62–69 (2012).

120. Green, D. M. & Swets, J. A. *Signal Detection Theory and Psychophysics*. xi, 455 (John Wiley, Oxford, England, 1966).

121. Skinner, B. F. *The Behavior of Organisms: An Experimental Analysis*. 457 (Appleton-Century, Oxford, England, 1938).

122. Pavlov, I. P. *Conditioned Reflexes: An Investigation of the Physiological Activity of the Cerebral Cortex*. xv, 430 (Oxford Univ. Press, Oxford, England, 1927).

123. Busse, L. *et al.* The Detection of Visual Contrast in the Behaving Mouse. *J. Neurosci.* **31**, 11351–11361 (2011).

124. Kurt, S. & Ehret, G. Auditory discrimination learning and knowledge transfer in mice depends on task difficulty. *Proceedings of the National Academy of Sciences* **107**, 8481–8485 (2010).

125. Wu, H.-P. P., Ioffe, J. C., Iverson, M. M., Boon, J. M. & Dyck, R. H. Novel, whisker-dependent texture discrimination task for mice. *Behavioural Brain Research* **237**, 238–242 (2013).

126. de Lafuente, V. & Romo, R. Neuronal correlates of subjective sensory experience. *Nat Neurosci* **8**, 1698–1703 (2005).

127. Brody, C. D., Hernández, A., Zainos, A., Lemus, L. & Romo, R. Analysing neuronal correlates of the comparison of two sequentially

presented sensory stimuli. *Philos Trans R Soc Lond B Biol Sci* **357**, 1843–1850 (2002).

128. de Lafuente, V. & Romo, R. Neural correlate of subjective sensory experience gradually builds up across cortical areas. *Proceedings of the National Academy of Sciences* **103**, 14266–14271 (2006).

129. Chen, J. L., Voigt, F. F., Javadzadeh, M., Krueppel, R. & Helmchen, F. Long-range population dynamics of anatomically defined neocortical networks. *eLife* **5**, e14679 (2016).

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

131. Andersen, R. A. & Mountcastle, V. B. The influence of the angle of gaze upon the excitability of the light- sensitive neurons of the posterior parietal cortex. *J. Neurosci.* **3**, 532–548 (1983).

132. Brotchie, P. R., Andersen, R. A., Snyder, L. H. & Goodman, S. J. Head position signals used by parietal neurons to encode locations of visual stimuli. *Nature* **375**, 232–235 (1995).

133. Mountcastle, V. B., Lynch, J. C., Georgopoulos, A., Sakata, H. & Acuna, C. Posterior parietal association cortex of the monkey: command functions for operations within extrapersonal space. *Journal of Neurophysiology* **38**, 871–908 (1975).

134. Musall, S., Kaufman, M. T., Juavinett, A. L., Gluf, S. & Churchland, A. K. Single-trial neural dynamics are dominated by richly varied movements. *Nat Neurosci* **22**, 1677–1686 (2019).

135. Steinmetz, M. A., Connor, C. E., Constantinidis, C. & McLaughlin, J. R. Covert attention suppresses neuronal responses in area 7a of the posterior parietal cortex. *Journal of Neurophysiology* **72**, 1020–1023 (1994).

136. Poulet, J. F. A. & Petersen, C. C. H. Internal brain state regulates membrane potential synchrony in barrel cortex of behaving mice. *Nature* **454**, 881–885 (2008).

137. Pancholi, R., Ryan, L. & Peron, S. Learning in a sensory cortical microstimulation task is associated with elevated representational stability. *Nat Commun* **14**, 3860 (2023).

138. Penfield, W. Activation of the Record of Human Experience. *Ann R Coll Surg Engl* **29**, 77–84 (1961).

139. Clark, K. L., Armstrong, K. M. & Moore, T. Probing neural circuitry and function with electrical microstimulation. *Proc Biol Sci* **278**, 1121–1130 (2011).

140. Histed, M. H., Ni, A. M. & Maunsell, J. H. R. Insights into cortical mechanisms of behavior from microstimulation experiments. *Progress in Neurobiology* **103**, 115–130 (2013).

141. Afraz, A. Behavioral optogenetics in nonhuman primates; a psychological perspective. *Current Research in Neurobiology* **5**, 100101 (2023).

142. Magiorkinis, E., Diamantis, A., Sidiropoulou, K. & Panteliadis, C. Highights in the History of Epilepsy: The Last 200 Years. *Epilepsy Research and Treatment* **2014**, 582039 (2014).

143. PENFIELD, W. EPILEPSY AND SURGICAL THERAPY. *Arch NeurPsych* **36**, 449–484 (1936).

144. Ladino, L. D., Rizvi, S. & Téllez-Zenteno, J. F. The Montreal procedure: The legacy of the great Wilder Penfield. *Epilepsy & Behavior* **83**, 151–161 (2018).

145. Romo, R., Hernández, A., Zainos, A. & Salinas, E. Somatosensory discrimination based on cortical microstimulation. *Nature* **392**, 387–390 (1998).

146. Romo, R., Hernández, A., Zainos, A., Brody, C. D. & Lemus, L. Sensing without Touching: Psychophysical Performance Based on Cortical Microstimulation. *Neuron* **26**, 273–278 (2000).

147. DeYoe, E. A., Lewine, J. D. & Doty, R. W. Laminar Variation in Threshold for Detection of Electrical Excitation of Striate Cortex by Macaques. *Journal of Neurophysiology* **94**, 3443–3450 (2005).

148. O'Doherty, J. E., Lebedev, M., Hanson, T. L., Fitzsimmons, N. & Nicolelis, M. A. L. A brain-machine interface instructed by direct intracortical microstimulation. *Front. Integr. Neurosci.* **3**, (2009).

149. Fox, K. C. R. *et al.* Intrinsic network architecture predicts the effects elicited by intracranial electrical stimulation of the human brain. *Nat Hum Behav* **4**, 1039–1052 (2020).

150. Houweling, A. R. & Brecht, M. Behavioural report of single neuron stimulation in somatosensory cortex. *Nature* **451**, 65–68 (2008).

151. Leal-Campanario, R., Delgado-García, J. M. & Gruart, A. Microstimulation of the somatosensory cortex can substitute for vibrissa stimulation during Pavlovian conditioning. *Proceedings of the National Academy of Sciences* **103**, 10052–10057 (2006).

152. Intracortical microstimulation of human somatosensory cortex. *Science Translational Medicine* https://www.science.org/doi/10.1126/scitranslmed.aaf8083.

153. Armenta Salas, M. *et al.* Proprioceptive and cutaneous sensations in humans elicited by intracortical microstimulation. *eLife* **7**, e32904 (2018).

154. Flesher, S. N. *et al.* A brain-computer interface that evokes tactile sensations improves robotic arm control. *Science* **372**, 831–836 (2021).

155. PENFIELD, W. & BOLDREY, E. SOMATIC MOTOR AND SENSORY REPRESENTATION IN THE CEREBRAL CORTEX OF MAN AS STUDIED BY ELECTRICAL STIMULATION1. *Brain* **60**, 389–443 (1937).

156. Histed, M. H., Bonin, V. & Reid, R. C. Direct Activation of Sparse, Distributed Populations of Cortical Neurons by Electrical Microstimulation. *Neuron* **63**, 508–522 (2009).

157. Nowak, L. G. & Bullier, J. Axons, but not cell bodies, are activated by electrical stimulation in cortical gray matterI. Evidence from chronaxie measurements. *Exp Brain Res* **118**, 477–488 (1998).

158. Nowak, L. G. & Bullier, J. Axons, but not cell bodies, are activated by electrical stimulation in cortical gray matterII. Evidence from selective inactivation of cell bodies and axon initial segments. *Exp Brain Res* **118**, 489–500 (1998).

159. Rattay, F. The basic mechanism for the electrical stimulation of the nervous system. *Neuroscience* **89**, 335–346 (1999).

160. Hubel, D. H. & Wiesel, T. N. Ferrier lecture - Functional architecture of macaque monkey visual cortex. *Proc Biol Sci* **198**, 1–59 (1977).

161. Mountcastle, V. B. Modality and topographic properties of single neurons of cat's somatic sensory cortex. *Journal of Neurophysiology* **20**, 408–434 (1957).

162. Schubert, D. *et al.* Layer-Specific Intracolumnar and Transcolumnar Functional Connectivity of Layer V Pyramidal Cells in Rat Barrel Cortex. *J. Neurosci.* **21**, 3580–3592 (2001).

163. Kim, U. & Ebner, F. F. Barrels and septa: Separate circuits in rat barrel field cortex. *Journal of Comparative Neurology* **408**, 489–505 (1999).

164. McGuire, B. A., Gilbert, C. D., Rivlin, P. K. & Wiesel, T. N. Targets of horizontal connections in macaque primary visual cortex. *Journal of Comparative Neurology* **305**, 370–392 (1991).

165. Butovas, S. & Schwarz, C. Spatiotemporal Effects of Microstimulation in Rat Neocortex: A Parametric Study Using Multielectrode Recordings. *Journal of Neurophysiology* **90**, 3024–3039 (2003).

166. Russell, B., Lees, R. M., Lak, A. & Packer, A. M. Desensitization of opsin responses during all-optical interrogation depends on imaging parameters. *NPh* **13**, 015007 (2026).

167. Logothetis, N. K. *et al.* The effects of electrical microstimulation on cortical signal propagation. *Nat Neurosci* **13**, 1283–1291 (2010).

168. Fan, L. Z. *et al.* All-Optical Electrophysiology Reveals the Role of Lateral Inhibition in Sensory Processing in Cortical Layer 1. *Cell* **180**, 521-535.e18 (2020).

169. Chettih, S. N. & Harvey, C. D. Single-neuron perturbations reveal feature-specific competition in V1. *Nature* **567**, 334–340 (2019).

170. Oldenburg, I. A. *et al.* The logic of recurrent circuits in the primary visual cortex. *Nat Neurosci* **27**, 137–147 (2024).

171. Youssef, D., Wittig, J. H., Jackson, S., Inati, S. K. & Zaghloul, K. A. Neuronal Spiking Responses to Direct Electrical Microstimulation in the Human Cortex. *J. Neurosci.* **43**, 4448–4460 (2023).

172. Steriade, M. & Contreras, D. Spike-Wave Complexes and Fast Components of Cortically Generated Seizures. I. Role of Neocortex and Thalamus. *Journal of Neurophysiology* **80**, 1439–1455 (1998).

173. Butovas, S. & Schwarz, C. Local Neuronal Responses to Intracortical Microstimulation in Rats' Barrel Cortex Are Dependent on Behavioral Context. *Front. Behav. Neurosci.* **16**, (2022).

174. Sadtler, P. T. *et al.* Neural constraints on learning. *Nature* **512**, 423–426 (2014).

175. Marshel, J. H. Cortical layer–specific critical dynamics triggering perception. *Science* **365**, 495 (2019).

176. Tehovnik, E. J. Electrical stimulation of neural tissue to evoke behavioral responses. *Journal of Neuroscience Methods* **65**, 1–17 (1996).

177. Kumaravelu, K., Sombeck, J., Miller, L. E., Bensmaia, S. J. & Grill, W. M. Stoney vs. Histed: Quantifying the spatial effects of intracortical microstimulation. *Brain Stimulation* **15**, 141–151 (2022).

178. Kumaravelu, K. & Grill, W. M. Neural mechanisms of the temporal response of cortical neurons to intracortical microstimulation. *Brain Stimulation: Basic, Translational, and Clinical Research in Neuromodulation* **17**, 365–381 (2024).

179. Hogan, M. K., Hamilton, G. F. & Horner, P. J. Neural Stimulation and Molecular Mechanisms of Plasticity and Regeneration: A Review. *Front. Cell. Neurosci.* **14**, (2020).

180. Hamani, C. *et al.* Insertional effect following electrode implantation: an underreported but important phenomenon. *Brain Commun* **6**, fcae093 (2024).

181. Crick, F. H. C. Thinking about the Brain. *Scientific American* **241**, 219–233 (1979).

182. Friedman, J. M. How the discovery of microbial opsins led to the development of optogenetics. *Cell* **184**, 5266–5270 (2021).

183. Josselyn, S. A. The past, present and future of light-gated ion channels and optogenetics. *eLife* **7**, e42367 (2018).

184. Emiliani, V. *et al.* Optogenetics for light control of biological systems. *Nat Rev Methods Primers* **2**, 55 (2022).

185. *Zeitschrift für wissenschaftliche Zoologie*. (1851).

186. Boll, F. Zur Anatomie und Physiologie der Retina. Aus dem Laboratorium f□r vergleichende Anatomie und Physiologie zu Rom, Achte Mittheilung [Offprint]. *Archiv f□r Anatomie und Physiologie, Physiologische Abtheilung* 4–35 (1877).

187. Gamgee, A. Kühne's Researches on Photo-Chemical Processes in the Retina. *Nature* **15**, 296–296 (1877).

188. Aubert, H. *Physiologie Der Netzhaut*. (Verlag von E. Morgenstern, Breslau, 1865).

189. Wald, G. CAROTENOIDS AND THE VISUAL CYCLE. *J Gen Physiol* **19**, 351–371 (1935).

190. Bortoff, A. Localization of slow potential responses in the *Necturus* retina. *Vision Research* **4**, 627-IN9 (1964).

191. Hodgkin, A. L., McNaughton, P. A. & Nunn, B. J. The ionic selectivity and calcium dependence of the light-sensitive pathway in toad rods. *The Journal of Physiology* **358**, 447–468 (1985).

192. Yau, K.-W. & Nakatani, K. Light-suppressible, cyclic GMP-sensitive conductance in the plasma membrane of a truncated rod outer segment. *Nature* **317**, 252–255 (1985).

193. Wigler, M. *et al.* Transfer of purified herpes virus thymidine kinase gene to cultured mouse cells. *Cell* **11**, 223–232 (1977).

194. Cohen, S. N., Chang, A. C. Y., Boyer, H. W. & Helling, R. B. Construction of Biologically Functional Bacterial Plasmids In Vitro. *Proceedings of the National Academy of Sciences* **70**, 3240–3244 (1973).

195. Zemelman, B. V., Lee, G. A., Ng, M. & Miesenböck, G. Selective Photostimulation of Genetically ChARGed Neurons. *Neuron* **33**, 15–22 (2002).

196. Miesenböck, G. The Optogenetic Catechism. *Science* **326**, 395–399 (2009).

197. Light- and drug-activated G-protein-coupled receptors to control intracellular signalling. https://doi.org/10.1113/expphysiol.2010.055517 doi:10.1113/expphysiol.2010.055517.

198. Oesterhelt, D. & Stoeckenius, W. Rhodopsin-like Protein from the Purple Membrane of Halobacterium halobium. *Nature New Biology* **233**, 149–152 (1971).

199. Bieszke, J. A. *et al.* The nop-1 gene of Neurospora crassa encodes a seven transmembrane helix retinal-binding protein homologous to archaeal rhodopsins. *Proceedings of the National Academy of Sciences* **96**, 8034–8039 (1999).

200. Racker, E. & Stoeckenius, W. Reconstitution of Purple Membrane Vesicles Catalyzing Light-driven Proton Uptake and Adenosine Triphosphate Formation. *Journal of Biological Chemistry* **249**, 662–663 (1974).

201. Chow, B. Y. *et al.* High-performance genetically targetable optical neural silencing by light-driven proton pumps. *Nature* **463**, 98–102 (2010).

202. Foster, K. W. *et al.* A rhodopsin is the functional photoreceptor for phototaxis in the unicellular eukaryote Chlamydomonas. *Nature* **311**, 756–759 (1984).

203. Holland, E. M., Braun, F. J., Nonnengässer, C., Harz, H. & Hegemann, P. The nature of rhodopsin-triggered photocurrents in Chlamydomonas. I. Kinetics and influence of divalent ions. *Biophysical Journal* **70**, 924–931 (1996).

204. Harz, H., Nonnengasser, C. & Hegemann, P. The photoreceptor current of the green alga Chlamydomonas. *Philos Trans R Soc Lond B Biol Sci* **338**, 39–52 (1992).

205. Nagel, G. *et al.* Channelrhodopsin-1: A Light-Gated Proton Channel in Green Algae. *Science* **296**, 2395–2398 (2002).

206. Nagel, G. *et al.* Channelrhodopsin-2, a directly light-gated cation-selective membrane channel. *Proceedings of the National Academy of Sciences* **100**, 13940–13945 (2003).

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

208. Zhang, F., Wang, L.-P., Boyden, E. S. & Deisseroth, K. Channelrhodopsin-2 and optical control of excitable cells. *Nat Methods* **3**, 785–792 (2006).

209. Miyashita, T., Shao, Y. R., Chung, J., Pourzia, O. & Feldman, D. Long-term channelrhodopsin-2 (ChR2) expression can induce abnormal axonal morphology and targeting in cerebral cortex. *Front. Neural Circuits* **7**, (2013).

210. Wiegert, J. S., Mahn, M., Prigge, M., Printz, Y. & Yizhar, O. Silencing Neurons: Tools, Applications, and Experimental Constraints. *Neuron* **95**, 504–529 (2017).

211. Owen, S. F., Liu, M. H. & Kreitzer, A. C. Thermal constraints on in vivo optogenetic manipulations. *Nat Neurosci* **22**, 1061–1065 (2019).

212. Stujenske, J. M., Spellman, T. & Gordon, J. A. Modeling the Spatiotemporal Dynamics of Light and Heat Propagation for In Vivo Optogenetics. *Cell Reports* **12**, 525–534 (2015).

213. Lima, S. Q. & Miesenböck, G. Remote Control of Behavior through Genetically Targeted Photostimulation of Neurons. *Cell* **121**, 141–152 (2005).

214. Aravanis, A. M. *et al.* An optical neural interface: in vivo control of rodent motor cortex with integrated fiberoptic and optogenetic technology. *J. Neural Eng.* **4**, S143 (2007).

215. Adamantidis, A. R., Zhang, F., Aravanis, A. M., Deisseroth, K. & de Lecea, L. Neural substrates of awakening probed with optogenetic control of hypocretin neurons. *Nature* **450**, 420–424 (2007).

216. Huber, D. *et al.* Sparse optical microstimulation in barrel cortex drives learned behaviour in freely moving mice. *Nature* **451**, 61–64 (2008).

217. Ranck, J. B. Which elements are excited in electrical stimulation of mammalian central nervous system: A review. *Brain Research* **98**, 417–440 (1975).

218. Histed, M. H. & Maunsell, J. H. R. Cortical neural populations can guide behavior by integrating inputs linearly, independent of synchrony. *Proceedings of the National Academy of Sciences* **111**, E178–E187 (2014).

219. Bruno, R. M. & Sakmann, B. Cortex Is Driven by Weak but Synchronously Active Thalamocortical Synapses. *Science* **312**, 1622–1627 (2006).

220. Fries, P. A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends in Cognitive Sciences* **9**, 474–480 (2005).

221. Xu, N. *et al.* Nonlinear dendritic integration of sensory and motor input during an active sensing task. *Nature* **492**, 247–251 (2012).

222. Tomko, G. J. & Crapper, D. R. Neuronal variability: non-stationary responses to identical visual stimuli. *Brain Research* **79**, 405–418 (1974).

223. van Vreeswijk, C. & Sompolinsky, H. Chaos in Neuronal Networks with Balanced Excitatory and Inhibitory Activity. *Science* **274**, 1724–1726 (1996).

224. Tiesinga, P., Fellous, J.-M. & Sejnowski, T. J. Regulation of spike timing in visual cortical circuits. *Nat Rev Neurosci* **9**, 97–107 (2008).

225. Shadlen, M. N. & Newsome, W. T. The Variable Discharge of Cortical Neurons: Implications for Connectivity, Computation, and Information Coding. *J. Neurosci.* **18**, 3870–3896 (1998).

226. Masquelier, T. Neural variability, or lack thereof. *Front. Comput. Neurosci.* **7**, (2013).

227. May, T. *et al.* Detection of Optogenetic Stimulation in Somatosensory Cortex by Non-Human Primates - Towards Artificial Tactile Sensation. *PLOS ONE* **9**, e114529 (2014).

228. Wu, C., Ivanova, E., Zhang, Y. & Pan, Z.-H. rAAV-Mediated Subcellular Targeting of Optogenetic Tools in Retinal Ganglion Cells In Vivo. *PLOS ONE* **8**, e66332 (2013).

229. Liu, X. *et al.* Optogenetic stimulation of a hippocampal engram activates fear memory recall. *Nature* **484**, 381–385 (2012).

230. O'Rawe, J. F. *et al.* Excitation creates a distributed pattern of cortical suppression due to varied recurrent input. *Neuron* **111**, 4086-4101.e5 (2023).

231. Wang, H. C., LeMessurier, A. M. & Feldman, D. E. Tuning instability of non-columnar neurons in the salt-and-pepper whisker map in somatosensory cortex. *Nat Commun* **13**, 6611 (2022).

232. Jendrichovsky, P. *et al.* Patchy harmonic functional connectivity of the mouse auditory cortex. *Proceedings of the National Academy of Sciences* **122**, e2510012122 (2025).

233. Peticolas, W. L., Goldsborough, J. P. & Rieckhoff, K. E. Double Photon Excitation in Organic Crystals. *Phys. Rev. Lett.* **10**, 43–45 (1963).

234. McNaughton, B. L., O'Keefe, J. & Barnes, C. A. The stereotrode: A new technique for simultaneous isolation of several single units in the central nervous system from multiple unit records. *Journal of Neuroscience Methods* **8**, 391–397 (1983).

235. Brakenhoff, G. J., van der Voort, H. T. M., van Spronsen, E. A., Linnemans, W. a. M. & Nanninga, N. Three-dimensional chromatin distribution in neuroblastoma nuclei shown by confocal scanning laser microscopy. *Nature* **317**, 748–749 (1985).

236. White, J. G., Amos, W. B. & Fordham, M. An evaluation of confocal versus conventional imaging of biological structures by fluorescence light microscopy. *J Cell Biol* **105**, 41–48 (1987).

237. Gu, M. & Sheppard, C. J. R. Comparison of three-dimensional imaging properties between two-photon and single-photon fluorescence microscopy. *Journal of Microscopy* **177**, 128–137 (1995).

238. Schilders, S. P. & Gu, M. Limiting Factors on Image Quality in Imaging through Turbid Media under Single-photon and Two-photon Excitation. *Microanal* **6**, 156–160 (2000).

239. Zhang, B. *et al.* Modified Golgi-Cox method for micrometer scale sectioning of the whole mouse brain. *Journal of Neuroscience Methods* **197**, 1–5 (2011).

240. Zipfel, W. R., Williams, R. M. & Webb, W. W. Nonlinear magic: multiphoton microscopy in the 404 biosciences. *Nat Biotechnol* **21**, 1369-1377, (2003).

241. Denk, W. & Svoboda, K. Photon Upmanship: Why Multiphoton Imaging Is More than a Gimmick. *Neuron* **18**, 351–357 (1997).

242. Göppert-Mayer, M. Über Elementarakte mit zwei Quantensprüngen. *Annalen der Physik* **401**, 273–294 (1931).

243. Denk, W. *et al.* Anatomical and functional imaging of neurons using 2-photon laser scanning microscopy. *Journal of Neuroscience Methods* **54**, 151–162 (1994).

244. Yuste, R. & Denk, W. Dendritic spines as basic functional units of neuronal integration. *Nature* **375**, 682–684 (1995).

245. Regehr, W. G. & Tank, D. W. Calcium concentration dynamics produced by synaptic activation of CA1 hippocampal pyramidal cells. *J. Neurosci.* **12**, 4202–4223 (1992).

246. Svoboda, K., Denk, W., Kleinfeld, D. & Tank, D. W. In vivo dendritic calcium dynamics in neocortical pyramidal neurons. *Nature* **385**, 161–165 (1997).

247. Peron, S. P., Freeman, J., Iyer, V., Guo, C. & Svoboda, K. A Cellular Resolution Map of Barrel Cortex Activity during Tactile Behavior. *Neuron* **86**, 783–799 (2015).

248. Kleinfeld, D. & Delaney, K. r. Distributed representation of vibrissa movement in the upper layers of somatosensory cortex revealed with voltage-sensitive dyes. *Journal of Comparative Neurology* **375**, 89–108 (1996).

249. Horton, N. G. *et al.* In vivo three-photon microscopy of subcortical structures within an intact mouse brain. *Nature Photon* **7**, 205–209 (2013).

250. Grienberger, C. & Konnerth, A. Imaging Calcium in Neurons. *Neuron* **73**, 862–885 (2012).

251. Nakai, J., Ohkura, M. & Imoto, K. A high signal-to-noise Ca2+ probe composed of a single green fluorescent protein. *Nat Biotechnol* **19**, 137–141 (2001).

252. Tallini, Y. N. *et al.* Imaging cellular signals in the heart in vivo: Cardiac expression of the high-signal Ca2+ indicator GCaMP2. *Proceedings of the National Academy of Sciences* **103**, 4753–4758 (2006).

253. Tian, L. *et al.* Imaging neural activity in worms, flies and mice with improved GCaMP calcium indicators. *Nat Methods* **6**, 875–881 (2009).

254. Akerboom, J. *et al.* Optimization of a GCaMP Calcium Indicator for Neural Activity Imaging. *J. Neurosci.* **32**, 13819–13840 (2012).

255. Chen, T.-W. *et al.* Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature* **499**, 295–300 (2013).

256. Dana, H. *et al.* High-performance calcium sensors for imaging activity in neuronal populations and microcompartments. *Nat Methods* **16**, 649–657 (2019).

257. Zhang, Y. *et al.* Fast and sensitive GCaMP calcium indicators for imaging neural populations. *Nature* **615**, 884–891 (2023).

258. Dana, H. *et al.* Sensitive red protein calcium indicators for imaging neural activity. *eLife* **5**, e12727 (2016).

259. Knöpfel, T. & Song, C. Optical voltage imaging in neurons: moving from technology development to practical tool. *Nat Rev Neurosci* **20**, 719–727 (2019).

260. Hao, Y. A. *et al.* A fast and responsive voltage indicator with enhanced sensitivity for unitary synaptic events. *Neuron* **112**, 3680-3696.e8 (2024).

261. Treptow, D., Bola, R., Martín-Badosa, E. & Montes-Usategui, M. Artifact-free holographic light shaping through moving acousto-optic holograms. *Sci Rep* **11**, 21261 (2021).

262. Rickgauer, J. P. & Tank, D. W. Two-photon excitation of channelrhodopsin-2 at saturation. *Proceedings of the National Academy of Sciences* **106**, 15025–15030 (2009).

263. Packer, A. M., Roska, B. & Häusser, M. Targeting neurons and photons for optogenetics. *Nat Neurosci* **16**, 805–815 (2013).

264. Rindner, D. J. & Lur, G. Practical considerations in an era of multicolor optogenetics. *Front. Cell. Neurosci.* **17**, (2023).

265. Li, L. *et al.* Colocalized, bidirectional optogenetic modulations in freely behaving mice with a wireless dual-color optoelectronic probe. *Nat Commun* **13**, 839 (2022).

266. Ernst, O. P. *et al.* Photoactivation of Channelrhodopsin *. *Journal of Biological Chemistry* **283**, 1637–1643 (2008).

267. Zhang, F. *et al.* Red-shifted optogenetic excitation: a tool for fast neural control derived from Volvox carteri. *Nat Neurosci* **11**, 631–633 (2008).

268. Kianianmomeni, A., Stehfest, K., Nematollahi, G., Hegemann, P. & Hallmann, A. Channelrhodopsins of Volvox carteri Are Photochromic Proteins That Are Specifically Expressed in Somatic Cells under Control of Light, Temperature, and the Sex Inducer. *Plant Physiol* **151**, 347–366 (2009).

269. Yizhar, O. *et al.* Neocortical excitation/inhibition balance in information processing and social dysfunction. *Nature* **477**, 171–178 (2011).

270. Prakash, R. *et al.* Two-photon optogenetic toolbox for fast inhibition, excitation and bistable modulation. *Nat Methods* **9**, 1171–1179 (2012).

271. Rickgauer, J. P., Deisseroth, K. & Tank, D. W. Simultaneous cellular-resolution optical perturbation and imaging of place cell firing fields. *Nat Neurosci* **17**, 1816–1824 (2014).

272. Packer, A. M., Russell, L. E., Dalgleish, H. W. P. & Häusser, M. Simultaneous all-optical manipulation and recording of neural circuit activity with cellular resolution in vivo. *Nat Methods* **12**, 140–146 (2015).

273. Carrillo-Reid, L., Yang, W., Bando, Y., Peterka, D. S. & Yuste, R. Imprinting and recalling cortical ensembles. *Science* **353**, 691–694 (2016).

274. Yang, W., Carrillo-Reid, L., Bando, Y., Peterka, D. S. & Yuste, R. Simultaneous two-photon imaging and two-photon optogenetics of cortical circuits in three dimensions. *eLife* **7**, e32671 (2018).

275. Forli, A. *et al.* Two-Photon Bidirectional Control and Imaging of Neuronal Excitability with High Spatial Resolution In Vivo. *Cell Reports* **22**, 3087–3098 (2018).

276. Forli, A. Two-Photon Bidirectional Control and Imaging of Neuronal Excitability with High 477 Spatial Resolution In Vivo. *Cell Reports* **22**, 3087–3098 (2018).

277. Rickgauer, J. P., Deisseroth, K. & Tank, D. W. Simultaneous cellular-resolution optical perturbation and imaging of place cell firing fields. *Nat Neurosci* **17**, 1816–1824 (2014).

278. Weiner, A. M., Leaird, D. E., Patel, J. S. & Wullert, J. R. Programmable shaping of femtosecond optical pulses by use of 128-element liquid crystal phase modulator. *IEEE Journal of Quantum Electronics* **28**, 908–920 (1992).

279. Wang, S. *et al.* All Optical Interface for Parallel, Remote, and Spatiotemporal Control of Neuronal Activity. *Nano Lett.* **7**, 3859–3863 (2007).

280. Nikolenko, V., Poskanzer, K. E. & Yuste, R. Two-photon photostimulation and imaging of neural circuits. *Nat Methods* **4**, 943–950 (2007).

281. Nikolenko, V. *et al.* SLM microscopy: scanless two-photon imaging and photostimulation using spatial light modulators. *Front. Neural Circuits* **2**, (2008).

282. Mayr, E. Cause and Effect in Biology. *Science* **134**, 1501–1506 (1961).

283. Carrillo-Reid, L., Han, S., Yang, W., Akrouh, A. & Yuste, R. Controlling Visually Guided Behavior by Holographic Recalling of Cortical Ensembles. *Cell* **178**, 447-457.e5 (2019).

284. Daie, K., Svoboda, K. & Druckmann, S. Targeted photostimulation uncovers circuit motifs supporting short-term memory. *Nat Neurosci* **24**, 259–265 (2021).

285. Robinson, N. T. M. *et al.* Targeted Activation of Hippocampal Place Cells Drives Memory-Guided Spatial Behavior. *Cell* **183**, 1586-1599.e10 (2020).

286. Rowland, J. M. *et al.* Propagation of activity through the cortical hierarchy and perception are determined by neural variability. *Nat Neurosci* **26**, 1584–1594 (2023).

287. Dalgleish, H. W. *et al.* How many neurons are sufficient for perception of cortical activity? *eLife* **9**, e58889 (2020).

288. Pancholi, R., Ryan, L. & Peron, S. Learning in a sensory cortical microstimulation task is associated with elevated representational stability. *Nat Commun* **14**, 3860 (2023).

289. Williams, Z. M., Elfar, J. C., Eskandar, E. N., Toth, L. J. & Assad, J. A. Parietal activity and the perceived direction of ambiguous apparent motion. *Nat Neurosci* **6**, 616–623 (2003).

290. Nienborg, H. & Cumming, B. G. Macaque V2 Neurons, But Not V1 Neurons, Show Choice-Related Activity. *J. Neurosci.* **26**, 9567–9578 (2006).

291. Janssen, P., Vogels, R., Liu, Y. & Orban, G. A. At Least at the Level of Inferior Temporal Cortex, the Stereo Correspondence Problem Is Solved. *Neuron* **37**, 693–701 (2003).

292. Murphey, D. K. & Maunsell, J. H. R. Behavioral Detection of Electrical Microstimulation in Different Cortical Visual Areas. *Current Biology* **17**, 862–867 (2007).

293. Singer, W. & Gray, C. M. Visual Feature Integration and the Temporal Correlation Hypothesis. *Annual Review of Neuroscience* **18**, 555–586 (1995).

294. Roelfsema, P. R., Engel, A. K., König, P. & Singer, W. The Role of Neuronal Synchronization in Response Selection: A Biologically Plausible Theory of Structured Representations in the Visual Cortex. *J Cogn Neurosci* **8**, 603–625 (1996).

295. Vinck, M., Batista-Brito, R., Knoblich, U. & Cardin, J. A. Arousal and Locomotion Make Distinct Contributions to Cortical Activity Patterns and Visual Encoding. *Neuron* **86**, 740–754 (2015).

296. McGinley, M. J., David, S. V. & McCormick, D. A. Cortical Membrane Potential Signature of Optimal States for Sensory Signal Detection. *Neuron* **87**, 179–192 (2015).

297. Gray, C. M., König, P., Engel, A. K. & Singer, W. Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature* **338**, 334–337 (1989).

298. Engel, A. K., König, P. & Singer, W. Direct physiological evidence for scene segmentation by temporal coding. *Proceedings of the National Academy of Sciences* **88**, 9136–9140 (1991).

299. Fries, P., Reynolds, J. H., Rorie, A. E. & Desimone, R. Modulation of Oscillatory Neuronal Synchronization by Selective Visual Attention. *Science* **291**, 1560–1563 (2001).

300. Rohenkohl, G., Bosman, C. A. & Fries, P. Gamma Synchronization between V1 and V4 Improves Behavioral Performance. *Neuron* **100**, 953-963.e3 (2018).

301. Fries, P. Rhythms for Cognition: Communication through Coherence. *Neuron* **88**, 220–235 (2015).

302. van Vugt, B. *et al.* The threshold for conscious report: Signal loss and response bias in visual and frontal cortex. *Science* **360**, 537–542 (2018).

303. O'Reilly, C., Iavarone, E., Yi, J. & Hill, S. L. Rodent somatosensory thalamocortical circuitry: Neurons, synapses, and connectivity. *Neuroscience & Biobehavioral Reviews* **126**, 213–235 (2021).

304. Nicolelis, M. A. L., Baccala, L. A., Lin, R. C. S. & Chapin, J. K. Sensorimotor Encoding by Synchronous Neural Ensemble Activity at Multiple Levels of the Somatosensory System. *Science* **268**, 1353–1358 (1995).

305. Temereanca, S., Brown, E. N. & Simons, D. J. Rapid Changes in Thalamic Firing Synchrony during Repetitive Whisker Stimulation. *J. Neurosci.* **28**, 11153–11164 (2008).

306. Bruno, R. M. & Simons, D. J. Feedforward Mechanisms of Excitatory and Inhibitory Cortical Receptive Fields. *J. Neurosci.* **22**, 10966–10975 (2002).

307. Sun, Q.-Q., Huguenard, J. R. & Prince, D. A. Barrel Cortex Microcircuits: Thalamocortical Feedforward Inhibition in Spiny Stellate Cells Is Mediated by a Small Number of Fast-Spiking Interneurons. *J. Neurosci.* **26**, 1219–1230 (2006).

308. Wang, Q., Webber, R. M. & Stanley, G. B. Thalamic synchrony and the adaptive gating of information flow to cortex. *Nat Neurosci* **13**, 1534–1541 (2010).

309. Crick, F. Function of the thalamic reticular complex: the searchlight hypothesis. *Proceedings of the National Academy of Sciences* **81**, 4586–4590 (1984).

310. S, H. *et al.* Thalamocortical rhythms during a vibrotactile detection task. *Proc Natl Acad Sci U S A* **111**, E1797-805 (2014).

311. Tauste Campo, A. *et al.* Feed-forward information and zero-lag synchronization in the sensory thalamocortical circuit are modulated during stimulus perception. *Proceedings of the National Academy of Sciences* **116**, 7513–7522 (2019).

312. Tauste Campo, A. *et al.* Thalamocortical interactions shape hierarchical neural variability during stimulus perception. *iScience* **27**, 110065 (2024).

313. Borden, P. Y. *et al.* Thalamic bursting and the role of timing and synchrony in thalamocortical signaling in the awake mouse. *Neuron* **110**, 2836-2853.e8 (2022).

314. Casas-Torremocha, D. *et al.* A Combinatorial Input Landscape in the "Higher-Order Relay" Posterior Thalamic Nucleus. *J. Neurosci.* **42**, 7757–7781 (2022).

315. Sobolewski, A., Kublik, E., Swiejkowski, D. A., Kamiński, J. & Wróbel, A. Alertness opens the effective flow of sensory information through rat thalamic posterior nucleus. *European Journal of Neuroscience* **41**, 1321–1331 (2015).

316. Suzuki, M. & Larkum, M. E. General Anesthesia Decouples Cortical Pyramidal Neurons. *Cell* **180**, 666-676.e13 (2020).

317. Diamond, M. E., Armstrong-James, M. & Ebner, F. F. Somatic sensory responses in the rostral sector of the posterior group (POm) and in the ventral posterior medial nucleus (VPM) of the rat thalamus. *Journal of Comparative Neurology* **318**, 462–476 (1992).

318. La Terra, D. *et al.* The role of higher-order thalamus during learning and correct performance in goal-directed behavior. *eLife* **11**, e77177 (2022).

319. El-Boustani, S. *et al.* Anatomically and functionally distinct thalamocortical inputs to primary and secondary mouse whisker somatosensory cortices. *Nat Commun* **11**, 3342 (2020).

320. Qi, J., Ye, C., Naskar, S., Inácio, A. R. & Lee, S. Posteromedial thalamic nucleus activity significantly contributes to perceptual discrimination. *PLOS Biology* **20**, e3001896 (2022).

321. Petty, G. H. & Bruno, R. M. Attentional modulation of secondary somatosensory and visual thalamus of mice. *eLife* **13**, RP97188 (2024).

322. Miller-Hansen, A. J. & Sherman, S. M. Conserved patterns of functional organization between cortex and thalamus in mice. *Proceedings of the National Academy of Sciences* **119**, e2201481119 (2022).

323. Petersen, C. C. H. Sensorimotor processing in the rodent barrel cortex. *Nat Rev Neurosci* **20**, 533–546 (2019).

324. Jurado-Parras, M.-T. *et al.* The Dorsal Striatum Energizes Motor Routines. *Current Biology* **30**, 4362-4372.e6 (2020).

325. Yttri, E. A. & Dudman, J. T. Opponent and bidirectional control of movement velocity in the basal ganglia. *Nature* **533**, 402–406 (2016).

326. Dhawale, A. K., Wolff, S. B. E., Ko, R. & Ölveczky, B. P. The basal ganglia control the detailed kinematics of learned motor skills. *Nat Neurosci* **24**, 1256–1269 (2021).

327. Sippy, T., Lapray, D., Crochet, S. & Petersen, C. C. H. Cell-Type-Specific Sensorimotor Processing in Striatal Projection Neurons during Goal-Directed Behavior. *Neuron* **88**, 298–305 (2015).

328. Lee, C. R. *et al.* Opposing Influence of Sensory and Motor Cortical Input on Striatal Circuitry and Choice Behavior. *Current Biology* **29**, 1313-1323.e5 (2019).

329. Sun, Z., Schneider, A., Alyahyay, M., Karvat, G. & Diester, I. Effects of Optogenetic Stimulation of Primary Somatosensory Cortex and Its Projections to Striatum on Vibrotactile Perception in Freely Moving Rats. *eNeuro* **8**, (2021).

330. Antzoulatos, E. G. & Miller, E. K. Differences between Neural Activity in Prefrontal Cortex and Striatum during Learning of Novel Abstract Categories. *Neuron* **71**, 243–249 (2011).

331. Antzoulatos, E. G. & Miller, E. K. Increases in Functional Connectivity between Prefrontal Cortex and Striatum during Category Learning. *Neuron* **83**, 216–225 (2014).

332. Koralek, A. C., Costa, R. M. & Carmena, J. M. Temporally Precise Cell-Specific Coherence Develops in Corticostriatal Networks during Learning. *Neuron* **79**, 865–872 (2013).

333. Ponvert, N. D. & Jaramillo, S. Auditory Thalamostriatal and Corticostriatal Pathways Convey Complementary Information about Sound Features. *J. Neurosci.* **39**, 271–280 (2019).

334. Uhlhaas, P. J. & Singer, W. Abnormal neural oscillations and synchrony in schizophrenia. *Nat Rev Neurosci* **11**, 100–113 (2010).

335. Cariani, P. & Baker, J. M. Survey of temporal coding of sensory information. *Front. Comput. Neurosci.* **19**, (2025).

336. Long, M. A., Jin, D. Z. & Fee, M. S. Support for a synaptic chain model of neuronal sequence generation. *Nature* **468**, 394–399 (2010).

337. Hahnloser, R. H. R., Kozhevnikov, A. A. & Fee, M. S. An ultra-sparse code underliesthe generation of neural sequences in a songbird. *Nature* **419**, 65–70 (2002).

338. Srivastava, K. H. *et al.* Motor control by precisely timed spike patterns. *Proceedings of the National Academy of Sciences* **114**, 1171–1176 (2017).

339. Diba, K. & Buzsáki, G. Forward and reverse hippocampal place-cell sequences during ripples. *Nat Neurosci* **10**, 1241–1242 (2007).

340. Skaggs, W. E. & McNaughton, B. L. Replay of Neuronal Firing Sequences in Rat Hippocampus During Sleep Following Spatial Experience. *Science* **271**, 1870–1873 (1996).

341. Chen, Y., Zhang, H., Cameron, M. & Sejnowski, T. Predictive sequence learning in the hippocampal formation. *Neuron* **112**, 2645-2658.e4 (2024).

342. Nolte, M., Reimann, M. W., King, J. G., Markram, H. & Muller, E. B. Cortical reliability amid noise and chaos. *Nat Commun* **10**, 3792 (2019).

343. London, M., Roth, A., Beeren, L., Häusser, M. & Latham, P. E. Sensitivity to perturbations in vivo implies high noise and suggests rate coding in cortex. *Nature* **466**, 123–127 (2010).

344. Schölvinck, M. L., Saleem, A. B., Benucci, A., Harris, K. D. & Carandini, M. Cortical State Determines Global Variability and Correlations in Visual Cortex. *J. Neurosci.* **35**, 170–178 (2015).

345. Oram, M. W. & Perrett, D. I. Time course of neural responses discriminating different views of the face and head. *Journal of Neurophysiology* **68**, 70–84 (1992).

346. Thorpe, S., Fize, D. & Marlot, C. Speed of processing in the human visual system. *Nature* **381**, 520–522 (1996).

347. Nowak, L. G. & Bullier, J. The Timing of Information Transfer in the Visual System. in *Extrastriate Cortex in Primates* (eds Rockland, K. S., Kaas, J. H. & Peters, A.) 205–241 (Springer US, Boston, MA, 1997). doi:10.1007/978-1-4757-9625-4_5.

348. Van Rullen, R., Gautrais, J., Delorme, A. & Thorpe, S. Face processing using one spike per neurone. *Biosystems* **48**, 229–239 (1998).

349. J, T. S. Spike arrival times: A highly efficient coding scheme for neural networks. *Parallel Processing in Neural Systems* 91–94 (1990).

350. Laurent, G., Wehr, M. & Davidowitz, H. Temporal Representations of Odors in an Olfactory Network. *J. Neurosci.* **16**, 3837–3847 (1996).

351. Wehr, M. & Laurent, G. Odour encoding by temporal sequences of firing in oscillating neural assemblies. *Nature* **384**, 162–166 (1996).

352. Stopfer, M., Bhagavan, S., Smith, B. H. & Laurent, G. Impaired odour discrimination on desynchronization of odour-encoding neural assemblies. *Nature* **390**, 70–74 (1997).

353. Dynamic Optimization of Odor Representations by Slow Temporal Patterning of Mitral Cell Activity. *Science* https://www.science.org/doi/10.1126/science.291.5505.889.

354. Luczak, A., Barthó, P., Marguet, S. L., Buzsáki, G. & Harris, K. D. Sequential structure of neocortical spontaneous activity in vivo. *Proceedings of the National Academy of Sciences* **104**, 347–352 (2007).

355. Luczak, A., Bartho, P. & Harris, K. D. Gating of Sensory Input by Spontaneous Cortical Activity. *J. Neurosci.* **33**, 1684–1695 (2013).

356. Montijn, J. S., Olcese, U. & Pennartz, C. M. A. Visual Stimulus Detection Correlates with the Consistency of Temporal Sequences within Stereotyped

Events of V1 Neuronal Population Activity. *J. Neurosci.* **36**, 8624–8640 (2016).

357. Abeles, M. *Local Cortical Circuits: An Electrophysiological Study*. (Springer Science & Business Media, 2012).

358. Abeles, M. *Corticonics: Neural Circuits of the Cerebral Cortex*. (Cambridge University Press, 1991).

359. Diesmann, M., Gewaltig, M.-O. & Aertsen, A. Stable propagation of synchronous spiking in cortical neural networks. *Nature* **402**, 529–533 (1999).

360. Gewaltig, M.-O., Diesmann, M. & Aertsen, A. Propagation of cortical synfire activity: survival probability in single trials and stability in the mean. *Neural Networks* **14**, 657–673 (2001).

361. Kumar, A., Rotter, S. & Aertsen, A. Spiking activity propagation in neuronal networks: reconciling different perspectives on neural coding. *Nat Rev Neurosci* **11**, 615–627 (2010).

362. Jahnke, S., Timme, M. & Memmesheimer, R.-M. Guiding Synchrony through Random Networks. *Phys. Rev. X* **2**, 041016 (2012).

363. Jahnke, S., Memmesheimer, R.-M. & Timme, M. Propagating synchrony in feed-forward networks. *Front. Comput. Neurosci.* **7**, (2013).

364. Carrillo-Reid, L., Miller, J. K., Hamm, J. P., Jackson, J. & Yuste, R. Endogenous Sequential Cortical Activity Evoked by Visual Stimuli. *J. Neurosci.* **35**, 8813–8828 (2015).

365. Gerstein, G. L., Williams, E. R., Diesmann, M., Grün, S. & Trengove, C. Detecting synfire chains in parallel spike data. *Journal of Neuroscience Methods* **206**, 54–64 (2012).

366. Mackevicius, E. L. *et al.* Unsupervised discovery of temporal sequences in high-dimensional datasets, with applications to neuroscience. *eLife* **8**, e38471 (2019).

367. Sotomayor-Gómez, B., Battaglia, F. P. & Vinck, M. SpikeShip: A method for fast, unsupervised discovery of high-dimensional neural spiking patterns. *PLOS Computational Biology* **19**, e1011335 (2023).

368. Agrawal, R., Imieliński, T. & Swami, A. Mining association rules between sets of items in large databases. in *Proceedings of the 1993 ACM SIGMOD international conference on Management of data* 207–216 (Association for Computing Machinery, New York, NY, USA, 1993). doi:10.1145/170035.170072.

369. Torre, E., Picado-Muino, D., Denker, M., Borgelt, C. & Grün, S. Statistical evaluation of synchronous spike patterns extracted by frequent item set mining. *Front. Comput. Neurosci.* **7**, (2013).

370. Quaglio, P., Yegenoglu, A., Torre, E., Endres, D. M. & Grün, S. Detection and Evaluation of Spatio-Temporal Spike Patterns in Massively Parallel Spike Train Data with SPADE. *Front. Comput. Neurosci.* **11**, (2017).

371. Yegenoglu, A., Quaglio, P., Torre, E., Grün, S. & Endres, D. Exploring the Usefulness of Formal Concept Analysis for Robust Detection of Spatio-

temporal Spike Patterns in Massively Parallel Spike Trains. in *Graph-Based Representation and Reasoning* (eds Haemmerlé, O., Stapleton, G. & Faron Zucker, C.) 3–16 (Springer International Publishing, Cham, 2016). doi:10.1007/978-3-319-40985-6_1.

372. Porrmann, F. *et al.* Acceleration of the SPADE Method Using a Custom-Tailored FP-Growth Implementation. *Front. Neuroinform.* **15**, (2021).

373. Burgess, C. P. *et al.* High-Yield Methods for Accurate Two-Alternative Visual Psychophysics in Head-Fixed Mice. *Cell Reports* **20**, 2513–2524 (2017).

374. Bhagat, J., Wells, M. J., Harris, K. D., Carandini, M. & Burgess, C. P. Rigbox: An Open-Source Toolbox for Probing Neurons and Behavior. *eNeuro* **7**, (2020).

375. cortex-lab.github.io/Rigbox/Burgess_hardware_setup.html. https://cortex-lab.github.io/Rigbox/Burgess_hardware_setup.html.

376. Instructions for steering wheel setup: software. *Burgess_setup.m* https://cortex-lab.github.io/Rigbox/Burgess_setup.html (2020).

377. Birman, D. *et al.* Pinpoint: trajectory planning for multi-probe electrophysiology and injections in an interactive web-based 3D environment. *eLife* **12**, (2023).

378. billkarsh. billkarsh/SpikeGLX. (2026).

379. billkarsh. billkarsh/CatGT. (2026).

380. Pachitariu, M., Sridhar, S., Pennington, J. & Stringer, C. Spike sorting with Kilosort4. *Nat Methods* **21**, 914–921 (2024).

381. Fabre, J. M. J., Beest, E. H. van, Peters, A. J., Carandini, M. & Harris, K. D. Bombcell: automated curation and cell classification of spike-sorted electrophysiology data. Zenodo https://doi.org/10.5281/zenodo.8172822 (2023).

382. Stella, A., Quaglio, P., Torre, E. & Grün, S. 3d-SPADE: Significance evaluation of spatio-temporal patterns of various temporal extents. *Biosystems* **185**, 104022 (2019).

383. Fuglstad, J. G., Saldanha, P., Paglia, J. & Whitlock, J. R. Histological E-data Registration in rodent Brain Spaces. *eLife* **12**, e83496 (2023).

384. Weiler, S., Vélez-Fort, M. & Margrie, T. W. Overcoming off-target optical stimulation-evoked cortical activity in the mouse brain in vivo. 2024.08.14.607152 Preprint at https://doi.org/10.1101/2024.08.14.607152 (2024).

385. Multiprobe rig for in vivo recordings. *Ultraprecise Micromanipulator – Sensapex* https://sensapex.com/products/multiprobe-rig/.

386. Siegle, J. H. *et al.* Survey of spiking in the mouse visual system reveals functional hierarchy. *Nature* **592**, 86–92 (2021).

387. Mathis, A. *et al.* DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nat Neurosci* **21**, 1281–1289 (2018).

388. Yang, W. *et al.* Selection of experience for memory by hippocampal sharp wave ripples. *Science* **383**, 1478–1483 (2024).

389. Zhang, Y. *et al.* Long-term mesoscale imaging of 3D intercellular dynamics across a mammalian organ. *Cell* **187**, 6104-6122.e25 (2024).

390. Hamid, A. A., Frank, M. J. & Moore, C. I. Wave-like dopamine dynamics as a mechanism for spatiotemporal credit assignment. *Cell* **184**, 2733-2749.e16 (2021).

391. MacDowell, C. J. *et al.* Multiplexed subspaces route neural activity across brain-wide networks. *Nat Commun* **16**, 3359 (2025).

392. Sotomayor-Gómez, B., Battaglia, F. P. & Vinck, M. Firing rates in visual cortex show representational drift, while temporal spike sequences remain stable. *Cell Reports* **44**, (2025).

393. García-Rosales, F., Sotomayor-Gómez, B., Poeppel, D. & Hechavarria, J. C. Vocal production differentially affects fast- and broad-spiking neurons in auditory cortex. 2026.01.15.699648 Preprint at https://doi.org/10.64898/2026.01.15.699648 (2026).

394. Galvan, A. & Wichmann, T. Basal ganglia neurons in healthy and parkinsonian primates generate recurring sequences of spikes. *Journal of Neurophysiology* **129**, 1010–1020 (2023).

395. Russell, L. E. All-optical interrogation of neural circuits in behaving mice. *Nat Protoc* **17**, 1579-1620, (2022).

396. Gill, J. V. Precise Holographic Manipulation of Olfactory Circuits Reveals Coding Features 425 Determining Perceptual Detection. *Neuron* **108**, 382-393 5 (2020).

397. Lees, R. M., Pichler, B., Packer, A. M., & S.P.I.E. Contribution of optical resolution to the spatial 401 precision of two-photon optogenetic photostimulation in vivo. *NPh* **11**, 015006, (2024).

398. Mattis, J. Principles for applying optogenetic tools derived from direct comparative 483 analysis of microbial opsins. *Nat Methods* **9**, 159-172, (2012).

399. Papaioannou, S. & Medini, P. Advantages, Pitfalls, and Developments of All Optical 428 Interrogation Strategies of Microcircuits in vivo. *Front. Neurosci* **16, Frontiers**, 429 (2022).

400. Packer, A. M. Two-photon optogenetics of dendritic spines and neural circuits. *Nat 431 Methods* **9**, 1202-1205, (2012).

401. Hegemann, P., Oesterhelt, D. & Bamberg, E. The transport activity of the light-driven chloride pump halorhodopsin is regulated by green and blue light. *Biochimica et Biophysica Acta (BBA) - Biomembranes* **819**, 195–205 (1985).

402. Friedrich, T. *et al.* Proteorhodopsin is a Light-driven Proton Pump with Variable Vectoriality. *Journal of Molecular Biology* **321**, 821–838 (2002).

403. Hegemann, P., Ehlenbeck, S. & Gradmann, D. Multiple Photocycles of Channelrhodopsin. *443 Biophysical Journal* **89**, 3911–3918 (2005).

404. Ishizuka, T., Kakuda, M., Araki, R. & Yawo, H. Kinetic evaluation of photosensitivity in genetically engineered neurons expressing green algae light-gated channels. *Neuroscience Research* **54**, 85–94 (2006).

405. Nikolic, K., Degenaar, P. & Toumazou, C. Modeling and Engineering aspects of ChannelRhodopsin2 System for Neural Photostimulation. in *2006 International Conference of the IEEE Engineering in Medicine and Biology Society* 1626–1629 (2006). doi:10.1109/IEMBS.2006.260766.

406. Bamann, C., Kirsch, T., Nagel, G. & Bamberg, E. Spectral Characteristics of the Photocycle of Channelrhodopsin-2 and Its Implication for Channel Function. *Journal of Molecular Biology* **375**, 686–694 (2008).

407. Berndt, A., Yizhar, O., Gunaydin, L. A., Hegemann, P. & Deisseroth, K. Bi-stable neural state switches. *Nat Neurosci* **12**, 229–234 (2009).

408. Stehfest, K. & Hegemann, P. Evolution of the Channelrhodopsin Photocycle Model. *459 ChemPhysChem* **11**, 1120–1126 (2010).

409. Schneider, F., Grimm, C. & Hegemann, P. Biophysics of Channelrhodopsin. *Annu. Rev. 486 Biophys* **44**, 167–186 (2015).

410. Ronzitti, E. *et al.* Submillisecond Optogenetic Control of Neuronal Firing with Two-Photon Holographic Photoactivation of Chronos. *J. Neurosci.* **37**, 10679–10689 (2017).

411. Watson, B. O., Yuste, R. & Packer, A. M. PackIO and EphysViewer: software tools for 497 acquisition and analysis of neuroscience data. 054080, (2016) doi:10.1101/054080.

412. Pachitariu, M. Suite2p: beyond 10,000 neurons with standard two-photon microscopy. 499 061507, (2017) doi:10.1101/061507.

413. Russell, L. E. *et al.* The influence of cortical activity on perception depends on behavioral state and sensory context. *Nat Commun* **15**, 2456 (2024).

414. Chung, S., Li, X. & Nelson, S. B. Short-Term Depression at Thalamocortical Synapses 461 Contributes to Rapid Adaptation of Cortical Sensory Responses In Vivo. *Neuron* **34**, 437-462 446 (2002).

415. LaFosse, P. K. Cellular-resolution optogenetics reveals attenuation-by-suppression in 464 visual cortical neurons. in *Proceedings of the National Academy of Sciences 121(45), 465 e2318837121, Proceedings of the National Academy of Sciences* 466 (2024). doi:10.1073/pnas.2318837121.

416. Venkatachalam, V. & Cohen, A. E. Imaging GFP-Based Reporters in Neurons with 511 Multiwavelength Optogenetic Control. *Biophysical Journal* **107**, 1554–1563 (2014).

417. Mardinly, A. R. Precise multimodal optical control of neural ensemble activity. *Nat* **21**, 881-893, (2018).

418. Sridharan, S. High-performance microbial opsins for spatially and temporally precise 516 perturbations of large neuronal networks. *Neuron* **110**, 1139-1155 6 (2022).

419. Maschio, M. Linking Neurons to Network Function and Behavior by Two-Photon 519 Holographic Optogenetics and Volumetric Imaging. *Neuron* **94**, 774-789 5 (2017).

420. Forli, A. Optogenetic strategies for high-efficiency all-optical interrogation using blue-light-522 sensitive opsins. *eLife* **10**, 524 (2021).

421. Zhang, Z., Russell, L. E., Packer, A. M., Gauld, O. M. & Häusser, M. Closed-loop all-optical interrogation of neural circuits in vivo. *Nat Methods* **15**, 1037–1040 (2018).

422. Buetfering, C. *et al.* Behaviorally relevant decision coding in primary somatosensory cortex neurons. *Nat Neurosci* **25**, 1225–1236 (2022).

423. Gauld, O. M. *et al.* A latent pool of neurons silenced by sensory-evoked inhibition can be recruited to enhance perception. *Neuron* **112**, 2386-2403.e6 (2024).

424. Znamenskiy, P. & Zador, A. M. Corticostriatal neurons in auditory cortex drive decisions during auditory discrimination. *Nature* **497**, 482–485 (2013).

425. Wang, Z.-Q. *et al.* Corticostriatal Neurons in the Anterior Auditory Field Regulate Frequency Discrimination Behavior. *Neurosci. Bull.* **39**, 962–972 (2023).

426. Guo, W., Clause, A. R., Barth-Maron, A. & Polley, D. B. A Corticothalamic Circuit for Dynamic Switching between Feature Detection and Discrimination. *Neuron* **95**, 180-194.e5 (2017).

427. Leben, R., Lindquist, R. L., Hauser, A. E., Niesner, R. & Rakhymzhan, A. Two-Photon Excitation Spectra of Various Fluorescent Proteins within a Broad Excitation Range. *International Journal of Molecular Sciences* **23**, 13407 (2022).

428. Valenta, H. *et al.* Per-pixel unmixing of spectrally overlapping fluorophores using intra-exposure excitation modulation. *Talanta* **269**, 125397 (2024).

429. Banerjee, A. *et al.* Value-guided remapping of sensory cortex by lateral orbitofrontal cortex. *Nature* **585**, 245–250 (2020).

430. Zakiewicz, I. M., Bjaalie, J. G. & Leergaard, T. B. Brain-wide map of efferent projections from rat barrel cortex. *Front. Neuroinform.* **8**, (2014).

431. Projection-specific Activity of Layer 2/3 Neurons Imaged in Mouse Primary Somatosensory Barrel Cortex During a Whisker Detection Task | Function | American Physiological Society. *Function* https://journals.physiology.org/doi/10.1093/function/zqaa008.

432. Peters, A. J., Chen, S. X. & Komiyama, T. Emergence of reproducible spatiotemporal activity during motor learning. *Nature* **510**, 263–267 (2014).

433. Vinograd, A., Nair, A., Kim, J. H., Linderman, S. W. & Anderson, D. J. Causal evidence of a line attractor encoding an affective state. *Nature* **634**, 910–918 (2024).

434. Kozai, T. D. Y. & Vazquez, A. L. Photoelectric artefact from optogenetics and imaging on microelectrodes and bioelectronics: new challenges and opportunities. *J. Mater. Chem. B* **3**, 4965–4978 (2015).

435. Sinnen, B. L. *et al.* Optogenetic Control of Synaptic Composition and Function. *Neuron* **93**, 646-660.e5 (2017).

436. Kvitsiani, D. *et al.* Distinct behavioural and network correlates of two interneuron types in prefrontal cortex. *Nature* **498**, 363–366 (2013).

437. Académie des sciences (France), sciences (France), A. des & scientifique (France), C. national de la recherche. *Comptes Rendus Hebdomadaires Des Séances de l'Académie Des Sciences*. vol. t.9 (1839) 1–920 (publiés avec le concours du Centre national de la recherche scientifique par MM. les secrétaires perpétuels, Paris, 1839).

438. Einstein, A. Über einen die Erzeugung und Verwandlung des Lichtes betreffenden heuristischen Gesichtspunkt. *Annalen der Physik* **322**, 132–148 (1905).

439. Amroune, K., Fontolan, L., Baude, A., Robbe, D. & Bureau, I. Sparse innervation and local heterogeneity in the vibrissal corticostriatal projection. *eLife* **14**, (2025).

440. Yamashita, T. *et al.* Diverse Long-Range Axonal Projections of Excitatory Layer 2/3 Neurons in Mouse Barrel Cortex. *Front. Neuroanat.* **12**, (2018).

441. Bourassa, J., Pinault, D. & Deschênes, M. Corticothalamic Projections from the Cortical Barrel Field to the Somatosensory Thalamus in Rats: A Single-fibre Study Using Biocytin as an Anterograde Tracer. *European Journal of Neuroscience* **7**, 19–30 (1995).

442. Veinante, P., Lavallée, P. & Deschênes, M. Corticothalamic projections from layer 5 of the vibrissal barrel cortex in the rat. *Journal of Comparative Neurology* **424**, 197–204 (2000).

443. Takasaki, K., Abbasi-Asl, R. & Waters, J. Superficial Bound of the Depth Limit of Two-Photon Imaging in Mouse Brain. *eNeuro* **7**, (2020).

444. Sermet, B. S. *et al.* Pathway-, layer- and cell-type-specific thalamic input to mouse barrel cortex. *eLife* **8**, e52665 (2019).

445. Gu, M. *et al.* Rabies virus-based labeling of layer 6 corticothalamic neurons for two-photon imaging in vivo. *iScience* **26**, (2023).

446. Augustinaite, S. & Kuhn, B. Chronic Cranial Window for Imaging Cortical Activity in Head-Fixed Mice. *STAR Protocols* **1**, 100194 (2020).

447. Kline, A. M. *et al.* Large cranial windows distort sensory maps and degrade feature integration in higher-order cortex. 2026.02.13.705795 Preprint at https://doi.org/10.64898/2026.02.13.705795 (2026).

448. Zhang, H. *et al.* Bright and Photostable Voltage Sensors Derived from mBaoJin. 2025.05.30.657123 Preprint at https://doi.org/10.1101/2025.05.30.657123 (2025).

449. Lakunina, A. *et al.* Neuropixels Opto: Combining high-resolution electrophysiology and optogenetics. 2025.02.04.636286 Preprint at https://doi.org/10.1101/2025.02.04.636286 (2025).

450. Sourmpis, C., Petersen, C. C., Gerstner, W. & Bellec, G. Biologically informed cortical models predict optogenetic perturbations. *eLife* **14**, RP106827 (2026).

451. Jazayeri, M. & Afraz, A. Navigating the Neural Space in Search of the Neural Code. *Neuron* **93**, 1003–1014 (2017).

452. Paz, J. T. *et al.* Closed-loop optogenetic control of thalamus as a tool for interrupting seizures after cortical injury. *Nat Neurosci* **16**, 64–70 (2013).

453. Vanduffel, W. & Li, X. Exciting inhibition in primates. *eLife* **9**, e59381 (2020).

454. Li, N. *et al.* Spatiotemporal constraints on optogenetic inactivation in cortical circuits. *eLife* https://elifesciences.org/articles/48622/peer-reviews (2019) doi:10.7554/eLife.48622.

455. Dimidschstein, J. *et al.* A viral strategy for targeting and manipulating interneurons across vertebrate species. *Nat Neurosci* **19**, 1743–1749 (2016).

456. Jager, P. *et al.* Dual midbrain and forebrain origins of thalamic inhibitory interneurons. *eLife* **10**, e59272 (2021).