Connectivity and integrative properties of layer 5 neurons in the mouse visual cortex

Alessandro Roberto Galloni

A dissertation submitted for the degree of

Doctor of Philosophy

University College London

and

The Francis Crick Institute

June 2020

Declaration

I, Alessandro Galloni, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Signed,

Alessandro Galloni
20th June, 2020

Abstract

A central question in neuroscience is how expectations shape sensory processing. In the cerebral cortex, feedforward connections convey sensory information, while feedback connections are thought to be crucial for directing attention, signaling contextual information, and enabling perceptual inference. Thick-tufted layer 5 (ttL5) pyramidal neurons play an important role in integrating sensory, internal, and motor information. They have large complex morphologies that receive thousands of synaptic inputs from across the brain, and express a diverse set of active conductances which support highly nonlinear dendritic computations. Their axons also form the largest output pathway from the cerebral cortex, making them ideally positioned to integrate and summarise cortical computations to drive behaviour.

In this thesis, I perform whole-cell patch clamp recordings to characterize the intrinsic properties of a population of ttL5 neurons in the medial secondary visual cortex (V2m) of mice, genetically labelled in the Glt25d2-Cre line. These neurons are found to form a homogeneous population with integrative properties that are broadly consistent with several known features of ttL5 neurons.

The inputs to Glt25d2-Cre neurons in V2m are found to originate in several brain regions, associated with both sensory and internal representations. Using optogenetics and spatially patterned optical stimulation, I map the spatial distributions of synapses from these regions on the dendritic trees of the ttL5 neurons. These inputs target distinct dendritic domains in a pattern that argues against classical notions of hierarchical connectivity.

By electrically stimulating the largest input populations, I show that these pathways display synaptic facilitation and sum linearly across a wide range of temporal intervals. Furthermore, I find that Ca^{2+}-mediated supralinearities, such as backpropagation-activated spike bursts, which have been extensively described in primary visual and somatosensory cortex, are absent in V2m neurons. Finally, through a combination of electrophysiological recordings, morphological reconstructions, and computational modelling, I show that dendritic excitability in ttL5 neurons is modulated by apical trunk length.

In summary, I thoroughly describe the properties of a new ttL5-specific mouse line and provide new evidence that ttL5 neuron properties and canonical notions of hierarchical connectivity are not universally applicable throughout the cortex.

Impact Statement

The research presented in this thesis has brought new insights into the connectivity and intrinsic properties of thick-tufted layer 5 pyramidal neurons in the secondary visual cortex of mice. My experiments have revealed surprising patterns of brain-wide connectivity to these cells and general principles by which morphology can influence the excitability of neurons. These results provide evidence against the canonical view of hierarchical cortical connectivity and computational uniformity of the cerebral cortex. In the process of these experiments, I have also discovered some previously unknown methodological limitations of optogenetic stimulation using multiple opsins.

Beyond the immediate scientific relevance to my field, these findings could also have broader implications to other fields. Layer 5 pyramidal neurons are generally considered a key component in the computations of the brain, and understanding them may thus play an important role in developing treatments for the many neurological conditions that involve the cerebral cortex. Furthermore, some areas of recent research in artificial intelligence have been heavily inspired by the connectivity and properties of this class of cells. A deeper understanding of these features could thus accelerate progress in this field towards developing truly intelligent artificial systems.

Overall, by thoroughly characterizing a novel genetically defined population of layer 5 pyramidal neurons, I have laid the groundwork for future experimental and computational investigations of these cells. This new knowledge brings us closer towards the ultimate goal of obtaining a complete mechanistic understanding of the functions and computational principles of the brain.

Acknowledgements

Throughout my PhD, I have been fortunate to have benefited from the generous help and support of many people. I wish to thank the scientific collaborators, colleagues, friends, and family who have directly or indirectly contributed to make this work possible.

I am grateful to the Francis Crick Institute and the Boehringer Ingelheim Fonds for generously funding me throughout this project and for providing me with many valuable training and career opportunities. In particular, the biological research facility and the light microscopy facility at the Crick were an indispensable resource.

Ede Rancz, my principal supervisor, provided invaluable support and advice throughout my PhD, helping me learn to design and perform rigorous experiments. Thank you for creating a collaborative and stimulating environment, for giving me the freedom to explore my ideas, and for giving me the honor of being the first member of the lab.

Andreas Schaefer, my co-primary supervisor, for his insightful questions, guidance, and enthusiasm. My thesis committee—Marco Beato, Vassilis Pachnis, and Denis Burdakov—for useful feedback throughout the development of this project and for helping me stay on track. Members of the Neuroscience Interest Group and the Division of Neurophysiology at the Francis Crick Institute, for providing feedback on my work and creating a supportive academic atmosphere. Molly Strom provided the Chronos Cre-OFF virus used in chapter 4. Arnd Roth, Lee Fletcher, and Alexandra Tran-Van-Minh gave me useful advice and comments on the experiments and modelling in chapter 6.

Anja Schmaltz and Sara Rolle, for showing me the ropes when I was learning how to patch-clamp, patiently putting up with all my questions, and keeping me company during the long days at the rig. Celia Garau, for always being by my side and letting me bounce off ideas and random thoughts.

Members of the Rancz lab: Zhiwen Ye, Alexandra Tran-Van-Minh, Aeron Laffere, and Xavier Kano, for making the lab atmosphere fun and relaxed, and for the many late-night discussions. Zhiwen provided data for some of the figures in this thesis and Aeron collaborated extensively with me on the computational modelling in chapter 6. Members of Gamma Rhythm, for giving me an excuse to get out of the lab and make some noise. Your support and friendship have made the last four years fly by. Special thanks to Zhiwen, who also helped proofread most of this thesis and who has been there for me through the highs and the lows.

Jason Rihel and Nicolò Musner, for shaping my first experiences as a scientist during my undergraduate years, teaching me many valuable lessons, and inspiring me to pursue a PhD.

Many people have also helped me develop and run public engagement activities, which have been a fun diversion from my research: Jenny Jopson, Carles Bosch, Andrew Erskine, Tobias Ackels, Alexandra Tran-Van-Minh, Cristina Marin, Tom Warner, and Justina Yeung.

Finally, I would like to thank my family and the people that are closest to my heart, for their spiritual and intellectual encouragement and companionship, for being there for me when I needed it, for always believing in me, and for shaping me into the person I am today: Marta, Roberto, Michela, Daniele, Catherine, and Yajing.

Contents

1 Introduction 19

1.1 Anatomy of the visual cortex: functional specialization 22

1.2 Anatomy of the visual cortex: hierarchical organization 23

1.3 The organization of visual areas in mice 25

1.4 Integrative properties of neurons 28

1.4.1 Electrical properties of biological membranes 29

1.4.2 Backpropagating action potentials 30

1.4.3 Synaptic potentials and dendritic excitability 32

1.5 Classification of L5 pyramidal neurons 34

1.6 Theories of cortical function 36

1.6.1 Deep neural networks 36

1.6.2 Canonical cortical columns 38

1.6.3 Predictive coding 39

1.7 Aims & approach 42

2 Materials & Methods 45

2.1 Animals 45

2.2 Stereotaxic surgeries 45

2.3 Serial-section two-photon tomography 48

2.4 Electrophysiology 49

2.4.1 Slice preparation 49

2.4.2 Whole-cell recordings 50

2.4.3 Data acquisition & analysis 52

2.4.4 Electrical stimulation 52
2.4.5 Optogenetics 53
Patterned illumination 53
Full-field perisomatic illumination 53
2.5 Immunohistochemistry, imaging & morphological reconstructions 54
2.6 Biophysical modelling 54

3 Intrinsic physiology 57
3.1 Introduction 57
3.2 Chapter Methods 59
3.2.1 Cell selection & recording conditions 59
3.2.2 Whole-brain cell mapping 59
3.2.3 Sample size determination 60
3.3 Results 62
3.3.1 Brain-wide distribution and projection targets of Glt neurons 62
3.3.2 Morphologies & recording location 64
3.3.3 Subthreshold properties 67
3.3.4 Suprathreshold properties 71
3.4 Discussion 75

4 Subcellular circuit mapping 79
4.1 Introduction 79
4.2 Chapter methods 81
4.2.1 Whole-brain input mapping 81
4.2.2 Cell selection 81
4.2.3 Recording conditions 83
4.2.4 sCRACM implementation and analysis 84
4.3 Results 88
4.3.1 Subcellular distribution of RSC inputs 88
4.3.2 Subcellular distribution of V1 inputs 92

4.3.3 Subcellular distribution of OFC inputs 95
4.3.4 Subcellular distribution of inputs from the anterior thalamic nuclei 98
4.3.5 Subcellular distribution of local V2m inputs 101
4.3.6 Comparison of presynaptic inputs maps 105
4.3.7 Axonal distributions of presynaptic pathways to Glt neurons 108
4.4 Discussion 110

5 Synaptic physiology 119

5.1 Introduction 119
5.2 Chapter methods 122
5.2.1 Stereotaxic surgeries 122
5.2.2 Cell selection & recording conditions 122
5.2.3 Electrical stimulation 123
5.3 Results 125
5.3.1 Electrical stimulation 125
Paired-pulse facilitation of individual inputs 125
Independent summation of RSC and V1 inputs 127
5.3.2 Full-field optogenetic stimulation 129
Single injection control experiments 129
Paired-pulse stimulation 132
Temporal interactions of Chronos and ChrimsonR 134
5.4 Discussion 137

6 Dendritic excitability & biophysical modelling 145

6.1 Introduction 145
6.2 Chapter methods 148
6.2.1 Cell selection 148
6.2.2 Recording conditions 149
6.2.3 Electrical stimulation 149

6.3 Results 152
    6.3.1 Thick-tufted L5 neurons in V2m lack BAC firing 152
    6.3.2 Thick-tufted L5 neurons in V2m lack a critical frequency ADP 154
    6.3.3 BAC firing is absent in short biophysical models 162
    6.3.4 Active propagation enhances tuft voltage in neurons with long apical dendrites 168
6.4 Discussion 175

7 General discussion 181

7.1 Hierarchical connectivity 183
7.2 Canonical cortical microcircuits & computations 186
7.3 Functions of dendritic plateaus 190
7.4 Conclusions 193

List of Figures

3.1 Brain-wide distribution of Glt25d2-Cre labelled neurons 61
3.2 Axon projection targets of Glt25d2-Cre labelled neurons 63
3.3 Anatomy of recorded Glt25d2-Cre neurons in V2m 65
3.4 Quantification of morphological parameters 66
3.5 Subthreshold properties 68
3.6 Frequency-dependent properties 70
3.7 Suprathreshold properties 72

4.1 Monosynaptic retrograde rabies labelling of inputs to Glt neurons in V2m 82
4.2 Subcellular optogenetic circuit mapping 85
4.3 Brain-wide distribution of rabies-labelled presynaptic neurons . 89
4.4 Subcellular distribution of RSC inputs 91
4.5 Influence of total RSC input on sCRACM maps 92
4.6 Subcellular distribution of V1 inputs 93
4.7 Influence of total V1 input on sCRACM maps 94
4.8 Subcellular distribution of OFC inputs 97
4.9 Influence of total OFC input on sCRACM maps 98
4.10 Subcellular distribution of anterior thalamus inputs 99
4.11 Influence of total anterior thalamus input on sCRACM maps . . 100
4.12 Control injection of Chronos Cre-OFF virus 102
4.13 Subcellular distribution of local inputs 103
4.14 Influence of total V2m input on sCRACM maps 104
4.15 Location of inputs relative to laminar position of neurons 106

4.16 Average sCRACM heatmap profiles 107
4.17 Average sCRACM heatmap profiles 109

5.1 Processing and analysis of voltage traces 123
5.2 Electrical paired-pulse stimulation of RSC and V1 126
5.3 Independent summation of RSC and V1 inputs 128
5.4 Dual opsin optogentic stimulation of RSC and V1 130
5.5 Single opsin control experiments 131
5.6 Optogenetic paired-pulse stimulation of Chronos and ChrimsonR 133
5.7 Temporal interactions of Chronos and ChrimsonR 134
5.8 Subthreshold suppression of ChrimsonR responses with blue light 136

6.1 Retrograde labelling of LP-projecting ttL5 neurons 150
6.2 Voltage responses to L1 stimulation 152
6.3 BAC firing in V1 and V2m 153
6.4 Critical frequency ADP 155
6.5 cfADP summary data 156
6.6 Clustering of maximum ADP integrals 157
6.7 Critical frequency ADP across $CaCl_2$ concentrations 158
6.8 Long current steps in V1 and V2m neurons 160
6.9 Correlation between dendritic length and excitability in V1 and V2m 161
6.10 BAC firing in morphologically detailed models 163
6.11 Influence of Ca^{2+} hot spot dimensions on BAC firing 164
6.12 Strong dendritic current can trigger Ca^{2+} plateaus in short morphologies 165
6.13 BAC firing in reduced compartmental models 166
6.14 Tuft voltage increases with trunk length 167
6.15 Length-dependent backpropagation depends on dendritic Na^+ channels 168

6.16 Backpropagation of single spikes in active and passive trunks of different length 169
6.17 Tuft voltage increases with trunk length independently of conductance gradients 170
6.18 Length-dependent backpropagation is independent of I_h 172
6.19 Passive effect of increasing trunk depolarization width 173
6.20 Effect of axial resistance on voltage propagation 174
6.21 Peak tuft voltage depends on trunk integral 176

Chapter 1

Introduction

Animals perform a wide range of behaviours and are remarkably flexible in adapting to rapidly changing environments. This is made possible by the nervous system, whose fundamental task is to consistently produce behaviours that maximize an animal's ability to survive and reproduce. Animals need to collect information about the environment, extract relevant features, and use these to guide their actions. From a computational perspective, a neural system needs to overcome many challenging problems to perform these tasks effectively.

Firstly, it must be able to accurately perceive the surrounding world. Aside from the problem of sensory transduction (i.e. the conversion of an external physical signal to an internal electrical signal) the nervous system needs to transform and interpret the sensory signals based on context and on the reliability of these signals. Any individual sensory organ gives at best an incomplete and often noisy estimate of an environmental variable. The brain must therefore infer the state of the surrounding world by making appropriate assumptions and by combining the evidence gleaned from different senses. For example, when searching for food or shelter an animal needs to navigate accurately and be able to estimate where it is within an environment. If it is relying on vision, it must be able to distinguish different objects and estimate their shapes, sizes and relative distances. As the eyes can only see two-dimensional projections of the actual three-dimensional surroundings, various types of inductive biases are required. For example, it might be useful to assume that objects that appear smaller are generally further away (which in turn requires some expectation about typical sizes for each type of object). As lighting conditions change throughout the day, the animal may need to dynamically adjust the weight it gives to different senses, relying less on vision and more on auditory or olfactory cues instead.

Secondly, the nervous system needs to determine long- and short-term objectives and choose actions that will maximise the probability of achieving them. This requires having detailed representations of the external environment, the internal state of the animal, and being able to simulate a range of possible decisions it could make. It needs to evaluate the expected outcomes of these decisions based on learned risks and rewards associated with different actions, and balance its priorities in a way that yields the greatest overall benefit. When traversing an unfamiliar environment, for instance, an animal must weigh up its reasons for exploring (such as needing to find food) with the anticipated risks of doing so (like the chance of encountering a predator).

Finally, there is the challenge of motor control, i.e. of physically executing the desired actions as accurately as possible. This task is far more difficult than one might intuitively think, as the environment we move through is highly dynamic and often unpredictable. Even in a controlled and predictable environment, to perform an action as simple as lifting a glass of water to drink we need to dynamically adjust the force with which we hold the glass based on its weight, texture and mechanical properties. Likewise, moving across an icy surface or unstable ground may require rapidly shifting body weight to avoid falling over.

Each of these three processing steps—sensing, deciding and acting—must be enacted while operating under both time and energy constraints. While arguably all behaviours on some level involve sensory-motor coupling, different amounts of contextual information are required depending on what the priorities are for a given objective. When rapid execution is critical and minimal decision making is required, the sensorimotor loop can take place entirely in peripheral circuits with little context modulation. This is the case for reflex arcs like the withdrawal reflex, where the action itself is highly stereotyped and is controlled entirely by circuits in the spinal cord. As speed may be important to prevent tissue damage, such sensorimotor loops can be executed in as little as 100–500 ms (Sandrini et al., 2005).

To perform more complex and demanding tasks effectively, such as longer sequences of motor actions or decisions that require some amount of contextual information, longer sensorimotor loops are recruited. These sacrifice processing time for increased flexibility and capability to smoothly integrate sensory, motor, and internal information to guide goal-directed behaviour. Many such tasks can be performed at the level of the midbrain and basal ganglia, where both innate and learned behaviours are processed. For example, these structures are involved in innate escape behaviours in rodents, where the decision to freeze or flee depends on whether there is a suitable shelter nearby (Evans et al., 2018).

To obtain even more nuanced control over decisions and actions that might require more detailed perceptual discrimination, sophisticated planning and reasoning abilities, or a finer level of motor control, the sensorimotor loops that are recruited may also require the cerebral cortex. This brain structure evolved in the first small mammals around 200 million years ago (Glenn Northcutt and Kaas, 1995; Rakic, 2009) and has undergone the greatest enlargement in humans compared to other animals. While the exact functions of cortex are still debated, it generally appears to play a crucial role in sensory perception, motor learning, memory, and higher cognitive function (Courtney et al., 1998; Glickfeld et al., 2013; Kawai et al., 2015).

1.1 Anatomy of the visual cortex: functional specialization

A core aspect of the cortex's ability to smoothly integrate sensory, motor, and internal information is the large number of recurrent connections, both within and across brain areas. To fully understand what role the computations within the cerebral cortex play in brain function, as well as how these computations are implemented, it is thus useful to study the organization of the cerebral cortex in terms of both function and of the biological components and circuitry that underlie it.

Functionally, cortex is highly diverse and has many subdivisions, each specializing in processing a distinct type of information, such as visual, auditory, or motor signals. On the surface, this fact is seemingly in contrast with the relative uniformity of the gray matter, which on first inspection appears fairly homogeneous in structure. The first evidence of structural subdivisions within cortex was found in 1782, when anatomical studies of the occipital lobe in fixed human brain tissue revealed a small but distinct white line (called the "stria of Gennari") through the middle of a region of cortex which became known as the striate cortex (Glickstein and Rizzolatti, 1984). This region was later found to be specialized for visual functions, based on the selective loss of vision in animals and human patients with lesions to the striate cortex (Glickstein and Whitteridge, 1987). Furthermore, the mapping of visual space onto this area was found to be retinotopically organized, meaning that adjacent locations on the retina (and thus in visual space) are mapped continuously onto adjacent locations within the striate cortex (Holmes, 1918). In parallel to the lesion studies, the advent of Nissl staining revealed a distinctive and stereotyped layered structure to the cerebral cortex, composed of approximately six horizontal layers of neurons, each containing a set of characteristic cell types with stereotyped morphologies, and patterns of connectivity and gene expression. Detailed characterization of the slight differences in the appearance of the layer across the broad cortical surface (including differences in cell size, packing density, and relative size of the different layers) made it possible for a large number of distinct cortical areas to be outlined on the basis of this cytoarchitecture (Brodmann, 1909; Amunts and Zilles, 2015). Within this new cytoarchitectonic mapping, the striate cortex became known as Brodmann area 17.

The fact that damage to the area 17 resulted in blindness, in conjunction with the prominent direct input to this area from the retina (via the thalamus) and its distinctive cytoarchitectonic appearance, led to the prevailing belief in the early 20th century that all visual processing occurred entirely within this region. While there existed early evidence that some forms of vision are still possible after ablation of area 17 (a phenomenon known as "blindsight"), this was largely dismissed by most physiologists at the time (Blythe et al., 1987; Zeki, 1990). However, as extracellular recordings of individual neurons in live animals became more common, it became increasingly clear that neurons outside this area could also be visually responsive. Specifically, two other cytoarchitectonic regions, called Brodmann area 18 and 19, form concentric rings along the cortical surface surrounding area 17. These areas were also found to contain visually responsive cells with retinotopic maps covering the full visual field (Hubel and Wiesel, 1965; Tusa et al., 1979; Albus and Beckmann, 1980). Brodmann area 17 thus started being referred to as "V1", indicating it as the primary, but not the only, visual area (Thompson et al., 1950; Hubel and Wiesel, 1998).

1.2 Anatomy of the visual cortex: hierarchical organization

The concept of hierarchy is a central feature of nearly all modern theories of visual processing. The early evidence for this was largely derived from a combination of anatomical tract tracing and single unit recordings in the different visual areas. The pathway connecting the retina to the striate cortex, via the optic nerve, lateral geniculate nucleus of the thalamus (LGN), and optic radiation, had been known since the anatomical work of Probst (1906) and Meyer (1907), based on stains of degenerating nerve fibres following lesions (Van Buren and Baldwin, 1958; Strich, 1968). However, the notion of hierarchical processing was originally suggested by Hubel and Wiesel (1962, 1965) to account for the progressive specialization of the receptive fields (i.e. the stimuli that individual neurons are most responsive to) of cells in the cat visual system as one moves further away from the retina. Indeed, while neurons in the retina and LGN have circular receptive fields with concentric regions of opposing light selectivity (Kuffler, 1953; Hubel and Wiesel, 1961), neurons in V1 displayed elongated receptive fields that responded strongly to oriented edges (Hubel and Wiesel, 1959, 1962), and more complex receptive fields were found rostral to V1 in Brodmann area 18 and 19 (which became known as V2 and V3) (Hubel and Wiesel, 1965). This organization was proposed to arise from a serial feedforward connectivity architecture that pooled responses of lower-level neurons to generate progressively more specialized responses (Serre et al., 2007).

Over time, improved methods for anatomical circuit tracing and functional recordings from the extrastriate visual areas revealed that Brodmann area 19 was in fact not a single visual area but contained multiple complete representations of the full visual field, and that visually responsive neurons could be found even beyond this area (Cragg, 1969; Zeki, 1969, 1973; van der Want et al., 1997). These visual areas were found to be highly selective for a variety of specialized visual features, such as color or motion. The view of visual processing as a serial sequence of transformations was thus replaced by the notion of parallel processing, routed from V1 and V2 to the large number of surrounding specialized visual areas (Zeki, 1978).

While this framework was still compatible with a hierarchical scheme, the relative order of different parallel areas is harder to define, particularly in the presence of recurrent connections between theses areas. To obtain an objective metric for estimating hierarchical relationships between different visual areas, one solution has been to delineate the hierarchy on the basis of anatomical criteria. In particular, tract tracing between different visual areas revealed asymmetrical connectivity patterns in the laminar projections, which appeared to accurately reflect the hierarchical relationships that had previously been defined by receptive field properties and distance from V1 (Tigges et al., 1977; Rockland and Pandya, 1979; Tigges et al., 1981; Maunsell and Essen, 1983; Cauller et al., 1998; Markov et al., 2014). The connections referred to as "feedforward" (i.e. directed from V1 towards more rostral cortical areas) were found to originate in pyramidal neurons of layers 2 and 3 (L2/3) and terminate broadly in L2/3 and L4. Meanwhile, the reciprocal "feedback" connections (i.e. projecting caudally in the direction of V1) originate mostly in deep layers (L5 and L6) and terminate in both superficial and deep layers of the hierarchically lower region, with the most superficial layer (L1) receiving particularly strong feedback projections. By extrapolating from the layer-specific connectivity patterns between early visual areas, whose hierarchical relationships had been defined based on criteria such as receptive field properties and rostro-caudal location, it was thus possible to establish a complete hierarchical map of the known visual areas, spanning multiple layers of depth along the parallel and recurrently connected processing streams (Van Essen and Maunsell, 1983; Felleman and Van Essen, 1991).

1.3 The organization of visual areas in mice

Most early research on the visual system, including the delineation of visual areas and hierarchies, was conducted on cats and primates due to their highly developed visual systems. In recent decades, however, rodents have also emerged as a prominent model for studying vision as well as general cortical function, largely because of the tractability of performing genetic manipulations and other sophisticated experiments in these animal to dissect the circuit mechanisms involved in visual processing.

As in other mammals, visual input in mice is transmitted from the retina to the LGN and then to V1, where receptive fields are retinotopically organized. Higher visual areas in mice were originally also defined based on cytoarchitecture and electrophysiological recordings (Caviness, 1975; Wagor et al., 1980). This revealed two separate secondary visual areas surrounding V1 in a medial band and in a region anterior and lateral to V1. The atlas created by Paxinos and Franklin (2007), which has long been one of the most commonly used mouse brain atlases, refers to these cytoarchitectonically defined medial and lateral secondary visual areas as V2m and V2l, respectively.

A more detailed subdivision of these regions was subsequently described by Wang and Burkhalter (2007), who combined extracellular microelectrode recordings with strategically positioned injections of three different anterograde fluorescent tracers across V1 to create a complete map of the mouse secondary visual areas. This revealed that the mouse secondary visual cortex is composed of nine distinct visual areas surrounding V1, each organized retinotopically and covering the full visual field. In particular, in this mapping two visual areas referred to as the anteromedial (AM) and posteromedial (PM) visual areas approximately cover the same area of cortex as the previously described medial secondary visual cortex (V2m). While not distinguishable in the cytoarchitecture, one method to see the boundary between these visual areas is through optical imaging of intrinsic signal across the cortex during presentation of visual stimuli, which allows simultaneous visualization of the full retinotopic mapping of multiple visual areas (Kalatsky and Stryker, 2003; Garrett et al., 2014). By combining this method with recordings of individual neuron responses to a variety of visual stimuli, the higher order visual areas of mice have been shown to be functionally specialized, albeit to a lesser extent than in the primate visual system (Andermann et al., 2011; Marshel et al., 2011; de Vries et al., 2020).

Using the same principle of asymmetrical laminar connectivity as in monkeys, a mapping of the visual hierarchy in rats has suggested that there are at least 3 levels, in which the areas homologous to the mouse AM and PM are found at the top (i.e. the deepest level) of the hierarchy (Coogan and Burkhalter, 1990, 1993). More recently, a partial mapping has also been done in mice, where a similar hierarchical ordering appears to exist (D'Souza et al., 2016).

Overall, the mouse visual cortex thus appears to be organized in parallel hierarchical streams, similar to what was observed in primates and cats, although with notably fewer visual areas with less specialized receptive fields and a shallower hierarchy. These differences may be due to constraints imposed by the smaller total size of the cortical sheet in mice, favouring fewer areas performing a wider range of functions and a more compressed hierarchy of receptive field properties (Huffman et al., 1999; Lindsey et al., 2019).

When it comes to delineating the different visual areas in mice, one source of ambiguity comes from the involvement of these same areas in other functions. Indeed, in recent years several papers have highlighted the many top-down inputs to the visual cortex, which drive neural activity along several non-visual dimensions, with features representing motor, spatial, and other internal state-related signals (Saleem et al., 2013, 2018; Stringer et al., 2019). These non-visual features are also strongly represented in a region of cortex known as the posterior parietal cortex (PPC), which borders the anterior side of several visual areas. The PPC is typically studied in relation to non-visual behaviours, such as spatial processing, decision making, and movement planning. This area was originally ascribed to the cytoarchitectonically defined area 7 in rodents (Kolb and Walkey, 1987), and is also defined by its anatomical location, between the visual and somatosensory cortex, and by its connectivity patterns (Hovde et al., 2019). However, due to variability in the exact labelling methodologies and features used to define the boundaries, both nomenclature and anatomical definitions of PPC are often inconsistent across different research groups.

A recent review of the anatomy and functions of the PPC (Lyamzin and Benucci, 2019) highlighted the discrepancies between the three most commonly used mouse brain atlases, namely the Paxinos and Franklin (2007) atlas, the Allen Reference Atlas (Dong, 2008), and the Allen Common Coordinate Framework (Wang et al., 2020). In particular, the Allen Reference Atlas places the boundary of PPC more posterior than either the Common Coordinate Framework or the Paxinos and Franklin (2007) atlas, making it overlap with the anterior portions of visual area AM and V2m, in the respective atlases. Thus, the associative and multimodal function of many neurons within V2m and AM in combination with the variability in anatomical definitions indicate that many studies targeting the PPC likely include recordings from these higher-order visual areas (Glickfeld and Olsen, 2017).

1.4 Integrative properties of neurons

Aside from the anatomy and connectivity of cortical circuits, an essential component to neural computation is how signals are integrated within individual neurons to drive action potentials (APs, also called spikes), which are ultimately the only output of the computations that take place. In the simplest possible case, neurons can be though of as point integrators, linearly summing the inputs they receive and producing a spike if the total input crosses a particular threshold or through some other nonlinear function of the membrane potential. This simplistic view of a neuron's operation is frequently used in computational neuroscience research and is often sufficient to model a variety of single-cell and network properties (Troyer and Miller, 1997; Izhikevich, 2003; Vogels and Abbott, 2009; Teeter et al., 2018; Li et al., 2019). It is also the basis of most artificial neural networks, which have been successful on specific tasks, such as computer vision (LeCun et al., 2015). However, since the work of Ramon y Cajal (1909) over a century ago, it has been well documented that biological neurons display huge morphological and functional diversity. While different pathways can target different populations of neurons, individual neurons themselves can receive multiple streams of sensory and internally generated information, which often impinge on functionally distinct dendritic compartments. Furthermore, most neurons have a diverse selection of voltage gated ion channels distributed throughout their dendrites. These endow dendrites with active properties that substantially increase the computational power of individual neurons (London and Hausser, 2005; Stuart and Spruston, 2015), allowing them for instance to operate as multilayer neural networks (Poirazi et al., 2003; Hausser and Mel, 2003; Harnett et al., 2013), and potentially enabling brain-wide learning algorithms that would otherwise be intractable (Guerguiev et al., 2017; Sacramento et al., 2018). It is therefore likely that neurons in the cortex exploit this structure to perform more nuanced transformations than would be possible by simple integrate-and-fire neurons alone.

1.4.1 Electrical properties of biological membranes

Without considering active conductances, the membrane of a neuron exhibits electrical properties similar to electronic resistor-capacitor (RC) circuits (Johnston and Wu, 1995). As the information processing capacity of neurons is generally believed to be primarily mediated by electric signals (as opposed to many non-neuronal cell types which communicate largely through diffusion of chemical signals), this abstraction of a neuron in terms of simple electronic components is very useful, allowing us to temporarily overlook the immense complexity of the underlying molecular dynamics. The specific membrane capacitance typically does not vary much and takes a value close to $1\,\mu F/cm^2$ in most neurons, making the overall capacitance (C_m) primarily determined by the total membrane (i.e. the size) of the cell. The membrane resistance (R_m) is dictated by the permeability of this membrane to ions, which depends on the presence of transmembrane ion channels. When an inward current flows across the membrane (typically carried by Na^+, K^+ or Ca^{2+} ions), this results in a change in the intracellular membrane potential (V_m) reative to the extracelluar medium that approximately obeys Ohm's law ($V = I{\cdot}R$). The time course of this change is described by an exponential curve with time constant $\tau = \mathrm{R_m} \cdot \mathrm{C_m}$. Beyond these passive properties of neurons, the main determinants of electrical signals in neurons are the voltage- and time-dependent changes in ion channels.

At resting state (i.e. in the absence of external stimuli) the membrane potential of a neuron is dictated by the intracellular and extracellular concentrations of ions and by their relative permeability across the cell membrane. Concentration gradients generated by transmembrane ion transporters (primarily the $\mathrm{Na^+/K^+}$ ATPase) are generally such that the voltage inside the cell is negative relative to the extracellular environment. If external stimuli cause the membrane potential to temporarily depolarize, this can trigger a rapid spike in the voltage. Such spikes are the main form through which the vast majority of neurons signal to each other. The most common form of spikes are those generated near the soma, in the axon initial segment, where the membrane contains a high density of voltage-gated $\mathrm{Na^+}$ channels. If the membrane potential becomes high enough to open these channels, a large influx of $\mathrm{Na^+}$ ions causes the intracellular potential to briefly become positive (generally on a sub-millisecond timescale). This depolarization in turn causes the opening of voltage-gated $\mathrm{K^+}$ channels, which flow out of the cell due to their concentration gradient, thus repolarizing the membrane potential. This process is repeated with voltage-gated $\mathrm{Na^+}$ and $\mathrm{K^+}$ channels throughout the axon, such that the spike is actively propagated along the axon and ultimately induces a release of neurotransmitters at the axon's synaptic terminals.

1.4.2 Backpropagating action potentials

In addition to propagating along the axon, voltage spikes in the soma also propagate backwards into the dendrites. In the absence of any other ion channels, this voltage propagates passively, resulting in an exponential attenuation of the voltage over distance. However, the dendrites of biological neurons are rarely entirely passive, and often contain many different types of ligand-gated and voltage-gated ion channels. These include $\mathrm{Na^+}$ and $\mathrm{K^+}$ channels (albeit at a lower concentration than in the axon) as well as Ca^{2+} channels, which can convey some particularly interesting properties to the neuron. This has been most carefully studied in a type of large pyramidal neurons which are found in the deeper portion of cortical layer 5, called thick-tufted layer 5 (ttL5) pyramidal neurons.

These cells have a large apical dendrite, which ends in a dendritic tuft in L1, and exhibit a high density of voltage-gated Ca^{2+} channels in a "hot zone" around the upper portion of this apical dendrite (Amitai et al., 1993; Yuste et al., 1994; Schiller et al., 1997). Unlike the Na^{+} and K^{+} channels mentioned above, the kinetics of these Ca^{2+} channels are slow and once activated they tend to stay open for prolonged periods of time (on the order of several milliseconds). This causes a large and sustained depolarization of the apical dendrite, called a Ca^{2+} plateau potential. Because of the voltage attenuation along the apical dendrite and the relatively high threshold to activate these channels, this response is not usually triggered by a single backpropagating AP or by local synaptic input alone. However, if a train of spikes occurs at sufficiently high frequency, the dendritic potentials sum and can thus reach a high enough voltage to elicit the Ca^{2+} plateau potential (Larkum et al., 1999a; Kampa and Stuart, 2006; Shai et al., 2015). Alternatively, the phenomenon could be elicited if a synaptic input near the top of the apical dendrite were to occur at the same time as a single backpropagating AP (Larkum et al., 1999b). This phenomenon is thus referred to as backpropagating AP activated Ca^{2+} spike firing (BAC firing).

The second possibility is particularly intriguing, as the same synaptic potential that can trigger BAC firing would otherwise have been strongly attenuated along the long apical dendrite, to the point that its influence on the somatic potential would be barely noticeable. Meanwhile, synaptic potentials in basal dendrites, which are closer to the soma and spike initiation zone, are not attenuated to the same degree and are thus more easily able to drive firing in the neuron. However, large Ca^{2+} plateau potentials in the apical dendrite do elicit large depolarizations at the soma, which can trigger a burst of somatic spikes. Since BAC firing only occurs if the backpropagating AP and synaptic input are closely matched in time, the presence of this mechanisms allows a neuron to act as a coincidence detector, amplifying the neuron's response to a stimulus if it receives simultaneous input to the basal and apical dendrites. This is thought to be particularly relevant in light of the layer specificity of long range projections in cortex, with L1 thought to receive a large number of top-down feedback signals.

1.4.3 Synaptic potentials and dendritic excitability

At most synapses, neurons do not directly transfer the presynaptic AP to the postsynaptic neuron (with the exception electrical synapses, where gap junctions allow currents to flow directly from one cell to the next). Instead, inputs from other neurons result in postsynaptic changes in the membrane potential through the binding of neurotransmitters to receptors on the postsynaptic membrane. The amplitude of an excitatory postsynaptic potential (EPSP) depends on the number of ion channels opened, the driving force (i.e. the concentration gradient) of the ions they are permeable to, and the local input resistance ($\mathrm{R_N = V_m/I}$). When multiple inputs arrive in brief succession, their summation depends on the location of the synapses, the type of ion channels activated, and the magnitude of the voltage change.

Small potentials originating on different dendrites tend to have little or no interaction, and therefore sum linearly at the level of the soma. However, if two stronger inputs impinge in brief succession nearby on the same dendrite, the reduced membrane resistance caused by open ion channels after the first EPSP can have a shunting effect on the second EPSP, resulting in an overall smaller depolarization (i.e. in sublinear summation). This effect may be counteracted if the combined EPSPs are sufficiently large to activate additional voltage sensitive ion channels that can increase the total amount of inward current and thus further depolarize the neuron. One example of such a supralinear interaction occurs when NMDA receptors are recruited. These ion channels are ligand-gated, but also require a strong depolarization to open, as at hyperpolarized potentials they are normally blocked by extracellular Mg^{2+} ions. The non-linear effects of NMDA receptor activation can make dendrites sensitive to the spatial and temporal sequence of input activation (Branco et al., 2010).

Another important mechanism for controlling excitability in neurons is the hyperpolarization-activated cyclic nucleotide-gated (HCN) channel, which mediates a voltage dependent cation current referred to as I_h (Wahl-Schott and Biel, 2009). These channels are active at resting membrane potentials and conduct mainly Na^+ ions under physiological conditions. The atypical nature of their activation results in channels closing when the membrane is depolarized, thus reducing the inward current and causing a counteracting hyperpolarization. I_h activation is also not instantaneous, increasing with an exponential time course that is faster at hyperpolarized potentials (McCormick and Pape, 1990). This conductance therefore acts as a stabilizing influence on the cell, allowing EPSPs to transiently depolarize the neuron while reducing the influence of sustained depolarizations or hyperpolarizations.

Notably, while HCN channels are strongly expressed in ttL5 neurons, they are not uniformly distributed but rather have a gradient along the apical dendrite, increasing by several orders of magnitude from the soma to the apical tuft dendrites (Stuart and Spruston, 1998; Williams and Stuart, 2000; Lorincz et al., 2002; Harnett et al., 2015). By dampening synaptic potentials evoked at distal tuft dendrites, this would further accentuate the electrical compartmentalization of these inputs. This may be a key feature in the integrative properties of L5 neurons, as compartmentalization may be required for them to separately process functionally distinct input streams. This may also be an important site for more diffuse arousal-mediated modulation, as the expression of HCN channels is itself influenced by the activation of adrenergic receptors (Carr et al., 2007; Barth et al., 2008; Phillips et al., 2016).

1.5 Classification of L5 pyramidal neurons

Classification of cell types is generally based on many different criteria, including location, dendritic morphology, intrinsic electrophysiological properties, gene expression patterns, projection targets, and *in vivo* response properties. Within L5, a distinction is often made between the upper (L5a) and lower (L5b) subdivisions due to the presence of two prominent excitatory cell types, which preferentially occupy one or the other layer (Manns et al., 2004; Groh et al., 2010; Oberlaender et al., 2011; Kim et al., 2015; Lur et al., 2016; Gouwens et al., 2019). Pyramidal neurons in L5a tend to have slender apical dendrites with few tuft branches in L1, and have axons that project primarily to other cortical neurons, both locally and to more distant cortical areas. Intrinsically they are also characterised by spike train adaptation and small I_h. Meanwhile, many L5b neurons are classified as ttL5 neurons, due to the highly branched nature of their apical tuft dendrites. Unlike the L5a neurons, these cells have few local connections and instead project mainly to subcortical targets such as the striatum, brainstem, and spinal cord. Their intrinsic physiology is characterised by large I_h and increased excitability, with bursts of spikes often observed at the beginning of depolarizing current steps.

Of the many classes of neurons found in cortex, ttL5 neurons are one of the most extensively studied (Larkman, 1991; Spruston, 2008; Ramaswamy and Markram, 2015). These cells have attracted particular attention with regard to synaptic integration in part due to their particularly large dendritic trees, which span all layers of cortex and integrate thousands of synapses from local and long-range pathways. They also form the main output pathway from cortex towards subcortical targets, including motor-related circuits, placing them in a prime position to summarize the computations that have occurred in the local cortical column. Furthermore, as mentioned above these cells exhibit several interesting integrative properties, including BAC firing, which may be key to understanding how different cortical areas interact with each other.

Recently, a large number of novel Cre-expressing mouse lines have been produced as part of the GENSAT project, with specificity for different subtypes of cortical pyramidal neurons (Gong et al., 2003, 2007; Gerfen et al., 2013). These lines were made following a large-scale screen of gene expression across the central nervous system to identify candidate genes that have some cell-type specificity and create a library of bacterial artificial chromosome vectors for manipulation of these different cell types. Of these, the gene for collagen beta(1-O)galactosyltransferase 2 (Glt25d2, also known as Colgalt2) was found to be selectively expressed in a subpopulation comprised primarily of ttL5 neurons. The promoter for this gene was therefore used to create a mouse line expressing enhanced green fluorescent protein (EGFP) in these cells. Groh et al. (2010) have recently used this EGFP line, as well as a different line specific for slender-tufted L5 neurons, to compare the properties of these two types of L5 neuron between primary visual cortex and somatosensory cortex. Subsequently, a Cre line with the Glt25d2 promoter was also generated, enabling the labelled neurons to express a wide range of transgenes for probing the connectivity and function of these cells. A notable difference of the Cre line compared to the EGFP line is that the bacterial artificial chromosome was randomly inserted into a different location in the mouse genome. Specifically, it was inserted into the Y chromosome, rendering all male mice transgenic. This has the advantage of eliminating the need to genotype animals, at the cost of not being able to record data on this cell line from female animals.

In general, the use of a labelled Cre line makes it possible to take a system-wide approach to studying the connectivity and function of these ttL5 neurons by combining *in vivo* and *in vitro* physiology and genetic methods to consistently record from the same well-defined population of cells in different experiments using a wide array of different experimental tools and techniques.

1.6 Theories of cortical function

The structural similarity of cortex across functionally different regions, along with the high level of recurrent connectivity, and the observation that the cerebral cortex is organized in parallel hierarchical processing streams has led to several hypotheses about the types of algorithms that might be implemented to support perception, behaviour, and flexible learning.

1.6.1 Deep neural networks

Advances in the field of machine learning over the last few decades, in particular the development of "deep learning" in artificial neural networks, have provided a practical demonstration of the power of hierarchical processing (LeCun et al., 2015). Deep artificial neural networks, which have been the driving force of recent progress in artificial intelligence, are able to form robust and invariant representations of input data that make them very useful for performing tasks, such as object recognition in images, which are easy for humans to do but very difficult to explicitly program into a computer. Their strong performance in such tasks is achieved despite the fact that these networks are often purely feedforward and are composed of neurons without dendrites, implementing extremely simple forms of non-linear point integration. Furthermore, the internal representations formed by individual artificial neurons optimized for visual tasks bear a strong resemblance to the stimulus selectivity and receptive field properties of biological neurons within the hierarchy of visual areas (Yamins et al., 2014; Yamins and DiCarlo, 2016). Unlike neural circuits in the brain, artificial neural networks are fully accessible to investigation, i.e. every neuron and connection can be monitored simultaneously. If their computations bear any resemblance to the computations in the brain, it would make them extremely useful and convenient from an experimental point of view to bring new insights into the functions of cortex (Richards et al., 2019). The main drawback to this approach is their lack of biological realism, which severely limits the insights we can extrapolate from them to biological neurons.

In addition to the similarity in receptive field properties of biological neurons in the visual cortex and artificial neural networks trained to recognise images, some evidence that pure feedforward processing underlies a large portion of vision comes from analysing the timescales in which perception occurs. The first visually evoked spikes in the primary visual cortex of mice occur within the first 80 ms from stimulus presentation and provide sufficient information to reliably perform some perceptual decisions (Resulaj et al., 2018). Similarly, humans can discriminate faces in as little as 100 ms (Thorpe et al., 2001; Calvo and Nummenmaa, 2011). These fast timescales in which perception takes place suggests that at least some elements of visual perception can take place before the activation of additional feedback loops through the recurrent circuitry.

Two important deviations from biological neurons are the lack of dendritic compartments and the learning rules that are used to find useful synaptic connection strengths (referred to as weights). Indeed the performance of these networks relies primarily on learning optimal weights between the artificial neurons, which is typically done in a supervised way by measuring the overall error of the network relative to a desired output (e.g. it's classification of an image) and backpropagating this error through the network to adjust all the individual weights that contributed to the error. This type of learning rule is generally considered biologically implausible, as it requires modifying synapses based on information that is not locally available to those synapses.

Recent alternatives have suggested that a similar learning rule could be implemented biologically by making use of signal compartmentalization within dendrites and the non-linear processing that can occur within these dendrites. In particular, the ability to drive firing rate and burst rate independently by sending inputs to different dendritic compartments can be leveraged to send multiplexed signals that simultaneously propagate feedback errors down the hierarchy while passing sensory signals forward through the network (Guerguiev et al., 2017; Sacramento et al., 2018; Naud and Sprekeler, 2018; Payeur et al., 2020). Biologically it has been suggested that ttL5 neurons could be the key components of this type of learning within cortex, as they have a large electrotonic separation between the apical tuft and basal dendrites and stimulation of the apical dendrites can trigger BAC firing and a burst of somatic spikes.

1.6.2 Canonical cortical columns

Closely related to the idea of hierarchical processing is the notion of modularity. The mostly uniform appearance of cortex, in terms of layers, cell types, and local connectivity rules, is often seen as evidence that cortex is modular and can be split into many adjacent cortical columns, each spanning all six layers of cortex and a few hundred micrometers in width (Mountcastle, 1997; Douglas and Martin, 2004). In this framework, each module would be identical in structure and circuitry and thus be performing the same set of fundamental computations, albeit on functionally different inputs. Different functions, such as auditory, visual, and somatosensory processing, could be seen as specific applications of the same universal algorithm, which learns to extract statistical regularities in any given set of input signals and generate an internal model of the environment through transformations that enable abstraction and prediction. The primary difference between brain regions would thus lie in the input information that they are processing, while the general computation and the circuits implementing it would be largely conserved.

The remarkable flexibility and potential universality of cortical computations is supported by the finding that experimentally rewiring sensory information pathways during development (e.g. by routing visual signals from the LGN to auditory cortex instead of the visual cortex) can induce a cortical area to remap to the new type of sensory input it is receiving (Frost and Metin, 1985; Sur et al., 1988; Metin and Frost, 1989; Roe et al., 1990, 1992; Pallas, 2001). This implies that there exists a canonical microcircuit, acting as the universal building block of a hierarchical computation that is repeated with little variation across all of cortex (Mountcastle, 1997; Douglas and Martin, 2004). While the cortex contains many different types of neurons, these would be connected together in a stereotyped way to support the same information processing mechanism across different regions.

On the basis of this hypothesis, much work has gone into understanding what these universal computations might be by studying the structure of individual cortical columns. A key component of this approach involves finding canonical cortical microcircuits, i.e. particular wiring motifs between defined cell types across different layers within a region of cortex small enough to encompass the hypothetical cortical column (Shipp, 2007; Bastos et al., 2012; Harris and Shepherd, 2015). These studies have led to the common understanding that within a cortical column feedforward information generally propagates from L4 to L2/3 and then to the deeper L5 and L6, and is modulated in this pathway by a variety of inhibitory interneuron cell types with cell-type specific influences in different layers of cortex.

While the exact nature of the computations and the theoretical basis for this circuit architecture are still unknown, there are several advantages to such a modular organization. Assuming that universal computations and canonical microcircuits indeed exist without substantial variations between cortical areas, other than in the exact synaptic weights which have been learned, modularity would render the process of constructing the required circuits very efficient, as potentially a small number of genetically programmed developmental steps could be applied to encode a fixed set of wiring rules broadly across all cortical areas. Moreover, modularity typically allows for easy scalability. If it were advantageous for a nervous system to have more processing power, this could thus be achieved by simply increasing the number of cortical modules.

1.6.3 Predictive coding

One possible form of universal computation which has attracted much attention is Bayesian inference, of which the most popular class of implementation algorithms is called predictive coding (Rao and Ballard, 1999; Friston and

Kiebel, 2009). These algorithms relie heavily on both hierarchical processing and reciprocal feedforward and feedback connections between cortical areas.

While limited forms of perception are possible with purely feedforward processing, in general sensation is not a passive phenomenon. The brain acquires information about the environment by actively choosing where to look and what features to direct attention towards. Visual scenes are often inherently ambiguous, and their interpretation therefore requires implicit knowledge of natural image statistics, and integration of noisy and contradictory sensory inputs with an internal model of the world. This model itself needs constant updating to maintain an accurate representation of the current state of the environment. Although different variations of the predictive coding algorithm exist, a common feature for enabling this type of model update is the use of top-down predictions from hierarchically higher areas towards lower areas. These predictions, based on contextual expectations and natural statistics, shape the lower-level responses by inhibiting predictable activity, allowing subsequent layers to process only the components of a sensory stimulus which were unexpected and are therefore most informative for updating the internal model.

Although the exact biological implementation of predictive coding algorithms is not well understood, there is substantial evidence that this type of inference is occurring in the brain. At the behavioural level, animals often seem to update their expectations about future events in a manner that is close to the theoretical optimum calculated through Bayes' rule, which is the general form of Bayesian inference that predictive coding algorithms are thought to implement (Kording and Wolpert, 2004). Furthermore, the effects of top-down suppression of predicted activity can be seen when recording from the cortex during presentation of either predictable or surprising stimuli. For example, the phenomenon of "repetition suppression" occurs when the same stimulus (e.g. a particular image or tone) is constantly present or repeated multiple times and results in reduced responses of single neurons (Gross et al., 1967; Sawamura et al., 2006; Vogels, 2016). This suppression is itself modulated by context, such as the perceived probability of a stimulus being repeated (Summerfield et al., 2008). Likewise, the sensory stimuli that occur during self-generated movements, such as sounds produced while running or the vestibular signals generated during orienting movements, are often suppressed at early sensory stages based on the efferent copy of the motor signal, which allows the sensory consequences to be predicted (Schneider et al., 2018; Cullen, 2019). Conversely, when an expected stimulus is omitted or if the visual scene changes in an unexpected way (such as when a part of the visual field freezes) mismatch responses have been recorded, which are consistent with the hypothesised responses of error units in a predictive coding framework (Bendixen et al., 2009; Keller et al., 2012; Keller and Mrsic-Flogel, 2018).

1.7 Aims & approach

The objective of the project I describe in this thesis was to characterize the connectivity and functional properties of a specific population of Glt25d2-Cre (referred to as Glt) neurons located in the medial secondary visual cortex (V2m). The visual cortex generally is a useful model system to study hierarchical processing in the brain as there are several visual areas, processing information both serially and in parallel.

Due to their location in V2m, the Glt neurons I have studied are highly associative, receiving a diverse set of multimodal sensory and internally generated signals, and would classically be thought of as hierarchically above neurons in the primary visual cortex. I therefore aim to elucidate the cellular mechanisms through which they integrate predictive feedback and feedforward streams of information. Specifically, I aim to directly address the following questions:

1. What are the morphological and intrinsic electrophysiological features of Glt25d2-Cre neurons in V2m?
2. Are these properties consistent with the known properties of ttL5 neurons?
3. How do local and long-range synaptic inputs map onto their intricate dendritic trees?
4. In what way are feedforward sensory signals and higher-order feedback signals integrated?
5. How does morphology affect intrinsic properties and the integration of synaptic inputs?

To address these questions, I have performed whole-cell patch-clamp recordings from ttL5 neurons labelled in the Glt25d2-Cre line or using other projection-specific markers. This intracellular recording technique allows the dynamic study of sub-threshold changes in current and voltage across the cell membrane. The recordings have been done in acute brain slices of adult mice, which allow easy access to the neurons for visually guided placement of the recording electrodes while preserving the local connectivity and cytoarchitecture of the tissue. Using optogenetics and spatially patterned optical stimulation, I have studied the axons linking distant areas of the cortex to the ttL5 neurons in V2m and mapped their connectivity at the subcellular level. I have also employed a combination of electrical and optogenetic stimulation of these axons in order to understand the logic of their inter-areal interactions. Finally, I have uncovered more general principles of integration in ttL5 neurons by recording cells across different brain regions and exploring the cellular mechanisms that govern their supralinearities and integrative properties in compartmental biophysical models. By adopting this combination of computational and experimental techniques, I have gained novel insight into the computations implemented by cortical circuits and how these computations emerge from the connectivity and intrinsic properties of the underlying neural network.

Chapter 2

Materials & Methods

2.1 Animals

All animal experiments were performed under project licence 70/8935 in accordance with the guidelines and regulations of the UK Home Office under the Animals (Scientific Procedures) Act 1986, and all surgical procedures were approved by the named veterinary surgeon at the Francis Crick Institute. The mice used were all male and were transgenic from either the Tg(Colgalt2-Cre)NF107Gsat (MGI:5311719, referred to as Glt) or Tg(Rbp4-Cre)KL100Gsat (MGI:4367067, referred to as Rbp4) lines created through the Gensat project (Gerfen et al., 2013; Groh et al., 2010). Transgenic animals were crossed with the Ai14 reporter line expressing tdTomato (MGI:3809524). Animals were housed in individually ventilated cages under a normal 12-hour light/dark cycle.

2.2 Stereotaxic surgeries

Surgeries were performed on mice aged 3–8 weeks using aseptic technique under isoflurane (2–4%) anaesthesia. Following induction of anaesthesia, each animal was subcutaneously injected with a mixture of meloxicam (2 mg/kg) and buprenorphine (0.1 mg/kg). While under anaesthesia, the fur on the animal's head was shaved and the animal was placed on a heated pad that maintained its body temperature at 36–37 °C, monitored by a rectal probe, and head-fixed in a stereotaxic frame. Liquid gel sterile eye drops were placed

Area	Anteroposterior distance (mm)	Mediolateral distance (mm)	Depth from pia (mm)
RSC	[−3.2, −2.7]	0.5	[0.5, 0.7]
V1	[−2.8, −3.5]	[−2.5, −2.7]	[0.5, 0.6]
V2m	[−3.1, −2.8]	1	0.25
LP	[−2.2, −1.7]	1	2.4
OFC	[2, 2.8]	1	[1.5, 2.25]

Table 2.1: stereotaxic coordinates used for injections. Brackets indicate a range of values. Abbreviations: RSC, retrosplenial cortex; V1, primary visual cortex; V2m, medial secondary visual cortex; LP, lateral posterior thalamus; OFC, orbitofrontal cortex.

over the animal's eyes to prevent them from drying during the surgery. The shaved skin on the top of the head was thoroughly cleaned with sterile gauze swabs dipped in chlorhexidine. After verifying that the animal was in a sufficiently deep anaesthesia, determined by the lack of a withdrawal reflex to a toe pinch, an incision (approximately 8 mm long) was made on the skin above the skull. A glass microinjection pipette (pulled to a tip diameter of around 20 µm), attached to a Nanoject II (Drummond Scientific) delivery system, was positioned above the skull and lowered until it contacted either the lambda or bregma sutures. The position of the head-fixed skull was then adjusted in order to level the skull such that the difference in vertical position of lambda and bregma sutures was less that 0.1 mm. This step helped ensure consistent and accurate targeting of stereotaxic coordinates for injections. A small hole (0.5–0.7 mm) was drilled in the bone at the stereotaxic coordinate of the injection site and the dura was carefully removed with the help of fine forceps and a lance-shaped dissecting knife. The injection pipette was then slowly lowered to the appropriate depth (based on the targeted stereotaxic coordinates) and pressure injections were made at a rate of 0.4 nL/s. To reduce backflow, the pipette was left in the brain approximately 5 minutes after completion of the injection before being slowly retracted. The skin was then closed with sutures and the animal was returned to its cage to recover. The cage was placed inside a heated chamber for the first 30–60 min of the recovery.

2.2. Stereotaxic surgeries

The stereotaxic coordinates and content of the injection varied depending on the experiment. The coordinates for each brain region were determined from Paxinos and Franklin (2007). Injections were targeted to either retrosplenial cortex (RSC), primary visual cortex (V1), medial secondary visual cortex (V2m), lateral posterior thalamus (LP), or orbitofrontal cortex (OFC). The stereotaxic injection coordinates used for these regions are listed in Table 2.1. In all animals, only the left hemisphere was injected. Injections contained either Alexa Fluor 488-conjugated cholera toxin subunit B (CTB, 0.8% w/v) or an adeno-associated virus (AAV, serotype 1). The viruses contained either Chronos (also called *Stigeoclonium helveticum*, ShChR; full name: rAAV1-Syn-Chronos-GFP) or ChrimsonR (also called *Chlamydomonas noctigama*, CnChR; full name: rAAV1-Syn-ChrimsonR-tdTomato), which were obtained from UNC Vector Core (https://www.med.unc.edu/genetherapy/vectorcore/) and are described in Klapoetke et al. (2014). A second Chronos virus, referred to as Chronos Cre-OFF (full name: AAV1-EF1-creOFF-Chronos-GFP), was also used to express Chronos-GFP only in neurons that did not express Cre recombinase. In each area, injections of 100 nL of CTB or viral suspension were made at 1–3 different sites within the given coordinate range of the area. Following the injections, I waited at least 3 weeks for AAV injections and 2 weeks for CTB injections before slicing the brain.

Rabies tracing experiments (Callaway and Luo, 2015; Reardon et al., 2016) were performed by Dr. Zhiwen Ye and involved two separate injections of virus. During the first surgery, 10–20 nL of a 1:2 mixture of Cre-dependent AAVs encoding the avian tumor virus receptor A (TVA) and N2c glycoprotein (AAV1-Syn-flex-H2B-N2CG), respectively, were injected in V2m. This was followed 5–7 days later by a second surgery, during which a 50–100 nL injection of N2c rabies virus (EnvA-CVS-N2CG-mCherry) was made at the same site.

Chronos Cre-OFF, EnvA-CVS-N2cG-mCherry rabies virus, TVA and N2c glycoprotein expressing AAVs (AAV1-Syn-flex-H2B-N2CG) were a generous gift of Dr. Molly Strom.

2.3 Serial-section two-photon tomography

In order to count and quantify the whole-brain distribution of cells and axons, labelled through either Cre-dependent tdTomato expression or retrogradely transported mCherry, I was assisted by Dr. Zhiwen Ye, who performed serial-section two-photon tomography experiments.

Adult Glt mice aged between 2 and 5 months (in the case of the injected mice, 2-4 weeks after the surgery) were transcardially perfused under terminal anaesthesia with cold phosphate-buffer (PB, 0.1 M) followed by 4% paraformaldehyde (PFA) in PB (0.1 M). The skull was post-fixed in 4% PFA at 4 °C for at least 24 hours before the brain was removed.

Brain samples were subsequently embedded in 4-5% agarose (in 0.1 M PB) and imaged using serial two-photon tomography (Ragan et al., 2012). This methods involves alternately cutting 40 µm thick coronal sections and imaging eight optical sections in depth (at 5 µm intervals with 1.2 µm x 1.2 µm resolution). Excitation was provided by a pulsed femto-second laser with 800 nm wavelength (MaiTai eHP, Spectraphysics). Images were acquired through a 16X/0.8 NA objective (Nikon MRP07220) in three channels (green, red, blue) using photomultiplier tubes. Image tiles for each channel and optical plane were stitched together with a custom-written MATLAB software.

Following acquisition, the two-photon image stacks were registered to the Allen CCF ((Wang et al., 2020)). Automated mouse atlas propagation (Niedworok et al., 2016) was used for anatomical registration and segmentation of brain samples. The average template brain from the CCF was first registered to a down-sampled version of each brain sample using affine and freeform registration. The resulting transformation parameters were then applied to the CCF itself, thus transforming it onto individual sample space and accurately segmenting each sample.

After segmenting the brains according to the brain areas defined in the CCF, the image stacks were processed to automatically detect and count individual cells. Full resolution images were filtered with a Gaussian blur (sigma = 1) using Fiji (ImageJ 1.52e) to reduce image noise. Cell candidates were over-detected using a dynamic threshold and a 30 voxel cube was extracted for each putative cell. These were supplied as input to a custom machine learning algorithm provided by Dr. Charlie Rousseau, which classified the voxels into noise and cell categories. The resulting cells were counted for each segmented area in the CCF.

To register and count each cell according to the segmentation in the Allen CCF, cell coordinates were reverse-transformed onto the CCF space using the open source registration tool, Elastix (Klein et al., 2010). All cells in each region of interest were projected onto a 2D matrix. In order to visualize cell density across the brain, for each unique cell position a cell density value was calculated as total cell counts in a nearby 9x9 pixel grid, normalized to the sum of all presynaptic cells. As the Allen CCF is segmented by both area and cortical layer, these boundaries intrinsic to the atlas were used when comparing cell counts in different areas and layers.

2.4 Electrophysiology

2.4.1 Slice preparation

Adult mice (5–12 weeks old) were deeply anaesthetised with isoflurane and decapitated. The brain was rapidly removed and placed in oxygenated ice-cold slicing artificial cerebrospinal fluid (ACSF) containing (in mM): 125 sucrose, 62.5 NaCl, 2.5 KCl, 1.25 NaH_2PO_4, 26 $NaHCO_3$, 2 $MgCl_2$, 1 $CaCl_2$, 25 dextrose; osmolarity 340–350 mOsm. The cerebellum and frontal cortex were then removed manually with a coronal cut using a single-edged razor blade and the rostral surface was affixed to a metal platform with cyanoacrylate glue. Coronal slices (300 µm thick) containing visual cortex were prepared using a vibrating blade microtome (Leica VT1200S or Campden 7000smz-2), and were sliced with the brain at a slight forward angle (approximately 10 degrees) to improve preservation of L5 apical dendrites, which extend in a direction that is approximately perpendicular to the pia. Slices used for recording were all between 2.3 and 4.3 mm posterior to bregma (the most rostral slice being the first to contain a fused corpus callosum). As the slices were cut at a slight angle from the coronal plane, estimated coordinates of recorded cells relative to bregma were based primarily on the shape of the the dorsal region of cortex, which was visually compared to coronal sections in Paxinos and Franklin (2007). Slices were transferred to a submerged holding chamber with regular ACSF containing (in mM): 125 NaCl, 2.5 KCl, 1.25 NaH_2PO_4, 26 $NaHCO_3$, 1 $MgCl_2$, 1.5 or 2 $CaCl_2$ (depending on the experiment), 25 dextrose; osmolarity 308-–312 mOsm. The holding chamber was held in a water bath at 35 °C for the first 30—60 min after slicing, and was kept at room temperature (22 °C) for the remaining time (up to 12 hours) after that. All solutions and chambers were continuously bubbled with carbogen (95% O_2 / 5% CO_2).

2.4.2 Whole-cell recordings

After the 35 °C incubation period, individual slices were transferred from the holding chamber to the recording chamber, where they were perfused at a rate of 6 mL/min with regular ACSF (see above) continuously bubbled with carbogen. The ACSF was heated to 35 ± 1 °C with an inline heater for all experiments except the optogenetic circuit mapping experiments in chapter chapter 4, which were performed at room temperature (22 ± 1 °C). Depending on the experiment, the recording ACSF also contained one or more of the following: kynurenic acid (1 mM), picrotoxin (50 µM), CPG52432 (1 µM), tetrodotoxin (TTX, 1 µM), 4-aminopyridine (4-AP, 100 µM). Filamented borosilicate thick-walled glass micropipettes were pulled and heat-polished using a two-stage horizontal puller (Zeitz DMZ Univeral Electrode Puller) to obtain an electrode resistance of 3–6 MΩ. The glass electrodes were filled with intracellular solution that was optimized for recording in either current clamp or voltage clamp configuration, depending on the experiment.

Internal solution used in current clamp experiments contained (in mM): 115 $KMeSO_3$ (CH_3KO_3S), 5 NaCl, 3 $MgCl_2$, 10 HEPES, 0.05 EGTA, 3 Na_2ATP, 0.4 NaGTP, 5 K_2-phosphocreatine, 0.5% (w/v) biocytin hydrochloride, 50 µM Alexa Fluor 488 hydrazide; osmolarity 290-295 mOsm; pH adjusted to 7.3 with KOH. Internal solution used in voltage clamp experiments contained (in mM): 120 $CsMeSO_3$ (CH_3O_3SCs), 3 CsCl, 10 HEPES, 1 EGTA, 4 Na_2ATP, 0.3 NaGTP, 5 Na_2-phosphoreatine, 3.5 QX-314 chloride, 0.5% (w/v) biocytin hydrochloride, 50 µM Alexa Fluor 488 hydrazide; osmolarity 290–295 mOsm; pH adjusted to 7.3 with CsOH.

Despite slicing the brain at an angle, because of the curvature of the caudal neocortex in mice, in most slices the L5 neuron apical dendrites still extended toward the pia at a shallow angle relative to the slice surface. The neurons on one side of each slice thus had apical dendrites directed towards the slice surface, which were typically cut before the apical bifurcation. To avoid recording from neurons with cut apical dendrites the slices were positioned in the recording chamber so that the apical dendrites could be seen to extend either parallel to the slice or at a slight downward angle into the slice.

Visually guided whole-cell patch-clamp recordings of ttL5 neurons were performed using a Scientifica SliceScope Pro 3000 microscope equipped with a 40x/0.8NA objective and an infrared (IR) Dodt Gradient Contrast system. Electrodes and recording headstages were mounted on motorized micromanipulators (Scientifica MicroStar). Electrode voltage was offset immediately above the slice and positive pressure was applied by blowing air through a stiff rubber tube connected to the back of the electrode before slowly lowering the pipette into the tissue. While lowering the pipette, the electrode was set in voltage-clamp mode and short 10 mV pulses were continuously applied in order to monitor the resistance through the pipette opening using an oscilloscope.

Once the electrode tip was in close proximity to the targeted cell, a small dimple could be seen in the cell membrane. At this point the pressure was released and the resistance was closely monitored while applying gentle suction to the pipette until a gigaohm seal was formed. Whole-cell configuration was then obtained with short suction pulses. Most recorded cells had access resistance of 10–25 MΩ, and recordings were excluded from analysis if the resistance exceeded 40 MΩ. When performing voltage clamp recordings, following whole-cell capacitance compensation, series resistance compensation was also applied and set to the highest value possible without inducing oscillations in the cell (typically 40–70 %).

2.4.3 Data acquisition & analysis

Recorded signals were amplified and low-pass filtered through an 8 KHz Bessel filter using a MultiClamp 700B amplifier (Molecular Devices). Filtered signals were then digitized at 20 kHz with a National Instruments DAQ board (PCIe-6323). Acquisition and stimulus generation were done with the Igor Pro (Wavemetrics) and NeuroMatic (Rothman and Silver, 2018) software packages. Further analysis and data visualization were performed with custom macros and scripts written in Igor Pro and MATLAB (Mathworks). Raincloud plots (consisting of a scatter plot, a box plot, and a kernel density plot) were generated in MATLAB using scripts written by Allen et al. (2019). All box plots presented show the median, interquartile range, 2nd and 98th percentile of the dataset. Unless otherwise specified, all reported data values refer to the mean ± standard error (SEM). When considering statistical significance, for experiments where multiple comparisons were made between the same groups the α significance threshold was adjusted using the Bonferroni correction. The liquid junction potential was not corrected for.

2.4.4 Electrical stimulation

Extracellular monopolar stimulation was carried out using a Digitimer DS3 isolated constant current stimulator, passing brief DC current pulses (0.1–1 ms, 20–320 µA) through a glass micropipette (20–50 µm tip diameter) filled with ACSF. The current was passed between two silver/silver chloride (Ag/AgCl) wires: one inside the pipette and the other coiled around the outside of the pipette. Pulse timing was controlled with 5 V TTL signals via the National Instruments board.

2.4.5 Optogenetics

Patterned illumination Patterend optogenetic stimulation was implemented using a digitial micromirror device (DMD) and a 463 nm laser (laser-coupled Polygon 400, Migtex Systems). The patterned light from the DMD was reflected by a mirror with 90/10 reflection/transmission ratio positioned above the optical filter turret in line with the objective. The pattern consisted of a $500 \times 1000\,\mu m$ grid composed of 12×24 square spots of light (each approximately $41.7\,\mu m$ wide) and was delivered through a 5×/0.15NA dry objective (Olympus MPlanFL N). As some variations in light intensity were observed across the illumination field, for each spot the laser power was measured at the specimen plane using a PM100D (Thorlabs) optical power meter and the laser output associated with that spot was adjusted to obtain a measured power of approximately $300\,\mu W$ ($173\,mW/mm^2$).

The spatial sequence of stimuli was designed to maximise the distance between consecutive spots. This was achieved by computing the distance between the nearest pair of consecutive spots for 10^7 randomly generated pattern sequences and selecting the sequence with the largest distance. To reduce the search space, each sequence was generated iteratively starting from a spot in the center of the grid, randomly adding one spot at a time from the set of remaining unselected spots. After each generated sequence, spots in subsequent sequence generations were only selected if their distance was greater than the shortest distance measured in the previous sequence.

Full-field perisomatic illumination Full-field optogenetic stimulation was implemented using a CoolLED pE-4000 illumination system, passing light through a 40×/0.8NA water-immersion objective (Olympus LUMPlanFL N) during whole-cell recordings. During stimulation the soma was centered in the field of view, resulting in a circular illumination field with a radius of approximately $100\,\mu m$ from the soma. Chronos and ChrimsonR were stimulated with 470 nm and 635 nm light, respectively. The range of illumination intensities was 0–37.3 mW/mm^2 for the 470 nm light and 0–32.5 mW/mm^2 for the 635 nm light. The upper values of these ranges were the measured light intensities at the highest power setting for this illumination system.

2.5 Immunohistochemistry, imaging & morphological reconstructions

After recording, slices were fixed overnight at 4 °C in a 4% formaldehyde solution and were subsequently kept in PBS. For immunohistochemical detection, the fixed slices were first incubated for 1–2 hours at room temperature in blocking solution containing 0.5% Triton X-100 and 5% Normal Goat Serum (NGS) in PBS. Slices were then washed twice (10 min each) in PBS and incubated overnight in a staining solution containing 0.05% Triton X-100, 0.5% NGS, DyLight 594-conjugated streptavidin (2 µg/ml). Slices were then washed in PBS (3 times, 5 min each) and stained with DAPI (5 µg/ml) for 10 min. After another wash (3 times, 5 min each), slices were mounted on glass slides and images were acquired with a confocal microscope (Leica multiphoton SP5; objective: 20x/0.7NA or 10x/0.4NA; pinhole size: 1 airy unit). Image processing was done with the FIJI software package (Schindelin et al., 2012). The confocal image stacks were used to manually reconstruct neurons using Neurolucida 360 (MBF bioscience).

2.6 Biophysical modelling

Simulations were performed in the NEURON simulation environment (7.7.1, Hines and Carnevale (2001)) embedded in Python 3.6. To model the consequences of morphological differences between V1 and V2m ttL5 cells, existing models of ttL5 pyramidal cells were used with either accurate morphological detail (biophysical model 3, cell #1 from Hay et al. (2011), referred to as detailed model) or simplified multicompartment morphologies (Ca^{2+} enriched model 2 from Bahl et al. (2012), referred to as reduced model). To study the effect of morphology in the detailed model, biophysical model 3 from Hay et al. (2011) was applied to the reconstructed morphology from one of the recorded ttL5 neurons in V2m (which had a substantially shorter apical trunk than the morphology used in the original model). Each morphology contained low-voltage-activated (LVA) and high-voltage-activated (HVA) Ca^{2+} channels located in a region around the main apical bifurcation. This region was 685–885 µm from the soma in the long morphology (cell #1 from Hay et al. (2011)) and 300–500 µm in the short morphology.

Subsequent simulations using the reduced model were done by modifying only selected parameters described in the results, such as the length of the apical trunk compartment, leaving all other parameters unchanged. Briefly, this reduced model (Bahl et al., 2012) is divided into sections representing the soma, axon (hillock and initial segment, AIS), basal dendrites, apical trunk, and apical tuft. Active conductances are present in all compartments and include the following: hyperpolarization-activated cation (HCN) channels (basal dendrite, apical trunk, tuft), transient voltage-activated Na^+ (Nat) channels (soma, axon hillock, AIS, apical trunk, tuft), persistent voltage-activated Na^+ (Nap) channels (soma), fast voltage-activated K^+ (Kfast) channels (soma, apical trunk, tuft), slow voltage-activated K^+ (Kslow) channels (soma, apical trunk, tuft), muscarinic K^+ (Km) channels (soma), slow Ca^{2+} (Cas) channels (tuft), Ca^{2+} dependent K^+(KCa) channels (tuft), and a Ca^{2+} pump (tuft). In this model the density of Kfast and Kslow channels decays exponentially from the soma to the tuft. The density of Nat channels decays linearly from the soma to the tuft, while HCN channels linearly increase in density. N.B., the tuft, but not the trunk, contains Ca^{2+} channels; consequently there is no hot spot similar to the apical bifurcation in the detailed model. When varying trunk length, Nat, Kfast, Kslow, and HCN conductances in each trunk segment were redistributed so as to take into account the new distance of each segment from the soma (thereby changing the total conductance in the trunk).

Chapter 3

Intrinsic physiology

3.1 Introduction

Neurons are a highly diverse class of cells with a wide range of genetic, morphological, and electrophysiological properties that enable different subpopulations to take on specialized roles within the computations that the network is performing. To understand these computations it is necessary to know the properties of different populations of neurons and the diversity within these populations, which determines how these neurons respond to any given set of inputs.

The most effective tool for measuring the subthreshold potentials in a neuron is the whole-cell patch clamp technique. This has enabled a large number of discoveries concerning how dendrites integrate information as well as the ion channels that are involved. Investigations of different cell types (Groh et al., 2010; Kim et al., 2015; Gouwens et al., 2019) have revealed a large diversity of integrative properties.

The morphology and biophysical properties of a neuron represent the foundation of single-cell computations. The effect of the thousands of inputs a neuron receives depends critically on how they interact and are integrated within the postsynaptic cell. Ultimately this rich subthreshold dynamic information is compressed into a single stream of spikes. Both the summation of individual synaptic potentials and the transformation of subthreshold voltages to output spikes can be highly non-linear depending on the how passive and active properties interact to shape the signal.

To fully understand the functions of a neuron, the inputs it is receiving, and the transformations it is performing on these inputs, we would ideally need to study the same cells with a variety of methods, both *in vivo* and *in vitro*. Genetic handles are useful for being able to consistently identify and study the same population of neurons across different experiments. One of the objectives of this thesis is to study the properties of one such genetically labelled population.

The Glt25d2-Cre labelled population of neurons in V2m (referred to as Glt) has not previously been described. These neurons are hypothesised to be ttL5 neurons, based primarily on a description of a mouse line driving enhanced green fluorescent protein (EGFP) under the same Glt promoter (Groh et al., 2010). In this chapter I will describe the basic anatomical and physiological properties of this Cre-labelled population and thereby directly address the first two aims described in the introduction. These experiments were primarily designed to determine if these cells can be considered ttL5 neurons by thoroughly characterizing their locations, projection patterns, morphologies, and intrinsic electrical properties. This comprehensive description constituted an important first step to validate the use of this Cre-driver line as a reliable marker of ttL5 neurons and enable me to subsequently explore the computations and functional properties of this population of Glt neurons in later experiments.

3.2 Chapter Methods

3.2.1 Cell selection & recording conditions

After preparing brain slices as described in chapter 2, whole-cell patch clamp recordings were made from Glt neurons in slices containing V2m (based on the approximate stereotaxic coordinates described in Paxinos and Franklin (2007)). The area to record from was first identified under a low magnification (10x/0.25NA) dry objective while illuminating the sample with IR light. To confirm the presence of Glt neurons at the desired recording location, the tdTomato was visualized with short pulses (lasting a few seconds) of epifluorescence stimulation with 550 nm light from the CoolLed pE-4000 illumination system. The pulses were kept short and the light intensity was set as low as possible in order to minimize photobleaching as well as any other possible side effects of strong illumination. After identifying the cluster of Glt cells approximately located in V2m, I switched to the high magnification (40x/0.8NA) water immersion objective.

Although Glt25d2-Cre occasionally labels neurons that are not in L5, these cells were avoided and only labelled neurons in L5 were targeted for recordings. This was easily done, as L5 is clearly distinguishable under IR illumination by the presence of many cells with very large pyramidal-shaped somas. Within the fluorescently labelled L5 neurons in V2m, cells were chosen at random. No consideration was thus given to criteria such as soma size and depth (provided they were within L5), which are sometimes used in existing literature to facilitate selection of ttL5 neurons.

All recordings in this chapter were made at near-physiological temperature (35 °C) in external ACSF containing 1.5 mM $CaCl_2$, 1 mM kynurenic acid, and 50 µM picrotoxin.

3.2.2 Whole-brain cell mapping

The brain-wide distribution of Glt neurons and their projection targets was measured and quantified from fixed brains, using serial-section two-photon tomography as described in chapter 2. While the location of cells could be accurately measured in the transgenic mice crossed with a tdTomato reporter line, the axons in these mice were not bright enough to accurately quantify. In order to increase the brightness of the labelling for axon tracing experiments, stereotaxic surgery was performed on Glt mice to inject 50–100 nL of Cre-dependent AAV expressing EGFP.

3.2.3 Sample size determination

An important consideration when making claims about the properties of a population based on limited experimental samples is the selection of an appropriate sample size. While ideally this would be as large as possible in order to obtain an accurate representation of the true population distribution of measured parameters, practical considerations (such as the time and cost of acquiring samples) limit the number of samples that can feasibly be acquire.

When making statistical comparisons between different populations or samples, the main criterion for determining sample size is the smallest effect size which we want to be able to distinguish, which in turn depends on either the expected magnitude of the difference or on the smallest difference that is could be considered to be biologically meaningful.

Although the data in this chapter is of a descriptive nature, with the main purpose of determining whether Glt cells in V2m can be considered ttL5 neurons, I aimed to acquire a sample size large enough to accurately represent the diversity within the recorded Glt population and to potentially compare it to other populations, should such a comparison be useful in the future. For this reason, I determined that a sample size of around n = 30 would be sufficient for my purposes, on the basis of previous literature involving similar descriptions of neuronal populations, where the range of samples sizes is often between 10 and 50 cells per group (Kasper et al., 1994; Pike et al., 2000; Groh et al., 2010; Vélez-Fort et al., 2014; Ueta et al., 2014; van Aerde and Feldmeyer, 2015; Kim et al., 2015; Naka et al., 2019; Fletcher and Williams, 2019).

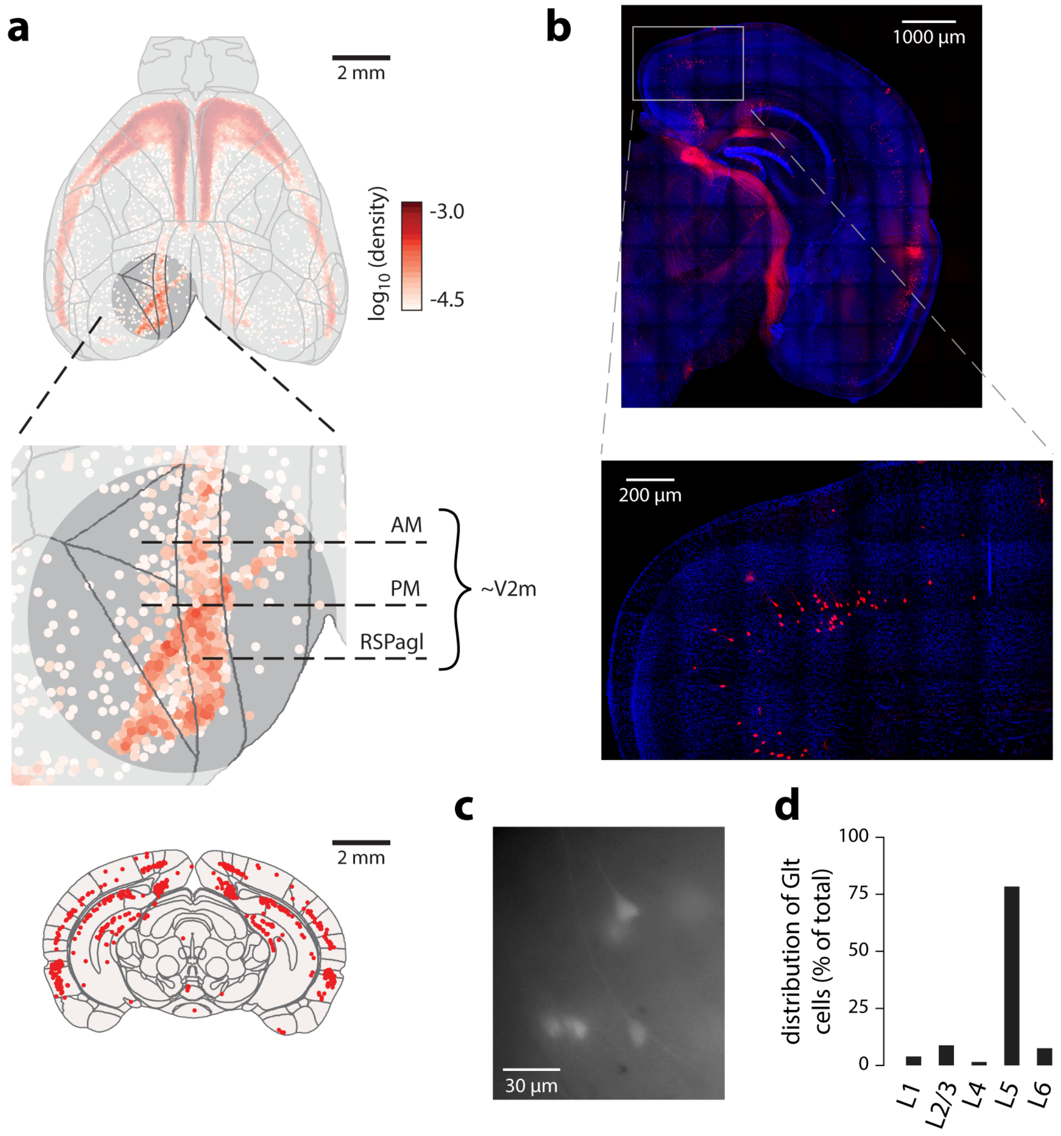

Figure 3.1: Brain-wide distribution of Glt25d2-Cre labelled neurons. a. *Top:* Top-down view of cell distribution throughout cortex relative to the Allen CCF in one Glt25d2-Cre×tdTomato animal. *Middle:* Highlighted and magnified circle illustrating the approximate region, between visual cortex and retrosplenial cortex, where recordings were made. *Bottom:* Coronal view with overlaid cells from a 200 µm thick section. **b.** Confocal image of one hemisphere from a coronal section of a Glt brain, highlighting the area of dense labelling that was targeted for recordings. Red: tdTomato expression in Glt neurons; Blue: DAPI nuclear stain. **c.** Epifluorescence image acquired during a whole-cell patch clamp recording from a Glt neuron in V2m. **d.** Percentage of Glt neurons across each cortical layer (averaged across all visual areas). The cell counts were assigned to layers based on the layer-specific segmentation in the Allen CCF. The data in this figure was partially collected by Dr. Zhiwen Ye.

3.3 Results

3.3.1 Brain-wide distribution and projection targets of Glt neurons

The Glt25d2-Cre line labels neurons in several areas throughout the brain. Compared to the EGFP line, with fluorescence under direct control of the Glt promoter (Groh et al., 2010), labelling in this Cre line was found to be substantially sparser, with relatively few labelled cells across most of the visual and somatosensory cortices, and denser clusters of cells spread throughout various brain regions (Figure 3.1a,b). One region of higher density was found on the dorsal side of cortex in a posteromedial band. In the brain coordinates segmented according to the Allen CCF, this region overlapped with the higher-order visual areas AM and PM, as well as the agranular retrosplenial cortex (RSPagl). This entire region approximately corresponds to area V2m in the Paxinos and Franklin (2007) brain atlas.

When performing patch-clamp recordings in acute brain slices (Figure 3.1c), it is not possible to distinguish the boundaries between AM, PM, and RSPagl, and stereotaxic coordinates are difficult to target accurately, due to the limited field of view during recordings and the fact that the brain is often not of identical dimensions to that of the reference atlas. For this reason, throughout this thesis I have chosen to refer to this area by the V2m terminology and segmentation described in Paxinos and Franklin (2007), which is defined cytoarchitectonically and corresponds well to the area of dense Glt labelling.

As expected from the EGFP line, Glt labelling was found to be highly specific for neurons in cortical layer 5 (L5), with very few cells in any other layer (Figure 3.1d). Within L5, an important distinguishing feature between slender-tufted and thick-tufted pyramidal neurons is the axon projection patter (Groh et al., 2010). Glt neurons were found to project strongly to ipsilateral subcortical targets, including the striatum and the superior colliculus (Figure 3.2), as is charateristic of ttL5 neurons.

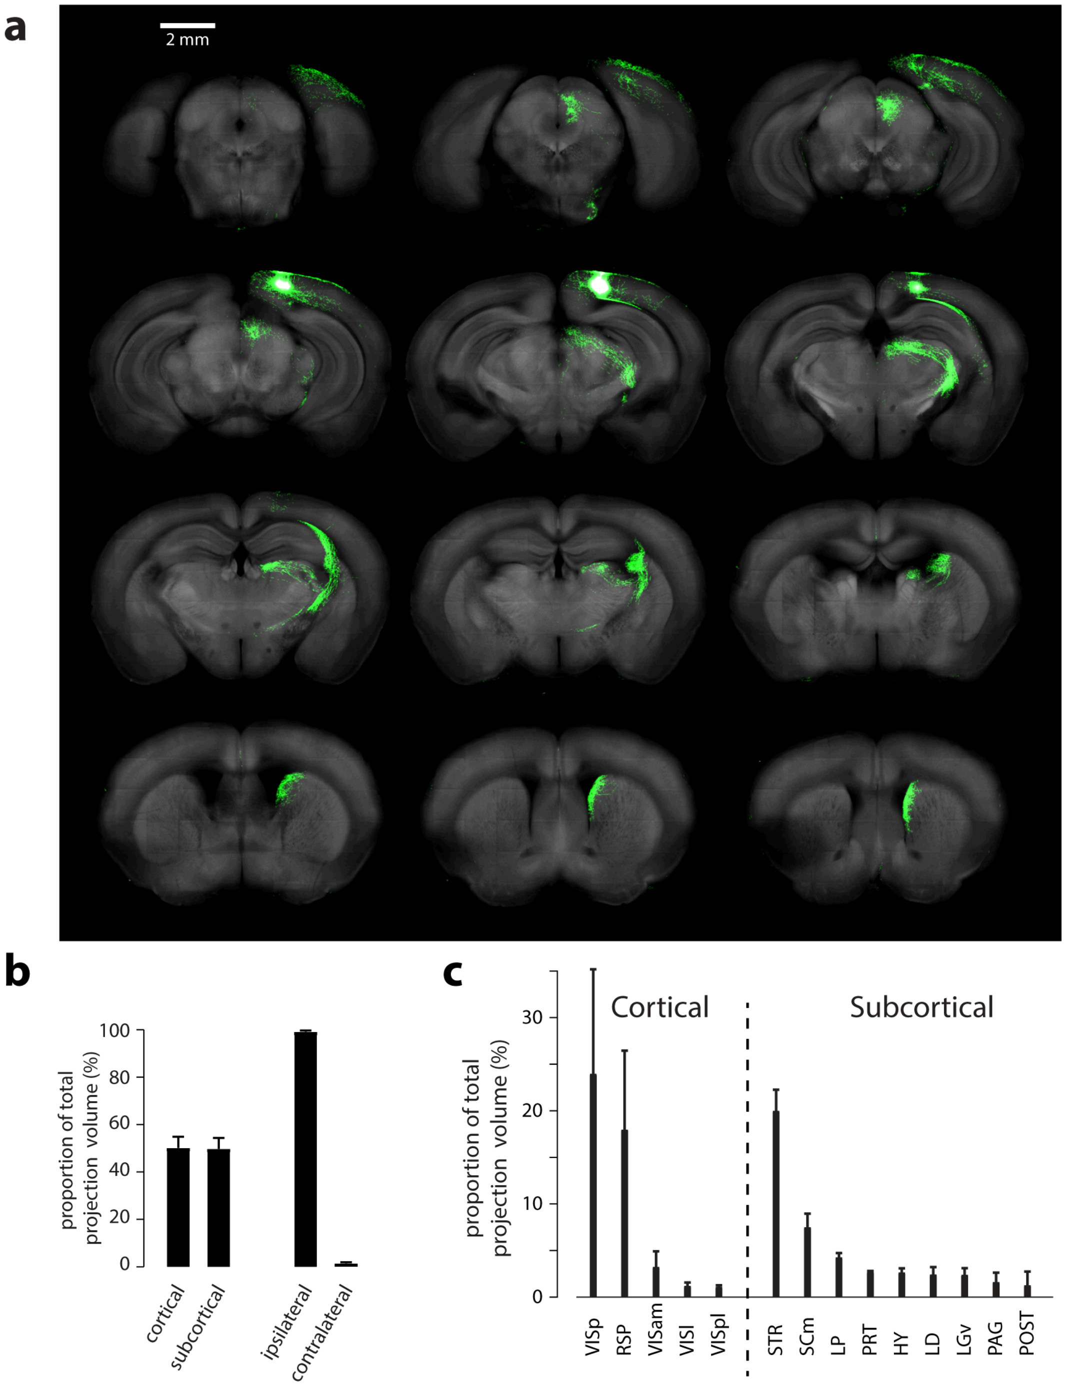

Figure 3.2: Axon projection targets of Glt25d2-Cre labelled neurons. **a.** Images from the two-photon serial tomography in a Glt mouse expressing Cre-dependent EGFP, which was stereotaxically injected in V2m. **b.** Proportion of axon projecting to cortical, subcortical, ipsilateral, and contralateral targets (n = 4 mice). **c.** Axon projections split by target region.
Abbreviations: VISp, primary visual area; RSP, retrosplenial area; VISam, anteromedial visual area; VISl, lateral visual area; VISpl, posterolateral visual area; STR, striatum; SCm, motor related superior colliculus; LP, lateral posterior thalamus; PRT, pretectal region; HY, hypothalamus; LD, lateral dorsal thalamus; LGv, ventral lateral geniculate thalamus; PAG, periaqueductal grey; POST, postsubiculum.

3.3.2 Morphologies & recording location

Recordings from the Glt population in V2m were targeted to L5, clearly distinguishable in acute brain slices by the presence of pyramidal neurons with large cell bodies and by the location approximately 500 µm below the pia, in order to avoid recording from the small percentage of labelled neurons located in other layers. The approximate location of the recorded cells is illustrated in Figure 3.3a. The recorded Glt neurons were broadly distributed along the anteroposterior axis, ranging from approximately 2.4 to 3.9 mm posterior to bregma. The anteroposterior location of all recorded cells is shown in Figure 3.3b and was estimated by visually comparing each slice to coronal sections in Paxinos and Franklin (2007).

During recordings, neurons were intracellularly filled with biocytin hydrochloride (0.5% w/v) and later stained with fluorescently tagged streptavidin. Of these, 11 Glt neurons were imaged and morphologically reconstructed (Figure 3.3b). These cells all presented morphological features typical of ttL5 neurons, including a large apical dendrite which extended to the pia and branched broadly in L1, and a large number of oblique dendrites along the apical dendrite (Figure 3.3c). By comparison, the other main class of pyramidal neuron residing in L5 are known as slender-tufted L5 neurons and are clearly distinguishable from ttL5, based on their smaller dendritic tufts and fewer oblique dendrites.

The main distinguishing morphological features are quantified in Figure 3.4. These features are all broadly consistent with similar values reported in the literature (Groh et al., 2010; Oswald et al., 2013; Ramaswamy and Markram, 2015; Jiang et al., 2020) and in the Allen Cell Types Database (http://celltypes.brain-map.org/). The morphological segregation of dendrites into the apical and basal dendritic compartments is reflected in the Sholl analysis, where two peaks of dendritic branching can be seen.

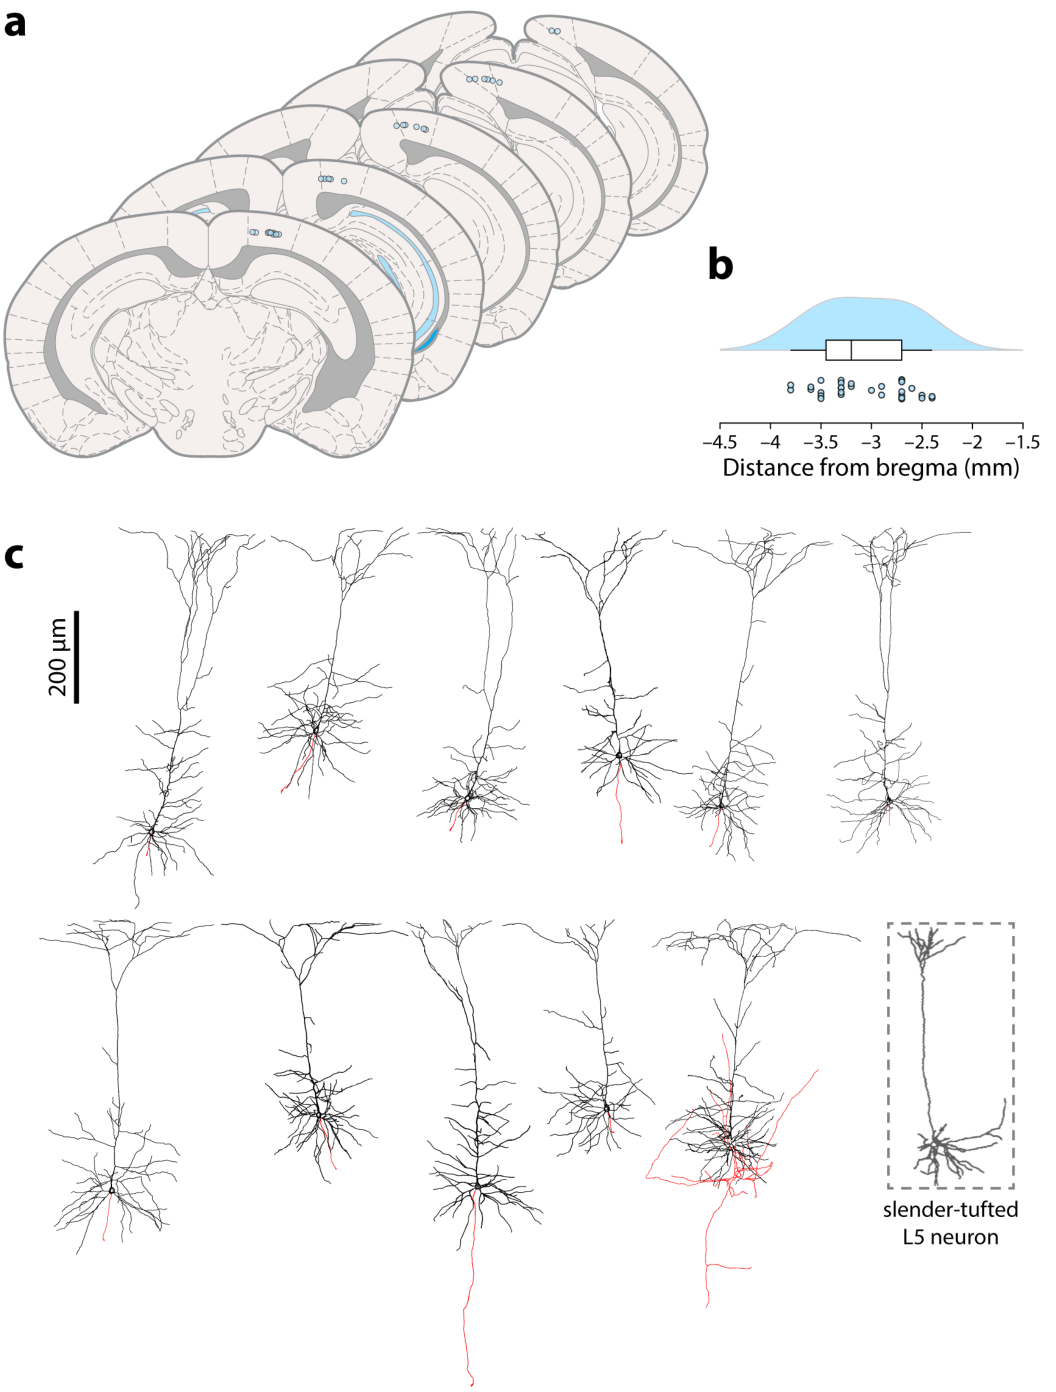

Figure 3.3: Anatomy of recorded Glt25d2-Cre neurons in V2m. **a.** Schematic of coronal sections, indicating approximate locations of recorded Glt neurons (each circle represents one neuron). **b.** Approximate anteroposterior locations of all recorded neurons relative to bregma. **c.** Morphologies of 11 reconstructed Glt neurons, with soma and dendrites in black and the axon labelled in red. For comparison, a slender-tufted L5 neuron from Groh et al. (2010) is shown in the dashed box (bottom right).

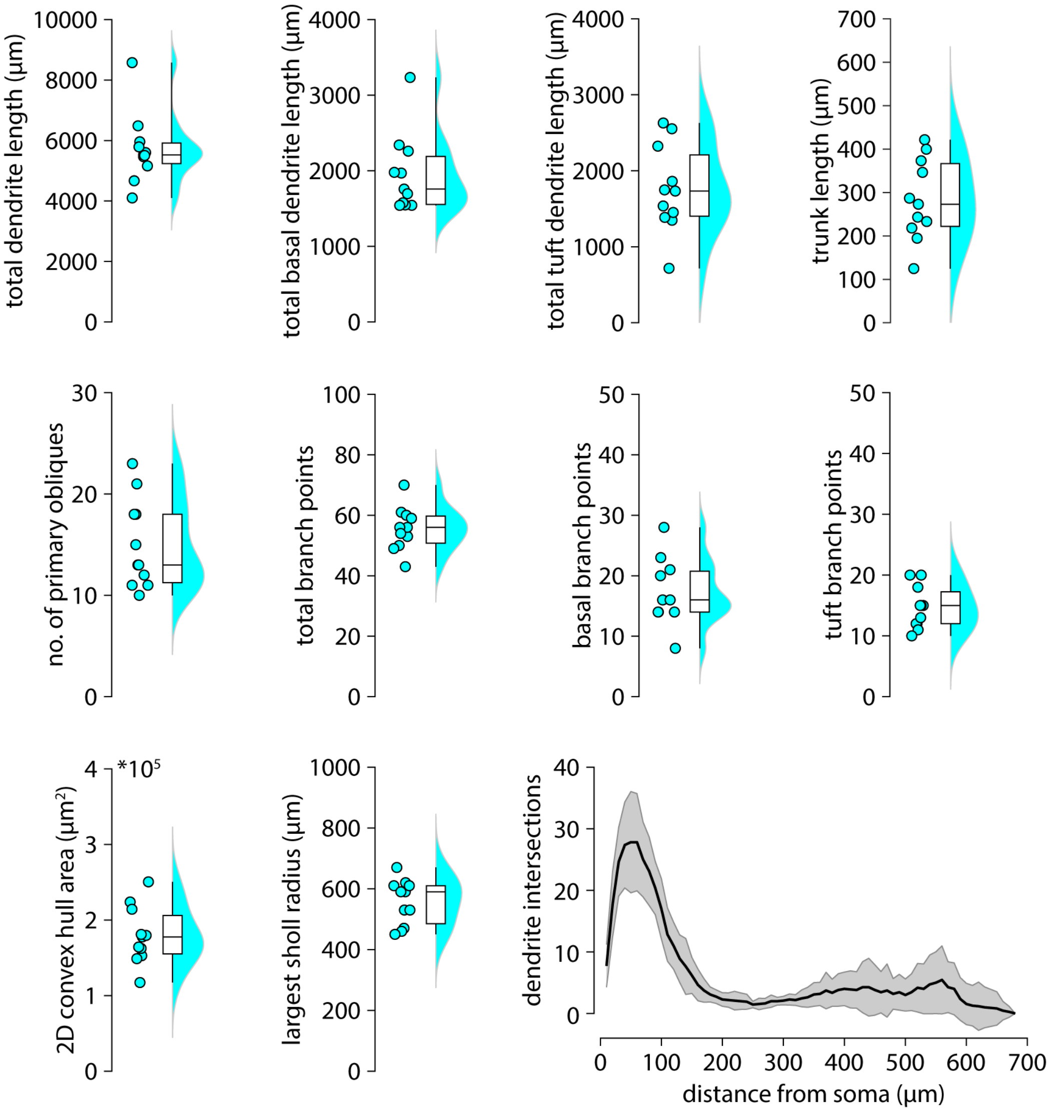

Figure 3.4: Quantification of morphological parameters. Morphological parameters of the 11 reconstructed Glt neurons, quantified using Neurolucida Explorer. Basal dendrites exclude the apical dendrite and any oblique dendrites originating from the apical dendrite (even if these branch off very close to the soma). Tuft length includes all dendrites after the main bifurcation on the apical dendrite. The apical trunk length was measured from the soma to this bifurcation. Any dendrites branching off from the apical trunk were considered oblique dendrites. The 2D convex hull is the smallest polygon that could be fit around the 2D projections shown in Figure 3.3c. The Sholl plot (bottom right) shows the number of intersections between dendrites and concentric spheres centred on the soma, with radius increasing in 10 µm increments.

3.3.3 Subthreshold properties

Whole-cell patch-clamp recordings were made in current clamp configuration from Glt neuron (n = 29) in V2m. To obtain a detailed characterisation that would capture a wide range of integrative properties, each neuron was subjected to several stimulation protocols designed to probe both subthreshold and suprathreshold properties of these cells. As these experiments were designed to probe intrinsic properties of each cell, independently of its inputs, pharmacological blockers of fast excitatory (kynurenic acid, 1 mM) and inhibitory (picrotoxin, 50 µM) synaptic transmission were added to the recording ACSF (containing 1.5 mM $CaCl_2$).

The passive membrane properties of neurons are important factors for the integration of inputs, determining their spatial and temporal summation. Shortly after breaking into the cell, the resting membrane potential (RMP) was recorded and a series of 500 ms long current steps (−400 pA to +600 pA in steps of 20 pA) were injected in the soma. Input resistance (R_N) of each cell was determined from the slope of a linear fit of the steady-state voltage during the hyperpolarizing current steps (Figure 3.5a,b). To maintain consistent parameters for comparison of properties between cells, for all subsequent recordings a holding current was set so as to maintain the resting potential of the neuron at −70 mV.

Both RMP (Figure 3.5c) and R_N (Figure 3.5d) were broadly similar across all recorded neurons (RMP = −67.3 ± 0.8 mV; R_N = 64.4 ± 3.7 MΩ). To measure the membrane time constant (τ) with minimal recruitment of active conductances, I injected 1 ms long 400 pA negative current steps and fit a double exponential to the voltage trace from 5 ms after the stimulus (to avoid measuring the pipette capacitance) to the point where the voltage returned to its resting value (Figure 3.5e). The value of τ resulting from this fit was calculated as the weighted average of the two exponential coefficients and was 18.8 ± 1.1 ms (Figure 3.5f).

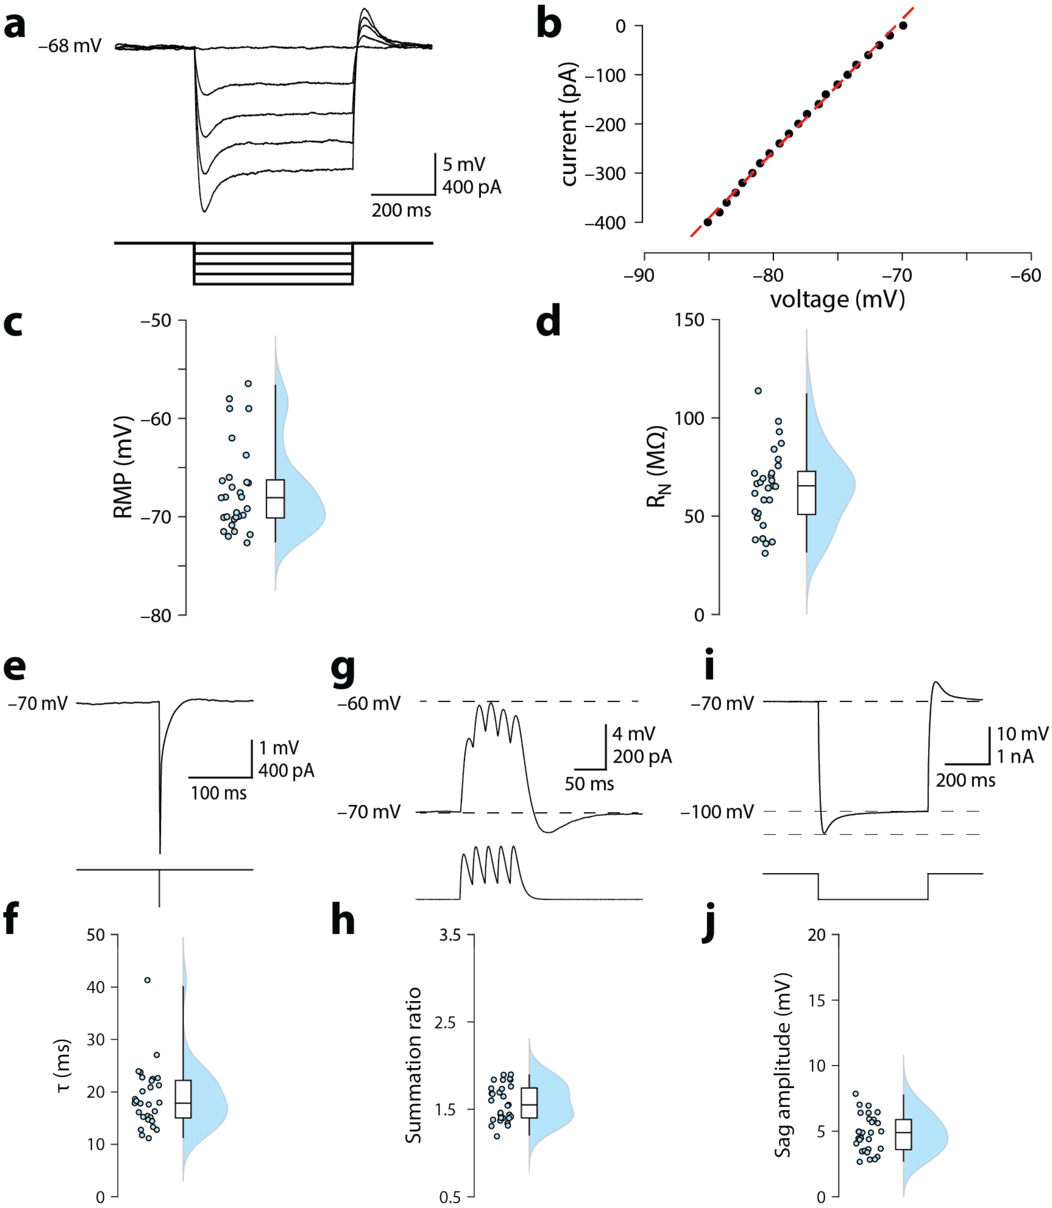

Figure 3.5: Subthreshold properties of ttL5 neurons in V2m. a. Recorded hyperpolarizing voltage (top) and injected current steps (bottom) used to calculate input resistance. For clarity, only a subset of the recorded traces is shown. Voltage was measured as the mean in the last 100 ms of each current step. **b.** Current-voltage (IV) relationship for all recorded hypoerpolarizing steps and linear least-squares fit (dashed line). **c.** Resting membrane potential recorded shortly after break-in. **d.** Input resistance for all cells, derived from the slope of the best fit line illustrated in *b*. **e.** Voltage induced by a 1 ms wide current step. **f.** Time constant measured from fitting a double exponential to the decay phase of the voltage shown in *e*. **g.** Voltage following somatic injection of 5 αEPSCs at 50 Hz. **h.** Ratio of peak voltages at the 5th and 1st αEPSC. **i.** Current step resulting in 30 mV steady-state hyperpolarization, from which sag amplitude was measured as the difference between peak and steady-state voltage. **j.** Sag amplitude for each recorded cell.

The most important consequence of the membrane time constant is its influence on temporal summation. Several other features, such as active conductances, can also have a substantial effect on summation. To quantify summation more directly I injected five αEPSC waveforms at 50 Hz and defined the summation ratio as the ratio between the peak voltages at the 5th and 1st current injection. The amplitude of the injected current was adjusted so as to obtain a maximum depolarization of 10 mV above baseline (Figure 3.5g). With this stimulus, the resulting depolarization exhibited a sharp reduction in amplitude after the first αEPSC. In 28/29 Glt cells the peak voltage was reached following either the third or fourth αEPSC. Measured summation ratios were 1.57 ± 0.04, indicating the presence of strong excitability-regulating mechanisms in these neurons that suppressed larger depolarizations (Figure 3.5h).

One known mechanisms that has a large influence on neuronal excitability is the H-current (I_H), driven by hyperpolarization-activated cyclic nucleotide–gated (HCN) ion channels. A fraction of these voltage-gated cation channels are open at resting potential and close following depolarization, thus reducing the total inward current to the cell. This negative-feedback control of neuronal excitability imparts a stabilizing effect on the membrane potential. Experimentally, the action of HCN channels appears as a rebound depolarization (known as sag) during hyperpolarizing current steps, which is often used as a proxy for I_H. To measure the sag, I injected 30 mV hyperpolarizing steps (from −70 to −100 mV) and measured the difference in voltage between the peak and steady-state hyperpolarization (Figure 3.5i). All recorded neurons exhibited substantial sag depolarizations (Figure 3.5j), measuring 4.82 ± 0.26 mV (16.1 ± 0.9 % of steady-state hyperpolarization).

Finally, the frequency response profile of the recorded neurons was tested by injecting a sinusoidal chirp stimulus (100 pA peak-to-peak) with frequency increasing linearly from 0 to 100 Hz over the course of 30 s. From this, the impedance profile as a function of frequency can be calculated by converting the current and voltage traces to the time domain using fast Fourier transforms

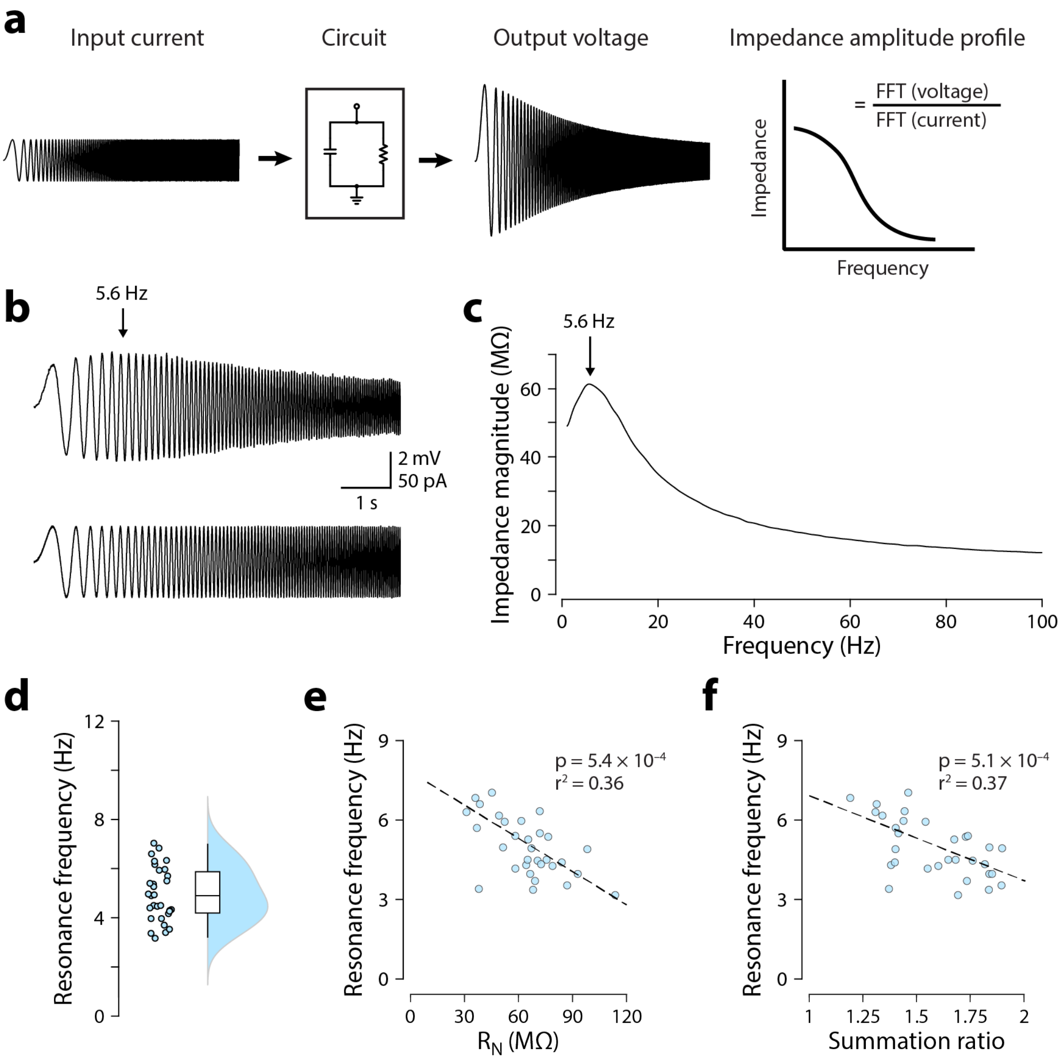

Figure 3.6: Frequency-dependent properties of ttL5 neurons in V2m. **a.** Diagram of typical voltage response to a 0–100 Hz chirp stimulus. Passive RC circuits function as low-pass filters with a monotonically decreasing impedance amplitude profile. Active conductances in neurons can additionally filter other frequency components. **b.** Voltage in Glt neuron during 0–100 Hz chirp stimulus. **c.** Impedance amplitude profile calculated from traces in *b*, indicating a resonance frequency at the impedance peak around 5.6 Hz. **d.** Resonance frequencies for all recorded neurons. **e,f,g.** Scatter plots showing correlation of resonance frequency with other intrinsic features. Dashed lines show least squares linear fits.

(FFT) and taking the ratio of these values (Figure 3.6a). Due to capacitive filtering, an entirely passive cell acting as an RC circuit inherently functions as a low-pass filter. However, in many cells active processes selectively enhance or dampen the signal in certain parts of the frequency spectrum, making the cell function more like a band-pass filter. The peak frequencies for any cell are termed resonance frequencies and have been previously observed in several cell types (Hutcheon et al., 1996; Hutcheon and Yarom, 2000; Pike et al., 2000; Gutfreund et al., 1995).

The current and voltage traces for a representative Glt neuron and its impedance amplitude profile are shown in Figure 3.6b,c. All recorded Glt cells had measurable resonance frequencies above 1 Hz (3.6d). Resonance frequencies were in the theta range, measuring 4.95 ± 0.21 Hz and were found to correlate with the R_N ($r^2 = 0.36$, $p = 5.4 \times 10^{-4}$, F-test, Figure 3.6e) and summation ratio ($r^2 = 0.37$, $p = 5.1 \times 10^{-4}$, F-test, Figure 3.6f).

3.3.4 Suprathreshold properties

Responses to depolarizing current injections were also highly stereotyped and characterized by several distinctive features. Each spike was followed by a prominent afterdepolarization (ADP), which was typically largest for the first spike in any depolarizing step (Figure 3.7a,b). During larger current steps, a spike doublet appeared at the beginning of the depolarization, with the second spike occurring approximately at the same location as the ADP (Figure 3.7a,b).

To measure spike features, for each cell I first determined the rheobase current, defined as the the smallest current step that could evoke at least one spike. Mean rheobase was 144 ± 9 pA (Figure 3.7c). Above rheobase, firing rate increased monotonically with current injection, although the slope of the firing rate (FI) curve can be seen to decrease somewhat at high current steps (Figure 3.7d). In order to obtain average spike shapes for each cell (minimizing any effects of variability there might be in individual spike shapes), I aimed to perform a detailed quantification of suprathreshold properties in traces containing at least 3 spikes. Since rheobase varied substantially across

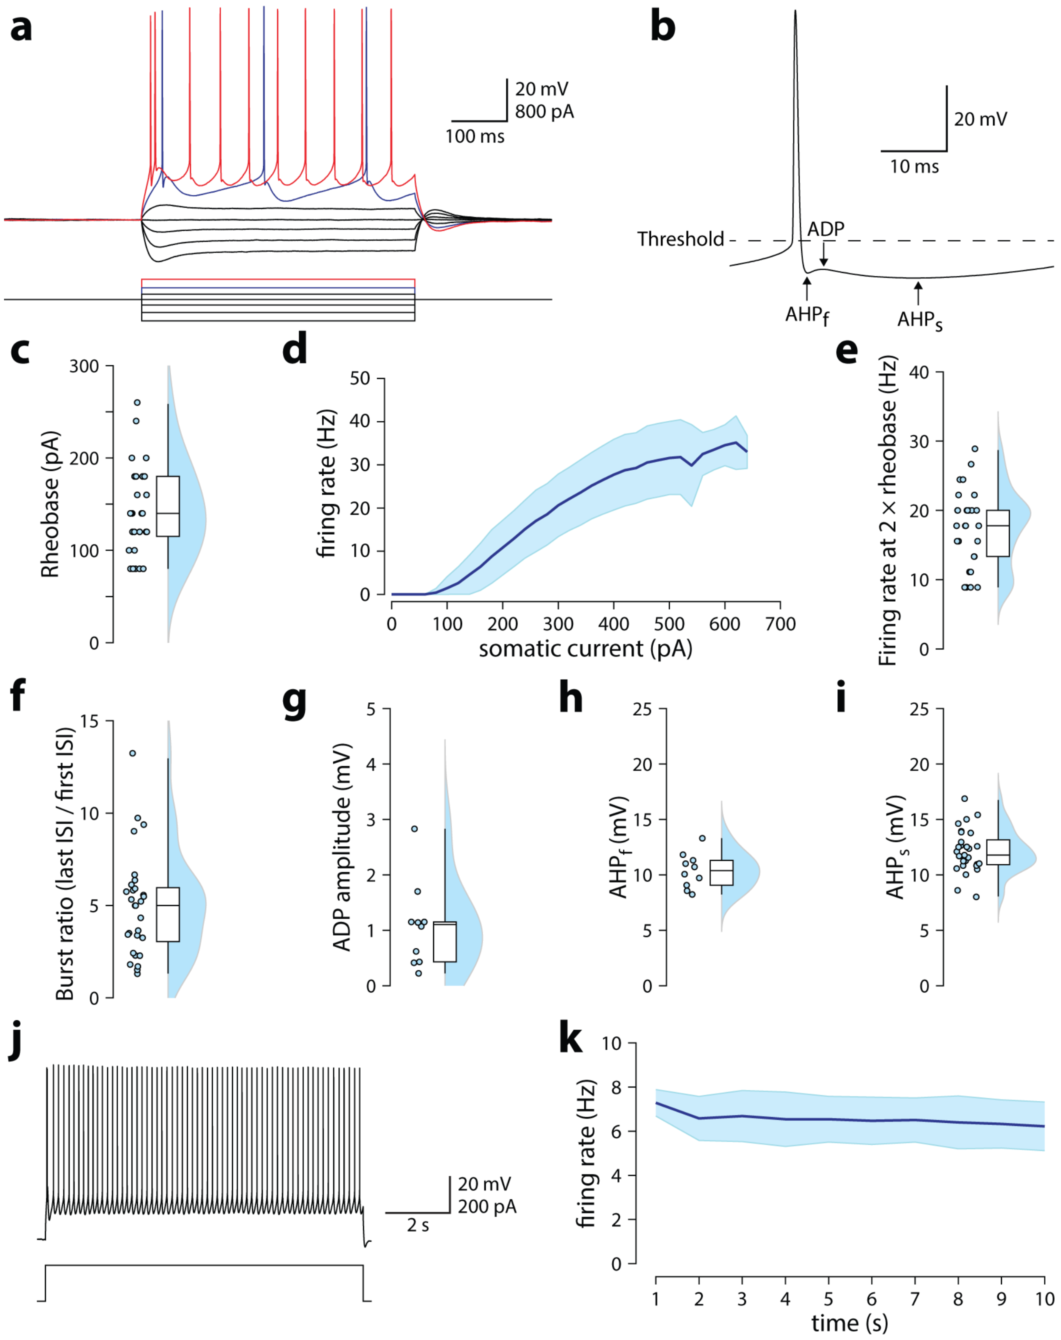

Figure 3.7: Suprathreshold properties of ttL5 neurons in V2m. a. Representative voltage during 500 ms current steps. Note initial spike doublet in the red trace. **b.** Average spike shape in 350 ms stimulus window at twice rheobase, indicating the main features that were quantified. **c.** Rheobase for recorded Glt neurons. **d.** FI curve, showing firing rates at current steps between 0 pA and 640 pA. **e–k.** Suprathreshold properties recorded at twice rheobase current injections showing firing rate (**e**), burst ratio (**f**), afterdepolarization amplitude (**g**), as well as fast (**h**) and slow (**i**) afterhyperpolarization relative to the threshold voltage. **j.** Voltage response during a 10 s long depolarizing current step. **k.** Mean firing rate computed in 1 s bins over the 10 s stimulus duration.

the population of Glt neurons (ranging from 80 pA to 360 pA), in order to be consistent when choosing which current trace to evaluate, I opted to perform the analyses for current steps at twice the recorded rheobase for each cell. At this stimulus amplitude, mean firing rate was 17.6 ± 1 Hz (Figure 3.7e), and was in all cases above 8 Hz.

In the current step at twice rheobase, the average spike shape was then computed in a 350 ms window (50–400 ms from stimulus onset). This quantification window was chosen so as to exclude the initial spike doublet, which had a firing rate that was several times higher and consequently a substantially different spike shape. This difference in firing rate is instead quantified separately through the "burst ratio", calculated as the ratio of the first and last inter-spike interval (ISI) in the trace, which was 4.98 ± 0.51 (Figure 3.7f). ADP amplitude was measured as the difference between the trough of the fast afterhyperpolarization ($\mathrm{AHP_f}$) and the following voltage peak. This value typically decreased as a function of current step size and was usually largest in the first spike (for current steps below the threshold for triggering a spike doublet). At twice rheobase current, the ADP amplitude (Figure 3.7g) was non-zero in 10/29 Glt neurons (1.1 ± 0.21 mV). In the same subset of neurons, the $\mathrm{AHP_f}$ was measured as the difference in voltage between the spike threshold (defined at the first peak in the 3^{rd} derivative of the voltage) and the voltage trough between the spike and the ADP. The mean $\mathrm{AHP_f}$ was 10.4 ± 0.5 mV (Figure 3.7h). In all recorded neurons the slow afterhyperpolarization ($\mathrm{AHP_s}$), measured between the spike threshold and the voltage trough following the ADP, was also quantified (12.1 ± 0.4 mV, Figure 3.7i).

From these 500 ms long current steps, spikes following the initial doublet appeared to be very regular. To verify if any spike-frequency adaptation occurred in these neurons, stimulus duration was increased to 10 s (Figure 3.7j), with a stimulus amplitude chosen to elicit 6–8 spikes in the first second. Even at this longer stimulus duration, recorded neurons displayed very little spike-frequency adaptation (not counting the spike doublet at the beginning of the spike train), with the firing rates decreasing by only 0.35 ± 0.72 Hz over the 10 s stimulus window ($p = 0.11$, repeated measures ANOVA test).

3.4 Discussion

Using whole-cell patch clamp recordings in acute brain slices from Glt25d2-Cre mice, I have obtained a general characterization of the properties of Glt neurons in V2m. The purpose of these initial experiments was determine if Glt25d2-Cre in V2m can be reliably used to target ttL5 neurons. My results confirm this hypothesis and provide a thorough description of these cells, which may be critical for interpreting future functional studies.

The most distinguishing morphological features of Glt neurons in V2m, and of ttL5 neurons more broadly, are a large apical dendrite extending from the soma in L5 to the pia, with many oblique dendrites along the apical trunk and a large and highly branched tuft of dendrites in L1. Physiologically, they are also easily recognisable by their spiking properties, with a short burst of spikes (typically consisting of two spikes) at the beginning of depolarizing current steps followed by a regular spike train with little or no adaptation. Each spike also has a characteristic shape, distinguished by the presence of a small ADP immediately following each spike.

One of the more noteworthy features that I observed while characterizing the ttL5 neurons in V2m was a strong sag potential, which is often considered a proxy for I_h. This is generally a distinctive property of ttL5 neurons that have previously been described, and serves an important role in maintaining a stable resting membrane potential. Neurons with very large morphologies receive thousands of excitatory inputs and thus require powerful regulatory mechanisms that maintain excitability within a useful dynamic range. Strong sag potentials are known, for instance, to be present in several other types of large excitatory neurons, such as cerebellar Purkinje cells (Fernandez et al., 2007) or hippocampal pyramidal neurons (Magee, 1998). The presence of I_h is likely to underlie other observed phenomena, such as the strong dampening of the voltage during injection of multiple EPSC waveforms. It is also known to be a contributing factor to the emergence of neuronal resonance, which helps control the oscillatory dynamics of a population of neurons (Hutcheon et al., 1996; Zemankovics et al., 2010; Schmidt et al., 2017; Labarrera et al., 2018). In individual cells this is visible as a frequency-dependent amplification of inputs, acting akin to a band-pass filter and predisposing the neuron to oscillatory firing patterns at that frequency (Pike et al., 2000; Hutcheon and Yarom, 2000). In a larger population of neurons, such oscillations may have an important role in coordinating different communication protocols and signalling feedback predictions through the temporal structure and coherence of activity across different brain regions (Engel et al., 2001; Fries, 2015).

Overall, Glt neurons in V2m appear to form a unimodal population of ttL5 neurons, with properties that are consistently representative and characteristic of ttL5 neurons. Historically, the majority of research on ttL5 neurons has described these cells in the cerebral cortex of rats (Ramaswamy and Markram, 2015). While there are many similarities in brain structure, and neuronal morphology and physiology between rats and mice, these remain fundamentally different species separated by 12 million years of evolutionary history (Kimura et al., 2015). The smaller number of papers which describe these cells in mice have generally focused on a very small subset of specific properties, such as soma size or burst propensity (Groh et al., 2010; Shai et al., 2015; Kim et al., 2015). In addition to providing the first description of Glt25d2-Cre labelled cells in V2m, the results in this chapter thus constitute one of the most complete and multifaceted descriptions of mouse ttL5 neurons to date, providing a useful benchmark for comparisons to other cell types or subpopulations of ttL5 neurons located in different cortical areas.

It is worth noting that, within the recorded Glt population, the apparent uniformity in intrinsic properties does not exclude the presence of further subdivisions, which might be differentiable with other features which were not quantified here. Individual subgroups of Glt neurons in V2m might, for instance, differ in expression of specific subsets of genes, or could project to different subsets of the full list of projection targets described in Figure 3.2

(e.g. some Glt neurons could conceivably project to the superior colliculus without also projecting to the striatum).

Generally, there are very many criteria that can be used to categorize cell types, and every population of neurons is likely to have some diversity in these properties. When a population is more narrowly defined, it is expected that both the diversity and the number of neurons that fit the definition would be smaller. The theoretical lower bound of this would be to create a definition that is specific enough to capture a single neuron. Such a narrow definition, however, would not be useful in mice where individual neurons are rarely critical to the overall function of the network. To be able to make general claims, the most useful definition would capture a population that is as large as possible while keeping the variance within the population as small as possible. Genetically labelled populations are particularly useful in this respect, as they allow for easy reproducibility and enable a wider range of experimental tools.

It has previously been shown that different genetically labelled populations of L5 neurons in V1 are able to define specific subpopulations with subtle but significant differences in intrinsic and functional properties (Kim et al., 2015). When considering all ttL5 neurons in V2m, it is thus plausible that these could be subdivided into several distinct or partially overlapping groups with small variability encompassing different ranges within the overall parameter range of all ttL5 neurons. Conversely, a label might be broader than the general ttL5 population, and also include other known neuron subtypes, such as slender-tufted L5 neurons. This is the case, for example, in the Rbp4-Cre line, which labels a broad selection of both thick-tufted and slender-tufted L5 pyramidal neurons (Gerfen et al., 2013; Baker et al., 2018).

Chapter 4

Subcellular circuit mapping

4.1 Introduction

Given a particular set of intrinsic properties, the inputs are what ultimately defines the functional properties of a neuron and its role in cortical computation. While the intrinsic physiology of neurons is key in shaping the computations that a network can perform, an essential requirement to fully understand a cortical circuit is to have a map of the connections between different populations of neurons. Indeed network architecture is a major determinant of the capabilities of any neural network (Zador, 2019). Using highly constrained wiring diagrams is, for instance, the basis for convolutional neural networks, which have been one of the biggest advances in artificial neural networks in the last three decades (LeCun et al., 1989, 1998).

Mapping anatomical pathways connecting neurons in the brain is one of the oldest approaches to studying brain function. Consequently, a wide variety of tracing techniques have been developed over the years to create wiring diagrams for different neural circuits (Lanciego and Wouterlood, 2011). While many of the methods that have been used historically can reveal broad connectivity patterns, recent technical developments have leveraged genetic tools to map connectivity in a cell-type specific way. In particular, glycoprotein-deleted rabies virus can be targeted to Cre-expressing neurons and be used to monosynaptically label presynaptic input populations across the brain (Wickersham et al., 2007; Callaway and Luo, 2015) and optogenetics allows highly targeted stimulation of specific pathways and neural populations to determine functional connectivity (Klapoetke et al., 2014).

Knowing which neurons are connected to each other and where these neurons are located is an important initial step to understanding the functional properties of a circuit. Recording the activity of neurons *in vivo* in both presynaptic and postsynaptic populations can inform us about both the functional properties of cells and what computations are required in order to transform the inputs into output signals. However, to understand the mechanisms by which these computations are implemented in the postsynaptic neurons at a single-cell level, it is also important to known how they connect at a subcellular scale by having a map of the dendritic distribution of the synapses. This is because synapses can have drastically different effects on the postsynaptic neuron depending on the dendritic locations where they occur (Larkum et al., 1999b; Hausser and Mel, 2003). In this chapter, I will address the third aim of my thesis (presented in chapter 1), by describing experiments in which I applied optogenetics and patterned optical stimulation in conjunction with patch-clamp recordings from Glt neurons to map the dendritic input distribution from the largest presynaptic clusters identified using retrograde rabies tracing.

4.2 Chapter methods

4.2.1 Whole-brain input mapping

Whole-brain mapping of presynaptic neurons was done by Dr. Zhiwen Ye in 3 Glt animals using the CVS-N2c strain of glycoprotein-deleted rabies virus (Callaway and Luo, 2015; Reardon et al., 2016). This virus was expressed in Glt neurons in V2m, and was retrogradely transported from these cells to label with mCherry the presynaptic neurons with monosynaptic connections to the targeted Glt population (Figure 4.1).

After allowing sufficient time for the virus to express, the brains of injected animals were perfused and fixed with formaldehyde (4%), and the brain-wide distribution of rabies-labelled neurons was measured using the same serial-section two-photon tomography as was used in chapter 3 (see chapter 2 for full description of method). In brief, each presynaptic neuron was accurately mapped by making serial sections of the fixed rabies-labelled brain and imaging each section under a two-photon microscope. These images were then used to locate and count presynaptic neurons with an automated registration and segmentation software that mapped each neuron's position to the Allen CCF. For each brain, area-wise cell numbers were normalized by the total presynaptic cell count.

4.2.2 Cell selection

After preparing brain slices as described in chapter 2, whole-cell patch clamp recordings were made from Glt neurons in slices containing V2m (based on the approximate stereotaxic coordinates described in Paxinos and Franklin (2007)). The area to record from was first identified under a low magnification (10x/0.25NA) dry objective while illuminating the sample with IR light. To confirm the presence of Glt neurons at the desired recording location, the tdTomato was visualized with short pulses (lasting a few seconds) of epifluorescence stimulation with 550 nm light from the CoolLed pE-4000 illumination system. The pulses were kept short and the light intensity was set as low as

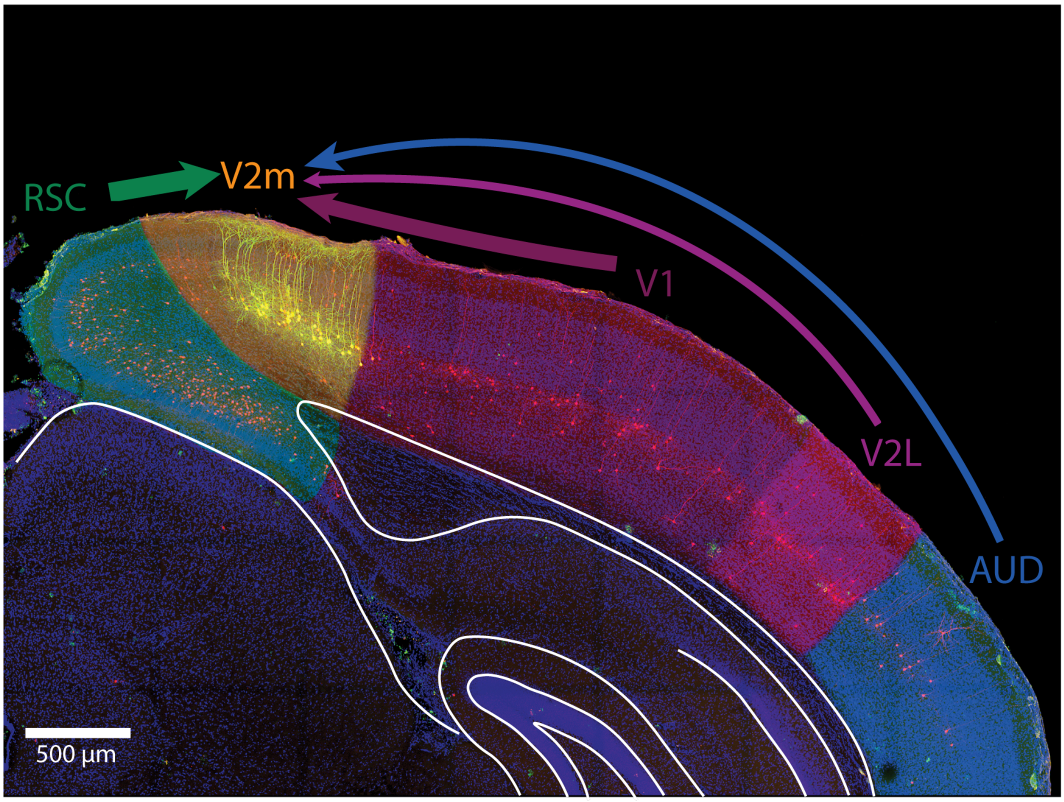

Figure 4.1: Monosynaptic retrograde rabies labelling of inputs to Glt neurons in V2m. Confocal image of a coronal section from a Glt25d2-Cre animal injected with glycoprotein-deleted rabies virus in V2m. Glt starter cells (yellow) were separately infected with adeno-associated virus containing Cre-dependent genes expressing the glycoprotein required for trans-synaptic spread and avian receptor protein (TVA) needed for uptake of the rabies virus. Presynaptic neurons (red) visible in this section have been labelled based on their approximate location in reference to the Paxinos and Franklin (2007) atlas. NB, most long-range inputs originate from L5 of the presynaptic region. This confocal image was acquired by Dr. Ede Rancz.

possible in order to minimize photobleaching as well as any other possible side effects of strong illumination. After identifying the cluster of Glt cells approximately located in V2m, I switched to the high magnification (40×/0.8NA) water immersion objective.

Although Glt25d2-Cre occasionally labels neurons that are not in L5, only labelled neurons in L5 were targeted for recordings while cells in other layers were avoided. This was easily done, as L5 is clearly distinguishable under IR illumination by the presence of many cells with very large pyramidal-shaped somas. Within the fluorescently labelled L5 neurons in V2m, cells were chosen at random. No consideration was thus given to criteria such as soma size and depth (provided they were within L5), which are sometimes used in existing L5 neuron literature to facilitate selection of ttL5 neurons.

4.2.3 Recording conditions

Since the experiments described in this chapter were designed to reveal anatomical features (namely the location of synapses from different presynaptic populations along the dendritic tree of Glt neurons) rather than the physiology of the Glt neurons (which was instead covered in chapter 3), all recordings presented here were performed at room temperature (22 °C). This was done to remain consistent with previous literature and because acute brain slices remain viable for longer at room temperature compared to warmer temperatures. The recording ACSF in these experiments contained 2 mM $CaCl_2$ as well as 1 µM TTX and 100 µM 4-AP. The choice of 2 mM $CaCl_2$ rather than the more physiological 1.5 mM $CaCl_2$ I used in chapter 3 was made both for consistency with previous literature using this technique and because these experiments were meant to reveal the presence of synapses, not their physiological parameters. As presynaptic vesicle release depends on Ca^{2+} ions entering the synaptic bouton, a higher concentration is likely to increase recorded synaptic currents and thus improve the signal-to-noise ratio in these experiments.

4.2.4 sCRACM implementation and analysis

To determine the dendritic distribution of defined inputs from different presynaptic populations onto Glt neurons in V2m, I adopted the method of subcellular channelrhodopsin-assisted circuit mapping (sCRACM), which was first described by Petreanu et al. (2009). This involved labelling neurons in individual presynaptic regions (identified from previous retrograde rabies experiments) with the optogenetic activator Chronos-GFP (Klapoetke et al., 2014). The optogenetic activator was delivered through injected AAVs, which were allowed to express for approximately 4 weeks (minimum 3 weeks, maximum 8 weeks) before performing the recordings (Figure 4.2a).

Following the expression period, whole-cell voltage clamp recordings were made from Glt neurons in V2m, with a holding potential of −70 mV to reveal excitatory postsynaptic currents (EPSCs). The addition of 1 µM TTX to the extracellular recording solution (containing 2 mM $CaCl_2$) blocked any axonal propagation of spikes and thereby ensured that any light-evoked responses were direct monosynaptic responses resulting from stimulation of Chronos-expressing axon terminals, rather than from passing axons terminating in unknown locations on the dendrites. I also blocked voltage-gated K^+ channels by adding 100 µM 4-AP in order to inhibit the repolarization of the axon and thereby increase the amplitude of evoked responses.

Before every recording, I first identified a suitable tdTomato positive Glt neuron in V2m. Since the slice contained light-sensitive axons expressing Chronos, during this identification phase much care was take to visualize the cells without stimulating the optogenetic activators. The epifluorescence system used to visualize fluorescent neurons was a CoolLED pE-4000 illumination system, set to stimulate with 550 nm light the full field-of-view (approximately 100 µm in diameter) through the high magnification (40×/0.8NA) water immersion objective. While this wavelength partially overlaps with the excitation spectrum of Chronos (Klapoetke et al., 2014), the power was set to the smallest value possible while still making out fluorescent neurons (approximately

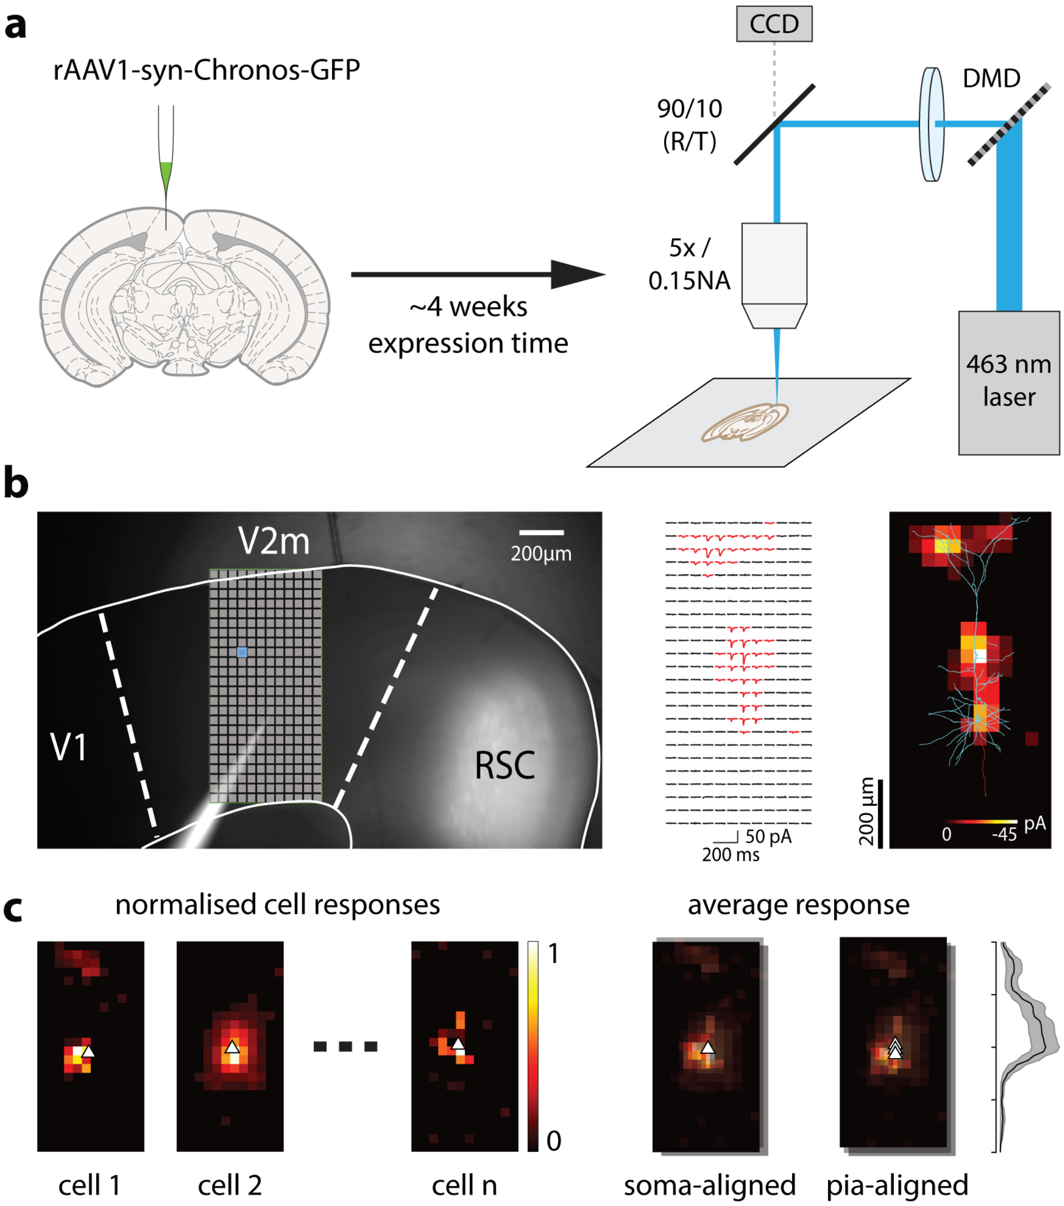

Figure 4.2: Subcellular optogenetic circuit mapping. a. Diagram of experimental setup, showing virus injection and optical path for patterned laser stimulation. **b.** Example sCRACM recording and analysed heatmap for one neuron. **c.** Illustration of analysis procedure by which each heatmap was normalised, aligned, and averaged. Shaded area in the profile plot represents the SEM.

$20\,\mu W$, or ($2.5\,mW/mm^2$) to minimise photoexcitation and bleaching of the Chronos-expressing axons. At this irradiance, postsynaptic responses were never observed when stimulating with this wavelength after obtaining whole-cell patch-clamp configuration. As a further precaution, the exposure to even 550 nm was minimised by keeping the illumination time as brief as possible, turning it on for only a few seconds at a time while identifying which cells to record from, and keeping the light off for the remaining time (including the entire duration of the recording).

After identifying a fluorescent Glt neuron to record from, I aligned the visible initial portion of the apical dendrite with the sCRACM laser stimulation grid, which had be pre-programmed into a digital micromirror device (DMD) coupled to the optical path of the microscope. This was done by manually rotating the recording chamber (containing the brain slice being continually perfused with ACSF) until the apical dendrite of the targeted neurons (approximately the first 50–100 µm) visible under the 40×/0.8NA water immersion objective) could be seen to extend vertically within the field-of-view. This alignment step ensured that the entire apical dendrite would fall within the sCRACM stimulation area, and facilitated the subsequent analysis when aligning and averaging the input maps of individual neurons (Figure 4.2c). During this process, care was taken to not disrupt the continuous flow of ACSF through the recording chamber.

The identified Glt neuron, rotated to have the apical dendrite aligned with the rectangular stimulation grid, was then targeted for patch-clamp recording. After achieving whole-cell configuration, the objective was changed to a 5×/0.15NA dry objective through which the sCRACM laser stimulus was delivered. This stimulus consisted of a $500 \times 1000\,\mu m$ grid divided into 12×24 spots of light (each spot being a square of side $\sim 41.7\,\mu m$) delivered through the objective using the DMD. The grid was approximately centered on the soma and aligned to the pia along one of the short edges of the stimulation rectangle (Figure 4.2b).

During recordings, each spot within the grid was stimulated for 1 ms, with 100 ms intervals between consecutive spots. Each recording trial consisted of a single repetition of all 288 stimuli, and was additionally followed by a “full-field” stimulus, in which all stimulation spots were illuminated simultaneously for 1 ms. For each cell, I recorded 5–20 trials with 30 s pauses between trials to allow the axons to recover and minimise any opsin desensitization. Recordings were baselined in a 40 ms window before each stimulus and averaged across trials, and the peak current was measured in a 20 ms window after the stimulus. Peak currents lower than seven times the standard deviation of the baseline noise were scored as zero. For each cell, I attempted to collect as many trials as possible (up to a maximum of 20) in order to improve the signal-to-noise ratio, and I did not analyse any cell for which I was not able to obtain at least 5 trials.

Following each recording, an image was taken in which the location of the recorded cell (filled with Alexa Fluor 488) relative to the stimulation grid could be seen. This allowed me to to later align the recorded responses to the location of the pia and soma within the stimulation grid. Because peak response amplitudes varied between cells and preparations, to obtain average input distributions from a presynaptic population, the heatmap for each cell was normalised to the peak EPSC value for that cell (Figure 4.2c). Heatmaps were then aligned horizontally by soma location and then vertically by either soma or pia location before averaging across cells. For visualization purposes, before averaging the heatmaps were aligned as accurately as the soma could be localised (which was typically to within one quadrant of a stimulation spot). The normalised profile of the average heatmap is also plotted so as to facilitate direct comparisons of the laminar locations of inputs (Figure 4.2c, right). The shaded region representing the SEM was obtained by calculating a matrix of SEM values for each pixel in the average heatmap, summing across each row, and normalising to the peak of the average heatmap profile.

4.3 Results

The monosynaptic retrograde rabies labelling revealed that Glt neurons in V2m receive a large number of inputs clustered in several distinct regions throughout the brain, primarily on the ipsilateral hemisphere (Figure 4.1). These inputs were subdivided into "proximal cortex" (defined here as the collection of visual areas and the adjacent retrosplenial cortex), "distal cortex" (referring to all other cortical areas that were not counted as proximal), and inputs from the thalamus. The few presynaptic neurons on the contralateral hemisphere originated primarily in the same region of V2m as the primary Glt neurons where the rabies virus was expressed (Figure 4.3).

The most prominent presynaptic regions were the proximal cortical areas (particularly the primary visual cortex and the retrosplenial cortex), with a smaller number of distal cortical inputs tightly clustered primarily in the orbitofrontal cortex and anterior cingulate cortex. Prominent thalamic inputs were also observed, originating in areas such as the anterior thalamus (AM), which has been linked to vestibular signalling (Rancz et al., 2015), and the lateral posterior thalamus (LP), which is often regarded as a higher-orded visual thalamic region (Bennett et al., 2019).

Given this list of the largest presynaptic clusters of neurons projecting to the Glt population in V2m, I next sought to determine which dendritic domains each of these input pathways was targeting. This was done through a series of experiments, each targeting an injection of Chronos (a blue-sensitive optogenetic activator) to only one presynaptic region and mapping the inputs from that region onto Glt dendrites using the sCRACM method.

4.3.1 Subcellular distribution of RSC inputs

By far the largest presynaptic population of cells identified from the rabies tracing experiment was located in RSC (labelled as RSP in the Allen CCF, Figure 4.3c). I therefore chose this area as the first to be mapped with sCRACM (Figure 4.4a). RSC inputs were found to primarily target the perisomatic dendrites, although a smaller amount of input was also observed in the superficial

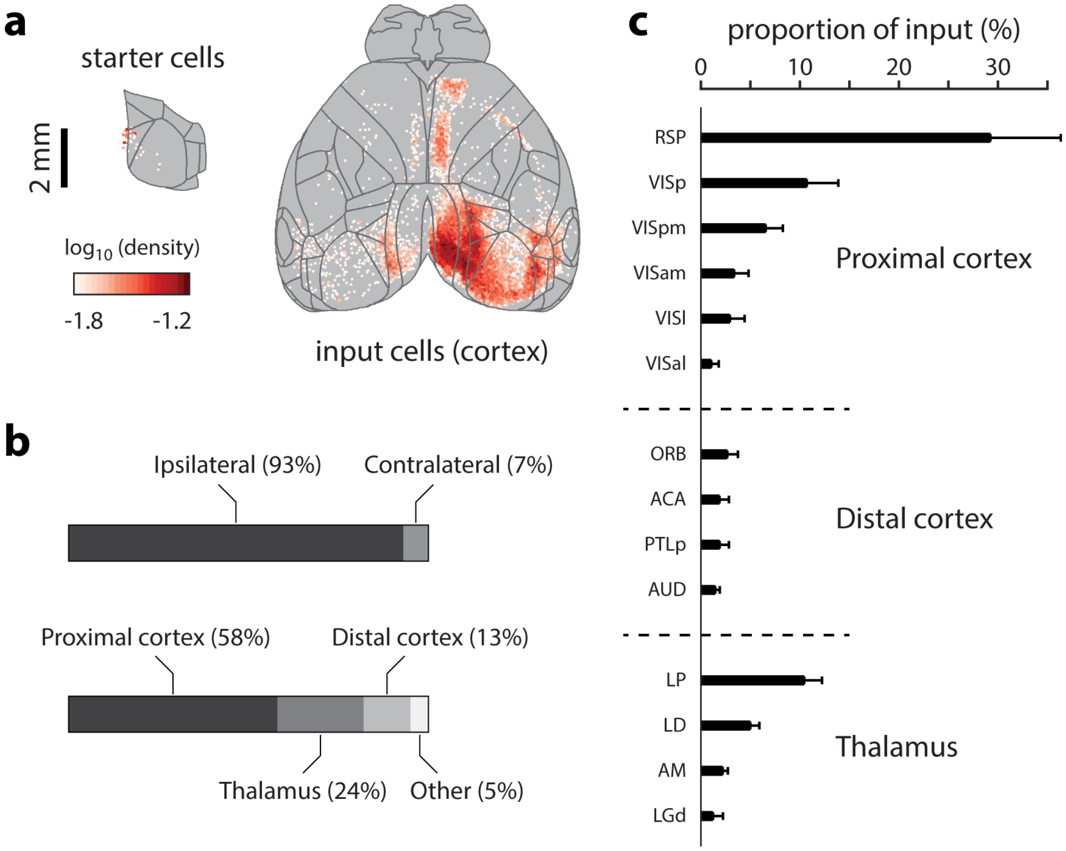

Figure 4.3: Brain-wide distribution of rabies-labelled presynaptic neurons. **a.** Top-down view of cell distribution throughout cortex (relative to the Allen Brain Atlas) for Glt starter cells (left) and presynaptic rabies-labelled cells (right). The data from 3 animals is overlayed. **b.** Proportion of presynaptic cells divided by hemisphere (top) or broad region of origin (bottom). **c.** Proportion of input subdivided by brain area (mean ± SD for n = 3 mice). Only ipsilateral inputs are shown for areas with at least 1% of the total presynaptic cells. The segmentation and area labels shown here are derived from the Allen Brain Atlas, which differs somewhat from the Paxinos and Franklin (2007) atlas. The data in this figure was collected by Dr. Zhiwen Ye.
Abbreviations: RSP, retrosplenial area; VISp, primary visual area; VISpm, posteromedial visual area; VISam, anteromedial visual area; VISl, lateral visual area; VISal, anterolateral visual area; ORB, orbital area; ACA, anterior cingulate area; PTLp, posterior parietal association areas; AUD, auditory areas; LP, lateral posterior thalamus; LD, lateral dorsal thalamus; AM, anteromedial thalamus; LGd, dorsal lateral geniculate thalamus.

layers of the cortex (Figure 4.4b,c). To visualise the diversity in responses across individual cells, each heatmap was horizontally projected into a single column by computing the sum of responses across each row of the full heatmap for that cell. Individual cells were then re-normalized to the highest value within the projected column and sorted by soma depth (Figure 4.4d). These individual responses were diverse, with most neurons receiving strong responses to the perisomatic region and a subset of neurons also receiving input to the apical and tuft dendrites.

The location of inputs relative to the soma and pia were further quantified by fitting a Gaussian function to each cell's projected heatmap. The average peak position of these Gaussian fits ($\bar{\mu}$) when soma-aligning the cells was $31 \pm 17\,\mu m$ above the soma, and $436 \pm 20\,\mu m$ below the pia when pia-aligning the cells. The Gaussian fits also had an average width (i.e. standard deviation, $\bar{\sigma}$) of $74 \pm 11\,\mu m$.

As most recorded cells were found to only receive perisomatic input, with little or no input to the apical dendrite, one concern is that apical input might be present but go undetected due to passive attenuation of the currents over the long distance between the soma and apical tuft. While strong distal currents would still be detected, weaker currents could thus fall below the signal-to-noise threshold. Since the intensity of the sCRACM light stimuli were calibrated to deliver equal light power at every spot on the sCRACM grid in all cells, there is some variability in the total amount of postsynaptic current in each cell, which is determined by the number and strength of the synapses to that cell.

If it were true that apical inputs are present but go undetected, one would expect these inputs to appear more prominently in neurons with very strong total input, and the sCRACM map to become progressively more biased towards basal inputs as the total current recorded in a cell decreases. This was tested by recording the postsynaptic response to a 1 ms full-field stimulus, i.e. turning on all spots of the sCRACM grid at the same time. With this stimulus, the total RSC input current to Glt neurons was on average $280 \pm 37\ pA$ (Fig-

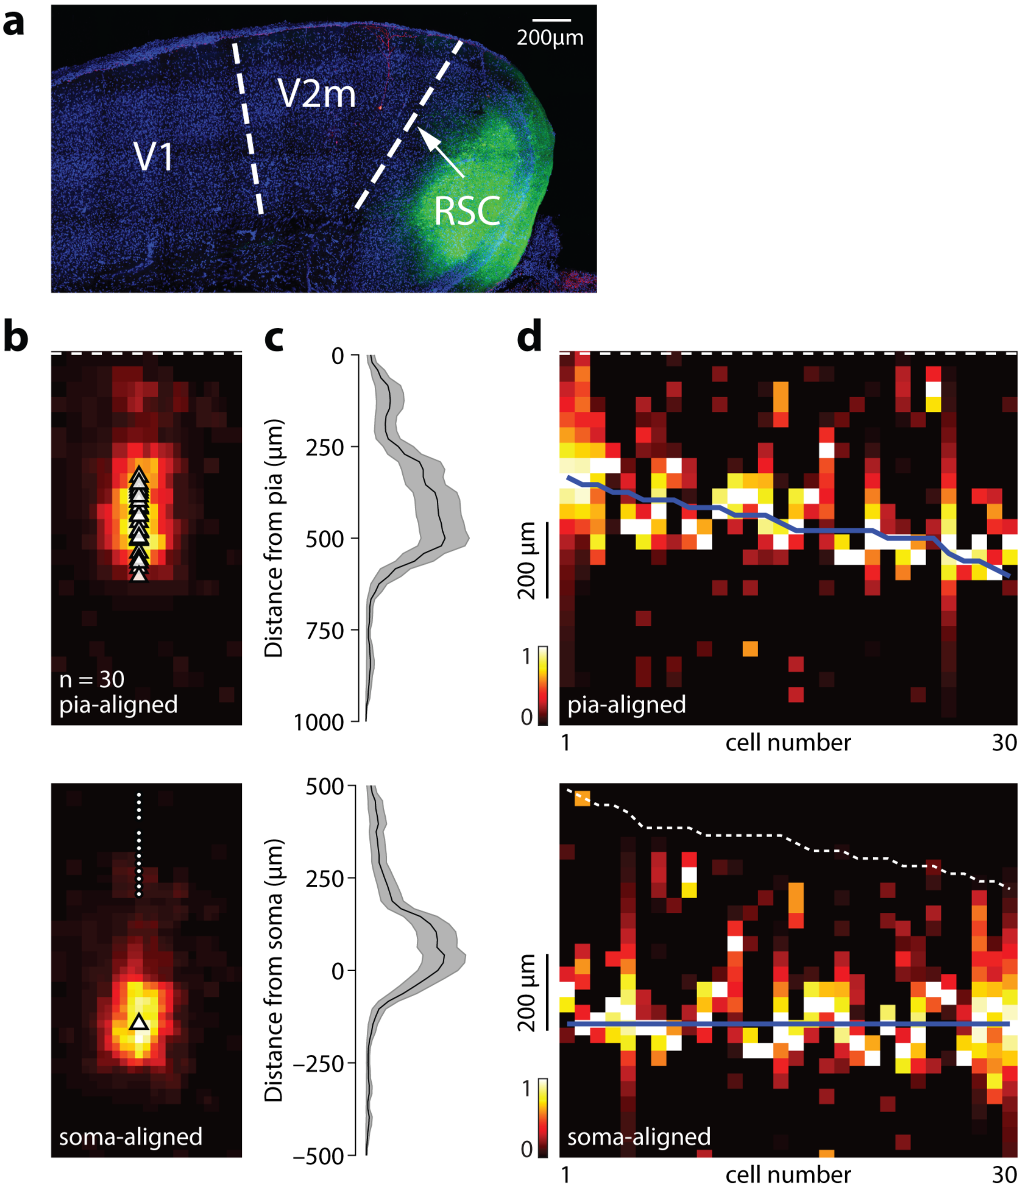

Figure 4.4: Subcellular distribution of RSC inputs. a. Confocal image of a representative brain slice containing recorded neurons in V2m and the injection site in RSC. **b.** *top:* Average pia-aligned input map based on 30 recorded neurons. *bottom:* the same heatmaps are shown aligned by the soma. **c.** Input profile of average pia-aligned (top) and soma-aligned (bottom) responses. **d.** Normalised horizontal projections of heatmaps for individual cells, aligned by the pia (top) or soma (bottom) and ordered by cortical depth of the soma. In all heatmaps, dashed white lines and white dots indicate the pia locations, while white triangles and blue lines indicate soma locations.

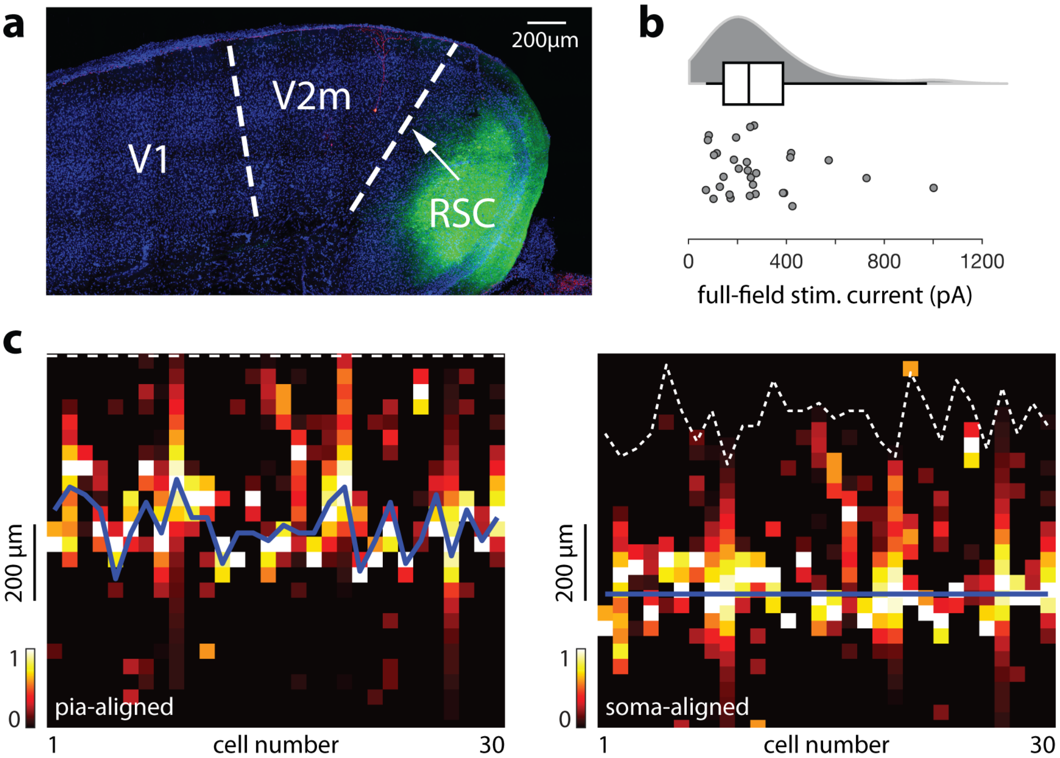

Figure 4.5: Influence of total RSC input on sCRACM maps. a. Confocal image of a representative brain slice containing recorded neurons in V2m and the injection site in RSC (same as in Figure 4.4). **b.** Distribution of peak currents in response to simultaneous 1 ms stimulation of the full sCRACM grid. **c.** Normalised horizontal projections of heatmaps for individual cells, aligned by the pia (left) or soma (right) and ordered by peak current during full-field stimulation (low to high current). Dashed white lines indicate pia location; blue lines indicate soma location.

ure 4.5b). When sorting cells according to this full-field current, there appears to be no clear relationship between location of inputs and total current to each cell (Figure 4.5c).

4.3.2 Subcellular distribution of V1 inputs

The second largest input population identified from the rabies tracing was in V1 (labelled as VISp in the Allen CCF, Figure 4.3c). Injections to V1 were targeted to the deeper layers (Figure 4.6a), as in the rabies tracing most inputs were found to originate in L5. However, to some extent the Chronos-GFP virus also spread to others layers, as is apparent from the fluorescence in Figure 4.6a, so it is possible that the recorded responses also reflect some amount of input from more superficial layers of V1. An exact quantification of the number of

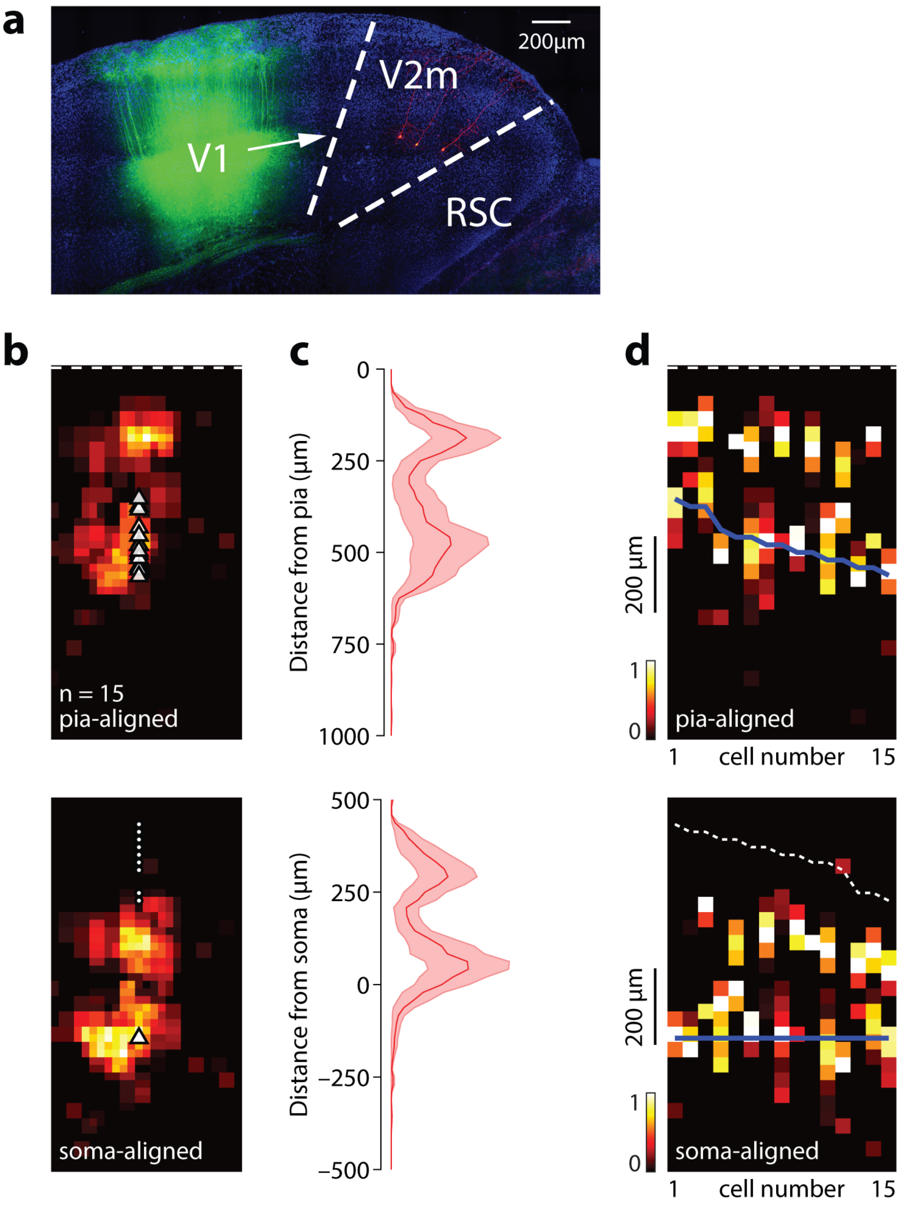

Figure 4.6: Subcellular distribution of V1 inputs. a. Confocal image of a representative brain slice containing recorded neurons in V2m and the injection site in V1. **b.** *top:* Average pia-aligned input map based on 15 recorded neurons. *bottom:* the same heatmaps are shown aligned by the soma. **c.** Input profile of average pia-aligned (top) and soma-aligned (bottom) responses. **d.** Normalised horizontal projections of heatmaps for individual cells, aligned by the pia (top) or soma (bottom) and ordered by cortical depth of the soma. In all heatmaps, dashed white lines and white dots indicate the pia locations, while white triangles and blue lines indicate soma locations.

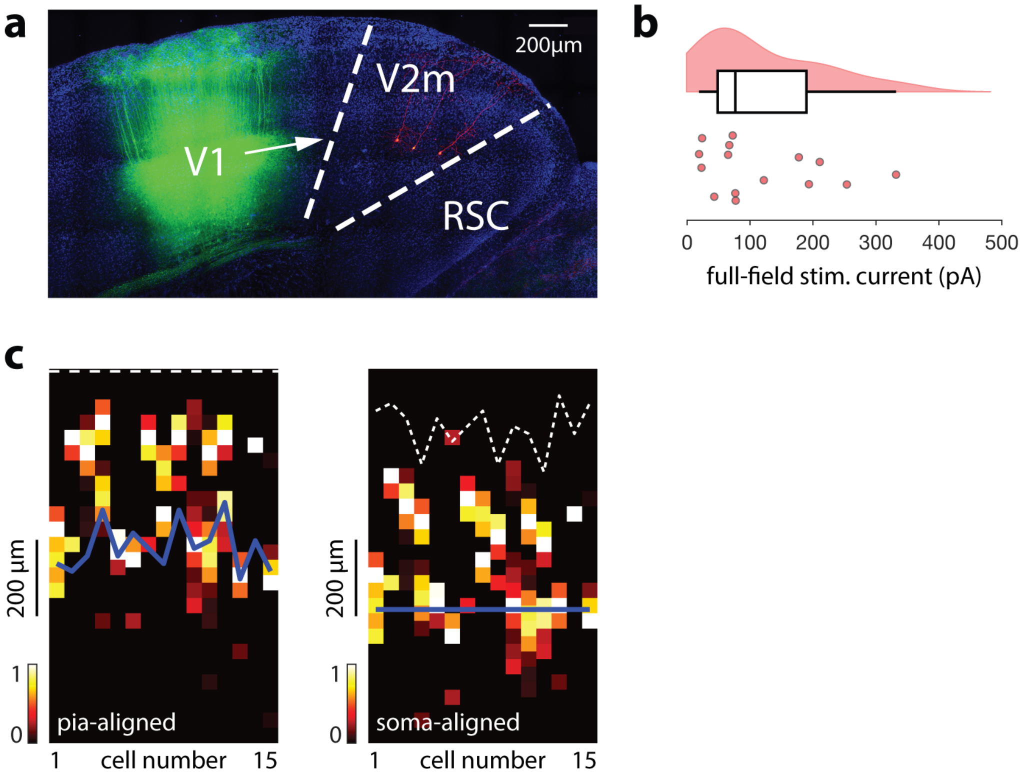

Figure 4.7: Influence of total V1 input on sCRACM maps. a. Confocal image of a representative brain slice containing recorded neurons in V2m and the injection site in V1 (same as in Figure 4.6). **b.** Distribution of peak currents in response to simultaneous 1 ms stimulation of the full sCRACM grid. **c.** Normalised horizontal projections of heatmaps for individual cells, aligned by the pia (left) or soma (right) and ordered by peak current during full-field stimulation (low to high current). Dashed white lines indicate pia location; blue lines indicate soma location.

infected cells in each layer was not possible in these samples due to extensive expression of fluorescence in the neuropil making it impossible to accurately count the cell bodies.

The V1 inputs were found to be bimodally distributed, with strong input to dendrites located in both deep and superficial layers (Figure 4.6b,c). Although several neurons only showed responses in either the apical or the perisomatic area, most received some input to both regions (Figure 4.6c). Notably, the apical input did not extend all the way to the apical tuft in L1, but rather appeared to target mostly the dendrites in L2/3.

To quantify the locations of these input domains, for each cell two Gaussian function were separately fitted to the superficial and the deep portions of the projected heatmaps. For the purposes of this quantification, the cutoff separating superficial and deep was defined as 300 µm from the pia. The mean input peak $\bar{\mu}$ for deep layer inputs was located 31 ± 13 µm above the soma and 451 ± 19 µm from the pia, with a standard deviation of the Gaussian fit $\bar{\sigma} = 46.8 \pm 8$ µm. On the other hand, superficial inputs targeted a region with $\bar{\mu} = 278 \pm 15$ µm from the soma and 183 ± 10 µm from the pia ($\bar{\sigma} = 44 \pm 4$ µm). Cells with only basal inputs (n = 5/15) were excluded from the quantification of the superficial input location and cells with no detectable input near the basal region (n = 1/15) were excluded when quantifying the location of deep layer inputs.

As with the inputs from RSC, full-field stimuli were used to measure the total current from to the Glt neurons from the presynaptic population in V1. Total currents for V1 inputs were 117 ± 25 *pA* (Figure 4.7b), which was significantly smaller than the inputs from RSC ($p = 6.01 \times 10^{-4}$, Mann–Whitney U test). Since these smaller currents were sufficient to reveal apical input to L2/3, this observation provides further support for the argument that the lack of apical input from RSC is not due to the size of currents from this pathway. Furthermore, when reordering the individual heatmaps from each recorded cell according to the total current during full-field stimulation, cells receiving apical input appear to be evenly distributed and are not biased towards cells with strong total current (Figure 4.7), as was the case for RSC inputs. This indicated that cells with stronger currents were not more likely to show apical inputs than cells with smaller currents.

4.3.3 Subcellular distribution of OFC inputs

I next investigated the circuit formed by the projection to Glt neurons from the orbitofrontal cortex (OFC, labelled as ORB in the Allen CCF), which contained the most distant cluster of input cells (Figure 4.3c). in the (Paxinos and Franklin, 2007) atlas, the OFC is split into three subdivisions: medial orbital (MO), ventral orbital (VO), and lateral orbital (LO). Of these, rabies labelled cells were located mainly on the ventral half of the LO and VO regions, immediately above the anterior olfactory areas. For the sCRACM experiments, injections of Chronos were thus targeted to these subdivisions of the OFC (Figure 4.8a).

The OFC has been implicated in a large range of sophisticated behaviours and cognitive processes, as well as top-down modulation of many brain regions, including visual areas, to implement functions such as attention. For this reason, it is generally regarded as being very high in the hierarchy of cortical regions, in the sense of being relatively remote from any direct source of sensory input (Rolls, 2004; Gilbert and Sigman, 2007; Zanto et al., 2011). This pathway therefore seems to be of particular interest in the context of cortical feedback to the visual cortex, as many would consider it to be a more canonical top-down projection compared to the input from RSC, which is not referenced as frequently in relation to cortical hierarchies.

The input from OFC was found to be highly localised to the perisomatic region (Figure 4.8b,c). Even across individual cells, the input appeared to target this area with remarkable specificity regardless of soma depth (Figure 4.8d). The mean of the Gaussian functions fitted to the individual input maps was located $9 \pm 11\,\mu\text{m}$ above the soma and $480 \pm 22\,\mu\text{m}$ below the pia, with a standard deviation $\bar{\sigma} = 85 \pm 7\,\mu\text{m}$.

While such focused targeting of the inputs to the perisomatic regions could be cause for even more caution than the RSC input, full-field stimulation of OFC inputs induced peak currents of $482 \pm 65\ pA$ (Figure 4.9), which is substantially larger than the currents from either RSC ($p = 0.0095$, Mann–Whitney U test) or V1 ($p = 7.11 \times 10^{-5}$, Mann–Whitney U test). As with the previous pathways, for OFC inputs there was also no meaningful difference in dendritic targeting between cells with lower and higher total currents (Figure 4.9c). Altogether, these factors support the notion that Glt neurons truly do not receive strong input beyond the basal dendrites and proximal portion of the apical and oblique dendrites, and that this result is not an artefact of low input current.

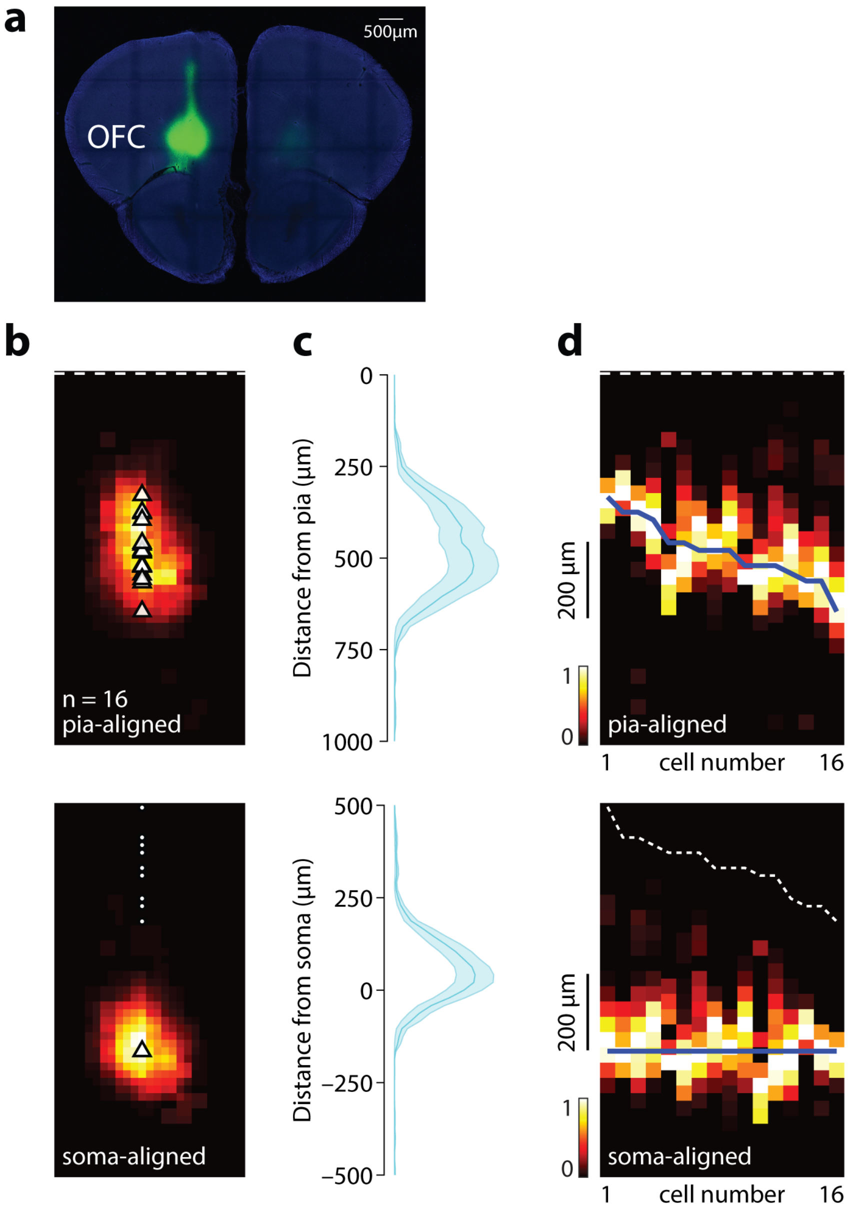

Figure 4.8: Subcellular distribution of OFC inputs. a. Confocal image of a representative brain slice containing the injection site in OFC. **b.** *top:* Average pia-aligned input map based on 16 recorded neurons. *bottom:* the same heatmaps are shown aligned by the soma. **c.** Input profile of average pia-aligned (top) and soma-aligned (bottom) responses. **d.** Normalised horizontal projections of heatmaps for individual cells, aligned by the pia (top) or soma (bottom) and ordered by cortical depth of the soma. In all heatmaps, dashed white lines and white dots indicate the pia locations, while white triangles and blue lines indicate soma locations.

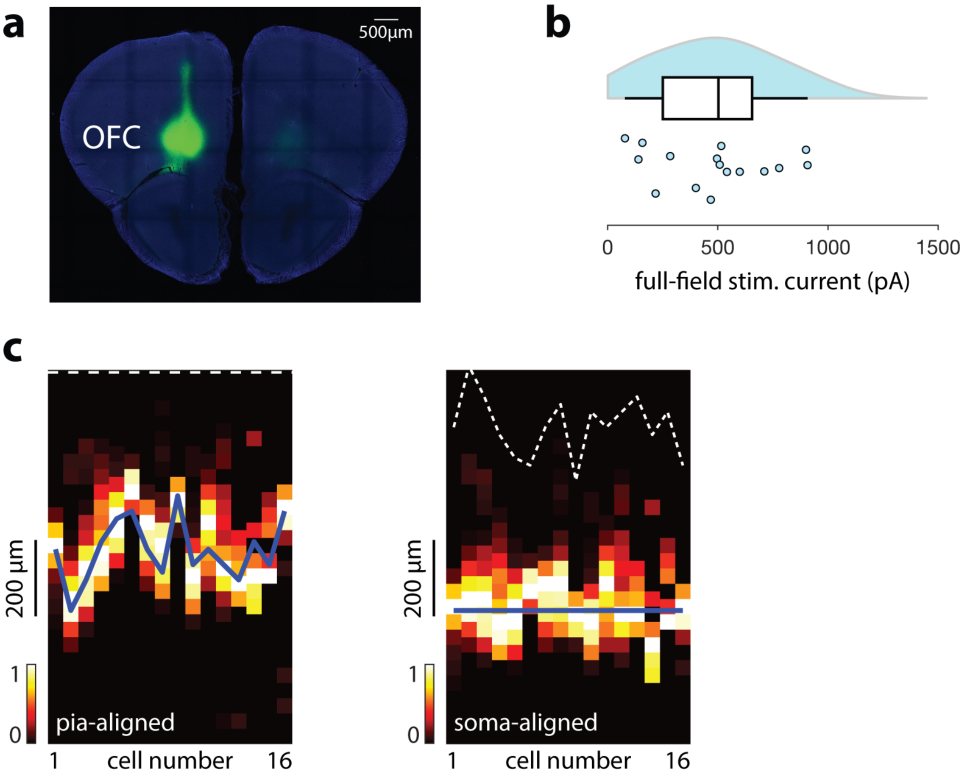

Figure 4.9: Influence of total OFC input on sCRACM maps. a. Confocal image of a representative brain slice containing the injection site in OFC (same as in Figure 4.8). **b.** Distribution of peak currents in response to simultaneous 1 ms stimulation of the full sCRACM grid. **c.** Normalised horizontal projections of heatmaps for individual cells, aligned by the pia (left) or soma (right) and ordered by peak current during full-field stimulation (low to high current). Dashed white lines indicate pia location; blue lines indicate soma location.

4.3.4 Subcellular distribution of inputs from the anterior thalamic nuclei

Inputs from the thalamus play an important role in cortical processing that is likely to be substantially different from the cortico-cortical connections described above, due to the diverse nature of the signals they carry. Of the thalamic inputs to Glt neurons identified from the rabies tracing, the anterior thalamus was investigated first as it is known to process both vestibular and spatial signals (Jankowski et al., 2013; Rancz et al., 2015) and could thus be an important source of multimodal signals with non-visual sensory components into V2m.

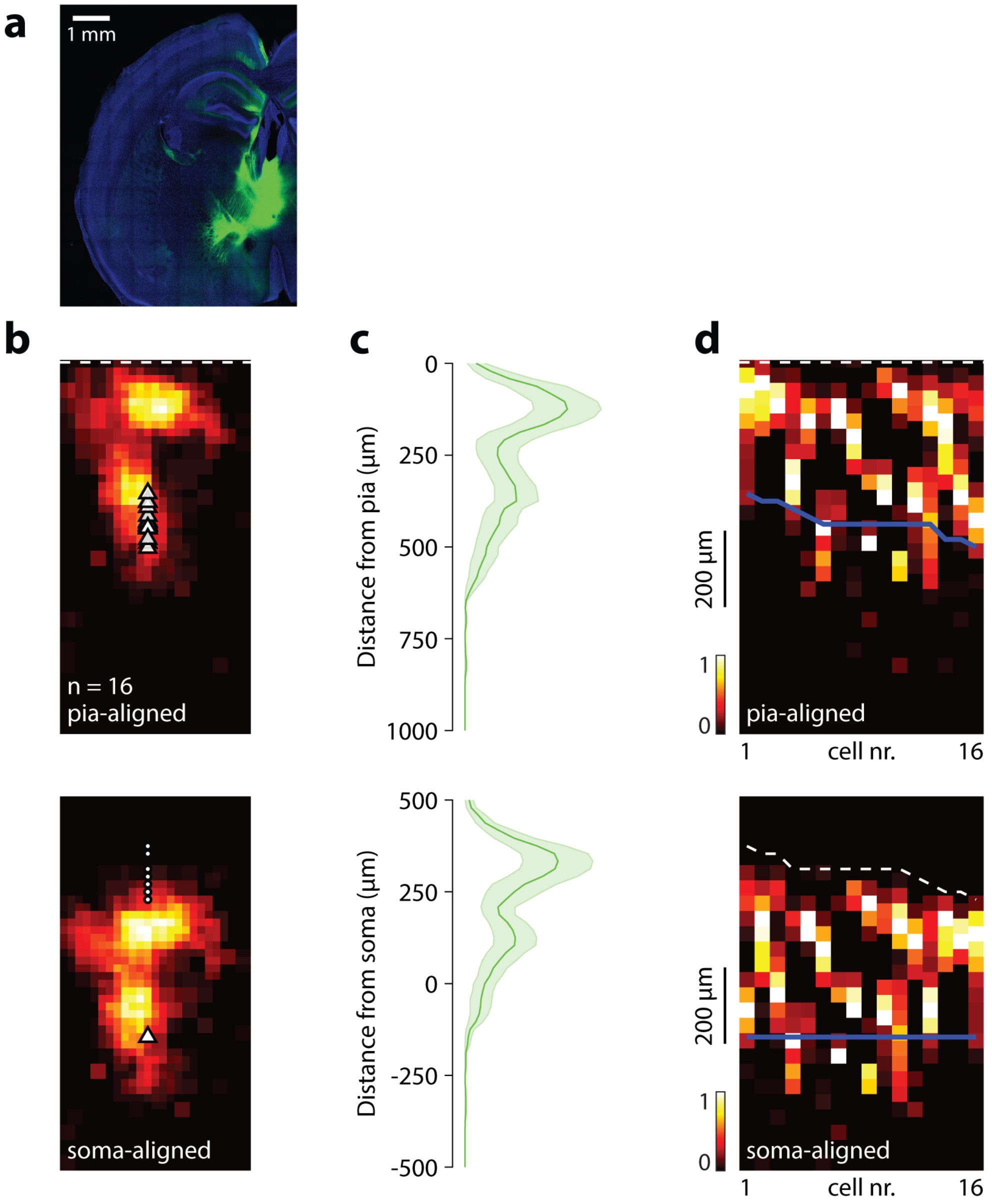

Figure 4.10: Subcellular distribution of anterior thalamus inputs. a. Confocal image of a representative brain slice containing the injection site in anterior thalamus. **b.** *top:* Average pia-aligned input map based on 16 recorded neurons. *bottom:* the same heatmaps are shown aligned by the soma. **c.** Input profile of average pia-aligned (top) and soma-aligned (bottom) responses. **d.** Normalised horizontal projections of heatmaps for individual cells, aligned by the pia (top) or soma (bottom) and ordered by cortical depth of the soma. In all heatmaps, dashed white lines and white dots indicate the pia locations, while white triangles and blue lines indicate soma locations.

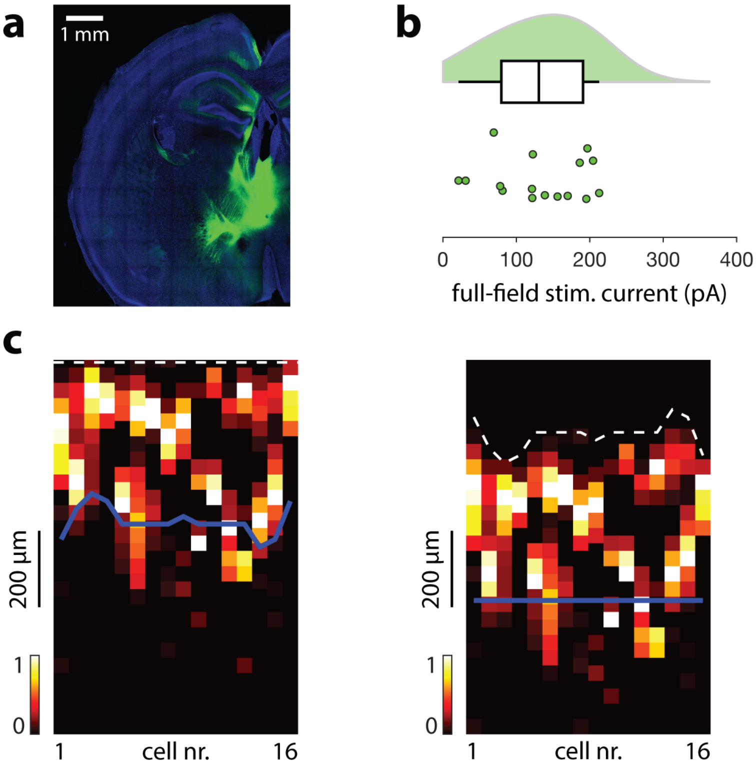

Figure 4.11: Influence of total anterior thalamus input on sCRACM maps. a. Confocal image of a representative brain slice containing the injection site in anterior thalamus (same as in Figure 4.10). **b.** Distribution of peak currents in response to simultaneous 1 ms stimulation of the full sCRACM grid. **c.** Normalised horizontal projections of heatmaps for individual cells, aligned by the pia (left) or soma (right) and ordered by peak current during full-field stimulation (low to high current). Dashed white lines indicate pia location; blue lines indicate soma location.

The spatially targeted stimulation of this pathway showed a bimodal connectivity pattern, with particularly strong input to tuft dendrites in L1 and a somewhat weaker connection to dendrites in the upper portion of the perisomatic region and extending into L2/3 (Figure 4.10). This patterns was found consistently across individual cells, with all recorded neurons receiving at least some input to the apical dendrites, and in many cases receiving little or no input to the perisomatic region.

As with the V1 inputs, the input map to each cell was fitted with two Gaussian functions around the superficial and deep regions of the heatmap (respectively less or more than 300 µm from the pia). The mean of the Gaussian functions $\bar{\mu}$ for apical inputs was located 27 ± 91 µm below the pia and

321 ± 12 µm above the soma, with a standard deviation $\bar{\sigma} = 60 \pm 6$ µm. Meanwhile, in cells with input to deeper layers this was targeted to 403 ± 19 µm from the pia and 44 ± 20 µm above the soma, with $\bar{\sigma} = 62 \pm 9$ µm

Given the strong mapping of anterior thalamus inputs to the apical tuft, this pathway did not face the same concern regarding detection of these inputs as the recordings of RSC and OFC inputs. Nonetheless, peak full-field currents from this pathway could provide a useful indication of the current amplitudes required to detect these inputs. The currents following full-field stimulation of anterior thalamus inputs were $132 \pm 15\, pA$ (Figure 4.11), which is similar to the full-field currents from V1 (p = 0.31, Mann–Whitney U test) but significantly smaller than either OFC or RSC currents ($p = 1.12 \times 10^{-4}$ and $p = 1.9 \times 10^{-3}$, respectively; Mann–Whitney U test). This is noteworthy, as it implies that full-field currents that are at least as large as this should be detectable throughout the full dendritic tree.

4.3.5 Subcellular distribution of local V2m inputs

In addition to the long-range inputs described above, I also measured the subcellular distribution of inputs originating locally from the superficial layers of V2m. To do this I used a Chronos-GFP virus for which expression was limited to neurons that did not express Cre recombinase. The efficacy of the virus was tested by injecting 100–150 nL of virus into both L2/3 and L5 of V2m in a mouse line with dense Cre-labelling in L5 neurons (Rbp4-Cre). Despite the large injection volume, Rbp4-Cre labelled and Chronos-GFP positive neurons were almost entirely non-overlapping, indicating that the virus was working as expected (Figure 4.12).

For subcellular input mapping of local connection, injections of 100 nL were targeted to L2/3 of V2m in Glt25d2-Cre mice (Figure 4.13a). Although GFP fluorescence was brightest in the superficial layers, it is likely that some virus was also taken up by non-Glt neurons in deeper layers, which could contribute to the resulting input maps. However, as for the other inputs, a complete quantification of the cells labelled in each layer of the injection site

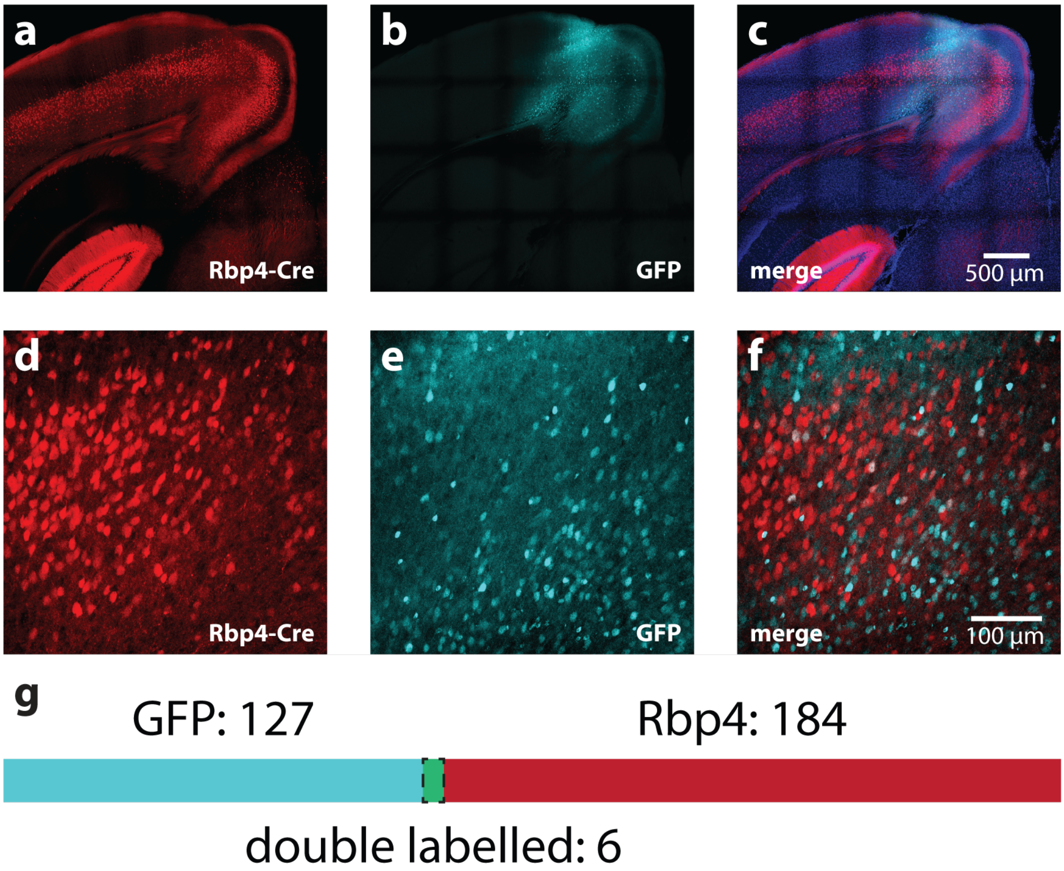

Figure 4.12: Control injection of Chronos Cre-OFF virus. a. tdTomato-labelled Rbp4-Cre neurons in L5 across cortex. **b.** GFP-tagged Chronos Cre-OFF virus injected across L2/3 and L5 of V2m. **c.** Composite image of tdTomato, GFP, and DAPI labelling. **d.** Magnified view of L5 neurons at the injection site in V2m. **e.** Chronos Cre-OFF expressing neurons in the same region. **f.** Composite image of the tdTomato and GFP labelling. **g.** Number of cells in the field of view around the injection site infected by labelled by either Chronos-GFP, tdTomato, or both.

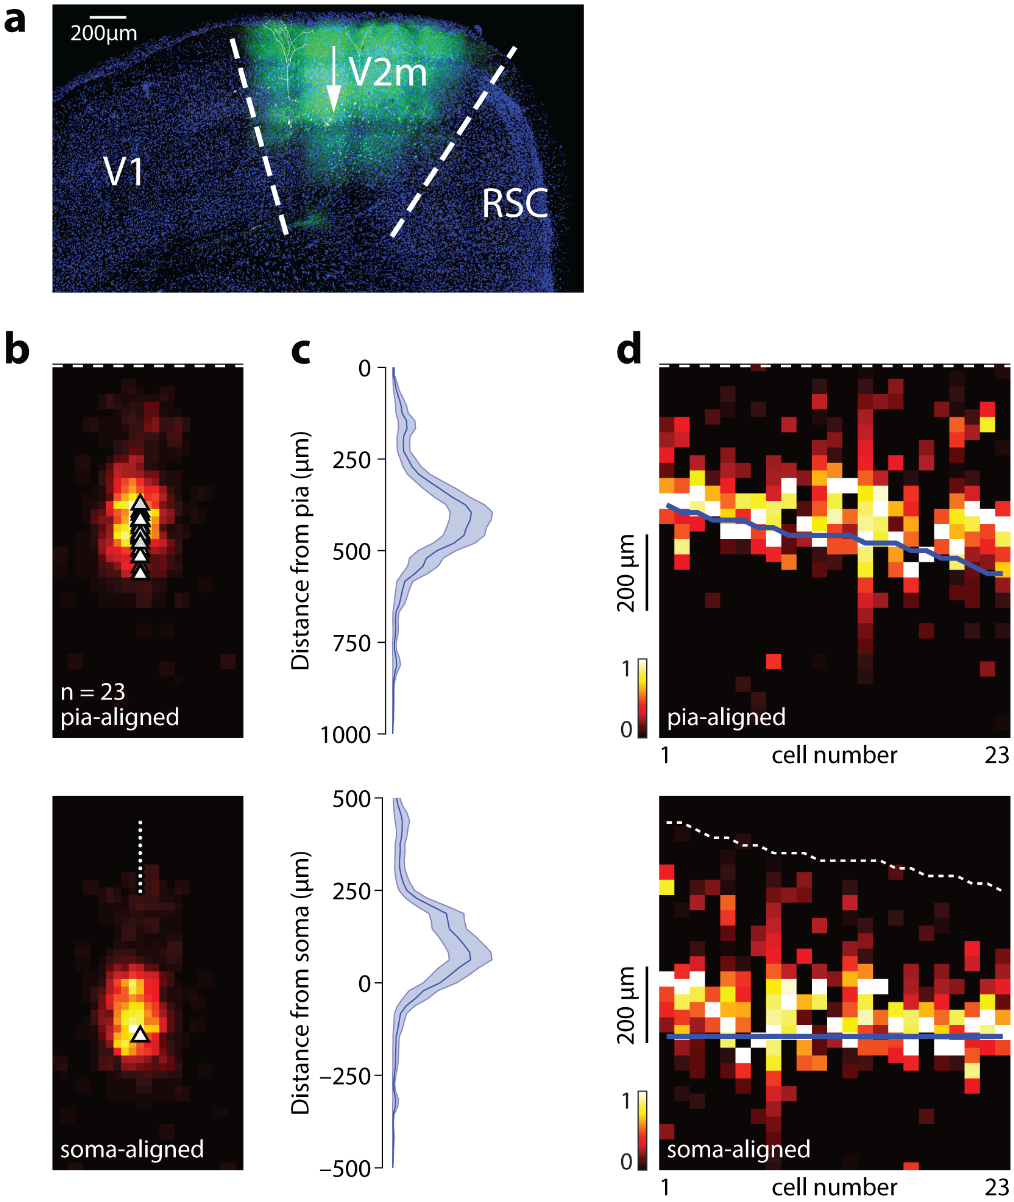

Figure 4.13: Subcellular distribution of inputs from local populations in superficial V2m. **a.** Confocal image of a representative brain slice containing recorded neurons in L5 and the injection site in superficial layers of V2m. **b.** *top:* Average pia-aligned input map based on 15 recorded neurons. *bottom:* the same heatmaps are shown aligned by the soma. **c.** Input profile of average pia-aligned (top) and soma-aligned (bottom) responses. **d.** Normalised horizontal projections of heatmaps for individual cells, aligned by the pia (top) or soma (bottom) and ordered by cortical depth of the soma. In all heatmaps, dashed white lines and white dots indicate the pia locations, while white triangles and blue lines indicate soma locations.

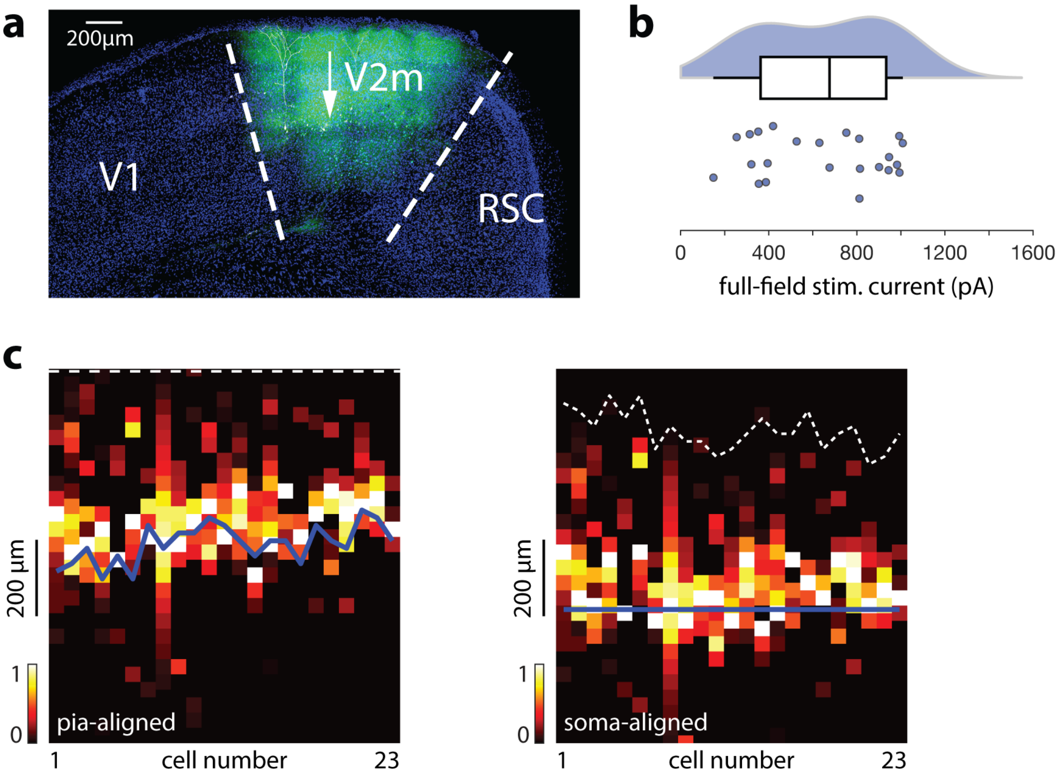

Figure 4.14: Influence of total V2m input on sCRACM maps. a. Confocal image of a representative brain slice containing recorded neurons in L5 and the injection site in superficial layers of V2m (same as in Figure 4.13). **b.** Distribution of peak currents in response to simultaneous 1 ms stimulation of the full sCRACM grid. **c.** Normalised horizontal projections of heatmaps for individual cells, aligned by the pia (left) or soma (right) and ordered by peak current during full-field stimulation (low to high current). Dashed white lines indicate pia location; blue lines indicate soma location.

was not possible due to extensive expression of Chronos-GFP in the surrounding neuropil masking the individual somas.

The input from local neurons was found to target primarily the perisomatic region, with little input to superficial layers (Figure 4.13b,c). Inputs in individual cells targeted a region on average 60 ± 10 µm above the soma and 414 ± 10 µm from the pia, and had a standard deviation of 84 ± 10 µm. Full-field stimulus currents in these cells were the largest of all the recorded inputs at 641 ± 60 pA (Figure 4.14), and as with previous areas the input maps appeared to be unaffected the current in each cell.

4.3.6 Comparison of presynaptic inputs maps

Overall, substantial differences were observed in the dendritic domains targeted by the different inputs. While OFC, RSC, and V2m inputs all predominantly targeted the perisomatic regions, the OFC inputs closely followed the soma location, targeting approximately the same perisomatic region in every cell, while the local V2m inputs appeared to be more tied to a particular laminar location, such that the inputs were biased towards the upper basal and oblique dendrites in deeper neurons (Figure 4.13c). Indeed for V2m inputs, soma depth was significantly correlated with the location of the input relative to the soma ($r^2 = 0.29$, $p = 0.0084$, F-test, Figure 4.15a) while the same was not true of the inputs to OFC ($p = 0.69$, F-test, Figure 4.15b), RSC ($p = 0.26$, F-test, Figure 4.15c), or the perisomatic V1 inputs ($p = 0.38$, F-test, Figure 4.15d). However, a similar correlation was observed in the superficial inputs from both V1 and the anterior thalamus, which also tended to maintain a fixed laminar position regardless of soma depth (V1: $r^2 = 0.67$, $p = 0.0039$; anterior thalamus: $r^2 = 0.42$, $p = 0.007$; F-test, Figure 4.15d,e).

The more specific perisomatic targeting of OFC inputs compared to V2m inputs is also apparent from a more direct comparison of the mean soma-aligned responses, which were significantly closer to the soma for OFC inputs ($\bar{\mu}_{\mathrm{OFC}} = 9 \pm 11\,\mu\mathrm{m}$, $\bar{\mu}_{\mathrm{V2m}} = 60 \pm 10\,\mu\mathrm{m}$, $p = 0.0017$, two-sample t-test, Figure 4.16).

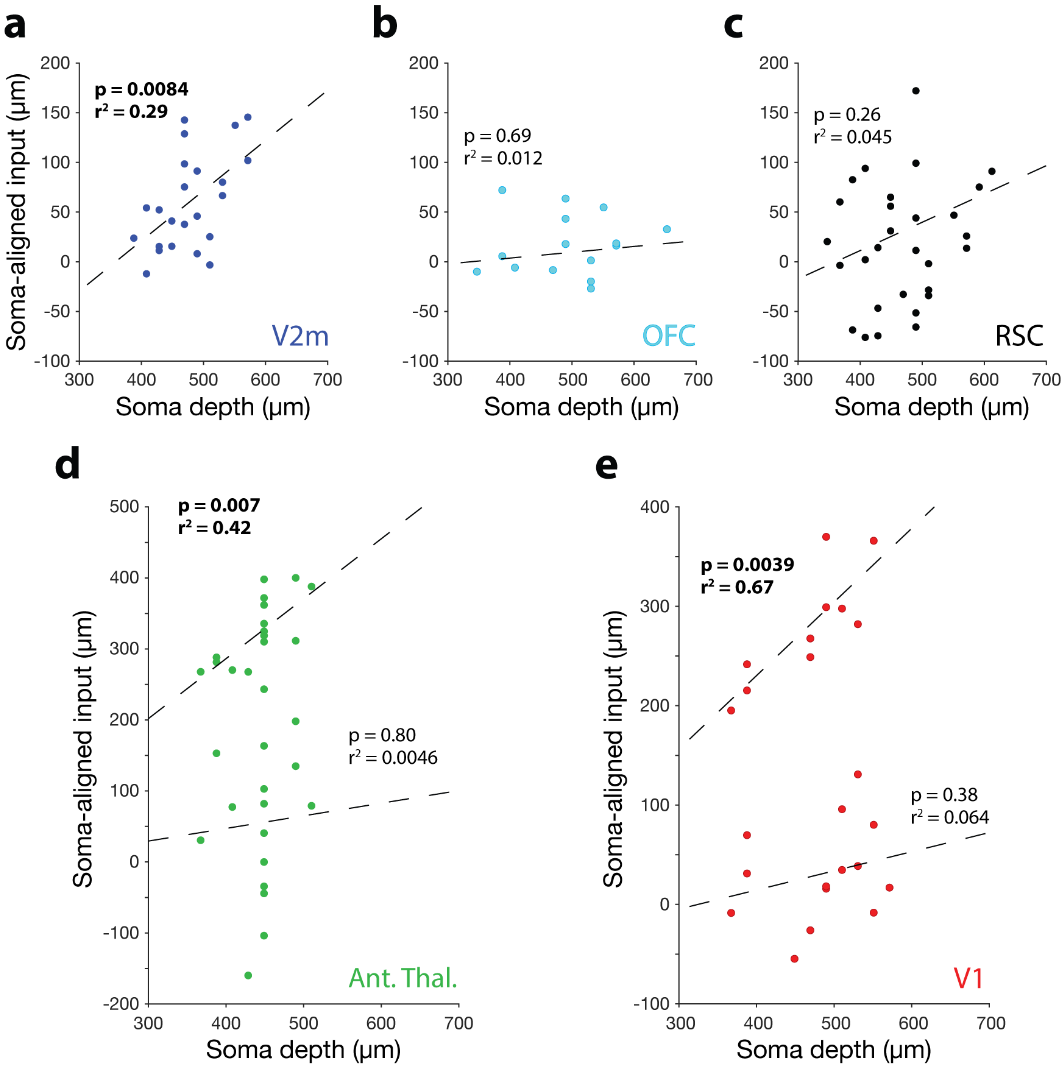

Figure 4.15: Location of inputs relative to laminar position of neurons. **a.** Scatter plot of input distance relative to the soma against somatic depth relative to the pia. V2m inputs had a significant positive correlation, indicating that the inputs to deeper neurons targeted more superficial dendrites. **b.** Same as in *a* for RSC inputs. **c.** Same as in *a* for OFC inputs. **d.** Same as in *a*, but with the correlation calculated separately for superficial and deep V1 inputs.

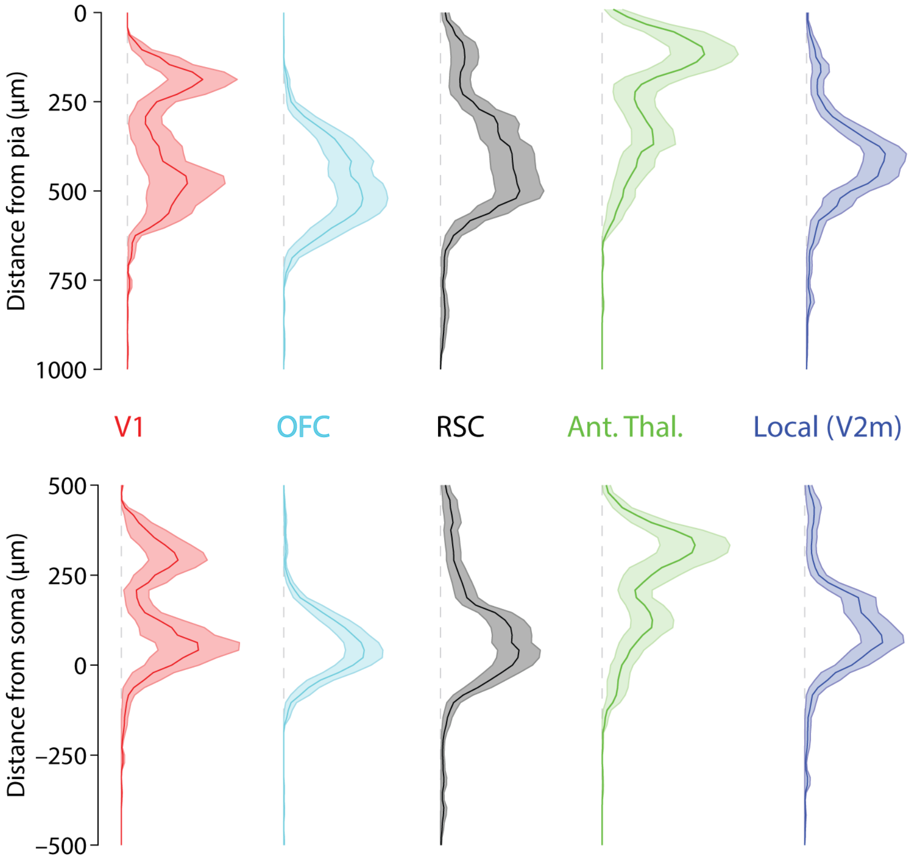

Figure 4.16: Direct comparison of average heatmap profiles from each injected region. a. Pia-aligned heatmap profiles. **b.** Soma-aligned heatmap profiles. Profiles are computed from the sum of rows in the average heatmap for one region. Errors are likewise computed as the sum of rows for the equivalent heatmap with each pixel encoding the SEM of responses at that location across cells.

4.3.7 Axonal distributions of presynaptic pathways to Glt neurons

Finally, an important additional consideration is the degree of overlap between the axons from any given region and the related sCRACM maps. Indeed, while the presence of axons is a prerequisite for there to be synapses in a region, the two do not necessarily have a perfect correspondence, since axons can pass through a region without synapsing or can make synapses with only a subset of cell types. Moreover, the axon distributions have historically been used as a basis for estimating hierarchical relationships, whereas relatively little data exists regarding the functional connections of these pathways.

The axons from each of the long-range inputs were thus imaged using a confocal microscope and normalised axon fluorescence was measured across different cortical layers (Figure 4.17). The V1 and RSC pathways exhibited axonal branching patterns consistent with classical feedforward and feedback pathways, with V1 projecting mostly in L2/3 and to a lesser extent in other layers except L1, while RSC axons showed the strongest fluorescence in L1 and L5. These axonal patterns, however, differed remarkably from the sCRACM input maps, which displayed a sharper preference for specific subsets of these axons. Despite the strong fluorescence in L1, RSC synapses thus seem biased towards the deeper layer axons, while synapses from V1 show a bimodal distribution that is not present in the axonal branching.

Perhaps even more surprising was the OFC pathway, which also branched extensively between L2/3 and L5 while avoiding L1. Furthermore, the synapses mapped with sCRACM were clustered in a perisomatic region that was considerably narrower than the broad area of axon fluorescence.

The anterior thalamus showed similar axonal branching as the RSC input, with strong fluorescence in L1 and a smaller amount in other layers. The corresponding sCRACM currents displayed a similar distribution, with strong input to the apical tuft of Glt neurons, but with notably stronger currents between L2/3 and L5 than would be expected based on the axonal pattern.

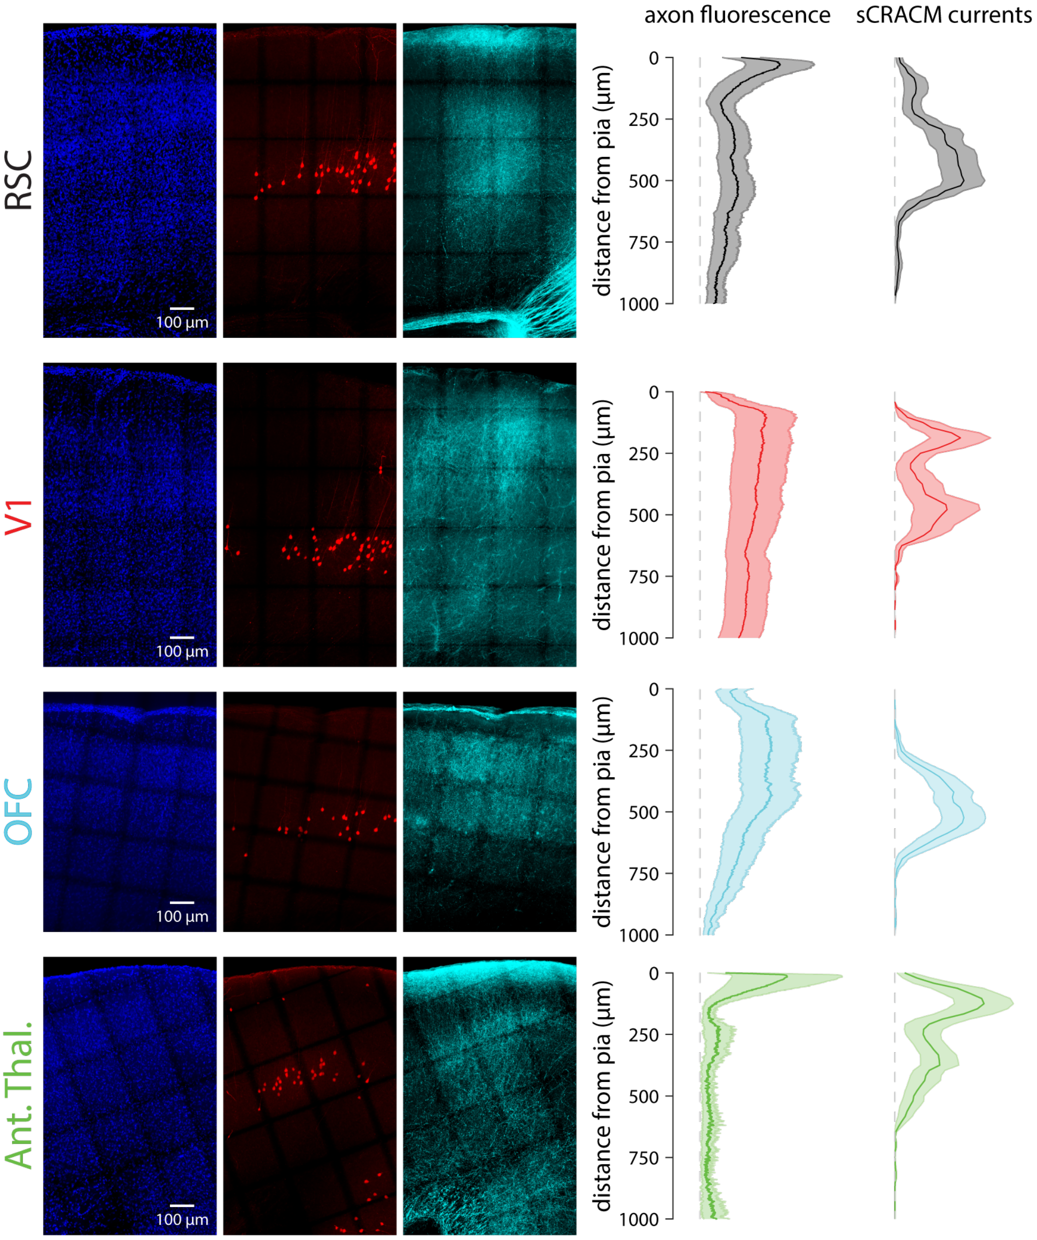

Figure 4.17: Direct comparison of average heatmap profiles from each injected region. *Left:* Confocal images of the region of V2m containing Glt neurons, showing a DAPI stain (blue), Glt neurons expressing tdTomato (red), and the axons expressing Chronos-GFP (cyan). *Right:* Comparison of the normalised GFP fluorescence intensity across different cortical layers with the corresponding sCRACM input distributions. Axon fluorescence was measured in rectangular areas approximately 300 µm wide spanning the full depth of cortex, and normalised to the minimum and maximum values within this region. For each injection, the mean and SEM were calculated based on sections from 3 different mice.

4.4 Discussion

One of the fundamental operational features of the cortex is the high level of recurrent connectivity. The precise structural architecture of this connectivity between and within brain regions is likely to be a key component of the computational power of the brain. Here I have described the subcellular connectivity of several local and long-range projections to Glt neurons in V2m. The locations of these upstream projection neurons were determined based on monosynaptic retrograde rabies labelling targeted to the Glt population. These locations were then used as the basis for injection of optogenetic activators that allowed nerve terminals to be stimulated in a spatially precise way along the dendritic trees of Glt neurons. The subcellular optogenetic input mapping has revealed several unexpected properties of these connections.

The largest identified presynaptic population came from the retrosplenial cortex. This area is involved in a wide variety of functions and behaviours that are associated with memory, imagination, spatial navigation, and planning future actions (Vann et al., 2009). Indeed it is a integral part of the Default Mode Network—a brain-wide network of highly interconnected areas that are active during resting wakefulness and are thought to be at the core of an animal's internal model of the world (Raichle et al., 2001; Greicius et al., 2003; Fox et al., 2005). Consistent with this high level of cognitive functions, the axons from RSC to V2m were found to be consistent with the classical top-down connectivity patterns, with strong targeting of L1, that are generally thought to have a modulatory influence on lower-level neurons (Felleman and Van Essen, 1991). On the basis of this type of axon distribution, studies of ttL5 neurons often implicityl assume that top-down pathways make functional connections with the apical tuft dendrites to drive non-linear dendritic operations (Larkum, 2013; Manita et al., 2015; Takahashi et al., 2016; Richards and Lillicrap, 2019). My results, however, indicate that the majority of Glt neurons receive RSC input directly to the basal dendrites, with relatively weak connections with the apical dendrite. This implies that the axons in L1 must be predominantly targeting other cells types, such as the apical dendrites of L2/3 pyramidal neurons, or perhaps the small population of interneurons that reside in L1 (Markram et al., 2004).

The second largest group of neurons was located in the primary visual cortex. As the first cortical area to receive visual information, V1 has been the focus of most research on vision and it is universally recognised as the lowest area in the hierarchy of visual processing. Indeed, the original justification for estimating the relative hierarchical order of cortical areas based on laminar connectivity patterns relies on the axiomatic assumption that V1 is the lowest region of the hierarchy (Felleman and Van Essen, 1991; Hubel and Wiesel, 1998). The foundation for this claim is the anatomical connectivity of V1, which is the primary recipient of visual signals originating in the retina and relayed through the LGN, as well as the *in vivo* functional properties of V1 neurons, which have receptive fields that are more complex than those in the LGN, yet simpler than those in other cortical areas. The connectivity from V1 to V2 is thus by definition the canonical example of a feedforward connection. In primates, this connectivity is known to project broadly from V1 to middle layers of V2, while the reciprocal connectivity strongly targets L1 and L5 (Rockland and Pandya, 1979). As this connectivity principle consistently appeared between other visual areas as well, it has become common to use connectivity as a proxy for judging hierarchy (Harris et al., 2019).

The axonal projections from V1 to V2m observed in my experiments were consistent with classical connectivity rules, projecting primarily to middle and superficial layers while avoiding L1. The functional synaptic input to Glt neurons in V2m also covered a similar region, but was substantially more localised to two narrow region in the apical and basal domains, with relatively weak input in between. This may be a reflection of the dendrites of ttL5 neurons, which have a higher density of dendrites close to the soma and in the upper region of the apical dendrite, with a smaller number of oblique dendrites along the trunk of the apical dendrite.

When comparing the sCRAM distributions across individual cells, one interesting observation is the high variability between cells. The connection from V1, in particular, often targeted either the apical or the basal dendrites, with only a minority of cells receiving input to both domains. This may be due to the presence of distinct subpopulations within either the presynaptic neurons in V1 or the Glt neurons in V2m, with different functional roles in visual processing. The inputs to the basal compartment could, for instance, be driving a more conventional form of visual receptive field response, while the apical inputs could be carrying more diverse contextual information regarding features outside the primary receptive field of the cell. Indeed, there is an increasing body of evidence highlighting that even V1 responses can be highly multimodal, encoding a variety of mixed sensory and motor signals that could be relevant for accurate visual perception (Stringer et al., 2019), although the extent to which these multimodal responses get passed forward to higher visual areas is still unknown.

To verify if different V1 neurons preferentially target apical or basal compartments, it would be ideal to have retrograde tracers that could be targeted to specific dendrites within genetically defined neurons. However, some knowledge could also be gleaned by making small injections of a non-specific retrograde tracer, such as cholera toxin subunit B (CTB), in deep and superficial layers of V2m and assess the extent of overlap between these projection-specific populations.

When considering the inputs from V1, it should be noted that the retrograde rabies tracing experiments primarily labelled neurons in L5 of V1. This pathway may, however, not be the primary route for visual receptive field information to reach the ttL5 neurons in V2m. Indeed, in addition to these direct long-range projections from L5, visual information can reach the Glt neurons in V2m indirectly through local visually responsive neurons in L2/3 , which also receive visual signals from V1 and are likely encoding a different set of visual signals than the L5 neurons in V1 (Vries et al., 2018).

In the experiments presented here, sCRACM inputs from a superficial injection of Chronos in V2m primarily targeted the perisomatic compartment, consistent with a more conventional feedforward driving sensory input to the Glt neurons. It should be noted, however, that these injections were not restricted to any specific cell type, and are likely to have labelled more than just the L2/3 pyramidal neurons, including potentially some deeper-layer neurons, thus limiting the extent of claims that can be made in this regard. While a rigorous quantification of the number of cells labelled in each layer was not possible due to the extensive background fluorescence in the neuropil, the result is largely consistent with the findings by (Petreanu et al., 2009) on the inputs to L5 neurons in mouse somatosensory cortex, in which presynaptic labelling was successfully restricted to L2/3 neurons.

One complicating factor in the experiments mapping local inputs is the need to avoid expressing Chronos directly in the Glt neurons I was recording from, as synaptic current would be confounded by direct Chronos-mediated currents in the dendrites. While this challenge was almost entirely overcome by using a Cre-OFF virus, a small number of double-labelled neurons (expressing both Cre and Chronos) were observed in the control injections. However, even if I were to record from such neurons, theses cells would be easily detectable by observing the latency of postsynaptic currents following the onset of the light stimulus, as neurons directly expressing Chronos would exhibit current almost instantaneously. Furthermore, a neuron directly expressing Chronos would likely have strong light-evoked currents throughout the entire dendritic tree. In all the neurons I recorded, the light-evoked responses were spatially restricted and had latencies of several milliseconds from stimulus onset, indicating that these were indeed postsynaptic responses.

Another potential source of sensory information to the Glt neurons comes through the inputs from the anterior thalamus. Among thalamic nuclei, a distinction is often made between primarily sensory “feedforward” nuclei, which process and relay the signals from sensory organs to primary sensory regions in cortex, and higher-order nuclei, which do not communicate directly with sensory organs and are instead involved in communication between different cortical areas. While the cluster of neurons in the anterior thalamus receives a large amount of sensory signals from the vestibular brainstem nucleus via the lateral mamillary body (Rancz et al., 2015), it has also been associated with a variety of spatial and memory functions (Jankowski et al., 2013), making it potentially somewhat ambiguous in this respect.

The projections from the anterior thalamus observed in my experiments strongly resembled those from RSC, with the axons branching widely in L1 and to a lesser extent in deeper layers. On the basis of connectivity, this input thus appears to be most consistent with a higher-order thalamic nucleus providing vestibular-related spatial information to the Glt neurons in V2m. Unlike the connection from RSC, this axon distribution closely matched the distribution of functional synapses seen in the sCRACM recordings, which also showed strong currents in the apical tuft.

One possible source of error in the mapping of the anterior thalamus connections is expression outside the main injection site. As the brain region is located deep in the brain, during injection a small amount of Chronos expression was visible around the injection tract, which passed through the most anterior portion of RSC. However, this most likely did not influence the main result, as in the rabies mapping very few presynaptic neurons were observed in this region of RSC. Furthermore, the distinctly different pattern of functional input maps between the anterior thalamus and RSC inputs implies that there is likely little to no overlap between the two presynaptic populations.

The final input population to be mapped in these experiments was the cluster of presynaptic neurons in the ventral OFC. This area exhibited perhaps the most noteworthy connectivity pattern of the areas examined here. Despite the fact that various forms of sensory input, including visual stimuli, can cause a change in the activity of OFC neurons, compared to narrow stimulus selectivity of neurons in the visual cortex these responses appear to be far more complex and invariant to the fine-scale details of the stimulus, responding instead to more abstract features, such as aesthetic properties and memory-related predictions of sensory signals. Thus, although this area was not originally described as part of the visual hierarchy in primates (Felleman and Van Essen, 1991), it is commonly regarded as an unambiguous source of top-down feedback to sensory areas (Cela-Conde et al., 2004; Rolls, 2004; Gilbert and Sigman, 2007; Ishizu and Zeki, 2011; Zanto et al., 2011; Chaumon et al., 2014).

It was thus particularly surprising that the axons from OFC were found to more closely resemble canonical feedforward pathways, like the connections from V1, projecting broadly throughout the middle and upper layers of cortex while avoiding L1. Moreover, of all the connections mapped with sCRACM, the OFC most consistently targeted a narrow perisomatic area, with almost no input to the apical and tuft dendrites. This is in contrast with V1 axons, which despite the similar physical spread, had a distinctly different pattern of functional connections.

Among the various dendritic targets of all these feedforward and feedback pathways, one interesting observation can be made by looking at the correlation between location of the soma and the functional inputs. A significant positive correlation indicates that neurons with somas on the deeper side of L5 receive inputs further up the apical dendrite than more superficial neurons. This is of significance because it implies that these inputs might be synapsing with all dendrites within a particular cortical layer, regardless of the specific identity and location of the neuron. This can be seen in the inputs from superficial V2m neurons, as well as in the apical portion of the input from V1 and the anterior thalamus. In contrast to this layer-specific input, the projection from OFC connection stood out for its remarkable perisomatic specificity, closely targeting the basal dendrites of individual Glt neurons independently of the somatic depth and in spite of the broad distribution of the OFC axons across several cortical layers.

Finally, When evaluating the results presented in this chapter it is worth keeping in mind a few general methodological limitations that could influence the observed responses. Regarding the stimulation, although the size of each light spot was set to be approximately 42×42 µm, the actual resolution of sCRACM maps, i.e. the smallest distance between synapses that can be detected, is likely to be larger. This is partially due to light scattering in the tissue, which increases the area that is illuminated by each stimulus. However, there is also a more fundamental limitation, which is the spread of voltage in the stimulated axons. While the presence of TTX in the extracellular solution prevents active propagation, the voltage still propagates passively beyond the stimulated area in proportion to the length constant of the axon. Because of this limitation, the effective minimum stimulation area limiting the sCRACM resolution has previously been estimated to be around 60 µm (Petreanu et al., 2009). Stimulus spots smaller than this, such as the 42×42 µm spots I have used in these experiments, are thus oversampling the input distributions, making the true resolution lower than it appears from the number of pixels in each heatmap. Consequently, while the recorded heatmaps approximately matched the dendritic span, some responses were also observed in adjacent spots that did not contain dendrites. Furthermore, as stimulation was not restricted in depth, each heatmap should be viewed as a projection of inputs throughout the thickness of the slice. Thus somatic responses, which often appear to be strong, are likely caused by synapses on dendrites that are perpendicular to the slice, as the soma itself is not known to directly receive excitatory connections.

Another limitation arises from the method of recording synaptic currents through a somatic patch-clamp electrode. As the recordings were all performed in voltage-clamp configuration, there is likely to be a substantial space-clamp error, wherein the synaptic currents are able to change the voltage at the synapses and thus reduce the driving force of ions flowing through them (Spruston et al., 1993; Bar-Yehuda and Korngreen, 2008). The severity of this is known to increase as a function of distance, to the point where distal tuft dendrites are likely to be effectively unclamped. To avoid introducing additional biases, no correction was applied to compensate for these errors. Therefore, any observed responses are likely to underestimate the true synaptic currents in proportion to their distance from the soma.

This limitation could potentially be compounded if the total synaptic currents from one pathway were substantially smaller than from other regions, as overall weaker synapses could result in only being able to measure the synaptic currents closest to the soma while more distant synapses fall below the detectability threshold. Meanwhile, stronger pathways, while still biased in favour of perisomatic inputs, would at least accurately represent the presence of some input to the apical dendritic domain. The simplest way to overcome this limitation would be to adjust the laser power used when stimulating each recorded cell so as to obtain constant current amplitudes, as was done by Petreanu et al. (2009), thus applying the same bias to all input distributions and facilitating the direct comparison.

One difficulty of implementing this solution is that a peak current amplitude has to be chosen *a priori* and set as a target for all future recordings, of which the synaptic strength is unknown. If this value is chosen too high, many cells for which the input is not strong enough to reach this arbitrary threshold but which still have informative patterns of input distribution might be unnecessarily discarded. For example, the bias in synapse detectability would not be an issue in cells receiving detectable input to the apical tuft, as is these inputs are detected other synapses closer to the soma would also presumably be seen. It would thus not be sensible to discard this data, regardless of the overall amplitude of the measured currents.

On the practical side, varying the laser power for each recorded cell also presents technical challenges when implementing the sCRACM stimulation grid with a digital micromirror device, as I have done in these experiments, that are not present when controlling a laser beam with mirror galvanometers. Specifically, DMDs split the incoming laser beam across the full field of the micromirrors, resulting in substantially lower irradiance in each stimulus spot. To stimulate with enough intensity to consistently record postsynaptic responses, it is thus necessary to maintain the laser powers close to the maximum available output, which severely limits the range in which this can be varied to obtain desired currents in the recorded neurons.

For these reasons, in the experiments presented here I have adopted the alternative approach of using constant laser power and allowing different cells to exhibit different current amplitudes, as was also done by Yamawaki et al. (2019). One potential advantage of this approach is that it enables a comparison of the total currents from different input region, although these should be interpreted with caution as they are likely to be sensitive to the size of the virus injection (particularly for large and distributed areas like V1). However, this comparison still constitutes an important control in order to estimate the reliability of synapse detection for different pathways. For example, if inputs from one pathway were only detected near the soma and the total current amplitudes were small, this input map would have to be deemed unreliable and unable to exclude the possibility of substantial undetected input to the apical tuft as well.

By this reasoning, the sCRACM distributions I recorded when stimulating V1 and the anterior thalamus are not a cause for concern, as input to the apical dendrite was observed despite any electrotonic filtering, indicating that the underestimate was not severe enough to entirely mask these distal inputs. On the other hand, the predominantly perisomatic heatmaps recorded for OFC, RSC, and local V2m inputs could potentially be missing synapses in the apical tuft that were to weak to be detected. However, the total current recorded during full-field stimulation of these pathways was substantially larger than for either V1 or anterior thalamus. This implies that the input maps for these regions can also be considered reliable, as even if there were a small amount of tuft input the overall distribution would still be heavily skewed towards the perisomatic region.

Chapter 5

Synaptic physiology

5.1 Introduction

In addition to the intrinsic properties of a neuron and the location of synaptic inputs, the important factors shaping how the activity of upstream neurons is translated into potentials in the postsynaptic cell also include the properties of presynaptic nerve terminals. When an action potential invades the synaptic bouton, an influx of Ca^{2+} ions through voltage-gated Ca^{2+} channels causes membrane-bound vesicles containing neurotransmitters to fuse with the presynaptic membrane and release their content into the synaptic cleft (Südhof, 2012). Aside from the type, configuration, and density of neurotransmitter receptors on the postsynaptic membrane, the quantity and time course of presynaptic neurotransmitter release ultimately determine the total transmembrane current in the postsynaptic cell.

Neurotransmitter release itself is a complex and stochastic process that is controlled by a large number of interrelated factors, which in combination determine the release probability in response to a presynaptic action potential. One such factor is the presynaptic Ca^{2+} concentration, which is the initial trigger that activates the release machinery. A higher Ca^{2+} concentration thus generally results in a larger number of vesicles being released and in a larger postsynaptic potential. If multiple spikes occur in brief succession, presynaptic Ca^{2+} can accumulate faster than it is cleared away, resulting in postsynaptic potentials that progressively increase in amplitude with each spike. This type of short-term plasticity is known as paired-pulse facilitation (Jackman and Regehr, 2017).

However, a prerequisite for Ca^{2+} to trigger neurotransmitter release is that there must already be vesicles of neurotransmitters that are ready to be released. As more vesicles are released, the readily releasable pool decreases and needs time to replenish. If vesicles are released faster than they are replenished, the postsynaptic potentials gradually decrease in amplitude as fewer vesicles are released with each successive spike, regardless of the presynaptic Ca^{2+} concentration. This is referred to as short-term depression.

In synapses which inherently have low initial release probability (due to factors such as the volume of the synaptic bouton and the density of presynaptic Ca^{2+} channels), vesicle depletion is generally not a limiting factor and the synapse displays facilitation. Conversely, synapses with high release probability are more likely to deplete their vesicle pool and display short-term depression.

Synapses exhibiting depression and facilitation are sometimes also referred to as type 1 “driver” and type 2 “modulator” synapses, respectively (Sherman and Guillery, 1998, 2011; Viaene et al., 2011; Mo et al., 2017). Type 1 synapses are frequently seen in the feedforward pathway between sensory thalamic nuclei and cortical neurons, whereas type 2 synapses are believed to be more common in pathways between higher-order thalamic nuclei and cortical regions, and are thought to be modulatory because they often activate both ionotropic and metabotropic postsynaptic receptors, resulting in a slower and more complex effect on the postsynaptic neuron.

It is not entirely clear what role short-term plasticity has in system-level neural computations, but one way in which these mechanisms can be leveraged is through spike multiplexing (Naud and Sprekeler, 2018). Individual neurons can simultaneously send their axons to multiple different across distant cortical regions (Han et al., 2018) and can exhibit a variety of spiking patterns, including high-frequency bursts (Lisman, 1997). Although axons transmit the same number of action potentials to all downstream targets, if the short-term plasticity varies across their synapses it may be possible for a neuron to multiplex the signal and simultaneously communicate distinct messages (i.e. have different postsynaptic effects) to different sets of downstream neurons (Naud and Sprekeler, 2018). Specifically, a burst of spikes would be most effectively transmitted across a synapse with short-term facilitation, whereas the additional spikes in the burst would have little effect at a synapse with short-term depression.

On the postsynaptic side, depending on the exact sequence, timing, and spatial pattern of inputs, postsynaptic potentials can also sum in various non-linear ways. For example, some dendrites have been shown to be sensitive to the direction in which synapses are sequentially stimulated relative to the location of the soma (Branco et al., 2010; Briggman et al., 2011). This is dues to a variety of active conductances in dendrites, such as NMDA receptors, which can amplify postsynaptic potentials in a non-linear way depending on voltage and the concentration of a variety of internal signalling molecules.

Given the complex way in which different pathways can be integrated, depending on a wide range of both pre- and postsynaptic parameters, in this chapter I will attempt to address the fourth aim presented in chapter 1 and describe how the inputs that I mapped anatomically in chapter 4 interact functionally. This was done using a combination of electrical and optogenetic stimulation strategies to activate feedforward and feedback pathways both individually and in various temporal combinations.

5.2 Chapter methods

5.2.1 Stereotaxic surgeries

For the optogenetic experiments, the stereotaxic surgeries required injecting two different viruses (encoding the optogenetic activators Chronos and ChrimsonR) in separate brain regions. In order to maximise the chances that any one Glt neuron would simultaneously receive synapses expressing both opsins, I targeted the injections in both RSC and V1 to cover as wide a region as possible. Since RSC is a long and narrow region running straight alongside the midline, two injections were made in this region, one on the more caudal portion of RSC and the second approximately 300–500 µm anterior to this.

As V1 is a substantially larger region, three injections were made here. The shape of V1 when seen from above is vaguely triangular, being wide on the most caudal side and more narrow on the rostral side. The three injections were this also arranged in a roughly triangular pattern with two caudal injections about 1 mm apart (in the mediolateral direction) and one rostral injection about 500 µm anterior to this at an intermediate mediolateral position. After making all the injection for one brain, the glass microinjection pipette was discarded and a new one was loaded and filled with the second virus. There was thus no cross-contamination between the injection sites containing different viruses. Which virus I injected in V1 or RSC was alternated between different mice, to enable me to distinguish opsin-specific effects from the properties of the two pathways.

5.2.2 Cell selection & recording conditions

The recordings in this chapter were aimed at understanding the functional interactions between synapses under physiological conditions. Whole-cell patch clamp recordings from Glt neurons in V2m were thus made at 35 °C in ACSF containing 1.5 mM $CaCl_2$, and no drugs were added to the solution, so as to maintain both axonal conductance and synaptic transmission as close to physiological as possible. As in previous chapters, tdTomato fluorescence in

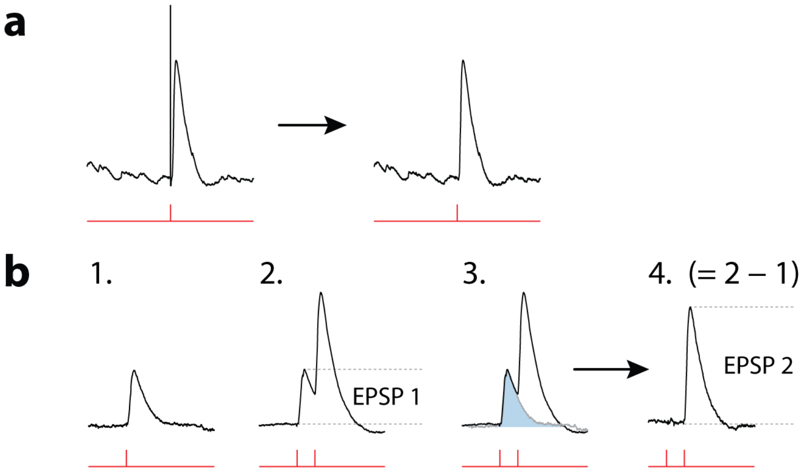

Figure 5.1: Processing and analysis of voltage traces. a. Example voltage trace before (left) and after (right) the removal of the stimulus artefact. **b.** Illustration of processing step to measure the first and second EPSPs in a voltage trace. (1.) The waveform of the first EPSP was initially recorded with a single stimulus pulse. (2.) During paired-pulse stimulation, the amplitude of the first EPSP was still measurable as the difference between peak an baseline. (3.) To accurately obtain the same measure for the second EPSP, the first EPSP waveform was subtracted from the trace. (4.) The amplitude of the second EPSP was then measured in this processed voltage trace.

the Glt neurons was visualized with short pulses of 550 nm light through the high magnification (40×/0.8NA) water immersion objective. Both the light intensity and duration were kept to a minimum in order to avoid activating or photobleaching the optogenetic activators expressed in presynaptic axons.

In electrical stimulation experiments, the recordings were done using 6–7 week old Glt mice. Since the optogenetic experiments required stereotaxic surgeries and time for the virus to express, in these experiments animals were a few weeks older, with injections occurring at an age of 4 weeks and recordings made when the animals were 8–9 weeks old.

5.2.3 Electrical stimulation

After positioning the slice within the recording chamber, two monopolar glass electrodes mounted to manipulators on either side of chamber were lowered under visual guidance through a 10x/0.25NA dry objective. The electrodes were placed approximately into the middle layers of V1 and RSC, which are located on either side of the Glt neurons in V2m. With the electrodes in place, neurons were targeted for recording under the high magnification (40×/0.8NA) water immersion objective, and were stimulated with short (0.1–1 ms) current pulses delivered through the electrodes to the presynaptic neurons in V1 and RSC.

During electrical stimulation, large stimulation artifacts occur in the recorded voltage trace. For visualization purposes, these stimulus artefacts were manually removed from the example traces shown throughout this chapter (Figure 5.1a). Removal of the stimulus artefact was done by deleting the voltage values around the stimulus time, in which the artefact is clearly distinguishable as a sharp voltage spike of very short duration. Following this deletion, the voltage trace was linearly interpolated between the voltage values on either side of the deleted stimulus artefact.

5.3 Results

5.3.1 Electrical stimulation

Having obtained a general description of the intrinsic physiology of Glt25d2-Cre neurons in V2m and mapped the subcellular dendritic targets of their inputs, I performed a series of experiments to determine how these inputs functionally interact to drive output in Glt neurons. To do this I used electrical stimulation to characterize the short-term plasticity intrinsic to the presynaptic nerve terminals and determine how summation is influenced by stimulation intervals between different pathways.

Since electrical stimulation is inherently non-selective to the cells being stimulated, the only way to restrict stimulation to a somewhat defined presynaptic population is to place the electrode in this region and stimulate the cell bodies directly. For the signal to then reach the postsynaptic neuron I was be recording from, it was necessary that the axons between the two still be intact. In a brain slice, the probability of this occurring decreases with distance between the stimulated and recorded regions. My initial experiments therefore focused on the inputs from V1 and RSC, which aside from being the two largest inputs identified in the retrograde rabies labelling (Figure 4.3) are also anatomically adjacent to V2m (Figure 5.2a). This proximity enables all three regions to be simultaneously present in the same coronal slice, with many of the axons connecting them still intact in slices of sufficient thickness.

The first experiments were designed to study the synaptic physiology of the connections from RSC and V1 and probe their functional interactions by electrically stimulating the inputs, both individually and in different temporal combinations, while performing whole-cell recordings from Glt neurons in V2m.

Paired-pulse facilitation of individual inputs

Each of the pathways was first stimulated individually using single electrical pulses to obtain an average waveform for excitatory postsynaptic potential

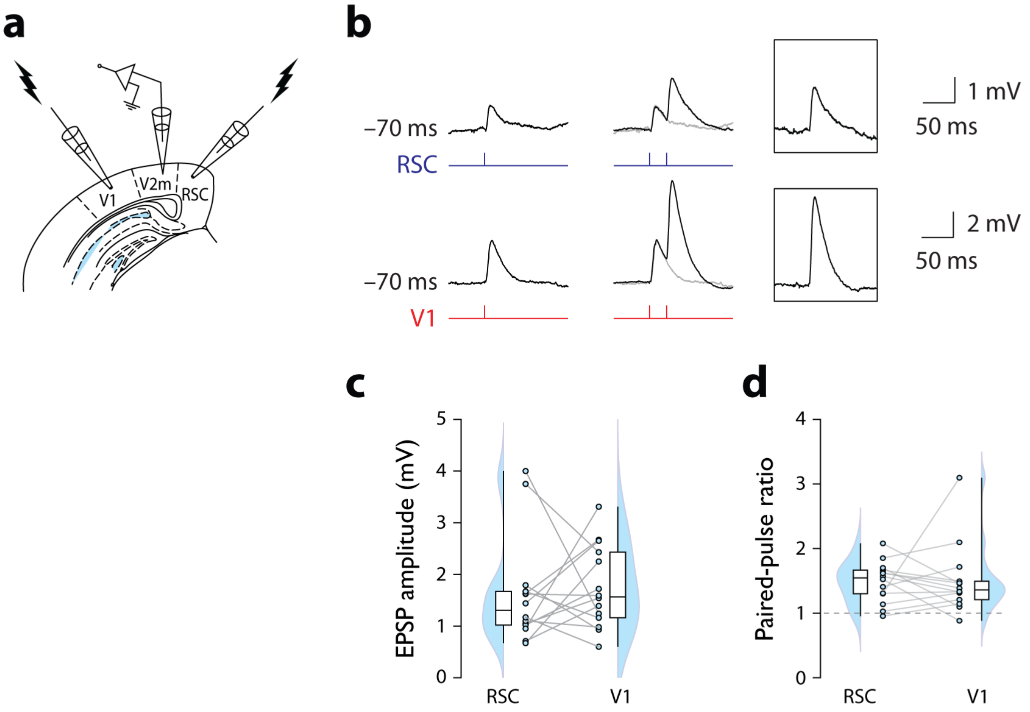

Figure 5.2: Electrical paired-pulse stimulation of RSC and V1. a. Diagram illustrating position of recording and stimulating electrodes. **b.** Example traces of single pulse and 40 Hz paired-pulse stimulation of RSC (top) and V1 (bottom). The gray traces show the EPSP waveforms from single pulse stimulation, which were subtracted from the paired-pulse trace. The boxed trace on the right displays the measured second EPSP after this subtraction. **c.** Summary data of EPSP amplitudes from single pulse stimulation of RSC and V1. **d.** Summary data of paired-pulse ratios for RSC and V1, computed as the ratio of the second and first EPSP amplitudes.

way, a full exploration of all inter-stimulus intervals was not deemed necessary here.

Stimulation current was set so as to obtain similar EPSP amplitudes when stimulating RSC and V1. The amplitude for the first EPSP, measured as the difference from the baseline voltage to the peak of the EPSP (Figure 5.1b), was 1.61 ± 0.27 mV for RSC stimulation and 1.75 ± 0.21 mV for V1 stimulation ($n = 14$, $p = 0.7$, paired-sample t-test, Figure 5.2c).

Since EPSPs evoked at such high frequency are subject to summation, with the second EPSP rising during the passive decaying phase of the first EPSP, when measuring the second EPSP a correction was applied in order to estimate its amplitude independently of the first EPSP (Figure 5.1b). To do this, the EPSP waveform obtained from single-pulse stimulation was scaled and subtracted from the first EPSP of the paired stimulation trace. Compared to simply measuring the second EPSP as the difference between the start and peak voltage of the EPSP, the result of this transformation is that the corrected amplitude of the second EPSP is effectively equal to the difference in voltage between the second EPSP peak and the expected voltage of the first EPSP at the same point in time (i.e. what the voltage would be at that time point if there had not been a second EPSP). This corrected value is thus somewhat larger than the uncorrected EPSP.

After subtracting the first EPSP from the voltage trace and measuring the second EPSP in this processed voltage trace, the paired-pulse ratio was measured as the ratio between the second and first EPSP amplitudes, with ratios above 1 indicating facilitation and ratios below 1 indicating short-term depression. Evoked EPSPs displayed paired-pulse facilitation for both V1 and RSC stimulation, with paired-pulse ratios of 1.48 ± 0.08 for RSC and 1.5 ± 0.14 for V1 ($n = 14$, $p = 0.91$, paired-sample t-test, Figure 5.2d).

Independent summation of RSC and V1 inputs

To study interactions between the two pathways, they were then stimulated in succession while varying the interval between the stimuli (Δt) from −75 ms to +75 ms in 5 ms steps (Figure 5.3a). Measured EPSP amplitudes were corrected in the same way as in Figure 5.1 and Figure 5.2 by fitting and subtracting the single pulse EPSP waveform of the first EPSP before measuring the second EPSP, to compensate for passive effects during the decaying phase of the first EPSP. The EPSPs were then normalised to the single pulse EPSP amplitude for each pathway and are plotted in Figure 5.3b. Across all recorded intervals, individual EPSP amplitudes were constant and unaffected by the presence of a preceding stimulus from the other pathway ($n = 9$; RSC: $p = 0.16$; V1: $p = 0.73$; repeated measures ANOVA test; Figure 5.3b).

At very short inter-stimulus intervals it was not possible to measure the individual EPSPs due to the EPSP peaks overlapping or because the second stimulus artefact obscured the first EPSP peak. The effect of combined stim-

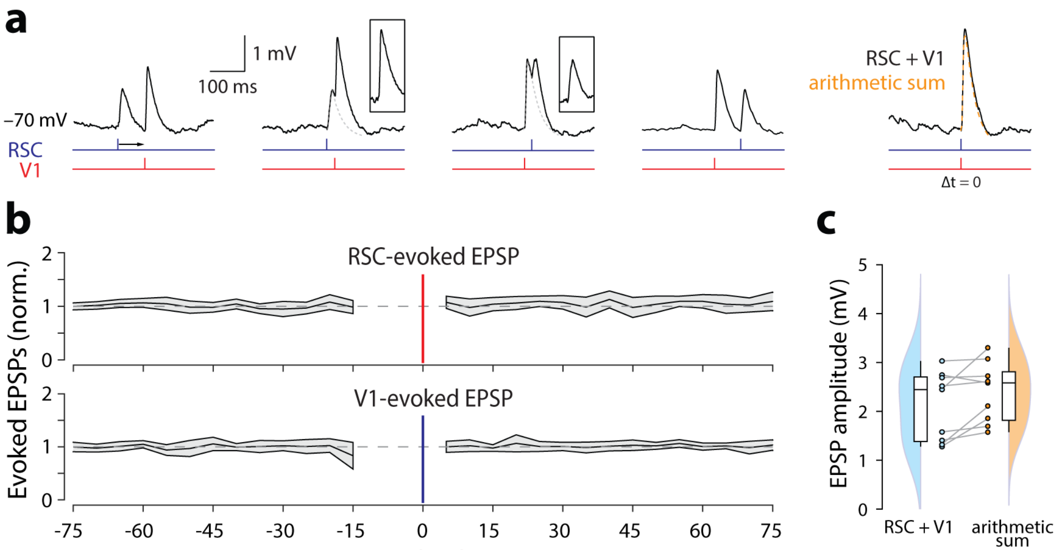

Figure 5.3: Independent summation of RSC and V1 inputs. a. Example traces from one cell showing 5 different inter-stimulus intervals (with the RSC stimulus occurring at $\Delta t = -75$ ms, -20 ms, $+20$ ms, $+75$ ms, or 0 ms relative to the V1 stimulus). The grey traces show fitted EPSP waveforms derived from single pulse stimulation, which were subtracted before measuring the second EPSP amplitude. The arithmetic sum was derived from the sum of independent EPSP waveforms for RSC and V1. **b.** Summary data of RSC-evoked (top) and V1-evoked (bottom) EPSP amplitudes at all time intervals. The vertical line at 0 ms indicates the position of the V1 stimulus (top) or the RSC stimulus (bottom) relative to the measured EPSP. Negative Δt values indicate that the measured EPSP precedes the other stimulus, with 0 ms indicating simultaneous stimulation. All amplitudes are normalised to the single pulse EPSP amplitude for each pathway. Shaded areas show standard deviation of responses. **c.** Summary data of total EPSP amplitude during simultaneous RSC and V1 stimulation ($\Delta t = 0$ ms) compared to the calculated linear sum based on summation of the independent waveforms.

ulation at the 0 ms time point is thus shown in Figure 5.3c. When stimulating the pathways simultaneously, EPSP summation was slightly sublinear. The average combined EPSP amplitude was 2.1 ± 0.2 mV, while the calculated linear sum of the independent EPSPs was 2.4 ± 0.2 mV, with the difference being on the boundary of statistical significance (n = 9, p = 0.05, paired-sample t-test).

5.3.2 Full-field optogenetic stimulation

The projections to V2m from RSC and V1 traverse a relatively short distance across the cortex, and enough axons are thus spared after slicing that EPSPs can be elicited in at least some V2m neurons. However, Glt neurons also receive inputs from many distal brain regions, such as OFC and thalamic nuclei (Figure 4.3), for which electrical stimulation would not work. To be able to study the functional interactions of these inputs as well, I attempted an alternative approach using optogenetics to directly stimulate the nerve terminals of these regions with light pulses. In particular, I implemented an approach described by Klapoetke et al. (2014) using two different optogenetic activators with blue- and red-shifted wavelength sensitivities—Chronos and ChrimsonR—to independently excite distinct sets of inputs.

Before studying the properties of distal inputs, this method was tested on the RSC and V1 inputs by injecting in each region AAVs containing either Chronos-GFP or ChrimsonR-tdTomato and allowing them to express for approximately 4 weeks before slicing the brain. Whole-cell current clamp recordings were then made and the RSC and V1 axons were stimulated with 1 ms pulses of full-field blue (470 nm) or red (635 nm) light delivered through the 40× water immersion objective centred on the soma. During stimulation, the illumination diaphragm on the microscope was adjusted to restrict the light to a circle of approximately 100 µm radius around the soma (Figure 5.4).

Single injection control experiments

Although Chronos and ChrimsonR have peak wavelength sensitivities in different parts of the visible spectrum, there is a substantial region of overlap

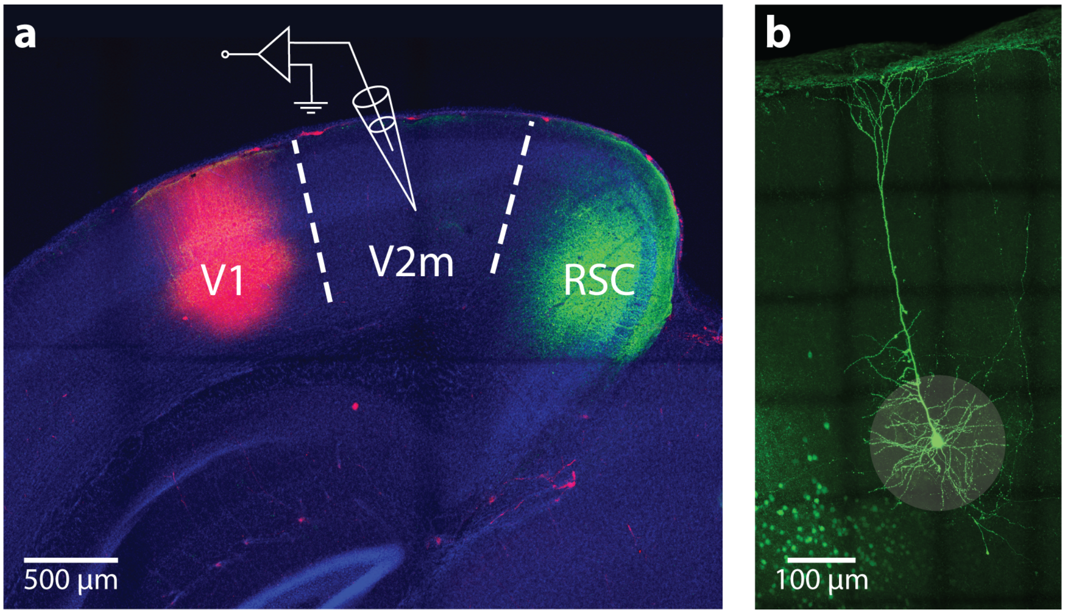

Figure 5.4: Dual opsin optogentic stimulation of RSC and V1. a. Coronal brain slice showing injection sites of Chronos-GFP in RSC and ChrimsonR-tdTomato in V1. Locations of Chronos and ChrimsonR injections were alternated in different mice. **b.** Confocal image of a Glt neuron in V2m filled with biocytin hydrochloride during whole-cell recording and stained with Alexa Fluor 488. The highlighted circle around the soma illustrates the approximately 100 µm radius that was directly stimulated by light during these optogenetics experiments.

between their excitation spectra. To obtain maximally targeted stimulation of ChrimsonR, Klapoetke et al. (2014) used 625 nm light which is relatively close to the 590 nm sensitivity peak and did not induce any photocurrent in Chronos-expressing neurons. At shorter wavelengths, however, both ChrimsonR and Chronos have some sensitivity. To independently stimulate Chronos without affecting ChrimsonR-expressing neurons they used 470 nm light, as the difference between the two excitation spectra is greatest around this wavelength. By carefully selecting the blue light power used to stimulate the cells, they were thus able to identify a narrow power range between approximately 0.1 and 0.5 mW/mm^2 in which spikes could be evoked with 100% probability in Chronos-expressing neurons without triggering any spikes in ChrimsonR-expressing neurons (see Fig. 5 in Klapoetke et al. (2014)).

The experiments described by Klapoetke et al. (2014) were carried out in L2/3 neurons of mice which had been transfected via *in utero* electroporation

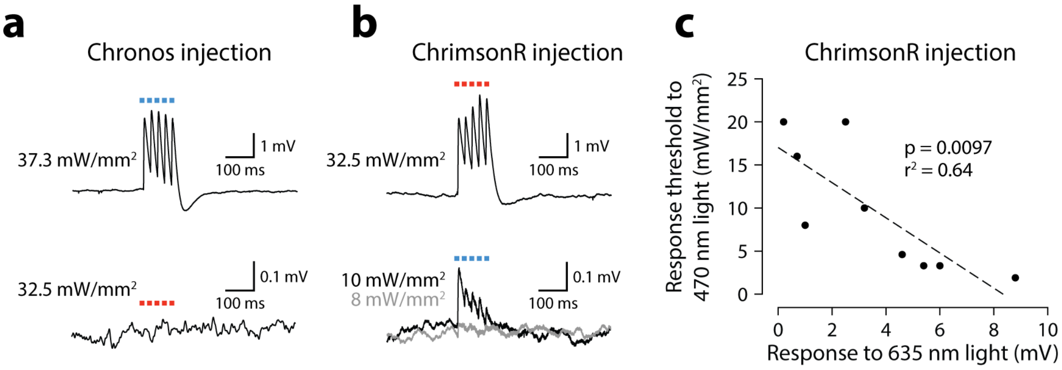

Figure 5.5: Single opsin control experiments. **a.** Recording of a Glt neuron during 1 ms pulses of 470 nm (top) or 635 nm (bottom) light in a mouse injected with Chronos-GFP in RSC. **b.** Recording of a Glt neuron during 1 ms pulses of 635 nm (top) or 470 nm (bottom) light in a mouse injected with ChrimsonR-tdTomato in RSC. ChrimsonR-evoked responses were observed when stimulating with 470 nm light above a certain light intensity threshold, defined as the lowest intensity to evoke average responses smaller than 0.1 mV. No EPSPs were observed below this threshold (grey line). **c.** ChrimsonR threshold to 470 nm stimulation for different cells plotted against EPSP amplitude of the same cell to 635 nm stimulation at maximum intensity (32.5 mW/mm^2).

and were recorded in ACSF containing 2 mM $CaCl_2$ while optically stimulating the soma and dendrites using 5 ms light pulses. As my experiments involved a different cell type and different recording and stimulation conditions, I tested the properties of Chronos and ChrimsonR under my experimental conditions by expressing a single opsin in RSC and stimulating the axons with different intensities of blue (470 nm) and red (635 nm) light while recording from Glt neurons in V2m.

For both opsins, response amplitudes were typically variable from cell to cell. When only Chronos was expressed, reliable postsynaptic responses could be recorded in response to blue light, while no responses were ever recorded in response to red light regardless of light power, consistent with the results in Klapoetke et al. (2014) (Figure 5.5a). However, the evoked responses when using light powers below 0.5 mW/mm^2 (which can reliably induce somatic spikes) were typically extremely small if present, and the responses scaled with higher light powers. This suggests that, unlike the somatic stimulation which has previously been described, the spike probability in the axons and nerve terminals was most likely far below 1.

When only ChrimsonR was expressed, EPSPs were observed in response to both blue and red light (Figure 5.5b). As my ultimate intention was to be able to independently stimulate Chronos in double-injected animals without also evoking EPSPs from ChrimsonR-expressing axons, I gradually decreased the blue light power (in the animals expressing only ChrimsonR) until a threshold was found for which the blue-evoked ChrimsonR responses were smaller than 0.1 mV. While this blue light threshold differed between cells, the threshold power was found to negatively correlate with the EPSP amplitude of the ChrimsonR response to red light delivered at 32.5 mW/mm^2 (n = 9, $r^2 = 0.64$, $p = 0.0097$, F-test; Figure 5.5c). Neurons with very strong responses to red light were thus also very sensitive to blue light, while in neurons with weaker red-evoked responses I was able to administer higher intensity blue light without evoking ChrimsonR responses.

As it would not be possible to directly measure the threshold in mice injected with both Chronos and ChrimsonR, to avoid evoking EPSPs from ChrimsonR-expressing terminals when stimulating with 470 nm light, in subsequent experiments the ChrimsonR threshold was estimated based on this linear fit and blue light powers for each cell were adjusted so as to be slightly below the fitted line, typically taking values around 4–8 mW/mm^2.

Paired-pulse stimulation

Once suitable light powers had been determined, I studied the temporal properties of Chronos- and ChrimsonR-evoked postsynaptic responses by stimulating the axons and nerve terminals in double-injected animals using the same stimulation protocols as with the electrical stimulation experiments (Figure 5.6a). Injection sites for the two opsins in RSC and V1 were alternated between mice to reduce the influence of any pathway-specific effects.

Paired-pulse responses for Chronos-expressing pathways appeared very similar to those recorded using electrical stimulation, with a similarly facilitat-

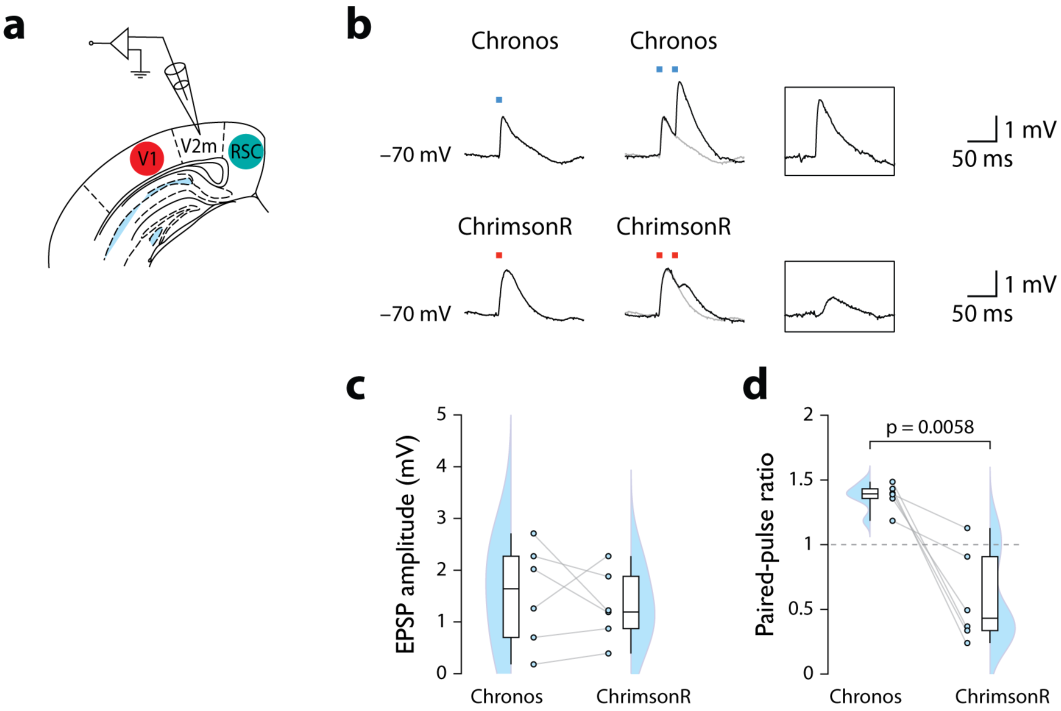

Figure 5.6: Optogenetic paired-pulse stimulation of Chronos and ChrimsonR. a. Diagram illustrating position of recording electrode and injection sites in RSC and V1. **b.** Example traces of single pulse and 40 Hz paired-pulse stimulation of pathways expressing Chronos (top) and ChrimsonR (bottom). The gray traces show the EPSP waveforms from single pulse stimulation, which were subtracted from the paired-pulse trace. The boxed trace on the right displays the measured second EPSP after this subtraction. **c.** Summary data of EPSP amplitudes from single pulse stimulation of Chronos and ChrimsonR. **d.** Summary data of paired-pulse ratios for Chronos and ChrimsonR stimulation, computed as the ratio of the second and first EPSP amplitudes.

ing paired-pulse ratio (Figure 5.6b). Chronos-evoked responses had a paired-pulse ratio of 1.37 ± 0.04 ($n = 6$), which was not statistically different from the mean ratio evoked during electrical stimulation of RSC and V1 (1.49 ± 0.08, $n = 28$; $p = 0.52$, two-sample t-test; Figure 5.6d). This is consistent with the data from Klapoetke et al. (2014), where no Chronos desensitization was present when stimulating at 40 Hz.

Paired-pulse stimulation of ChrimsonR, on the other hand, showed strong paired-pulse depression, despite being in the same pathways that were facilitating when stimulated electrically (Figure 5.6d). Mean paired-pulse ratio (0.58 ± 0.15, $n = 6$) was significantly smaller than in the Chronos-expressing pathways ($p = 0.0058$, paired-sample t-test). This was surprising as 40 Hz

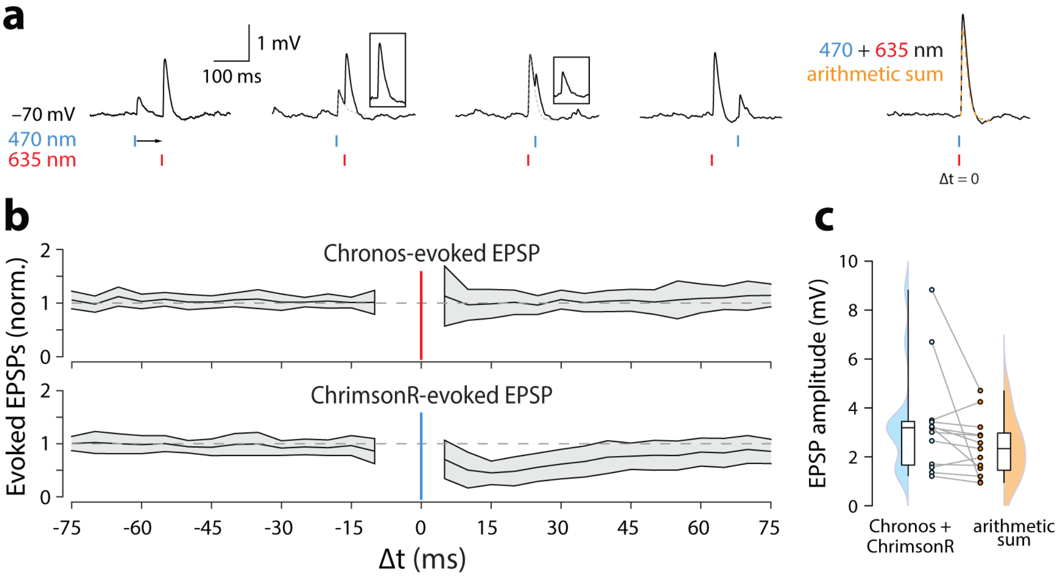

Figure 5.7: Temporal interactions of Chronos and ChrimsonR. a. Example traces from one cell showing 5 different inter-stimulus intervals (with the 470 nm stimulus occurring at $\Delta t = -75$ ms, -20 ms, $+20$ ms, $+75$ ms, or 0 ms relative to the 635 nm stimulus). The grey traces show fitted EPSP waveforms derived from single pulse stimulation, which were subtracted before measuring the second EPSP amplitude. The arithmetic sum was derived from the sum of independent EPSP waveforms for Chronos (470 nm) and ChrimsonR (635 nm) stimulation. **b.** Summary data of Chronos-evoked (top) and ChrimsonR-evoked (bottom) EPSP amplitudes at all time intervals. The vertical line at 0 ms indicates the position of the 635 nm stimulus (top) or the 470 nm stimulus (bottom) relative to the measured EPSP. Negative Δt values indicate that the measured EPSP precedes the other stimulus, with 0 ms indicating simultaneous stimulation. All amplitudes are normalised to the single pulse EPSP amplitude for each opsin. Shaded areas show standard deviation of responses. **c.** Summary data of total EPSP amplitude during simultaneous Chronos and ChrimsonR stimulation ($\Delta t = 0$ ms) compared to the calculated linear sum based on summation of the independent waveforms.

stimulation was expected to induce minimal ChrimsonR desensitization, with less than 15 % reduction in spike probability, based on the data from somatic expression in Klapoetke et al. (2014).

Temporal interactions of Chronos and ChrimsonR

While the paired-pulse stimulation results suggest that ChrimsonR would not be suitable for stimulation at high temporal frequencies, the observed desensitization would not preclude its use for assessing responses to single light pulses in conjunction with other stimuli. To determine the feasibility of combining Chronos and ChrimsonR stimulation, I therefore delivered 1 ms pulses of

470 nm and 635 nm light with inter-stimulus interval varying between −75 ms and +75 ms in 5 ms steps, as was previously done with electrical stimulation (Figure 5.7a).

Chronos-evoked responses remained constant at all stimulus intervals and were unaffected by the 635 nm stimulus ($n = 13$, $p = 0.80$, repeated measures ANOVA test), as was observed with electrical stimulation. ChrimsonR-evoked responses, however, were strongly suppressed following the 470 nm stimulus ($n = 13$, $p = 2.47\times10^{-12}$, repeated measures ANOVA test), despite the fact that the intensity of this stimulus was set below the estimated threshold for directly evoking ChrimsonR responses. Maximal suppression occurred when the red pulse followed the blue pulse by 15 ms. At this time interval the ChrimsonR response was reduced to 46 ± 6 % of its independently stimulated amplitude. When stimulated simultaneously, evoked Chronos and ChrimsonR responses had a combined amplitude of 3.35 ± 0.6 mV, which was not significantly different from the calculated arithmetic sum of the individual EPSPs (2.39 ± 0.33 mV; $n = 13$, $p = 0.089$, paired-sample t-test; Figure 5.7c).

As the response suppression was only present for ChrimsonR stimulation, to unambiguously investigate the response characteristics of ChrimsonR when preceded by a pulse of blue light I performed an additional control experiment in which a similar stimulus was delivered in an animal that was injected only with ChrimsonR. Blue and red light stimuli were pulsed both individually and in brief succession, with the blue pulse preceding the red pulse by 25 ms. At high powers of blue light, the ChrimsonR response was nearly abolished altogether (Figure 5.8a). Surprisingly, however, even when blue light power was reduced to being subthreshold for evoking a ChrimsonR response, it still had a dramatic effect on the following ChrimsonR response, reducing it by 75.3 ± 3.5 % ($n = 4$) compared to the independent red pulse (Figure 5.8b). When blue light power was reduced further to only 0.2 mW/mm^2, which was expected to have a 0 % probability of inducing any currents in ChrimsonR-expressing neurons (Klapoetke et al., 2014), a postsynaptic ChrimsonR am-

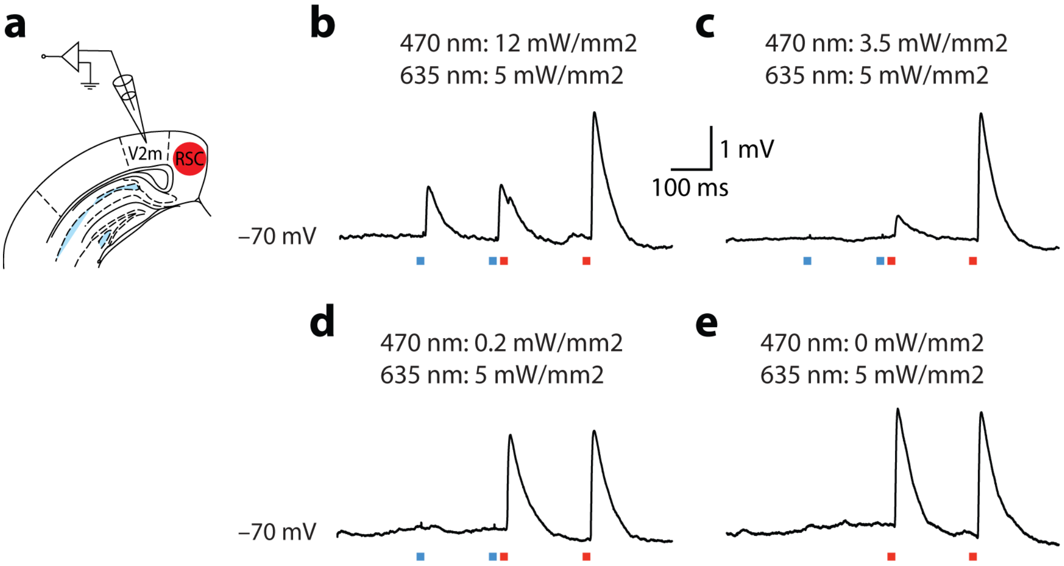

Figure 5.8: Subthreshold suppression of ChrimsonR responses with blue light. Example traces from a recorded Glt neuron in a mouse injected with ChrimsonR in RSC. **a.** Diagram illustrating ChrimsonR injection site in RSC and recording site in V2m. **b.** Responses to individual and paired light pulses with high blue light power. **c.** As in *b* but with subthreshold blue light power. **d.** As in *b* but with extremely low blue light power. **e.** Responses to only red stimuli, with blue light power set to 0.

plitude reduction of 9.5 ±1.6 % was still observed (Figure 5.8c). Only when the blue light power was reduced to 0 did the ChrimsonR-evoked EPSP maintain the same amplitude across both red light pulses. This final experiment provided conclusive evidence that it is not possible to independently optically excite distinct neural populations expressing Chronos and ChrimsonR within intervals shorter that approximately 75 ms while maintaining consistent stimulus amplitude.

5.4 Discussion

Here I have shown that the synapses from RSC and V1 evoke EPSPs that show surprisingly linear summation, despite the many known sources of non-linearity in ttL5 neurons. Indeed, the amplitudes of corrected EPSP were found to be approximately constant regardless of the relative timing of stimuli, with only marginally sublinear integration occurring when the two pathways were stimulated simultaneously. This suggests that synaptic currents for each input remain unchanged and do not actively interact, and that the observed EPSPs likely result from passive linear summation of the two independent pathways.

Part of this linearity may be explained by differences in dendrite targeting. I previously showed that RSC projections mainly target the basal dendrites, while V1 inputs target both basal and apical dendrites. Even among the basal inputs, it's possible that the synapses could be clustered on separate dendrites on the level of individual neurons in a way that is not necessarily detectable when averaging the sCRACM heatmaps across cells. This could be relevant because the location of synapses on dendrites is known to influence the way signals are filtered and the arithmetic operations the dendrites produce on these signals (Cash and Yuste, 1998, 1999; Silver, 2010).

In general, when synaptic inputs are broadly distributed across the dendritic tree they tend to sum linearly by nature of the passive cable properties of dendrites. On the other hand, when inputs are spatially clustered the depolarization caused by one synapse could reduce the ionic driving force of neighboring synapses and the local increase in membrane conductance could have a shunting effect, leading to sublinear summation. However, many dendrites also express ion channels that required a depolarized state to open, such as NMDA receptors and voltage-gated Ca^{2+} and Na^{+} channels. If the clustered synapses cause a sufficient local depolarization of the dendrite, these channels can boost the synaptic potentials and induce supralinear summation of the inputs.

The EPSP amplitudes I recorded, being marginally sublinear, are consistent with a mostly independent distribution of synapses between the two pathways, with possibly a small amount of clustering in the perisomatic area. However, as the evoked potentials were relatively small it is possible that supralinear mechanisms could have been recruited given sufficient stimulation if I had explored a larger range of evoked EPSP amplitudes. Large potentials, however, were generally difficult to obtain, most likely due to only a small fraction of the total axons being preserved across such large distances in a 300 µm thick slice. This limitation could potentially be overcome by increasing the slice thickness. However, this would substantially reduce the throughput of such experiments, because of the reduced visibility of neurons (due to light scattering and lower illumination intensity) and possibly also due to poorer oxygenation of deep neurons, which could influence the health of neurons throughout the slice.

Individual stimulation of both V1 and RSC pathways revealed substantial paired-pulse facilitation. This is sometimes considered an indication of a pathway having a more modulatory role, rather than being a main driver of activity (Sherman and Guillery, 1998, 2011). While a more modulatory role might be anticipated in the top-down connection from RSC, it is somewhat unexpected that the V1 pathway would also share this property. However, this may be consistent with my previous observation that the direct long-range input from V1 neurons is biased towards the apical compartment, which would also imply a modulatory role. This also suggests that, rather than coming directly from V1, the primary feedforward drive of visual input to Glt neurons in V2m comes indirectly through either the superficial layers of V2m or through a higher-order visual thalamic nucleus such as the lateral posterior thalamus (LP).

One aspect to note in these responses to paired-pulse stimulation is that presynaptic short-term plasticity is more commonly studied in voltage clamp configuration, because this is believed to give more accurate estimates of synaptic currents without recruiting active processes or varying the driving force of any ionic currents. While this may be true in isopotential compartments, voltage clamp is generally ineffective even small distances along the dendrites (Spruston et al., 1993; Bar-Yehuda and Korngreen, 2008; Jackman and Regehr, 2017). As my main objective was to determine the individual and combined influence of different pathways on the activity of Glt neurons, I therefore opted to perform the recordings in current clamp configuration, recording the potentials that are likely to occur within these neurons under normal conditions. The drawback of this is that the measured facilitation can not be said to depend entirely on presynaptic factors, as the EPSPs are likely to be influenced by interactions and active conductances in the postsynaptic cell, such as the recruitment of NMDA receptors which might be enhancing the amplitude of the second EPSP.

Another major limitation of the electrical stimulation experiments is that the effects of this stimulation method are more difficult to determined due to the lack of specificity in the neurons and pathways that are stimulated. Each electrode is thus likely to be simultaneously stimulating many different cell types across different cortical layers, which might have opposing influences on the Glt neurons that get masked when stimulating them together. In V1, for example, if the projections to the apical and basal dendrites that I discovered in chapter 4 were to originate from separate populations of L5 neurons, these two populations would likely have different synaptic properties and be integrated differently in the Glt neurons. Furthermore, during electrical stimulation passing fibers originating in other areas are also recruited, potentially adding to the ambiguity in interpreting the effects of stimulating each region.

To preserve physiological axonal and synaptic transmission, there experiments described here were all made in regular ACSF without any pharmacological agents. Unlike in chapter 4, I am thus not able to guarantee that all recorded postsynaptic potentials arise from direct monosynaptic connections between V1 or RSC and V2m. It is not inconceivable, for instance that stimulating RSC could trigger spikes in downstream V1 neurons, which in turn pass the signal forward to the recorded Glt neuron. This particular case is not very likely, as the probability of an axon being severed in the slice increases with distance travelled though the slice. The joint probability of an RSC axon travelling intact to V1, having a strong enough effect to trigger a spike, and for that downstream neuron to also have its axon fully intact between V1 and V2m is thus fairly low. However, it is quite plausible that an RSC axon travelling to the closer V2m could stimulate a local neuron, and that the observed EPSP (and related synaptic parameters) thus reflect the properties of this local neuron rather than of the actual RSC pathway.

Even without the numerous drawbacks in the specificity of electrical stimulation, by requiring intact axons these experiments are inherently only possible if the stimulated brain regions are in the same slice as the recorded V2m neurons. A large number of long-range inputs identified in the rabies tracing, such as those from OFC or anterior thalamus), would thus not be amenable to this type of investigation anyway, as there is no straightforward way to slice the brain while preserving these axons all the way to V2m.

For all of the reasons stated above, in order to enable experiments on the functional interactions of inputs from distant presynaptic population, after the electrical stimulation experiments I assessed the feasibility of an optogenetic approach. This was inspired by newly developed optogenetic activators with blue- and red-shifted excitation spectra, that would allow me to separately express the opsins in different presynaptic populations and thereby independently stimulate any two input pathways using different wavelengths of light (Klapoetke et al., 2014).

While this method initially showed promise, in performing the optogenetic stimulation experiments I uncovered several limitations that ultimately made this technique unsuitable for the experiments that I wanted to perform. The main problem stemmed from the the substantial overlap in spectral sensitivity between the two opsins in the blue range of the visible spectrum. While I could be confident that red light stimuli were only activating ChrimsonR-expressing axons, the broad sensitivity of this opsin made it difficult to stimulate Chronos-expressing axons with blue light without also recruiting ChrimsonR.

As I show, this problem alone could possibly be overcome by carefully setting the blue light power in a range that is subthreshold for ChrimsonR-expressing axons while still evoking Chronos-mediated responses. However, ChrimsonR was found to also suffer from a second drawback in the temporal domain. When presented with repeated light stimuli, while Chronos appeared to be largely stable even at high frequencies, ChrimsonR rapidly desensitized. My data is insufficient to conclude whether this is due to direct desensitization of the opsin, or to a consequence of the ChrimsonR current, such as Ca^{2+} influx altering the release probability of the ChrimsonR-expressing synapses. Regardless of the cause, this desensitization resulted in drastically reduced amplitudes to the second stimulus. Moreover, this desensitization occurred not only when delivering paired red stimuli, but critically also when a single red stimulus was preceded by a blue stimulus, even if this blue stimulus did not itself evoke any EPSP. As my aim was to perform experiments similar to those done using electrical stimulation, with the two pathways stimulated in rapid succession, overall this combination of spectral overlap and rapid desensitization thus rendered ChrimsonR ineffective for use in the dual-opsin experiments that I had originally planned.

These results do not preclude the combined use of Chronos and ChrimsonR altogether. However, if response amplitude is a relevant parameter for a particular experiment, my experiments indicate that Chronos and ChrimsonR can only be used in conjunction if the temporal interval between them is at least 75 ms (beyond which ChrimsonR amplitudes appeared to have fully recovered). Even under this circumstance, caution should be taken, particularly when stimulating axons and nerve terminals rather than opsin-expressing somas, as the maximum intensity of blue light that can be used without triggering ChrimsonR responses is highly variable and can not be determined unambiguously in preparations where both opsins are expressed simultaneously. The interaction may thus not be a problem if stimulation is restricted to the soma and only all-or-none spike responses are quantified. However, the inability to interpret responses when stimulating Chronos and ChrimsonR together at high temporal frequencies precludes the study of any synaptic interactions.

Similar challenges of combining optogenetic activators with different spectral sensitivity have also previously been described by Hooks et al. (2015) when attempting to independently photostimulate synaptic inputs expressing the blue channelrhodopsin ChR2 and the red-shifted ReaChR. They presented a solution to overcome the problem of spectral overlap by deliberately desensitizing the red-shifted opsin with long pulses of orange (590 nm) light. This creates a window of time after the offset of the orange light when blue light is no longer effective on ReaChR and any remaining responses to blue light can be confidently claimed to arise entirely from the ChR2-expressing axons.

Considering my data, it is very likely that a similar approach would also work to independently stimulate Chronos and ChrimsonR, at least for experiments in which the relative timing of the stimuli is not important. One potential future experiment that would be feasible with this approach is the simultaneous subcellular mapping of two input pathways using the sCRACM method. The synaptic distribution of ChrimsonR-expressing axons could first be mapped using conventional sCRACM methodologies with a red laser, which does not evoke any Chronos activity. The Chronos-expressing pathway could then also be mapped in the same neuron by isolating Chronos-evoked EPSCs with a strong and wide pulse of red light preceding each blue stimulus, which would eliminate ChrimsonR-dependent effects from the blue light response.

In addition to increased throughput (from simultaneously measuring the distribution of two different input pathways), a major advantage of a dual-opsin approach to sCRACM would be that different pathways could be determined in individual cells, thus revealing patterns that are not detectable from the average heatmap. For example, although tuft and basal inputs from different pathways can easily be distinguished, inputs distributions that are largely overlapping (such as the perisomatic components of the RSC, V1, and OFC inputs to Glt neurons) may correlate at the single-cell level, by either clustering on the same dendrites or avoiding each other and targeting separate dendrites. If there is substantial variation between cells, such spatial correlations within the dendrites of individual neurons would not be detectable in average responses, but would be revealed by a dual-opsin approach.

Chapter 6

Dendritic excitability & biophysical modelling

6.1 Introduction

One of the best know integrative properties of ttL5 neurons is their ability to fire in bursts and exhibit dendritic supralinearities. One such supralinearity is due to a high concentration of Ca^{2+} channels in their apical dendrites (Amitai et al., 1993; Yuste et al., 1994; Schiller et al., 1997), which can induce sustained depolarizations (known as Ca^{2+} plateaus) that trigger high-frequency bursts of somatic spikes.

This property is exemplified by their propensity for producing dendritic Ca^{2+} plateaus when subject to coincident stimulation at the soma and apical dendrite (Larkum et al., 1999b)—a phenomenon known as BAC firing. Dendritic Ca^{2+} plateaus can also be triggered by high-frequency trains of somatic spikes, which backpropagate into the dendrites and combine to reach the necessary voltage for Ca^{2+} channel activation (Larkum et al., 1999a; Kampa and Stuart, 2006; Shai et al., 2015). This dendritic depolarization is visible in the somatic voltage as an after-depolarization (ADP) following the somatic spike train. The frequency response is typically sharp, with the ADP visible only above a critical frequency of approximately 100 Hz.

While Ca^{2+} channel expression is clearly a necessary element in BAC firing, morphology is also known to have a large influence on the bursting properties of ttL5 neurons. In particular the location of oblique dendrites and the electrical coupling between the soma and dendrites is thought to be crucial for determining a neuron's integrative properties (Mainen and Sejnowski, 1996; Vetter et al., 2001; van Ooyen et al., 2002; Schaefer et al., 2003). Recent experimental work has further shown that even within the class of ttL5 neurons there can be substantial variation in intrinsic properties depending on their location within V1 or on the species from which they are recorded (Beaulieu-Laroche et al., 2018; Fletcher and Williams, 2019).

Our current understanding of ttL5 neurons has been summarized in a general model that explains their operation as multi-stage integrators of tuft and basal input (Larkum et al., 2009; Larkum, 2013; Stuart and Spruston, 2015). This view has been derived mainly from recordings of ttL5 neurons in V1 and primary somatosensory cortex of rats. These recordings have enabled the generation of detailed conductance-based models that reproduce the experimentally recorded somatic and dendritic spike properties (Hay et al., 2011). It has also been shown that even reduced models with simplified morphology are capable of capturing most of the relevant dynamics (Bahl et al., 2012).

BAC firing in ttL5 neurons has been suggested to have wide-ranging implications for cortical computation. As top-down connectivity between cortical areas often targets cortical L1 (Rockland and Pandya, 1979; Coogan and Burkhalter, 1990; Felleman and Van Essen, 1991; Cauller et al., 1998; Markov et al., 2014), where the apical tuft dendrites are located, BAC firing is believed to play a major role in integrating feedforward and feedback pathways in the brain to modulate sensory perception (Takahashi et al., 2016).

Although in previous chapters I have explored both the basic integrative properties of Glt neurons, as well as the connectivity of different inputs to their apical and basal dendrites, in this chapter I specifically investigate the supralinearities linked to dendritic Ca^{2+} plateaus. This was done using experimental configurations and stimulation protocols that have previously been shown to be optimal for these types of neuronal behaviours. In order to address the final aim from chapter 1, I have analysed some possible links between these types of supralinearities and dendritic morphology, through both experiments and biophysical models.

In performing these experiments, preliminary results while attempting to evoke BAC firing in V2m neurons gave me reason to hypothesise that these neurons might not conform to the classical models of ttL5 responses, which are based on recordings from primary sensory cortices (Larkum et al., 1999b; Hay et al., 2011). For this reason, I decided to systematically compare the bursting properties of ttL5 neurons in both V1 and V2m.

While previous chapters have focused entirely on Glt neurons in V2m, the analysis in chapter 3 of whole-brain distribution of these cells across cortex revealed that they occur very sparsely in V1 (Figure 3.1), making a comparison of Glt neurons across V1 and V2m very difficult to achieve. In order to make the comparison of ttL5 neurons across V1 and V2m fair and ensure that all the neurons recorded from V1 also belonged to a well-defined group of ttL5 neurons, in this chapter I therefore primarily used a different strategy to identify ttL5 neurons.

Aside from morphology and intrinsic physiology, a defining characteristic of all ttL5 neurons is their projection to subcortical targets. In addition to identifying ttL5 neurons based on Glt labelling (which is an effective strategy in V2m, but not in V1), I therefore used a retrograde tracer injected into the thalamus to uniformly label ttL5 neurons across both V1 and V2m. I thus demonstrate that any observed differences are true, not only in the genetically defined Glt population in V2m, but also in a more conventionally defined set of ttL5 neurons, selected according to the same unbiased criteria in both V1 and V2m.

6.2 Chapter methods

6.2.1 Cell selection

Recordings were made from ttL5 cells in both V1 and V2m. Some cells in V2m were selected based on their Glt labelling, and were thus identified through their expression of tdTomato and visualized with 550 nm light, as described in previous chapters. To visualize other ttL5 neurons across both V1 and V2m, I injected Alexa Fluor 488-conjugated cholera toxin subunit B (CTB) in the lateral posterior (LP) nucleus of the thalamus (Figure 6.1). This retrograde tracer produced fluorescence in the LP-projecting cells, which had a granular appearance around the soma characteristic of CTB labelling.

CTB injections were made in either Glt or Rbp4-Cre mice crossed with tdTomator reporter mice. In experiments using Rbp4 mice, the tdTomato fluorescence was not a primary selection criterion, but made it easier to distinguish L5 from L6, since this Cre line is specific for L5 neurons (but labels both ttL5 and slender-tufted neurons). After 2–4 weeks, dense labelling could be seen in both L5 and L6 of V1 and V2m. As the projections of other cell types in L5 (such as slender-tufted neurons) are primarily intracortical, by recording from the L5 population of CTB labelled neurons I could ensure that these all consisted of ttL5 neurons and that I was introducing minimal selection bias when sampling cells.

To minimize photobleaching, when exciting the alexa fluorophores (using 470 nm light from the CoolLED pE-4000 system) the illumination power was kept as low as possible and was only turned on in brief pulses (lasting at most a few seconds). As retrograde injections in thalamus label neurons in both L5 and L6, care was also taken to ensure that the recorded neurons were only in L5, which can be determined based on both the distance from the pia and white matter and the different appearance of the two layers when observed under regular IR illumination.

The distinction between V1 and V2m was largely based on approximate stereotaxic coordinates. However, CTB labelling was also frequently found to differ in intensity between V1 and other visual areas (both V2m and lateral visual areas). In a few cases where the recorded neurons were close to the border between V1 and V2m, this difference in CTB fluorescence was therefore also used to decide the identity of the neuron.

6.2.2 Recording conditions

Recordings presented here were all made in ACSF heated to 35 °C, but some variations were made in the ACSF composition in different experiments. In the literature there is considerable ambiguity regarding the extracellular Ca^{2+} *in vivo* (Lopes and Cunha, 2019). This parameter is potentially of particular relevance to the dendritic excitability experiments described in this chapter. While in chapter 3 I performed recordings in 1.5 mM $CaCl_2$ in an attempt to be somewhat close to *in vivo* physiology, nearly all previous literature specifically exploring dendritic Ca^{2+} plateaus has been done in 2 mM $CaCl_2$. To be consistent with my previous experiments while also producing results that can be compared to previous literature, I have thus performed some of these experiments in both 1.5 mM $CaCl_2$ and 2 mM $CaCl_2$.

Likewise, to obtain a broader exploration of the parameters space, experiments using extracellular stimulation were done both in clean ACSF and in ACSF containing 1 µM CGP52432, which was added to increase the chance of eliciting BAC firing and thus strengthen the evidence for the claims being made.

6.2.3 Electrical stimulation

After positioning the slice within the recording chamber, a single monopolar glass electrodes mounted to a manipulator on the opposite side of chamber from the recording electrode was slowly lowered and placed in L1 of the area I intended to record from (either V1 or V2m). With the electrode in place, I selected a fluorescently labelled neuron from the underlying L5 for whole-cell patch clamp recordings. For both V1 and V2m recordings, the electrode was located in approximately the same position relative to the ttL5 neurons, i.e. in

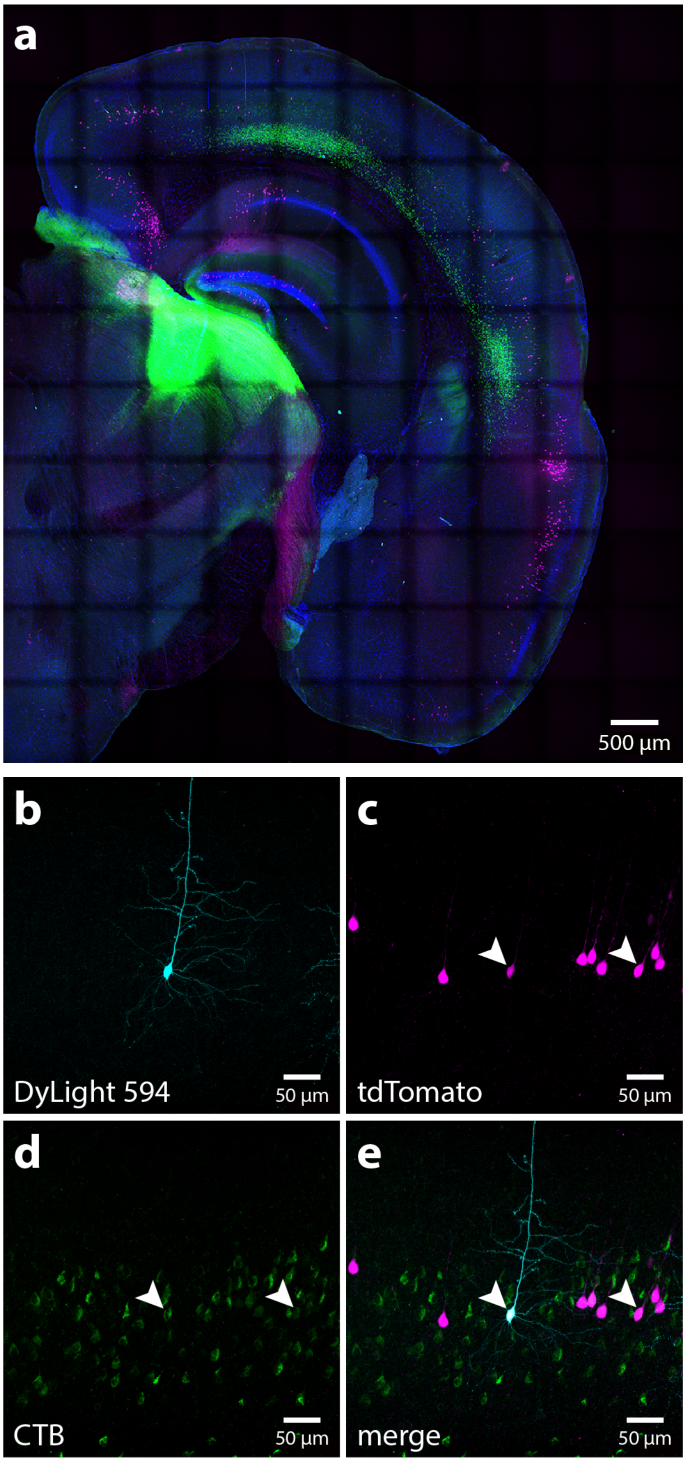

Figure 6.1: Retrograde labelling of LP-projecting ttL5 neurons. Confocal images from a Glt25d2-Cre mouse injected with CTB-Alexa Fluor 488 in LP. **a.** Coronal slice of one hemisphere containing visual cortex, showing CTB labelling in cortex and at the injection site (green), tdTomato-expressing Glt neurons (magenta), DAPI (blue), and neurons that were filled with biocytin hydrochloride during intracellular recording and stained with DyLight 594 (cyan). **b.** Biocytin-filled ttL5 neuron in V2m. **c.** Neighbouring tdTomato-expressing Glt neurons. **d.** CTB-labelled L5 neurons projecting to LP. **e.** Composite image of b–d.

L1 (the boundary of which is clearly visible under standard IR illumination) and within approximately 200 µm of the apical dendrite in the lateral direction.

The experiments using electrical stimulation of L1 were designed to determine is a given ttL5 neuron was capable of BAC firing, given the appropriate conditions. For each cell, I therefore applied short short (0.1–1 ms) current pulses through the stimulating electrode while exploring the full range of stimulus intensities. Specifically, I began each recording at the smallest stimulus intensity (20 µA) and after a few stimulus pulses I slowly increased the stimulus current by manually turning the analog dial controlling this on the Digitimer DS3 constant current stimulator. The current was increased by a small amount between each recording, but was left constant during the actual recording. The maximal current for each cell was 320 µA, but in most cells I ended the recordings before reaching this current when the L1 stimulus alone was strong enough to trigger a spike burst (consisting of 2 spikes) in the recorded cell. For visualization, stimulus artefacts occurring in the recorded voltage trace at the time of stimulation were manually deleted in the example traces in the same way as I did in chapter 5.

6.3 Results

To study the integrative properties of ttL5 neurons in V1 and V2m, I made current clamp recordings from Glt cells in V2m as well as from neurons in both V1 and V2m that projected to LP in the thalamus, identified though retrograde fluorescent CTB labelling. This alternative labelleing was done in order for recordings to be more easily comparable between brain regions and to ensure consistency in cell type when sampling from these regions.

6.3.1 Thick-tufted L5 neurons in V2m lack BAC firing

To reproduce the conditions required for triggering BAC firing in ttL5 neurons, I stimulated synaptic inputs near the apical tuft in L1 using an extracellular electrode in conjunction with somatic stimulation through the recording electrode. A single spike was generated in the soma by injecting a 5 ms depolarizing step through the recording electrode. At the offset of this current step, an extracellular current pulse was triggered to stimulate L1 inputs. The intensity of this extracellular pulse was adjusted to evoke either a subthreshold EPSP or at most a single spike at the soma.

Specifically, the stimulus current was initially set to 20 µA and was gradually increased between trials (Figure 6.2). As I increased the current, the evoked postsynaptic potentials gradually increased in amplitude. At interme-

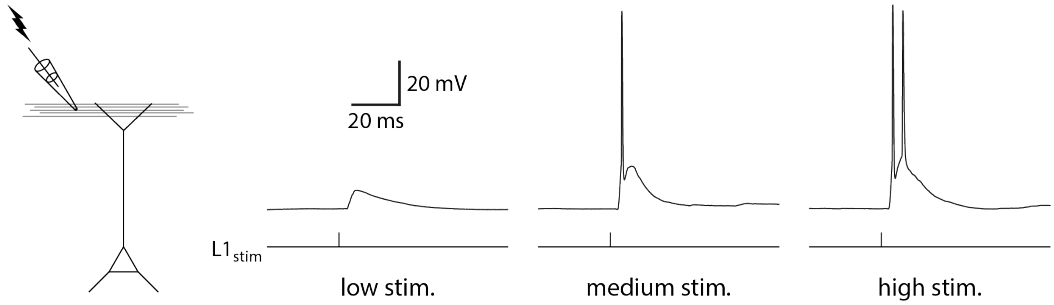

Figure 6.2: Voltage responses to L1 stimulation. Example postsynaptic potentials following L1 stimulation with gradually increasing current. For each recorded cell in both V1 and V2m, the full range of low and medium strength currents were explored in the attempt to evoke BAC firing. L1 currents strong enough to evoke a double spike, however, were excluded. The stimuli used were thus only strong enough to evoke either an EPSP or at most a single spike, but never a double spike.

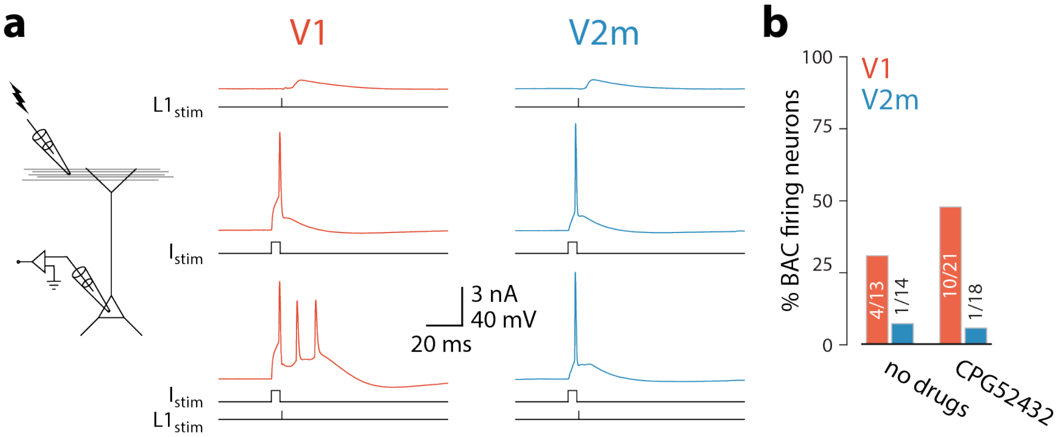

Figure 6.3: BAC firing in V1 and V2m. **a.** Left: diagram of experimental configuration with somatic whole-cell patch-clamp recording electrode and extracellular stimulation electrode in L1. Right: example traces from V1 (red) and V2m (blue) ttL5 neurons during BAC firing stimulation paradigm. **b.** Proportion of BAC firing (i.e. bursting) cells in V1 and V2m recorded with and without a GABA$_B$ antagonist in the ACSF.

diate current amplitudes, these were sufficient to trigger an action potential in the recorded neuron. As I approached the largest stimulus intensity (320 µA), the dendritic stimulus usually triggered a burst of 2 spikes. Any such recordings in which 2 spikes occurred during L1 stimulation alone were not used in any subsequent analyses to assess BAC firing.

In V1 neurons, combined stimulation occasionally evoked a prolonged plateau potential resulting in a burst of 3 or more spikes, as is typical of BAC firing. On the other hand, when I repeated these experiments in ttL5 neurons in V2m (both Glt and CTB-labelled) under the same recording conditions, such supralinear integration was rarely observed (Figure 6.3a).

To quantify these observations, I defined as “bursting” any cell in which three or more spikes could be evoked following combined somatic and L1 stimulation (each evoking no more than one spike). With this definition, when recording in normal ACSF with no added drugs, bursting was present in 4/13 V1 neurons but only in 1/14 V2m neurons (Figure 6.3b).

When stimulating extracellularly in L1, it is likely that recruited fibers are composed of both excitatory and inhibitory inputs, which may have induced opposing effects on the apical dendrites. In particular, GABA$_B$ receptors have been specifically shown to inhibit dendritic Ca^{2+} plateaus in ttL5 neurons (Perez-Garci et al., 2006). To avoid recruiting inhibitory inputs, I therefore repeated these experiments after adding the competitive $GABA_B$ receptor antagonist CGP52432 (1 µM) to the extracellular solution. With reduced inhibition, the proportion of bursting neurons in V1 increased to 10/21, while in V2m this proportion stayed at only 1/18 ($p = 4.6 \times 10^{-3}$, Fisher's exact test, Figure 6.3b). These results suggest a much-reduced propensity for BAC firing in V2m neurons compared to V1 neurons.

6.3.2 Thick-tufted L5 neurons in V2m lack a critical frequency ADP

To further investigate the prevalence of dendritic supralinearities in ttL5 neurons across the different visual areas, I recorded another hallmark of dendritic Ca^{2+} plateaus which is the somatic ADP following a high-frequency train of somatic spikes. This was first observed by Larkum et al. (1999a) in rat somatosensory cortex, and has since also been described in mouse primary visual cortex (Shai et al., 2015).

In these experiments, I recorded the somatic membrane potential from ttL5 neurons in V1 and V2m and evoked three action potentials using 3 ms wide pulses of somatic current injection at frequencies ranging from 50 Hz to 200 Hz in 10 Hz increments (Figure 6.4). In V1 neurons, increasing the somatic spike frequency above a critical value typically resulted in a sudden increase in the somatic ADP. However, when recording in V2m under the same experimental conditions, there was usually no change in ADP, even at stimulation frequencies as high as 200 Hz.

To quantify this effect, I aligned the peaks of the last spike for each frequency and measured the area of the ADP difference between the 50 Hz trace and the higher frequency traces in a 20 ms window (4–24 ms) following the last spike (Figure 6.4a, inset). This measure of ADP increased sharply above the critical frequency and was often largest around the value of this frequency (Figure 6.4b). Critical frequency was defined here as the lowest frequency at

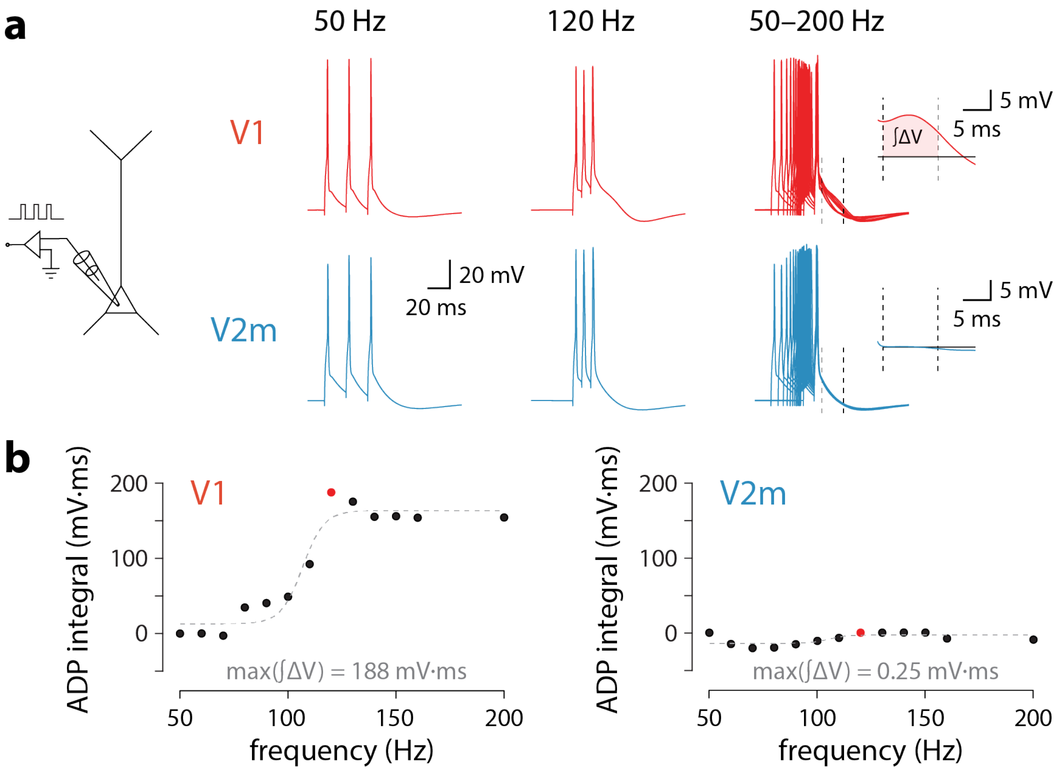

Figure 6.4: Critical frequency ADP. a. Left: diagram of experimental configuration. Right: example traces of V1 (red) and V2m (blue) ttL5 neurons stimulated with spike trains between 50 Hz and 200 Hz. Note the sustained after-depolarization following the 120 Hz spike train in the V1 neuron. Inset: ADP measured as the area between the 50 Hz and 120 Hz traces following the last spike. **b.** Quantification of ADP area at each measured frequency for the example neurons in *a*. The peak integral value is highlighted in red.

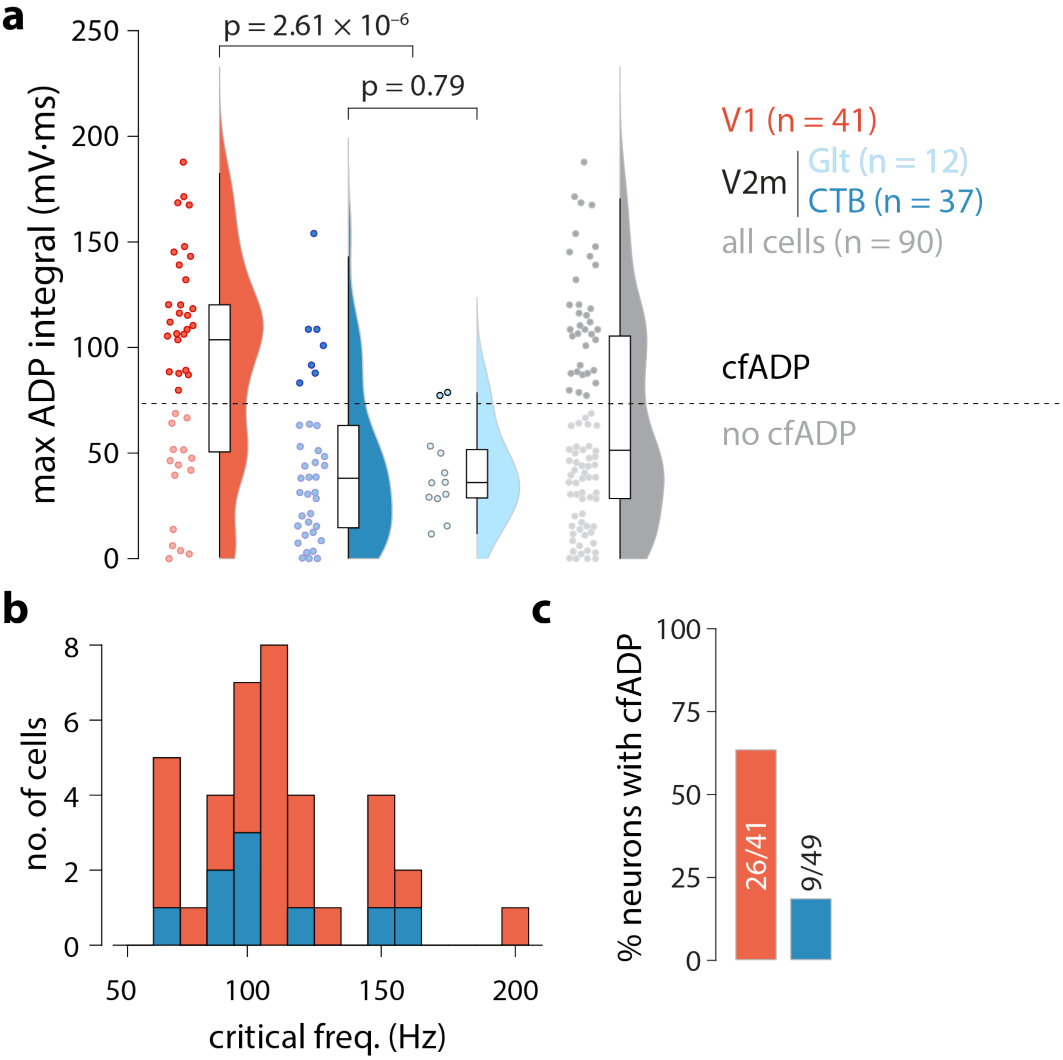

Figure 6.5: cfADP summary data. a. Summary data of peak ADP integral values for every recorded neuron in V1 and V2m. The dashed line indicates the division between the two groups of cells classified through k-means clustering. **b.** Histogram of critical frequencies across all recorded neurons, coloured by area (red: V1; blue: V2m). Cells with no identifiable critical frequency were excluded from this plot. **c.** Proportion of cells with cfADP in V1 (red) and V2m (blue).

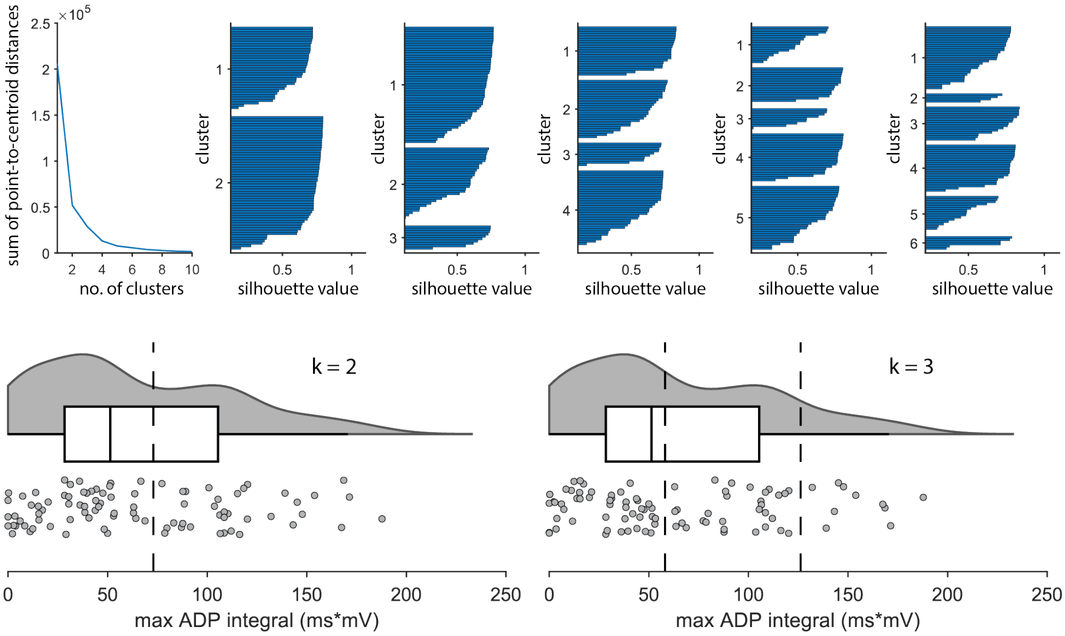

Figure 6.6: Clustering of maximum ADP integrals. *Top:* Analysis of cluster consistency using the elbow and silhouette methods. *Bottom:* Classification boundaries for the k-means algorithm applied to the pooled maximum ADP integral values, when defining either k = 2 or k = 3 clusters.

which a substantial increase in the ADP could be observed. The mean critical frequency cross all cells in both V1 and V2m (excluding cells that did not have a critical frequency) was 110.8 ± 4.9 Hz ($n = 37$, Figure 6.5b). The maximal value of the ADP measure across all frequencies for each cell is shown in Figure 6.5a. Neurons in V2m, including both Glt and CTB labelled cells, had significantly smaller ADP area (V1: mean = 90.9 ± 7.8 mV·ms, $n = 41$; V2m: mean = 42.2 ± 4.8 mV·ms, $n = 49$; $p = 2.61\times10^{-6}$, Mann–Whitney U test), reflecting that most of these cells lacked a critical frequency altogether. Within V2m, Glt and CTB-labelled ttL5 neurons did not differ significantly (Glt mean = 40.6 ± 6.1 mV·ms, $n = 12$; CTB mean = 42.7 ± 6.0 mV·ms, $n = 37$; $p = 0.79$, Mann–Whitney U test).

In order to make it possible to directly compare the number of neurons that had a critical frequency ADP (cfADP) to the number of neurons that displayed BAC firing, I split the recorded cells into two groups, representing cell with and without cfADP. In order to minimize any subjective bias in this categorization (particularly for ambiguous cells, with small but nonzero ADP),

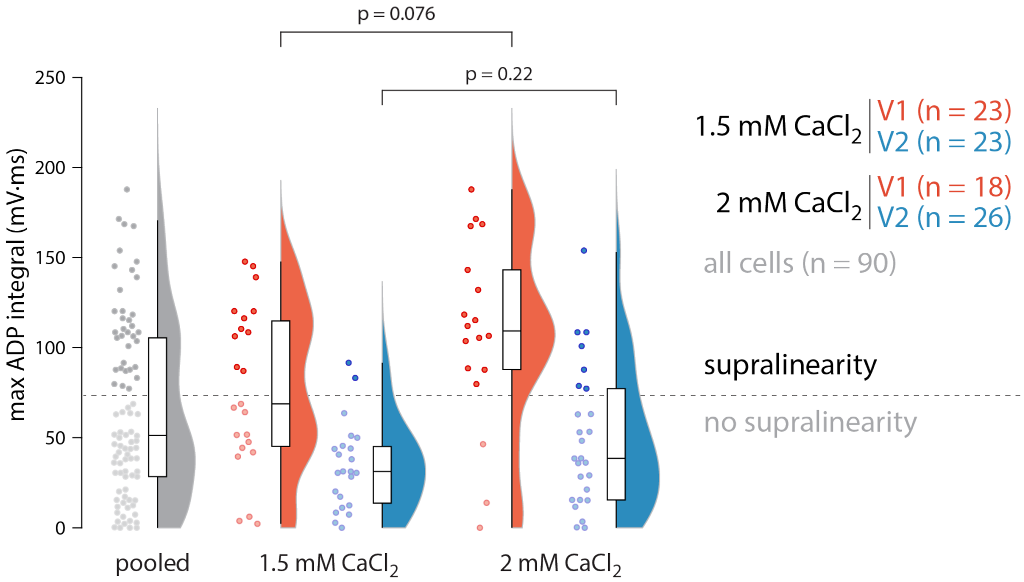

Figure 6.7: Critical frequency ADP across $CaCl_2$ concentrations. Maximum ADP integral for all cells split by recording ACSF containing either 1.5 or 2 mM $CaCl_2$. The difference between the two conditions was not statistically significant in either V1 or V2m ($p > 0.05$, Mann–Whitney U test).

I pooled the unlabelled ADP values from all cells (across both V1 and V2m) into a single list of number. I then applied the k-means clustering algorithm (using the built-in MATLAB function kmeans(X,k) with k = 2) to this list of numbers, and thereby obtained the same list of number split by their belonging to the group with either high or low ADP (i.e. with or without cfADP).

Although the purpose of this analysis was to obtain an unbiased split of the cells into two groups, with and without cfADP, the underlying data was also consistent with there being only k = 2 clusters in the ADP values. This was evaluated using the two most common metrics for determining number of clusters, namely the "elbow method" and the "silhouette value" (Figure 6.6). According to the first method, the number of clusters can be found at the sharpest inflection point when plotting the sum-of-distances to the cluster centroids. The silhouete value is instead evaluated based on the grouping that gives the greatest uniformity in silhouette value and cluster size. In addition to both of these metrics pointing to k = 2 as the optimal number of clusters, visual inspection of the distribution of ADP values also confirms that this split gives a boundary located on the trough of the bimodal distribution.

With this classification, the proportion of neurons with cfADP in V1 was 26/41, while in V2m this was more than three times lower at 9/49 ($p = 2.6\times10^{-5}$, Fisher's exact test). For consistency with previous experiments described in Chapter 3 and Chapter 5 and to facilitate comparisons to similar previously published experiments, the critical frequency ADP experiments were done with extracellular ACSF containing both 1.5 and 2 mM $CaCl_2$. Although some change in the magnitude of Ca^{2+} currents is expected, at the level of maximum ADP integral values there was no statistically significant difference between the two conditions in either V1 or V2m (Figure 6.7). As both p-values are close to significance, it is plausible that a small effect of $CaCl_2$ concentration would be detectable given sufficiently large sample size. However, as no statistically significant difference was found between these recording conditions with this sample size, and given the relatively small effect size of

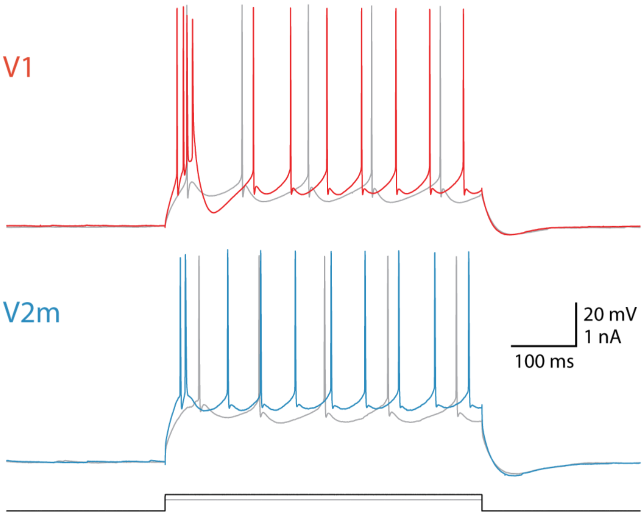

Figure 6.8: Long current steps in V1 and V2m neurons. Representative voltage traces for V1 and V2m neurons in response to 500 ms wide depolarizing current steps. The two example neurons were recorded from the same animal and both had the same rheobase (200 pA). For each cell the traces show the responses to stimulation at 60 pA (grey) and 180 pA (coloured) above rheobase.

$CaCl_2$ concentration compared to the much larger difference between areas, the data presented in Figure 6.5 contains cells pooled from both concentration conditions.

In both the BAC firing and ADP experiments described above, bursting was typically also apparent in the spiking response to a long (500 ms) depolarization at the soma. While all ttL5 neurons are generally characterized by a spike doublet at the beginning of the current step, in bursting neurons there is also a critical current step above which the initial spike burst is substantially larger, usually with 3 or 4 spikes and a deeper afterhyperpolarization (Figure 6.8).

Overall, these results show a much-diminished dendritic excitability, and as such different integrative properties, in V2m ttL5 neurons compared to V1 ttL5 neurons under the same conditions and in the same operational ranges. Previous research has indicated the length of the apical trunk as a possible factor involved in determining the dendritic excitability of ttL5 neurons in V1 (Fletcher and Williams, 2019). I therefore reconstructed the apical trunk of 22 neurons in V1 and 26 neurons in V2m from those recorded. Apical trunk

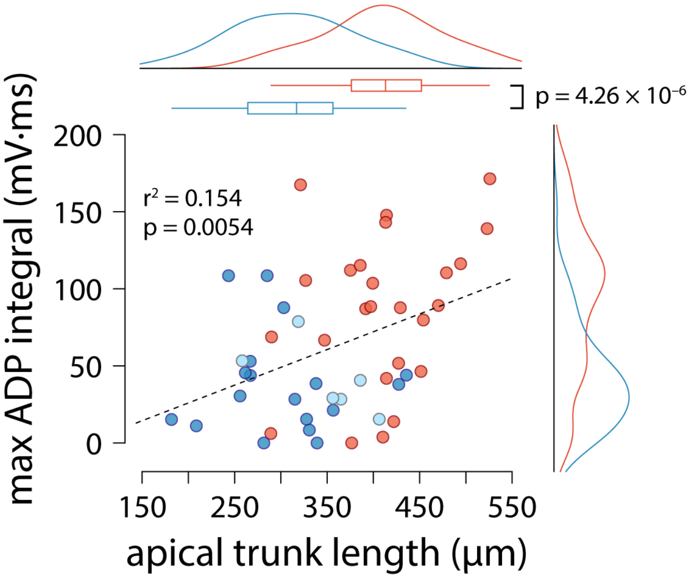

Figure 6.9: Correlation between dendritic length and excitability in V1 and V2m. Length of the apical trunk (soma to main bifurcation) plotted against the corresponding maximum ADP integral values for neurons in V1 (red) and V2m (Glt in cyan, CTB in blue). Dashed line shows the least squares linear fit; curves at the top and right are kernel density plots of the two variables in V1 and V2m (pooled across Glt and CTB neurons).

lengths were significantly shorter in V2m than in V1 (V1: mean = 409 ± 13 µm, n = 25; V2m: mean = 313 ± 13 µm, n = 24; $p = 4.26 \times 10^{-6}$, two-sample t-test, Figure 6.9). Additionally, there was a correlation between maximum ADP integral values and apical trunk length across the two populations, which explained approximately 15% of the variance in integral values ($r^2 = 0.154$, $p = 0.0054$; F-test). However, this correlation was not present within the individual populations (V1: $r^2 = 0.087$, $p = 0.15$; V2m: $r^2 = 0.02$, $p = 0.51$; F-test).

By observing data points at intermediate lengths (around 350 µm), it is clear that this correlation only explains some of the variance in maximum ADP integrals, and that among neurons with similar trunk length the cells in V2m still have smaller maximum ADP integrals. This could be due to a variety of other unmeasured variables associated with the identity of neurons in V1 and V2m, such as differences in ion channel expression. Indeed, when fitting a linear model to determine the contribution of cell identity (i.e. V1 or V2m) to the maximum ADP integral, the explained variance was even higher than for trunk length ($r^2 = 0.257$, $p = 3.42\times10^{-7}$; F-test).

To determine if there is any interaction between these two predictors on the maximum ADP integrals, I carried out a two-way ANOVA (using the MATLAb function "anovan()") with factors for both cell identity (labelled V1 or V2m) and trunk length. While both factors individually correlated significantly with maximum ADP integral, this analysis showed no significant interaction between the two predictors ($p = 0.11$, two-way ANOVA).

6.3.3 BAC firing is absent in short biophysical models

Since the experimental results suggested that there may be a weak interaction between apical trunk length and dendritic excitability, with longer trunks potentially resulting in increased excitability, I explored this interaction further through *in silico* simulation. To investigate the possible mechanisms underlying the dependence of bursting on apical trunk length, I collaborated with Aeron Laffere to run numerical simulations in conductance based compartmental models of ttL5 neurons. We first probed BAC firing in a morphologically detailed model published by Hay et al. (2011), using the model parameters (biophysical model 3) and morphology (cell #1) favoured for reproducing BAC firing. As in the original paper, BAC firing was triggered by injecting a 0.5 nA current at the apical bifurcation coupled to a somatic action potential evoked by square-pulse current injection at the soma. Mirroring the responses seen in the subset of bursting ttL5 neurons, coincident stimulation triggered BAC firing in this detailed model (Figure 6.10a, left), as has previously been shown. We then applied the same model to a different ttL5 morphology with a shorter apical dendrite, reconstructed from a recorded Glt neuron in V2m.

The only parameter adjustment made to the detailed model when applying it to this shorter morphology was to re-specify the location of the Ca^{2+} channel hot spot around the new apical branch point (300–500 µm from the soma). The amplitude of the dendritic current injection in the short morphology (0.194 nA) was also scaled so as to obtain the same depolarization

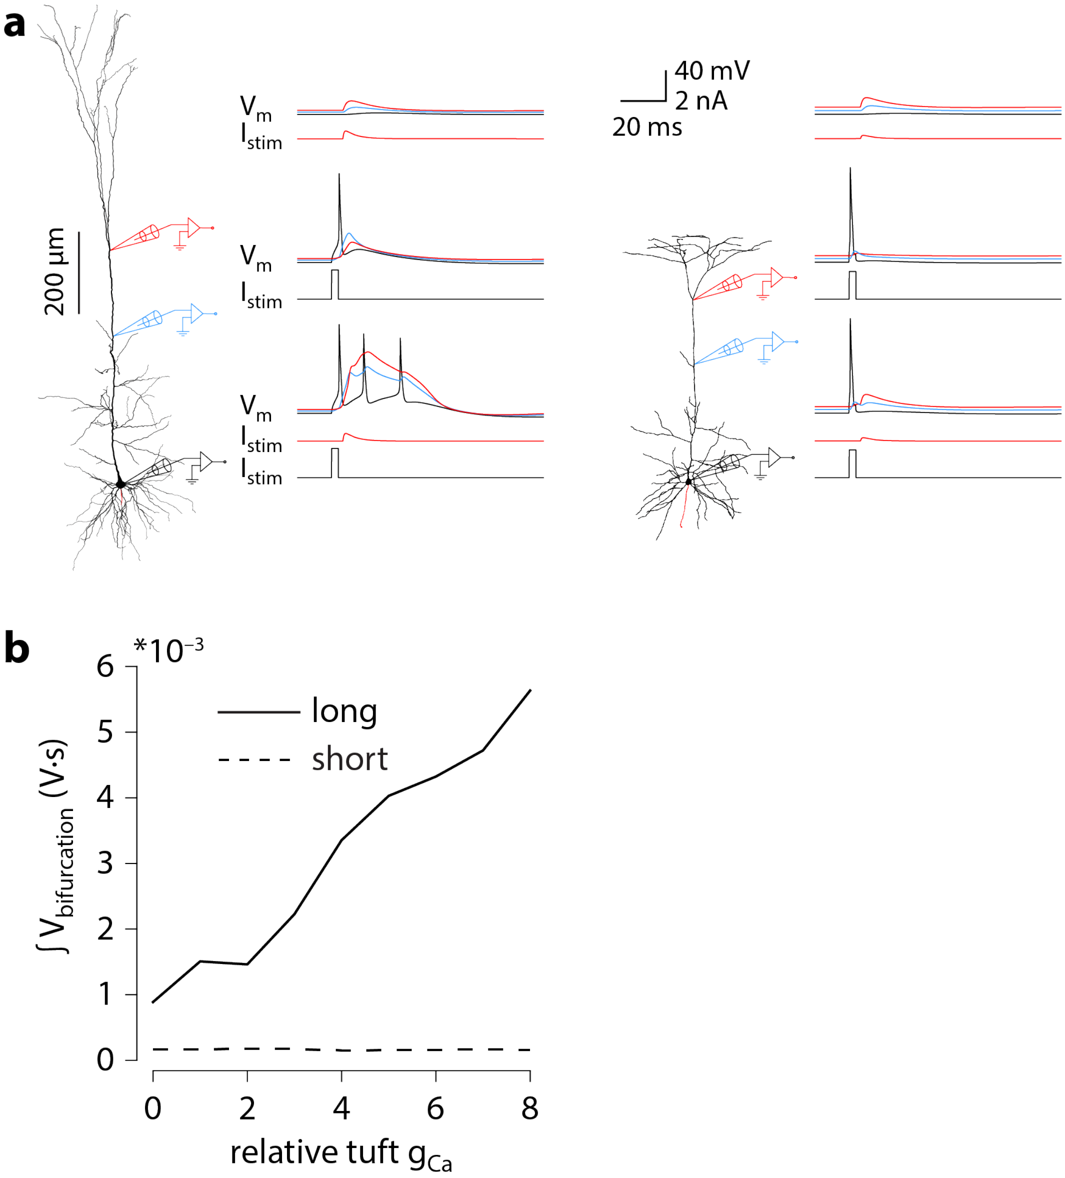

Figure 6.10: BAC firing in morphologically detailed models. a. *Left:* detailed morphology of a ttL5 pyramidal neuron from the model favoured by Hay et al. (2011) for reproducing BAC firing. *Right:* reconstructed morphology from a recorded Glt neuron in V2m. For each cell, injected current and voltage traces are shown for the soma (black), the apical trunk (blue, 400 or 200 µm from the soma), and the main bifurcation (red, 620 or 370 µm from the soma) under three different stimulation paradigms. The Ca^{2+} channel hot spot was located 685–885 µm from the soma in the long morphology and 300–500 µm in the short morphology. **b.** Integral of voltage at the branch point in the long and short model neuron during combined somatic and branch point stimulation, plotted against the relative g_{Ca}.

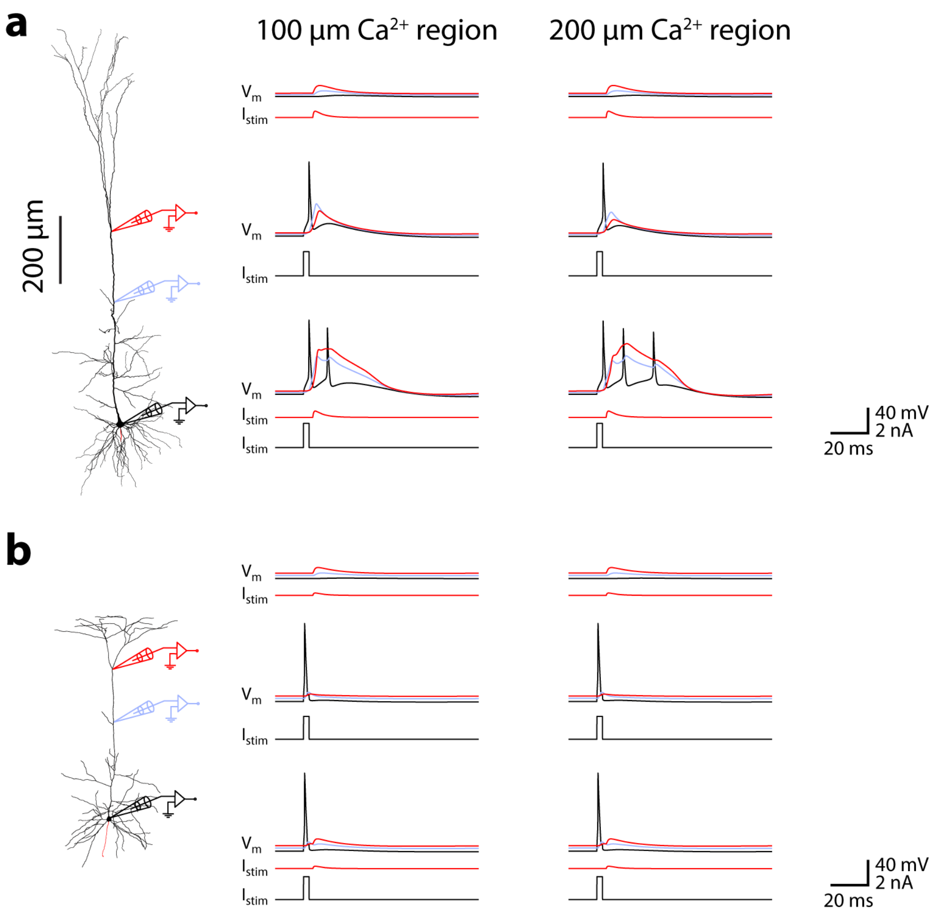

Figure 6.11: Influence of Ca^{2+} hot spot dimensions on BAC firing. a. BAC firing protocol in a model neuron with long morphology and Ca^{2+} hot spot located either 735–835 µm (left) or 685–885 µm (right) from the soma. **b.** BAC firing protocol in a model neuron with short morphology and Ca^{2+} hot spot located either 350–450 µm (left) or 300–500 µm (right) from the soma. Injected current and voltage traces shown for the soma (black), the apical trunk (blue), and the main bifurcation (red).

amplitude at the bifurcation in both model cells. With this morphology, coincident tuft and somatic stimulation evoked only a single somatic spike and did not trigger a dendritic Ca^{2+} plateau (Figure 6.10a, right).

The morphologically detailed models contained two Ca^{2+} conductances—low-voltage activated and high-voltage activated—clustered in an area around the apical branch point, which are responsible for the dendritic Ca^{2+} plateau potential. To explore the sensitivity of Ca^{2+} plateaus to the dendritic Ca^{2+} channel density in the long and short neurons, we scaled the Ca^{2+} conductance (g_{Ca}) between 0 and 8 times the original values. To minimize the number of

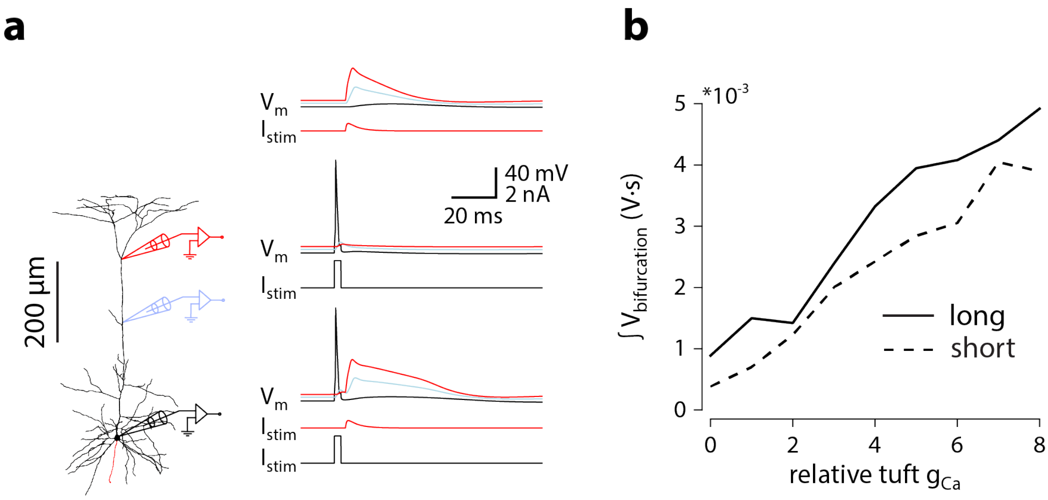

Figure 6.12: Strong dendritic current can trigger Ca^{2+} plateaus in short morphologies. a. BAC firing protocol in the model neuron with short morphology using the same dendritic current injection as in the long morphology (0.5 nA). **b.** Integral of voltage at the branch point in the long and short model neuron during combined somatic and branch point stimulation, plotted against the relative g_{Ca}.

variables, when scaling the relative g_{Ca} we kept the ratio of the two channels constant. In the long morphology the integral of the distal dendritic voltage, acting as an indicator of the large and sustained depolarization during a Ca^{2+} plateau, increased proportionally to g_{Ca}. In the short morphology, however, this value stayed constant across all g_{Ca} values (Figure 6.10b). This indicates that, although the size of a Ca^{2+} plateau depends on g_{Ca} in long neurons, in short neurons there is no Ca^{2+} channel activation and the magnitude of the voltage integral therefore does not depend on g_{Ca}.

For both the long and short morphology, BAC firing was probed with a Ca^{2+} channel hot spot size of either 100 µm or 200 µm. The size was found to not affect the response in the short neuron and had a modest effect on the size of the Ca^{2+} plateau in the long neuron (Figure 6.11). To test if the short neuron was capable of generating Ca^{2+} plateaus, we stimulated the short neuron with 0.5 nA at the dendritic electrode. The resulting dendritic potential was substantially larger and scaled with g_{Ca}, indicating that such large dendritic current injection could trigger a Ca^{2+} plateau (Figure 6.12).

However, this resulted in only a small depolarization at the soma, and even when combining it with a somatic spike it did not trigger a burst.

Biophysical models with detailed long and short morphologies thus were able to reproduce the results from my experiments. To more easily manipulate the dendritic length across a continuous range of values, we investigated the same phenomenon in a reduced ttL5 model based on Bahl et al. (2012). The simplicity of this model had the added benefit of reducing the number of

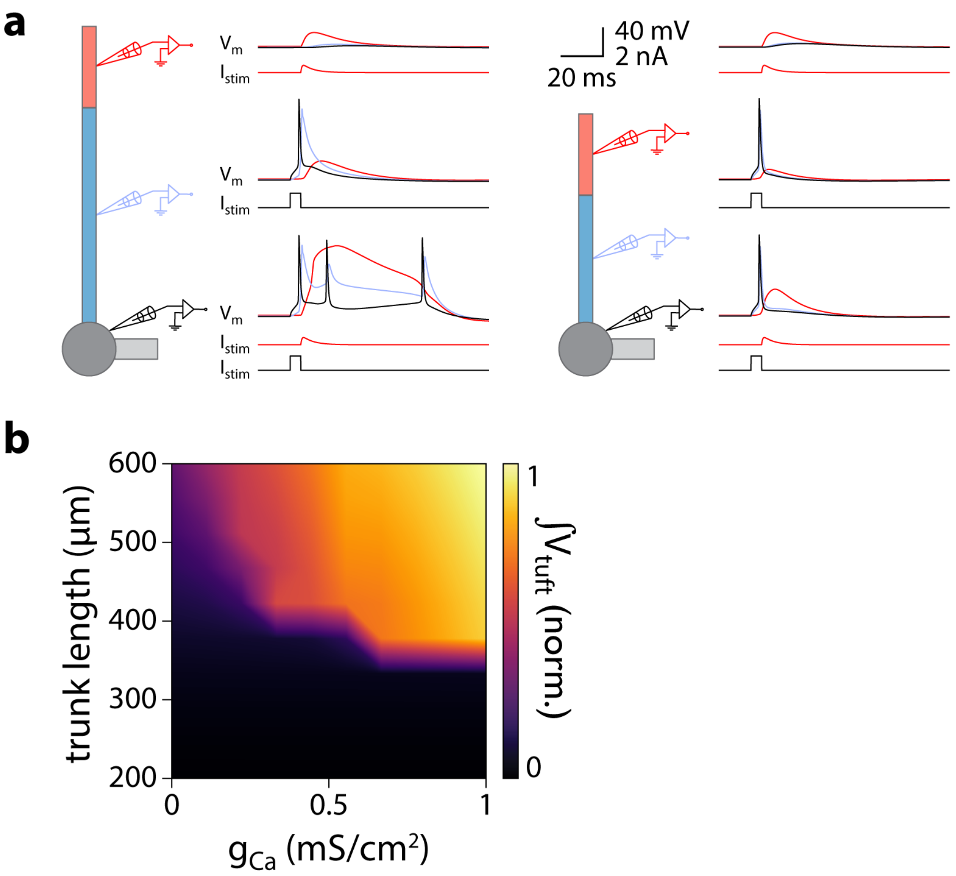

Figure 6.13: BAC firing in reduced compartmental models. a. *Left:* diagram of the reduced model neuron. The length of the apical trunk in this model was 600 µm. Voltage traces are given for three different stimulation paradigms, with recordings made at the centres of the somatic compartment (black), apical trunk compartment (blue), and tuft compartment (red). *Right:* the same current and voltage traces are shown for a version of the Bahl model modified to have an apical trunk length of 200 µm. **b.** Heatmap representing the normalised tuft voltage integral during combined somatic and tuft stimulation in the reduced model, plotted against the absolute density of Ca^{2+} channels in the tuft compartment and the length of the apical trunk compartment. Default $g_{Ca} \approx 0.45$ mS/cm^2.

variables, allowing us to more easily identify general principles of dendritic voltage propagation. One notable difference in this simplified model is that Ca^{2+} channels are all located in the tuft compartment, rather than being in a hot spot around the bifurcation. The tuft in this model should therefore be considered equivalent to a combination of the bifurcation and apical tuft dendrites.

As with the morphologically detailed model, the reduced model with the original published parameters displayed BAC firing triggered by coincident tuft and somatic stimulation (Figure 6.13a, left). Shortening the apical trunk was sufficient to eliminate this response (Figure 6.13a, right). We explored the dependence of BAC firing on apical trunk length and calcium conductance while

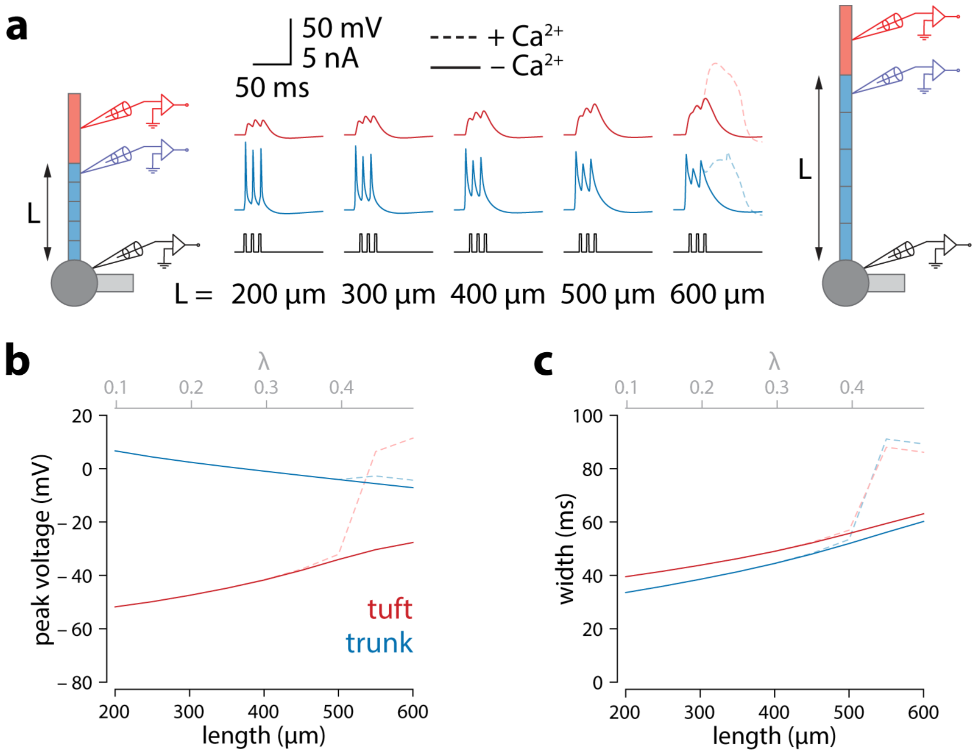

Figure 6.14: Tuft voltage increases with trunk length. a. Schematic of the simulation: stimulation site in the somatic compartment, and recording sites at the distal end of the apical trunk (blue) and in the centre of the tuft (red). Stimulus shown in black. Solid lines: $g_{Ca} = 0$; dashed lines: original g_{Ca}. **b.** Peak voltage reached in trunk (blue) and tuft (red) for a range of simulations with different trunk lengths, using the same stimulation protocol as in *a*. **c.** Same as in *b*, but plotting the width of the depolarization, measured 2 mV above baseline. Length constant $\lambda = 1009\,\mu m$.

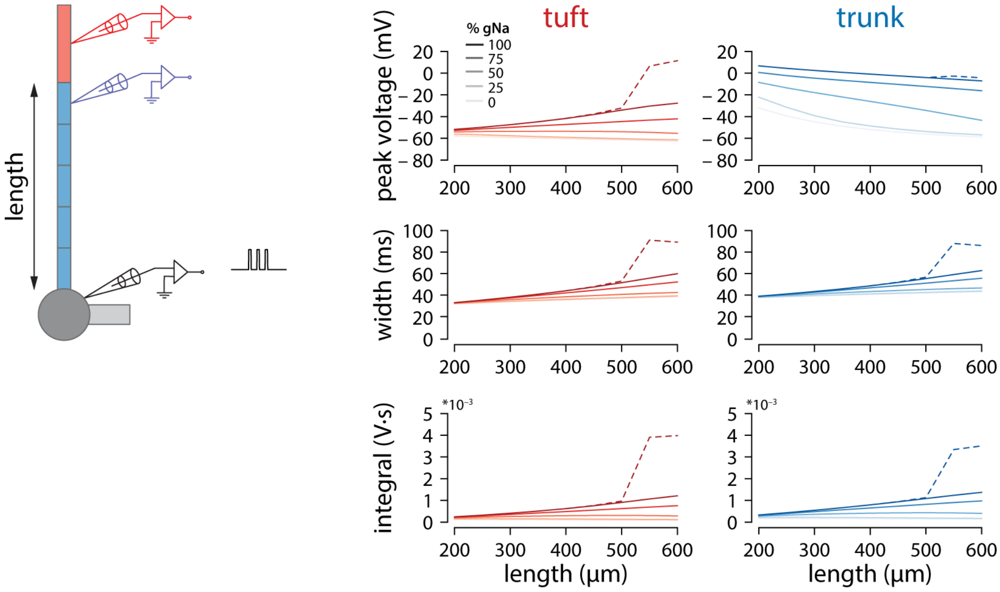

Figure 6.15: Length-dependent backpropagation depends on dendritic Na^+ channels. *Left:* Schematic of the simulation: stimulation site in the somatic compartment, and recording sites at the distal end of the apical trunk (blue) and in the centre of the tuft (red). Stimulus shown in black. *Right:* Peak voltage, width and integral values measured in the trunk and tuft for dendrites containing different Na^+ channel densities in the apical trunk. The colour saturation of different lines indicates channel density as a percentage of the original values (from 100 to 0 %).

measuring the time-integral of tuft voltage as an indicator of Ca^{2+} plateau potentials (Figure 6.13b). The presence of a Ca^{2+} plateau depended strongly on apical trunk length and was only sensitive to g_{Ca} above a critical length of approximately 350 µm ($\approx 0.35\ \lambda$). Below this length, no Ca^{2+} plateaus were triggered regardless of how high g_{Ca} was set to. These experiments show that a simplified model can also reproduce my experimental results, allowing us to explore and dissect the underlying parameters in more detail.

6.3.4 Active propagation enhances tuft voltage in neurons with long apical dendrites

To obtain a mechanistic understanding of what causes the length dependence of bursting, we studied the voltage in the final segment of the apical trunk as well as the apical tuft in the reduced model during backpropagation from the soma. We recreated the experimental paradigm used in Figure 6.4 by triggering 3 spikes at 100 Hz through somatic current injection. As with co-

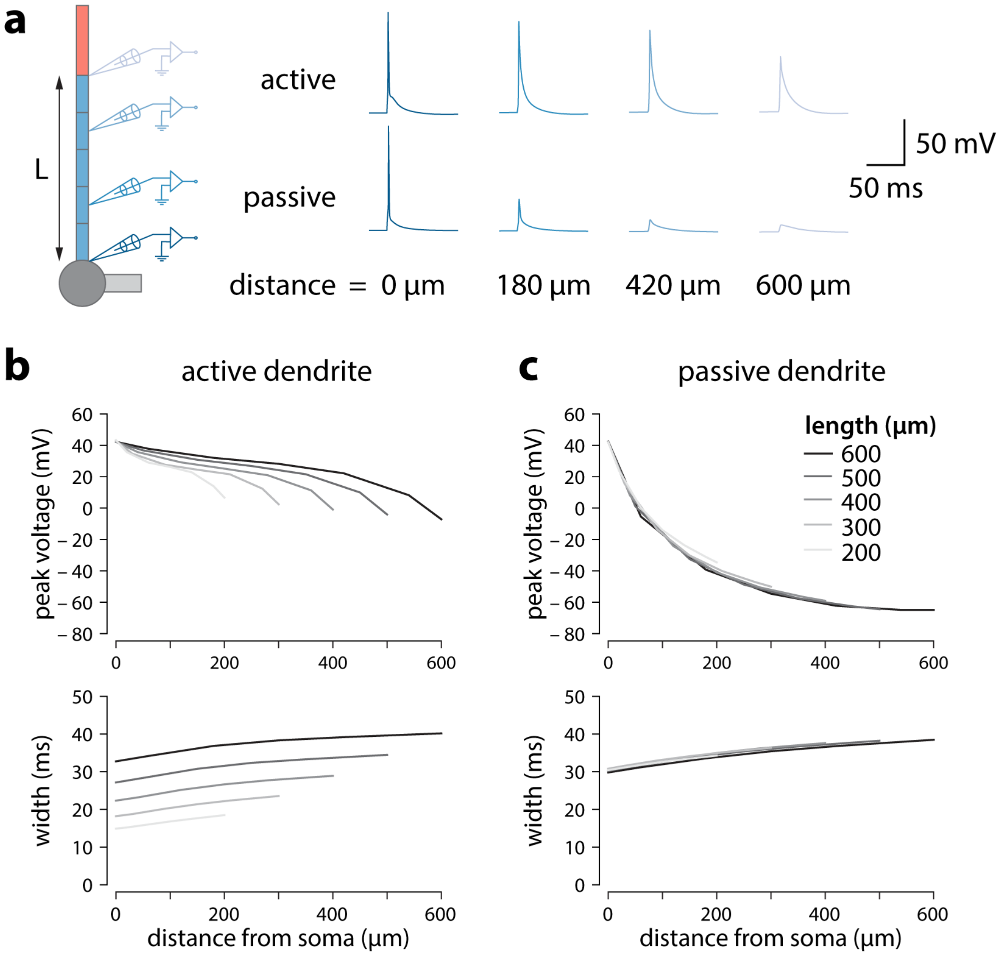

Figure 6.16: Backpropagation of single spikes in active and passive trunks of different length. a. Backpropagation of a somatic spike elicited through a single 3 ms wide 2 nA square current step at the soma in a model neuron with 600 µm apical trunk length. Voltage recordings were made at different distances along the trunk. **b.** Peak voltage and width measured at different absolute distances (same relative) for active model neurons. Width was measured as the interval between the voltage values 2 mV above baseline membrane potential. Length (L) refers to the length of the apical trunk. N.B., at any given absolute distance from the soma, peak voltage and width of the backpropagating spike are larger when the apical trunk is longer. **c.** Same as in *b* but with all voltage-dependent conductances removed from the trunk and tuft compartments. N.B., voltage attenuation is independent of trunk length.

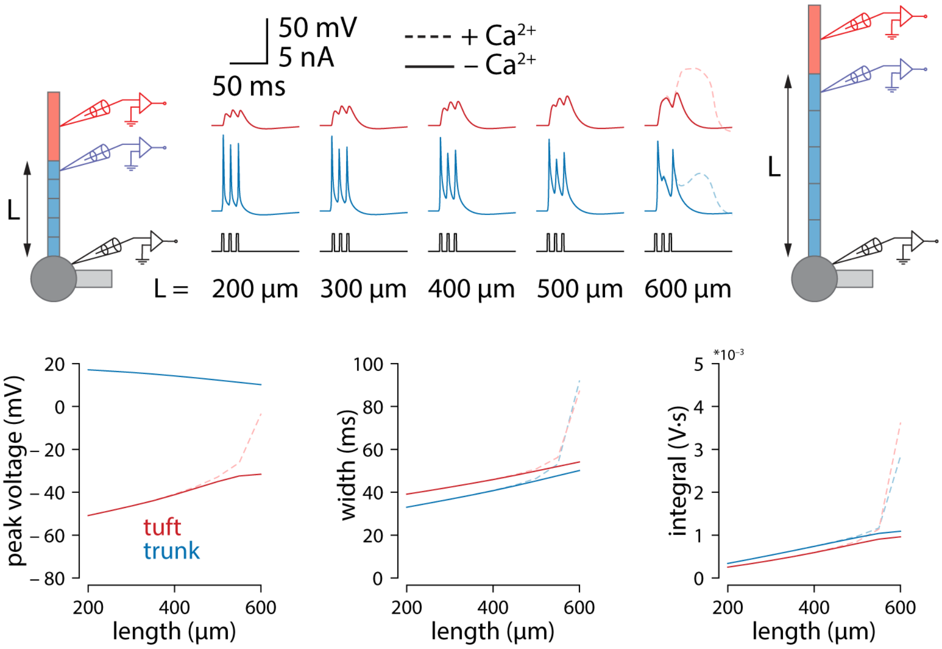

Figure 6.17: Tuft voltage increases with trunk length independently of conductance gradients. Same experiment as in Figure 6.14a but with uniform distribution of all active conductances in the apical trunk. Total conductance was kept constant for each channel.

incident backpropagating spike and tuft input, increasing the length of the apical trunk facilitated dendritic Ca^{2+} plateau initiation (Figure 6.14a). Upon closer inspection, the width and peak voltage in the tuft steadily increased with dendritic length (Figure 6.14b,c), even in the absence of Ca^{2+} currents ($g_{Ca} = 0$). In the presence of voltage-gated Ca^{2+} channels, this increased amplitude of backpropagating spikes triggered a larger all-or-none Ca^{2+} plateau above a certain threshold length.

We found that backpropagating spike amplitude in the tuft increased as a function of apical trunk length despite a decreasing backpropagating spike amplitude in the distal segment of the trunk (Figure 6.14b). We also observed that the width of the backpropagating spike (measured 2 mV above baseline) increased with length in both the tuft and trunk (Figure 6.14c). While waveform broadening is a natural consequence of passive filtering along dendrites, the sustained voltage in the distal trunk required active dendritic propagation. In the reduced model, this active propagation in the apical trunk was mediated primarily by voltage-gated Na^+ channels. Removing these channels caused a substantial reduction in peak voltage and width of the depolarization in the distal trunk, and importantly also abolished the trend of increasing tuft voltages with longer dendritic trunks (Figure 6.15).

More generally, active propagation caused backpropagating spikes to be larger and broader at all distances along a long dendrite compared to the same absolute distances in shorter dendrites (Figure 6.16). Because of this, when comparing the final positions along the trunk, the peak voltage is only marginally smaller in long dendrites despite the larger distance from the soma. This is not the case in a passive dendrite, where voltage attenuation depends only on distance from the soma and the attenuation at any give location is not sensitive to the total length of trunk beyond this distance.

The general phenomenon of enhanced voltage propagation in longer dendrites resulting in amplification of tuft voltage was not sensitive to the specific distribution of active conductances implemented in this model. Indeed, when all conductances were uniformly distributed along the apical trunk, the waveforms did not substantially change and the enhanced voltage continued to trigger Ca^{2+} plateaus only in neurons with long apical trunks (Figure 6.17). Enhanced voltage propagation was also insensitive to I_h, which is a major regulator of dendritic excitability, indicating that the effect likely requires regenerative excitatory conductances such as the Na^+ current (Figure 6.18).

While it might seem counter-intuitive that peak voltage in the tuft increases when the peak trunk voltage is decreasing, we propose that the increasing width of the depolarization can at least partially account for this via a passive mechanism. Wider depolarizations allow the tuft compartment to charge to a higher voltage. The rate and peak value of tuft charging depends on the passive properties of the tuft. The peak value of depolarization reached and the rate of voltage change are proportional to membrane resistance (R_m) and membrane capacitance (C_m), respectively. The product of these two parameters gives the membrane time constant (τ).

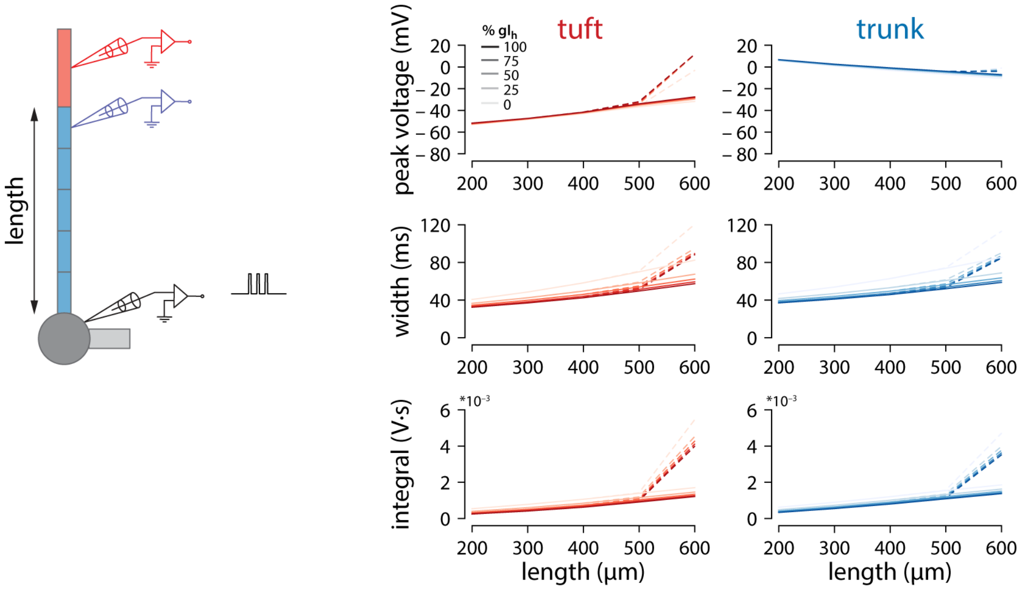

Figure 6.18: Length-dependent backpropagation is independent of I_h. *Left:* Schematic of the simulation: stimulation site in the somatic compartment, and recording sites at the distal end of the apical trunk (blue) and in the centre of the tuft (red). Stimulus shown in black. *Right:* Peak voltage, width and integral values measured in the trunk and tuft for dendrites containing different HCN (I_h) channel densities in the apical trunk (as a percentage of the original values).

To illustrate this, we applied voltage clamp to the end of the distal segment of the trunk and delivered 30 mV square voltage pulses of increasing width (Figure 6.19a). Due to the capacitive filtering of the tuft, short voltage steps did not fully charge the tuft while wide voltage steps allowed the tuft voltage to reach the steady-state values commanded by R_m. To directly test the hypothesis that the relationship between trunk depolarization width and tuft membrane time constant caused the backpropagating spike amplitude in the tuft to increase with length, one could vary R_m by changing g_{leak}. However, this would affect resting membrane potential and consequently alter voltage-dependent properties of the tuft. We therefore chose to vary C_m instead, in order not to affect other variables in the model. For a given value of depolarization amplitude and width, increasing C_m (and therefore τ) in the tuft caused a reduction in the peak tuft voltage (Figure 6.19b). These simulations show that the tuft time constant and the width of the backpropagating spike interact to create a higher tuft depolarization with longer apical trunks.

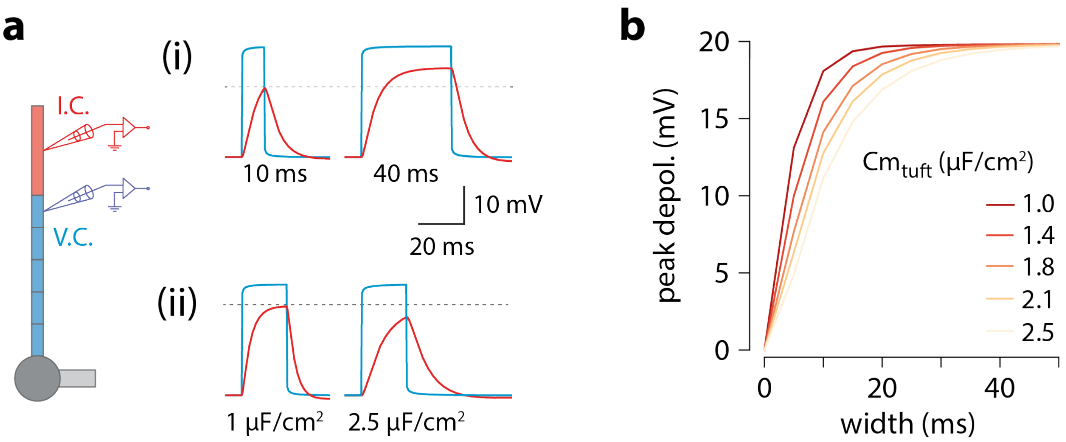

Figure 6.19: Passive effect of increasing trunk depolarization width. a. Example traces during voltage-clamp of distal trunk. Different width voltage steps (i) or same width under differing tuft membrane capacitance conditions (ii) were injected into the trunk (blue). The recorded tuft voltage is also shown in red, Vrest = -65 mV. **b.** Peak voltage values reached in the tuft for a range of trunk step widths and tuft capacitances. In the original model, tuft $C_m \approx 1.75\,\mu F/cm^2$.

It has previously been suggested that axial resistance (R_a) in the apical dendrite may influence the backpropagation efficacy in dendrites and burstiness of ttL5 neurons (Fletcher and Williams, 2019). To test this hypothesis, we measured peak voltage and width in the trunk and tuft for different trunk lengths under different R_a conditions. We found that peak tuft voltage (and therefore burstiness) increased with increasing trunk R_a, reaching the highest voltage near the reduced model's original value of R_a, and decreasing again for higher values (Figure 6.20). However, in these simulations burstiness always increased with trunk length regardless of R_a. This indicates that, although important, it was not the primary determinant for generating the length-dependent effect.

Overall, the combination of increased width and a relatively small reduction in amplitude resulted in a trunk voltage integral that increased with trunk length, thereby passing more charge to the adjacent tuft compartment. However, if active backpropagation was reduced or absent, the trunk integral and resulting tuft voltage decreased with length (Figure 6.15). The peak tuft voltage approximately followed the integral of voltage in the distal trunk. To illustrate this, we applied voltage-clamp to the end of the trunk and injected

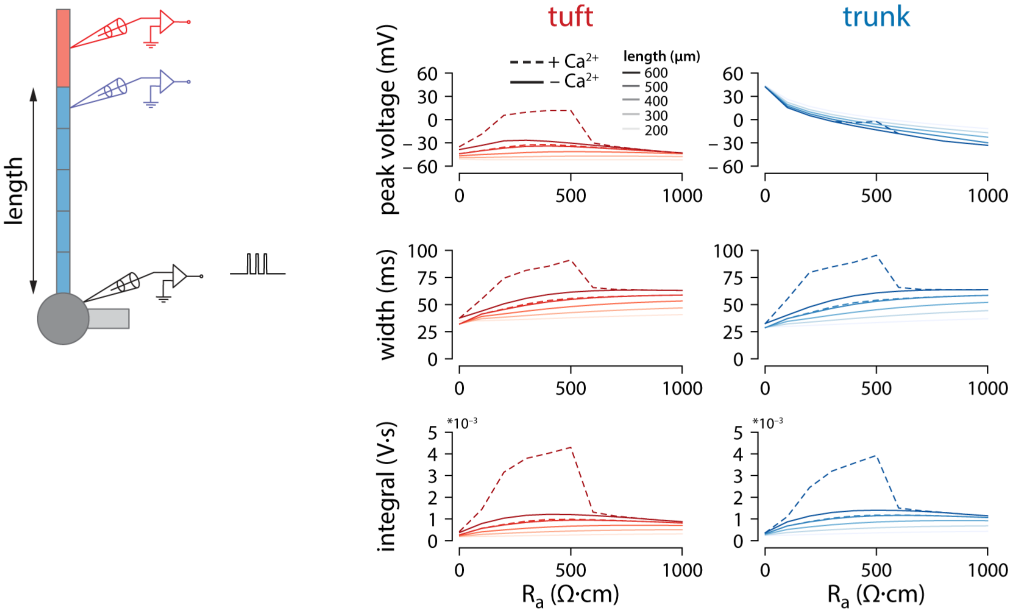

Figure 6.20: Effect of axial resistance on voltage propagation. Plots of peak voltage, width, and voltage integral reached in the tuft and trunk for varying values of trunk length and axial resistance (R_a). Default $R_a \approx 382.22\,\Omega{\cdot}\text{cm}$. The stimulus and recording conditions were the same as in Figure 6.14a, with 3 APs at 100 Hz triggered in the somatic compartment. Solid lines show simulations with $g_{Ca} = 0$, while dashed lines show the same simulations with the original $g_{Ca} = 0.45$ mS/cm^2.

square steps with a range of integrals obtained through various combinations of width and amplitude (Figure 6.21a). This revealed a zone above a critical trunk integral for which many different combinations of width and depolarization amplitude were sufficient to evoke a Ca^{2+} plateau in the tuft (Figure 6.21b).

6.4 Discussion

While much is known about the Ca^{2+}-mediated supralinearities of ttL5 neurons (Groh et al., 2010; Kasper et al., 1994; Kim et al., 2015; Larkum, 2013; Ramaswamy and Markram, 2015), our existing knowledge is mainly based on experiments done in V1 and primary somatosensory cortex of rats. Previous studies often selected cells based primarily on soma size and approximate shape of the dendrites, potentially introducing biases in the selection which could skew the results towards specific subtypes of L5 neurons. This may be particularly relevant in light of recent evidence showing that there are functional differences between different projection-defined populations in L5 (Kim et al., 2015; Lur et al., 2016). I sought to overcome some of these limitations by selecting cells using more easily reproducible criteria, selecting cells in L5 uniformly based on either their Glt-positive labelling or on their projection to LP, which I determined using retrograde CTB labelling. In the subset of the recorded neurons which were successfully filled with biocytin, I was able to confirm that these cells also had the characteristic morphological features of ttL5 neurons. I was thus able to maintain cortical area as the primary variant when comparing V1 and V2m neurons.

My results show that, contrary to common assumptions, there are considerable differences in the properties of ttL5 neurons across different brain regions. Indeed, both BAC firing and critical frequency ADP, which are considered hallmarks of dendritic Ca^{2+} plateaus, were found to be less common in V1 than I anticipated based on the results from rat somatosensory cortex, and were almost completely absent in V2m. This is most likely due to intrinsic properties of the cells in these areas, rather than on differences in local circuitry or tonic inhibition, as the extracellular stimulation to evoke BAC firing was done in the presence of a $GABA_B$ antagonist and in the critical frequency experiments Ca^{2+} plateaus were evoked solely by intracellular stimulation through the soma.

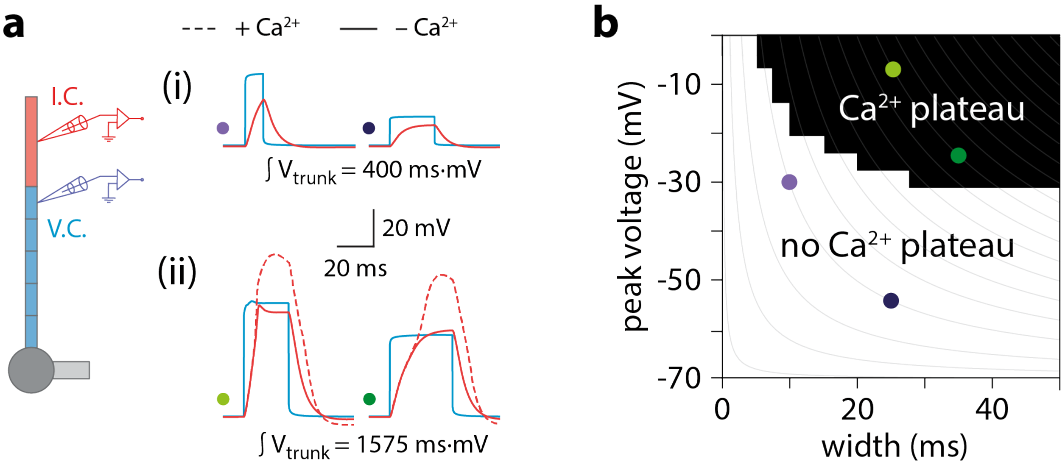

Figure 6.21: Peak tuft voltage depends on trunk integral. **a.** Example traces during voltage-clamp of distal trunk (blue), showing voltage steps of different width and amplitude but with the same voltage integral in the trunk. For smaller integrals (i) the voltage remains subthreshold while for larger integrals (ii), the tuft voltage crosses the threshold for a Ca^{2+} plateau. **b.** Voltage and width combinations for square voltage steps in the distal trunk which result in a Ca^{2+} plateau in the tuft. Coloured dots represent the value combinations illustrated in **a**. Grey lines indicate width and amplitude combinations with equal integral.

My experiments thus provide evidence against the commonly cited notion that the neocortex is composed of canonical circuits performing the same fundamental computations on different sets of inputs across different brain areas (Mountcastle, 1997; Douglas and Martin, 2004; Harris and Shepherd, 2015; Hawkins et al., 2019; Jiang et al., 2015; Markram et al., 2015; Miller, 2016). Instead, it implies that the non-linear computations performed through Ca^{2+} plateaus and BAC firing is not required in more associative areas outside of the primary sensory cortices. This may have to do with the cortical hierarchy being less clearly defined in these regions, making it more useful to maintain equal weighting between different sensory modalities and rely on other mechanisms to change the weights according to the reliability of each input.

Although all cells were selected according to the same criteria, I found that the ttL5 neurons in V2m had significantly shorter apical trunks than the V1 neurons. This is consistent with recent structural MRI data showing reduced neocortical thickness in the most caudal and medial portions of the rodent brain (Fletcher and Williams, 2019). It is worth noting, however, that the correlation observed between apical length and dendritic excitability (measured though the critical freqiency ADP integral), only explained a relatively small percentage of the observed variance, and differences between V1 and V2m neurons are still present when only considering neurons with similar trunk lengths. Other factors that I didn't measure in these experiments, but which could influence dendritic excitability, include the thickness of the apical dendrite, the location and size of oblique dendrites along the trunk, and differences in gene expression (particularly controlling the production and trafficking of transmembrane ion channels, such as voltage-gated Na^+ and Ca^{2+} channels).

In the BAC firing experiments using extracellular stimulation, the neurons in V2m did not produce spike bursts despite making every attempt to make the conditions for BAC firing as favourable as possible, by reducing tonic inhibition and turning the stimulation currents as high as I could without directly eliciting a spike doublet through L1 stimulation alone. However, although I equalized the conditions between recorded neuron in V1 and V2m as much as possible, it is still plausible that the difference could be partially explained by factors that are not intrinsic to the neurons. For example, if the inhibitory receptor subtypes responsible for inhibiting the apical dendrite differ between V1 and V2m, the $GABA_B$ antagonist I used may have been less effective in V2m, and the extracellular stimulation in V2m may thus have recruited stronger inhibitory circuits that prevented BAC firing.

Furthermore, even with the intrinsic difference between the properties of ttL5 neurons in V1 and V2m, the actual computations that these neurons perform *in vivo* could still be approximately the same. For instance, differences in neuromodulatory circuits between the two areas could compensate for the intrinsic differences in the neurons, such that BAC firing could be enabled in V2m neurons given the right conditions.

Despite all the possible unexplored parameters which could in theory partially account for differences between V1 and V2m, the simplest explanatory feature that was apparent from the experiments I did was the length of the apical trunk. While directly manipulating dendritic length is experimentally unfeasible, the mechanistic role of this parameter could easily be explored in a compartmental model. Applying an existing biophysical ttL5 model, which was designed to reproduce classic ttL5 properties such as BAC firing (Hay et al., 2011), to a morphology with a shorter apical trunk resulted in a loss of BAC firing independently of Ca^{2+} channel density. To determine the subcellular dynamics causing this effect, we used a reduced ttL5 model with a simplified morphological structure (Bahl et al., 2012) that could be more easily manipulated and where the effects would be more interpretable. As in the morphologically detailed model, the BAC firing response which was present in the original model disappeared when shortening the apical trunk. Importantly, in these simulation the only physical change that we made in the morphology was in trunk length. Other variables, such as dendrite diameter, were kept constant and can thus not explain the observed effects (although this does not exclude the possibility that dendrite diameter could itself be important for dendritic excitability). Because of the structural simplicity of this model, we were able to study this effect across a wide range of trunk lengths and thereby identify a sharp length cut-off at around 0.35 λ ($\approx 350\,\mu m$ in model space), below which BAC firing was abolished. A notable simplification in the reduced model is that the apical trunk and oblique dendrites are compressed into the same compartment. The model parameter fits were also based on recordings from rat neurons. Therefore, the numerical values of model length do not translate directly into apical trunk lengths for real mouse neurons. It is also worth noting that, as the reduced model does not have a distinct compartment to represent the apical bifurcation (where the Ca^{2+} channel hot spot is located in the morphologically detailed model), all Ca^{2+} channels are placed in the tuft compartment.

As backpropagation from the soma into the dendrites is a key aspect of BAC firing, this aspect was studied more closely by recording the dendritic voltage in response to backpropagating somatic spikes in models with different apical trunk lengths. Surprisingly, the voltage in the tuft did not attenuate but rather increased as a function of trunk length. This was in contrast to the trunk voltage, whose peak did indeed attenuate with increased distance from the soma. We were able to demonstrate that the tuft amplification depended on a sustained broadening depolarization in the distal trunk which was generated by voltage gated Na^+ channels present along the trunk. While this broadening was proportional to trunk length, the Na^+ channel activation reduced the attenuation of spike amplitude over distance. The combined effect of these changes caused neurons with a longer apical trunk to have a greater voltage integral in the trunk, leading to greater charging of the tuft. We hypothesise that, above a minimal threshold for peak trunk voltage, the primary determinant of peak tuft voltage is the time-averaged voltage in the trunk. Supporting this view, many different combinations of depolarization width and amplitude in the trunk were able to trigger Ca^{2+} plateaus. It is interesting to note that, in the presence of Na^+ channels, the backpropagating spike at any given absolute distance from the soma was larger and broader in neurons with longer trunks. This may be due to a cooperative effect of each trunk section on the sections both up- and downstream, with the voltage at each location decaying slower because of the more depolarised state of the remaining dendrite. This observation could potentially be verified experimentally with dendritic recordings by comparing the backpropagating spike width and amplitude at equivalent absolute distances from the soma in neurons with long and short apical dendrites. It would also be possible to measure the voltage at the apical bifurcation in response to a train of backpropagating spike in the presence of Ca^{2+} channel blockers, in order to determine if the voltage profile at the Ca^{2+} hot spot is indeed different in long neurons.

Axial resistance in the apical dendrite is another important factor influencing the electrical coupling to the soma and has been suggested to be involved in determining a difference in BAC firing between rostral and caudal

V1 neurons (Fletcher and Williams, 2019). While the tuft voltage did decrease when reducing the trunk R_a, the voltage was always larger in the neurons with longer trunk, indicating that the effects of variations in R_a and trunk length are independent. If R_a indeed correlates with length, these effects may combine to further enhance the tuft voltage in long neurons.

One notable counterexample to the principle that long neurons should be more prone to burst through enhanced backpropagation is the human ttL5 neuron, which was recently shown to have greater compartmentalization and reduced excitability compared to rat neurons, despite being substantially longer (Beaulieu-Laroche et al., 2018). However, this may still be consistent with our predictions as the paper also reported reduced ion channel densities in human ttL5 neurons, which we show to be crucial for the length-dependent enhancement. Furthermore, as the boosting effects of a broader depolarization are subject to saturating (when the depolarization is wide enough to fully charge the tuft), I would predict that the positive effect of length on tuft voltage does not increase monotonically, as above a certain apical trunk length the trunk voltage would attenuate to the point where it is no longer sufficient to trigger a Ca^{2+} plateau.

In conclusion, I have characterised previously undescribed key differences in the intrinsic integrative properties of ttL5 neurons in V1 and V2m. My results contribute new insights on the diversity of ttL5 neurons and on the possible mechanisms through which morphology can influence their integrative properties, adding to the growing body of evidence that the properties and canonical computations of cortical circuits across the brain may not be as stereotyped as is commonly believed. This cellular heterogeneity may functionally expand the ability of cortical areas to specialize in the computations that are required for processing their particular set of inputs, at the cost of reduced flexibility in generalizing to other types of input.

Chapter 7

General discussion

The cerebral cortex is a complex brain structure that is functionally diverse while also having evident architectural principles that are common throughout its different subdivisions. This is most apparent in the categories of cells that are universally found throughout the cortex (Harris and Shepherd, 2015), which form the building blocks of cortical computation and which seem to be arranged according to tightly controlled rules governing their location and connectivity, endowing each part of cortex with a complete set of the computational tools provided by the various cellular components.

While many parts of the nervous system have hard-wired circuits that are specialised for performing well-defined input-output transformations with high fidelity and low variability, such as the many reflex arcs and circuits dedicated to innate behaviours, one of the unique aspects of the cerebral cortex is its generality. This is possibly the main reason why this part of the brain has experienced the greatest expansion in volume throughout our evolutionary history (Herculano-Houzel et al., 2007; Kaas, 2008; Hofman, 2014). The ability to learn flexibly that comes with this generality enables knowledge to be passed between generations and thus overcome the information bottleneck of the genetic code (Zador, 2019). For this reason, understanding the structure and functions of the cortex and distilling the computational elements that allow it to adapt to such a wide variety of tasks may be the key to developing general artificial intelligence (Kumaran et al., 2016; Hassabis et al., 2017).

In this thesis, I have contributed to our understanding of the functional properties and computational diversity of the cerebral cortex by describing the properties of a genetically defined population of thick-tufted layer 5 pyramidal neurons in the medial secondary visual cortex of mice. While the visual cortex has historically been one of the most-studied cortical areas, the use of mice as a model system for these studies has only recently become common and the functions of higher-order visual areas are only beginning to be understood (de Vries et al., 2020). A large driver of the interest in mice for studying the cortex in general, and vision in particular, is the increasing availability of mouse lines that facilitate the targeting of specific cell types using a broad repertoire of molecular and genetic tools (Gong et al., 2003, 2007; Gerfen et al., 2013; Huang and Zeng, 2013). This has made it possible to reveal the connectivity and properties of neural circuits with unprecedented precision, and study well-defined groups of cells using complementary experimental approaches to reveal both structural and functional properties.

I have focused my experiments primarily on characterizing a population of neurons located in V2m and labelled in the Glt25d2-Cre mouse line. Using whole-cell patch clamp recordings in combination with electrical and optogenetic stimulation, I have provided a thorough characterization of these neurons, including their input and output connectivity, morphology, and a wide range of intrinsic properties. In addition to revealing several unexpected features of the connectivity in V2m and the diversity in the properties of neurons across brain regions, this multifaceted dataset constitutes a resource that future experiments can build on to understand the functional operations that occur in V2m *in vivo* and the interactions of different cell types and circuits in this area.

My results show that in V2m the Glt line predominantly labels L5 pyramidal neurons, and that these neurons specifically belong to the class of thick-tufted L5 neurons. These are characterized by large and highly arborised dendritic trees with wide tufts of dendrites in L1, axons projecting prominently to subcortical areas related to motor functions, and distinctive spiking patterns, which are most recognisable by the presence of spike ADPs, double spiking at the start of current steps, very little spike frequency adaptation, and large "sag" currents.

7.1 Hierarchical connectivity

The large dendritic trees of ttL5 neurons, spanning all layers or cortex, put them in a unique position to function as primary integrators of a wide range of local and long-range signals from across the brain, and act as the output layer of cortex by compressing and transforming these inputs into a signal that can be used to drive actions. The long-range connectivity between different brain regions is thought to be particularly important, as it establishes a hierarchy of parallel distributed processing pathways that has huge computational benefits, as is demonstrated by the advanced capabilities of comparatively simple artificial neural networks which rely heavily on this connectivity principle (Felleman and Van Essen, 1991; McClelland and Rogers, 2003; LeCun et al., 2015).

Accurately defining hierarchy between the different cortical areas is difficult, as recurrent connectivity abounds and no complete wiring diagrams exist to analyse brain-wide connectivity at a subcellular scale, due to the immense technical challenges of mapping large connectomes at high resolution (Bohland et al., 2009; Leergaard et al., 2012). While hierarchical location is commonly understood to mean processing depth in terms of distance from external inputs (such as sensory inputs from the retina), historical studies of hierarchy within the cortex have revealed that a useful proxy for this is the laminar pattern of connections between different cortical areas. Specifically, at least within the visual pathways in the brain the general principle that is consistently seen is that lower visual areas send their axons broadly to the middle layers of higher visual areas. These higher areas in turn are reciprocally connected to the lower areas and send their axons mostly into L1 and L5 of the lower area.

Aside from making it possible to assign a hierarchical position to “deep” cortical areas, whose inputs cannot be easily traced back to any one sensory input (Harris et al., 2019), this heuristic also has important implications for how the long-range signals connect to and are processed by the large number of different cell types that exist within a given cortical area.

Identifying functional synapses is generally challenging compared to the simpler task of characterizing dendritic and axonal morphology. Therefore, while direct knowledge of the dendrite-level connectivity of different pathways and cell types is limited, another useful rule that is frequently employed is axon-dendrite co-localization as a proxy for actual connectivity—a principle commonly known as “Peters’ rule” (Peters and Feldman, 1976; Rees et al., 2017).

For example, this rule is thought to be largely accurate in describing connections to L1, which predominantly receives long-range input from other cortical areas despite having a very low density of neurons, consisting mainly of a small number of interneurons (Markram et al., 2004). Projections to L1 thus largely connect with the apical dendrites of pyramidal neurons in deeper layers, and particularly cells such as ttL5 neurons, whose apical dendrites branch extensively in L1 (Spratling, 2002; Larkum, 2013).

To directly confirm or refute if Peters’ rule holds true for any given connection, it is necessary to test that connection by either stimulating the axons and recording postsynaptic responses or imaging the synapses using electron microscopy. While early experiments using electrical stimulation were able to confirm that some of the axons in L1 made functional synapses with ttL5 dendrites, it is only in recent years that it has been possible to refine this view with specific labelling of identified pathways (Cauller, 1995; Petreanu et al., 2009; Ohno et al., 2015; Rees et al., 2017).

Through subcellular optogenetic stimulation of axons projecting to Glt neurons in V2m, my experiments have provided additional data regarding both the general principle of laminar connectivity between brain regions and the extent to which this follows Peters' rule. Specifically, while most of the inputs to V2m that I studied conformed to the general laminar feedforward and feedback projection patterns that are characteristic in the cortical hierarchy, the functional synaptic connectivity was only partially accounted for by Peters' rule. For example, inputs from the retrosplenial cortex, which strongly branched in L1 as it typical of feedback connections, received synapses mostly on deeper dendrites and had very few synapses on the apical tuft dendrites. On the other hand, input from the anterior thalamus followed Peters' rule more closely, with axonal branching in L1 matched by equally strong synaptic input to this layer.

The most surprising connection that I observed, from the point of view of both Peters' rule and cortical hierarchy, was from the orbitofrontal cortex. The functional connections from OFC generally did not conform to Peters' rule, and instead appeared to target the basal dendrites with remarkable specificity, even in the face of variability in the exact depth of individual neurons. Furthermore this region, which is often placed at a high level of the cortical hierarchy (Rolls, 2004; Harris et al., 2019), was found to project to V2m in a pattern that is more typical of a feedforward pathway. This suggests that either OFC should be re-positioned to a lower level of the cortical hierarchy than is commonly thought, or that the rule for laminar-specific feedback should not be considered a universal rule but rather a convenient heuristic that applies to most, but perhaps not all, pathways in the brain. I would argue that the latter possibility is more sensible, given that there are still many functional reasons for considering OFC to be a higher-order area (such as its involvement in memory and complex cognitive behaviours) and there is nothing inherent to the concept of laminar connectivity that guarantees any particular depth in the neural network (relative to the input layer, i.e. the sensory organs).

Overall, these new observations are potentially relevant to understanding long-range circuit computations, as theories that consider the role of ttL5 neurons often implicitly assume that Peters' rule is accurate, and that feedback connections thus primarily target the apical tuft dendrites while feedforward pathways target basal dendrites. In recent years, this simplified view of ttL5 connectivity has progressively become the standard model for L5 neurons (Larkum, 2013; Richards et al., 2019; Richards and Lillicrap, 2019; Doron et al., 2019; Aru et al., 2019; Payeur et al., 2020; Suzuki and Larkum, 2020). However, in light of the evidence that I have presented, this simplified model may need re-visiting, and should perhaps be abandoned in favour of more flexible models that account for convergence of feedback and feedforward connections in both basal and apical dendrites.

7.2 Canonical cortical microcircuits & computations

The theory of canonical cortical microcircuits has been highly influential in neuroscience, and was originally developed as a distillation and generalization of a large body of anatomical and electrophysiological data (Douglas et al., 1989; Mountcastle, 1997; Douglas and Martin, 2004). It posits that the cerebral cortex is both modular and stereotyped, composed of cortical columns in which each cell type contributes to the microcircuit computation in a unique way that is consistent across all cortical areas. In this view, each column is implementing a universal computation that enables it to dynamically learn a particular range of input-output transformations, and the overall cortical computation emerges as a result of the interactions of many such modules (Hawkins et al., 2019).

The great power of such a modular framework is that it is inherently well-suited to being expanded to arbitrary sizes, increasing computational power in proportion to the number of modules, which is thus mostly limited by the physical and energetic constraints of the brain. If true, it would provide a good explanation for the differences in cognitive abilities between species with different amounts of cortical matter and gyration, which could have evolved to maximise the cortical surface that can fit in a given volume (Lui et al., 2011).

On the other hand, were this theory not true, the alternative is that cortex could be composed of a very large number of uniquely specialized circuits performing fundamentally different computations that can only be understood by studying each individual module separately. While most would agree that the true answer likely lies somewhere in between these two extremes, the question of where cortex lies on this spectrum of generality is important, as it constrains the theories that can explain cortical function and outlines which assumptions can be made about one region of cortex on the basis of discoveries made in another region.

A somewhat related question is how many classes of neurons exist in cortex. The concept of cell types is inherently born out of a need to simplify the vast diversity in connectivity, functional responses, morphology, and intrinsic properties of cells. At its core, this is a problem of clustering, since true neuronal properties occupy a continuous space of variables (Hodge et al., 2019). Thus, while specific differences between any two neurons can always be found, the relevant question is to what extent these differences are meaningful in terms of the general computations of the cortical circuit they are part of.

On a basic level, the cortex undeniable has some broad commonality in the neuronal properties between different cortical areas, and to some extent even between different species (Thomson and Lamy, 2007; Harris and Mrsic-Flogel, 2013; Harris and Shepherd, 2015; Calabrese and Woolley, 2015). However, different cortical regions have different functional requirements and on a structural level gradients have been described in cortical thickness, which in turn influence neuronal properties (Jacobs et al., 2001; Elston, 2002; Elston et al., 2006; Fletcher and Williams, 2019).

When it comes to ttL5 neurons in particular, although our knowledge of their physiology is largely based on recordings from primary visual and somatosensory cortices of rats (Larkum et al., 1999b; Hay et al., 2011; Ramaswamy and Markram, 2015), recent evidence has shown that even within V1 there are substantial variations in their integrative properties across the rostro-caudal axis. Specifically, Fletcher and Williams (2019) showed that the thickness of cortex is several times larger in the rostral end of V1 compared to the caudal end, causing neurons in the thinner portions to have substantially shorter dendrites. This was also associated with a much more compact electrotonic structure, which resulted in dendritic inputs to these neurons experiencing less attenuation over distance and having a stronger driving force on the neuron's firing rate. Interestingly, although these shorter neurons would be driven more strongly by their inputs (assuming equal synaptic potentials), they had relatively low variability in firing rate and were less prone to firing in bursts than their longer counterparts in rostral V1. This less excitable variant of ttL5 neurons has not previously received attention, possibly because neurons in the caudal portion of V1 are cut when making coronal sections, as is most commonly done, and require slicing angles that are up to 50 degrees off from the coronal plane.

My data from recordings across V1 and V2m strongly aligns and is complementary to these findings in many respects. Fletcher and Williams (2019) performed direct whole-cell recordings from the apical dendrites, and were thus able to measure voltage propagation and attenuation between the apical tuft and the soma in both directions along the apical dendrite. On the other hand, my experiments directly tested the phenomenon of BAC firing and the associated dendritic Ca^{2+} plateau, which is one of the most well-known forms of dendritic supralinearities that can drive bursting (Larkum et al., 1999b; Larkum, 2013).

As my slicing angles were only slightly angled relative to the coronal plane, the recordings I made from V1 were mainly sampled from the middle and rostral portions of V1. While neurons in V2m likely have different functions than the neurons in caudal V1, the principle of cortical thinning and reduced apical trunk length appears to apply similarly to both regions. This similarity can be clearly seen in the structural MRI data that Fletcher and Williams

(2019) present. The difference I observed between V1 and V2m is thus largely analogous to the difference they found between rostral and caudal V1.

One important difference to highlight between this data and my experiments is that my recordings were made in mice, while Fletcher and Williams (2019) recorded from rats. Aside from the many potential differences in gene expression patterns across neurons in these two species, one obvious factor is that the cortex is generally thicker in rats than in mice. Thus even the longer V1 neurons I recorded in mice had apical trunk lengths in the same range as the caudal V1 neurons in rats.

While caution is warranted, this by itself does not invalidate the possibility of a large causal influence of apical length on dendritic excitability and burstiness. However it does clearly indicate that there must be other equally important factors that are capable of compensating for these changes. The computational models that Aeron Laffere and I created show that trunk length alone is at least theoretically capable of causing large changes in dendritic excitability, voltage propagation, and propensity for BAC firing in a way that closely mirrors what was observed experimentally. However, this model also showed sensitivity to the density of various ion channels, most notably voltag-gated Na^+ and Ca^{2+} channels in the apical dendrite.

Given the difference in apical length between rat and mouse neurons, it is thus likely that changes in ion channel densities and in other relevant parameters that we did not explore (such as dendrite diameter and oblique dendrite distribution) are present in mice neurons to compensate for the shorter dendrites and enable V1 neurons in mice and rats to behave similarly.

If neurons are equally able to control dendritic excitability through changes in ion channel expression as through morphology, an interesting implication is that the differences between neurons in rostral V1, caudal V1, and V2m may have been important enough to preserve throughout the evolutionary history that separates the two species. This suggests that there might be specific computational advantages to having highly excitable and nonlinear neurons in rostral V1, while neurons in caudal V1 and V2m integrate their inputs more linearly.

7.3 Functions of dendritic plateaus

What role dendritic plateau potentials and BAC firing could play in cortical computations is still an open question in the field. Since BAC firing was first discovered, the prevalent hypothesis has been that its primary role is to signal the coincidence of feedforward and feedback pathways in the brain (Larkum et al., 1999b).

Supporting this view, one line of evidence comes from *in vivo* Ca^{2+} imaging of the apical dendrites during sensory stimulation. Using this method, Takahashi et al. (2016) showed that the presence of Ca^{2+} signals in the apical dendrite strongly correlates with the psychometric function of mice trained to detect weak sensory stimuli. In other words, the mice generally only reported perceiving a stimulus on trials when there was Ca^{2+} influx into the apical dendrites, and failed to report the stimulus when a Ca^{2+} signal was not recorded.

This result suggests that BAC firing could be performing some form of sensory gating, with signals only being passed on to downstream targets in the striatum and brainstem when the sensory input is matched with appropriate top-down signals, such as from brain regions controlling attention.

Supporting such a role in gating perception through the combination of sensory and top-down signals, it has recently been shown that one of the actions of anaesthetics is to decouple the apical dendrites from the somas of ttL5 neurons by suppressing the propagation of signals along the dendrite (Suzuki and Larkum, 2020).

Another somewhat related role for dendritic plateau potentials has been suggested in theories of cortical function that are inspired by “deep learning” in artificial neural networks (Guerguiev et al., 2017; Sacramento et al., 2018; Richards et al., 2019; Payeur et al., 2020). Instead of focusing on the coincidence detection enabled by BAC firing, these theories rely on the ability of apical input to drive burst probability independently of any input to the basal dendrites.

This mechanism is shown to enable networks of ttL5 neurons to implement "supervised learning" (the most common form of learning in artificial neural networks) by multiplexing the spike code and simultaneously passing sensory signals forward through the network while transmitting error signals backward through the network to act as teaching signals and change the synaptic strengths in a way that improves the overall output of the network.

While these theories are intriguing and may contain elements of truth, one issue with most existing theories concerning the role of BAC firing and ttL5 neurons in cortex is that they implicitly or explicitly assume that ttL5 are directly involved in hierarchical processing pathways, with ttL5 neurons projecting to other cortical areas both higher and lower in the cortical hierarchy. However, this is not fully consistent with the anatomical data, according to which ttL5 neurons primarily project to subcortical targets that drive motor activity.

While they undoubtedly receive and integrate long-range signals from across the brain, which may be considered as feedforward and feedback, they are thus unlikely to directly contribute to the processing hierarchy itself, which drives perception and conscious awareness. Instead, they may have a role more analogous to output motor neurons, reading out information that has already been computed in the hierarchical computations that are driven by intracortically projecting neurons.

This still leaves open the question of why dendritic plateau potentials might be required for computations in V1 but not in V2m. If the electrotonic compactness that was measured in ttL5 neurons in caudal V1 is also true of Glt neurons in V2m, synapses in the apical tufts of these neurons would have a much stronger driving effect than they do in the tufts of rostral V1 neurons. Ca^{2+} plateaus may thus serve to amplify voltage signals from the distal tuft inputs in V1 neurons in a manner that is not necessary in the more compact Glt neurons.

However, given the evidence that Ca^{2+} plateaus are likely involved in more sophisticated processes than just signal amplification, the reason for different sensitivities to input in the apical tuft of V1 and V2m neurons may depend on what the behavioural function of the ttL5 neurons in V2m is and exactly which motor centers they project to.

Glt neurons likely receive vestibular signals from the anterior thalamus (Rancz et al., 2015), which synapses with their apical tufts. If this input is used to process visual-motion computations, such as controlling eye muscles to stabilise images during head movements, it may be useful to have this drive the neurons directly and be integrated in a linear way with visual signals from V1, so as to smoothly control eye position.

Conversely, depending on exactly which inputs target the apical tufts in V1 and where these neurons project to, a more nonlinear modulatory influence may be more beneficial. For example, ttL5 neurons are known to project to the superior colliculus, where they can drive escape behaviours in response to looming stimuli (Yilmaz and Meister, 2013; Lur et al., 2016; Zingg et al., 2017; Evans et al., 2019). The synaptic input required for such binary escape decisions may benefit from the strong drive that comes from a burst of spikes, without requiring the precision that is conferred by more regular spiking and linear integration.

Since L5 neurons in V1 projecting to different targets are known to have different functional properties (Kim et al., 2015), it would be interesting to explore in more depth if there are differences in dendritic excitability between ttL5 neurons projecting to the superior colliculus compared to brainstem-projecting neurons. If such differences exist, this could potentially also explain why there is substantial variability in dendritic excitability even within V1.

7.4 Conclusions

The cerebral cortex is an intricate brain structure that holds much promise for casting light on fundamental questions regarding the nature of intelligence and how we might be able to abstract the operational principles of the cortex to create artificial systems with general intelligence.

While much is still unknown about the structure and function of the cortex and the computations that cortical circuits implement, several principles have emerged over the years which have been highly influential and provided an indispensable stepping stone for building more complete theories of the brain.

Chief among these are the theory of canonical cortical circuits and the related principles of long-range connectivity between cortical areas. While many aspects of cortical function are likely to be common between different cortical areas, in light of several lines of recent evidence, including data that I have presented in this thesis, stereotypy should not be assumed when studing a new cortical area and should be rigorously examined.

Although different cortical circuits may arise through common developmental mechanisms, the need for functional specialization may result in a loss of generality in these circuits in adults. It therefore seems increasingly likely that few generalizations can truly be made that apply universally across all of cortex, and a full understanding will require studying the unique computations and solutions that every cortical region has optimized for.

Understanding the manner in which diverse streams of feedforward and feedback information are integrated within individual ttL5 neurons will likely be an important part of this process. Ultimately this may provide key insights into both the universal and the specialized computations performed by cortical circuits, and may have widespread implications for our understanding of both artificial and biological intelligence.

References

Albus, K. and Beckmann, R. (1980). Second and third visual areas of the cat: interindividual variability in retinotopic arrangement and cortical location. *The Journal of Physiology*, 299:247–276.

Allen, M., Poggiali, D., Whitaker, K., Marshall, T. R., and Kievit, R. A. (2019). Raincloud plots: a multi-platform tool for robust data visualization. *Wellcome Open Research*, 4:63.

Amitai, Y., Friedman, A., Connors, B. W., and Gutnick, M. J. (1993). Regenerative Activity in Apical Dendrites of Pyramidal Cells in Neocortex. *Cerebral Cortex*, 3(1):26–38.

Amunts, K. and Zilles, K. (2015). Architectonic Mapping of the Human Brain beyond Brodmann. *Neuron*, 88(6):1086–1107.

Andermann, M. L., Kerlin, A. M., Roumis, D. K., Glickfeld, L. L., and Reid, R. C. (2011). Functional specialization of mouse higher visual cortical areas. *Neuron*, 72(6):1025–1039.

Aru, J., Suzuki, M., Rutiku, R., Larkum, M. E., and Bachmann, T. (2019). Coupling the State and Contents of Consciousness. *Frontiers in Systems Neuroscience*, 13.

Bahl, A., Stemmler, M. B., Herz, A. V. M., and Roth, A. (2012). Automated optimization of a reduced layer 5 pyramidal cell model based on experimental data. *Journal of Neuroscience Methods*, 210(1):22–34.

Baker, A., Kalmbach, B., Morishima, M., Kim, J., Juavinett, A., Li, N., and Dembrow, N. (2018). Specialized Subpopulations of Deep-Layer Pyramidal Neurons in the Neocortex: Bridging Cellular Properties to Functional Consequences. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 38(24):5441–5455.

Bar-Yehuda, D. and Korngreen, A. (2008). Space-clamp problems when voltage clamping neurons expressing voltage-gated conductances. *Journal of Neurophysiology*, 99(3):1127–1136.

Barth, A. M. I., Vizi, E. S., Zelles, T., and Lendvai, B. (2008). Alpha2-adrenergic receptors modify dendritic spike generation via HCN channels in the prefrontal cortex. *Journal of Neurophysiology*, 99(1):394–401.

Bastos, A. M., Usrey, W. M., Adams, R. A., Mangun, G. R., Fries, P., and Friston, K. J. (2012). Canonical microcircuits for predictive coding. *Neuron*, 76(4):695–711.

Beaulieu-Laroche, L., Toloza, E. H. S., van der Goes, M.-S., Lafourcade, M., Barnagian, D., Williams, Z. M., Eskandar, E. N., Frosch, M. P., Cash, S. S., and Harnett, M. T. (2018). Enhanced Dendritic Compartmentalization in Human Cortical Neurons. *Cell*, 175(3):643–651.e14.

Bendixen, A., Schroger, E., and Winkler, I. (2009). I Heard That Coming: Event-Related Potential Evidence for Stimulus-Driven Prediction in the Auditory System. *Journal of Neuroscience*, 29(26):8447–8451.

Bennett, C., Gale, S. D., Garrett, M. E., Newton, M. L., Callaway, E. M., Murphy, G. J., and Olsen, S. R. (2019). Higher-Order Thalamic Circuits Channel Parallel Streams of Visual Information in Mice. *Neuron*, 102(2):477–492.e5.

Blythe, I. M., Kennard, C., and Ruddock, K. H. (1987). Residual vision in patients with retrogeniculate lesions of the visual pathways. *Brain: A Journal of Neurology*, 110 (Pt 4):887–905.

Bohland, J. W., Wu, C., Barbas, H., Bokil, H., Bota, M., Breiter, H. C., Cline, H. T., Doyle, J. C., Freed, P. J., Greenspan, R. J., Haber, S. N., Hawrylycz, M., Herrera, D. G., Hilgetag, C. C., Huang, Z. J., Jones, A., Jones, E. G., Karten, H. J., Kleinfeld, D., Kötter, R., Lester, H. A., Lin, J. M., Mensh, B. D., Mikula, S., Panksepp, J., Price, J. L., Safdieh, J., Saper, C. B., Schiff, N. D., Schmahmann, J. D., Stillman, B. W., Svoboda, K., Swanson, L. W., Toga, A. W., Essen, D. C. V., Watson, J. D., and Mitra, P. P. (2009). A Proposal for a Coordinated Effort for the Determination of Brainwide Neuroanatomical Connectivity in Model Organisms at a Mesoscopic Scale. *PLOS Computational Biology*, 5(3):e1000334. Publisher: Public Library of Science.

Branco, T., Clark, B. A., and Hausser, M. (2010). Dendritic Discrimination of Temporal Input Sequences in Cortical Neurons. *Science*, 329(5999):1671–1675.

Briggman, K. L., Helmstaedter, M., and Denk, W. (2011). Wiring specificity in the direction-selectivity circuit of the retina. *Nature*, 471(7337):183–188. Number: 7337 Publisher: Nature Publishing Group.

Brodmann, K. (1909). *Vergleichende Lokalisationslehre der Grosshirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues von Dr. K. Brodmann, ...* J.A. Barth. Google-Books-ID: Qw5KQwAACAAJ.

Calabrese, A. and Woolley, S. M. N. (2015). Coding principles of the canonical cortical microcircuit in the avian brain. *Proceedings of the National Academy of Sciences of the United States of America*, 112(11):3517–3522.

Callaway, E. M. and Luo, L. (2015). Monosynaptic Circuit Tracing with Glycoprotein-Deleted Rabies Viruses. *Journal of Neuroscience*, 35(24):8979–8985. Publisher: Society for Neuroscience Section: Toolbox.

Calvo, M. G. and Nummenmaa, L. (2011). Time course of discrimination between emotional facial expressions: The role of visual saliency. *Vision Research*, 51(15):1751–1759.

Carr, D. B., Andrews, G. D., Glen, W. B., and Lavin, A. (2007). alpha2-Noradrenergic receptors activation enhances excitability and synaptic integration in rat prefrontal cortex pyramidal neurons via inhibition of HCN currents. *The Journal of Physiology*, 584(Pt 2):437–450.

Cash, S. and Yuste, R. (1998). Input Summation by Cultured Pyramidal Neurons Is Linear and Position-Independent. *Journal of Neuroscience*, 18(1):10–15. Publisher: Society for Neuroscience Section: ARTICLE.

Cash, S. and Yuste, R. (1999). Linear summation of excitatory inputs by CA1 pyramidal neurons. *Neuron*, 22(2):383–394.

Cauller, L. (1995). Layer I of primary sensory neocortex: where top-down converges upon bottom-up. *Behavioural Brain Research*, 71(1):163–170.

Cauller, L. J., Clancy, B., and Connors, B. W. (1998). Backward cortical projections to primary somatosensory cortex in rats extend long horizontal axons in layer I. *Journal of Comparative Neurology*, 390(2):297–310.

Caviness, V. S. (1975). Architectonic map of neocortex of the normal mouse. *The Journal of Comparative Neurology*, 164(2):247–263.

Cela-Conde, C. J., Marty, G., Maestú, F., Ortiz, T., Munar, E., Fernández, A., Roca, M., Rosselló, J., and Quesney, F. (2004). Activation of the prefrontal cortex in the human visual aesthetic perception. *Proceedings of the National Academy of Sciences*, 101(16):6321–6325. Publisher: National Academy of Sciences Section: Biological Sciences.

Chaumon, M., Kveraga, K., Barrett, L. F., and Bar, M. (2014). Visual Predictions in the Orbitofrontal Cortex Rely on Associative Content. *Cerebral Cortex (New York, NY)*, 24(11):2899–2907.

Coogan, T. A. and Burkhalter, A. (1990). Conserved patterns of cortico-cortical connections define areal hierarchy in rat visual cortex. *Experimental Brain Research*, 80(1):49–53.

Coogan, T. A. and Burkhalter, A. (1993). Hierarchical organization of areas in rat visual cortex. *Journal of Neuroscience*, 13(9):3749–3772. Publisher: Society for Neuroscience Section: Articles.

Courtney, S. M., Petit, L., Haxby, J. V., and Ungerleider, L. G. (1998). The role of prefrontal cortex in working memory: examining the contents of consciousness. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 353(1377):1819–1828.

Cragg, B. G. (1969). The topography of the afferent projections in the circumstriate visual cortex of the monkey studied by the Nauta method. *Vision Research*, 9(7):733–747.

Cullen, K. E. (2019). Vestibular processing during natural self-motion: implications for perception and action. *Nature Reviews. Neuroscience*, 20(6):346–363.

de Vries, S. E. J., Lecoq, J. A., Buice, M. A., Groblewski, P. A., Ocker, G. K., Oliver, M., Feng, D., Cain, N., Ledochowitsch, P., Millman, D., Roll, K., Garrett, M., Keenan, T., Kuan, L., Mihalas, S., Olsen, S., Thompson, C., Wakeman, W., Waters, J., Williams, D., Barber, C., Berbesque, N., Blanchard, B., Bowles, N., Caldejon, S. D., Casal, L., Cho, A., Cross, S., Dang, C., Dolbeare, T., Edwards, M., Galbraith, J., Gaudreault, N., Gilbert, T. L., Griffin, F., Hargrave, P., Howard, R., Huang, L., Jewell, S., Keller, N., Knoblich, U., Larkin, J. D., Larsen, R., Lau, C., Lee, E., Lee, F., Leon, A., Li, L., Long, F., Luviano, J., Mace, K., Nguyen, T., Perkins, J., Robertson, M., Seid, S., Shea-Brown, E., Shi, J., Sjoquist, N., Slaughterbeck, C., Sullivan, D., Valenza, R., White, C., Williford, A., Witten, D. M., Zhuang, J., Zeng, H., Farrell, C., Ng, L., Bernard, A., Phillips, J. W., Reid, R. C., and

Koch, C. (2020). A large-scale standardized physiological survey reveals functional organization of the mouse visual cortex. *Nature Neuroscience*, 23(1):138–151. Number: 1 Publisher: Nature Publishing Group.

Dong, H. W. (2008). *The Allen reference atlas: A digital color brain atlas of the C57Bl/6J male mouse*. The Allen reference atlas: A digital color brain atlas of the C57Bl/6J male mouse. John Wiley & Sons Inc, Hoboken, NJ, US. Pages: ix, 366.

Doron, G., Shin, J. N., Takahashi, N., Bocklisch, C., Skenderi, S., Drüke, M., Mont, L. d., Toumazo, M., Heimendahl, M. v., Brecht, M., Naud, R., and Larkum, M. E. (2019). Perirhinal input to neocortical layer 1 controls learning. *bioRxiv*, page 713883. Publisher: Cold Spring Harbor Laboratory Section: New Results.

Douglas, R. J. and Martin, K. A. (2004). Neuronal Circuits of the Neocortex. *Annual Review of Neuroscience*, 27(1):419–451.

Douglas, R. J., Martin, K. A., and Whitteridge, D. (1989). A Canonical Microcircuit for Neocortex. *Neural Computation*, 1(4):480–488. Publisher: MIT Press.

D'Souza, R. D., Meier, A. M., Bista, P., Wang, Q., and Burkhalter, A. (2016). Recruitment of inhibition and excitation across mouse visual cortex depends on the hierarchy of interconnecting areas. *eLife*, 5:e19332. Publisher: eLife Sciences Publications, Ltd.

Elston, G. N. (2002). Cortical heterogeneity: implications for visual processing and polysensory integration. *Journal of Neurocytology*, 31(3-5):317–335.

Elston, G. N., Benavides-Piccione, R., Elston, A., Zietsch, B., Defelipe, J., Manger, P., Casagrande, V., and Kaas, J. H. (2006). Specializations of the granular prefrontal cortex of primates: implications for cognitive processing. *The Anatomical Record. Part A, Discoveries in Molecular, Cellular, and Evolutionary Biology*, 288(1):26–35.

Engel, A. K., Fries, P., and Singer, W. (2001). Dynamic predictions: oscillations and synchrony in top-down processing. *Nature Reviews. Neuroscience*, 2(10):704–716.

Evans, D. A., Stempel, A. V., Vale, R., and Branco, T. (2019). Cognitive Control of Escape Behaviour. *Trends in Cognitive Sciences*, 23(4):334–348.

Evans, D. A., Stempel, A. V., Vale, R., Ruehle, S., Lefler, Y., and Branco, T. (2018). A synaptic threshold mechanism for computing escape decisions. *Nature*, 558(7711):590–594.

Felleman, D. J. and Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex (New York, N.Y.: 1991)*, 1(1):1–47.

Fernandez, F. R., Engbers, J. D. T., and Turner, R. W. (2007). Firing dynamics of cerebellar purkinje cells. *Journal of Neurophysiology*, 98(1):278–294.

Fletcher, L. N. and Williams, S. R. (2019). Neocortical Topology Governs the Dendritic Integrative Capacity of Layer 5 Pyramidal Neurons. *Neuron*, 101(1):76–90.e4.

Fox, M. D., Snyder, A. Z., Vincent, J. L., Corbetta, M., Essen, D. C. V., and Raichle, M. E. (2005). The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proceedings of the National Academy of Sciences*, 102(27):9673–9678.

Fries, P. (2015). Rhythms For Cognition: Communication Through Coherence. *Neuron*, 88(1):220–235.

Friston, K. and Kiebel, S. (2009). Predictive coding under the free-energy principle. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 364(1521):1211–1221.

Frost, D. O. and Metin, C. (1985). Induction of functional retinal projections to the somatosensory system. *Nature*, 317(6033):162–164.

Garrett, M. E., Nauhaus, I., Marshel, J. H., and Callaway, E. M. (2014). Topography and areal organization of mouse visual cortex. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 34(37):12587–12600.

Gerfen, C., Paletzki, R., and Heintz, N. (2013). GENSAT BAC Cre-Recombinase Driver Lines to Study the Functional Organization of Cerebral Cortical and Basal Ganglia Circuits. *Neuron*, 80(6):1368–1383.

Gilbert, C. D. and Sigman, M. (2007). Brain States: Top-Down Influences in Sensory Processing. *Neuron*, 54(5):677–696.

Glenn Northcutt, R. and Kaas, J. H. (1995). The emergence and evolution of mammalian neocortex. *Trends in Neurosciences*, 18(9):373–379.

Glickfeld, L. L., Histed, M. H., and Maunsell, J. H. R. (2013). Mouse Primary Visual Cortex Is Used to Detect Both Orientation and Contrast Changes. *Journal of Neuroscience*, 33(50):19416–19422.

Glickfeld, L. L. and Olsen, S. R. (2017). Higher-Order Areas of the Mouse Visual Cortex. *Annual Review of Vision Science*, 3:251–273.

Glickstein, M. and Rizzolatti, G. (1984). Francesco Gennari and the structure of the cerebral cortex. *Trends in Neurosciences*, 7(12):464–467.

Glickstein, M. and Whitteridge, D. (1987). Tatsuji Inouye and the mapping of the visual fields on the human cerebral cortex. *Trends in Neurosciences*, 10(9):350–353.

Gong, S., Doughty, M., Harbaugh, C. R., Cummins, A., Hatten, M. E., Heintz, N., and Gerfen, C. R. (2007). Targeting Cre recombinase to specific neuron populations with bacterial artificial chromosome constructs. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 27(37):9817–9823.

Gong, S., Zheng, C., Doughty, M. L., Losos, K., Didkovsky, N., Schambra, U. B., Nowak, N. J., Joyner, A., Leblanc, G., Hatten, M. E., and Heintz, N. (2003). A gene expression atlas of the central nervous system based on bacterial artificial chromosomes. *Nature*, 425(6961):917–925.

Gouwens, N. W., Sorensen, S. A., Berg, J., Lee, C., Jarsky, T., Ting, J., Sunkin, S. M., Feng, D., Anastassiou, C. A., Barkan, E., Bickley, K., Blesie, N., Braun, T., Brouner, K., Budzillo, A., Caldejon, S., Casper, T., Castelli, D., Chong, P., Crichton, K., Cuhaciyan, C., Daigle, T. L., Dalley, R., Dee, N., Desta, T., Ding, S.-L., Dingman, S., Doperalski, A., Dotson, N., Egdorf, T., Fisher, M., Frates, R. A. d., Garren, E., Garwood, M., Gary, A., Gaudreault, N., Godfrey, K., Gorham, M., Gu, H., Habel, C., Hadley, K., Harrington, J., Harris, J. A., Henry, A., Hill, D., Josephsen, S., Kebede, S., Kim, L., Kroll, M., Lee, B., Lemon, T., Link, K. E., Liu, X., Long, B., Mann, R., McGraw, M., Mihalas, S., Mukora, A., Murphy, G. J., Ng, L., Ngo, K., Nguyen, T. N., Nicovich, P. R., Oldre, A., Park, D., Parry, S., Perkins, J., Potekhina, L., Reid, D., Robertson, M., Sandman, D., Schroedter, M., Slaughterbeck, C., Soler-Llavina, G., Sulc, J., Szafer, A., Tasic, B., Taskin, N., Teeter, C., Thatra, N., Tung, H., Wakeman, W., Williams, G., Young, R., Zhou, Z., Farrell, C., Peng, H., Hawrylycz, M. J., Lein, E., Ng, L., Arkhipov, A., Bernard, A., Phillips, J. W., Zeng, H., and Koch, C. (2019). Classification of electrophysiological and morphological neuron types in the mouse visual cortex. *Nature Neuroscience*, 22(7):1182–1195.

Greicius, M. D., Krasnow, B., Reiss, A. L., and Menon, V. (2003). Functional connectivity in the resting brain: A network analysis of the default mode hypothesis. *Proceedings of the National Academy of Sciences*, 100(1):253–258.

Groh, A., Meyer, H. S., Schmidt, E. F., Heintz, N., Sakmann, B., and Krieger, P. (2010). Cell-Type Specific Properties of Pyramidal Neurons in Neocortex

Underlying a Layout that Is Modifiable Depending on the Cortical Area. *Cerebral Cortex*, 20(4):826–836.

Gross, C. G., Schiller, P. H., Wells, C., and Gerstein, G. L. (1967). Single-unit activity in temporal association cortex of the monkey. *Journal of Neurophysiology*, 30(4):833–843.

Guerguiev, J., Lillicrap, T. P., and Richards, B. A. (2017). Towards deep learning with segregated dendrites. *eLife*, 6:e22901.

Gutfreund, Y., Yarom, Y., and Segev, I. (1995). Subthreshold oscillations and resonant frequency in guinea-pig cortical neurons: physiology and modelling. *The Journal of Physiology*, 483(3):621–640.

Han, Y., Kebschull, J. M., Campbell, R. A. A., Cowan, D., Imhof, F., Zador, A. M., and Mrsic-Flogel, T. D. (2018). The logic of single-cell projections from visual cortex. *Nature*, 556(7699):51–56.

Harnett, M., Xu, N.-L., Magee, J., and Williams, S. (2013). Potassium Channels Control the Interaction between Active Dendritic Integration Compartments in Layer 5 Cortical Pyramidal Neurons. *Neuron*, 79(3):516–529.

Harnett, M. T., Magee, J. C., and Williams, S. R. (2015). Distribution and Function of HCN Channels in the Apical Dendritic Tuft of Neocortical Pyramidal Neurons. *The Journal of Neuroscience*, 35(3):1024–1037.

Harris, J. A., Mihalas, S., Hirokawa, K. E., Whitesell, J. D., Choi, H., Bernard, A., Bohn, P., Caldejon, S., Casal, L., Cho, A., Feiner, A., Feng, D., Gaudreault, N., Gerfen, C. R., Graddis, N., Groblewski, P. A., Henry, A. M., Ho, A., Howard, R., Knox, J. E., Kuan, L., Kuang, X., Lecoq, J., Lesnar, P., Li, Y., Luviano, J., McConoughey, S., Mortrud, M. T., Naeemi, M., Ng, L., Oh, S. W., Ouellette, B., Shen, E., Sorensen, S. A., Wakeman, W., Wang, Q., Wang, Y., Williford, A., Phillips, J. W., Jones, A. R., Koch, C., and Zeng, H. (2019). Hierarchical organization of cortical and thalamic connectivity.

Nature, 575(7781):195–202. Number: 7781 Publisher: Nature Publishing Group.

Harris, K. D. and Mrsic-Flogel, T. D. (2013). Cortical connectivity and sensory coding. *Nature*, 503(7474):51–58. Number: 7474 Publisher: Nature Publishing Group.

Harris, K. D. and Shepherd, G. M. G. (2015). The neocortical circuit: themes and variations. *Nature Neuroscience*, 18(2):170–181.

Hassabis, D., Kumaran, D., Summerfield, C., and Botvinick, M. (2017). Neuroscience-Inspired Artificial Intelligence. *Neuron*, 95(2):245–258.

Hausser, M. and Mel, B. (2003). Dendrites: bug or feature? *Current Opinion in Neurobiology*, 13(3):372–383.

Hawkins, J., Lewis, M., Klukas, M., Purdy, S., and Ahmad, S. (2019). A Framework for Intelligence and Cortical Function Based on Grid Cells in the Neocortex. *Frontiers in Neural Circuits*, 12.

Hay, E., Hill, S., Schurmann, F., Markram, H., and Segev, I. (2011). Models of Neocortical Layer 5b Pyramidal Cells Capturing a Wide Range of Dendritic and Perisomatic Active Properties. *PLoS Computational Biology*, 7(7).

Herculano-Houzel, S., Collins, C. E., Wong, P., and Kaas, J. H. (2007). Cellular scaling rules for primate brains. *Proceedings of the National Academy of Sciences*, 104(9):3562–3567. Publisher: National Academy of Sciences Section: Biological Sciences.

Hines, M. L. and Carnevale, N. T. (2001). Neuron: A Tool for Neuroscientists. *The Neuroscientist*, 7(2):123–135.

Hodge, R. D., Bakken, T. E., Miller, J. A., Smith, K. A., Barkan, E. R., Graybuck, L. T., Close, J. L., Long, B., Johansen, N., Penn, O., Yao, Z., Eggermont, J., Höllt, T., Levi, B. P., Shehata, S. I., Aevermann, B., Beller, A., Bertagnolli, D., Brouner, K., Casper, T., Cobbs, C., Dalley, R., Dee, N., Ding, S.-L., Ellenbogen, R. G., Fong, O., Garren, E., Goldy, J., Gwinn, R. P., Hirschstein, D., Keene, C. D., Keshk, M., Ko, A. L., Lathia, K., Mahfouz, A., Maltzer, Z., McGraw, M., Nguyen, T. N., Nyhus, J., Ojemann, J. G., Oldre, A., Parry, S., Reynolds, S., Rimorin, C., Shapovalova, N. V., Somasundaram, S., Szafer, A., Thomsen, E. R., Tieu, M., Quon, G., Scheuermann, R. H., Yuste, R., Sunkin, S. M., Lelieveldt, B., Feng, D., Ng, L., Bernard, A., Hawrylycz, M., Phillips, J. W., Tasic, B., Zeng, H., Jones, A. R., Koch, C., and Lein, E. S. (2019). Conserved cell types with divergent features in human versus mouse cortex. *Nature*, 573(7772):61–68. Number: 7772 Publisher: Nature Publishing Group.

Hofman, M. A. (2014). Evolution of the human brain: when bigger is better. *Frontiers in Neuroanatomy*, 8. Publisher: Frontiers.

Holmes, G. (1918). Disturbances of vision by cerebral lesions. *The British Journal of Ophthalmology*, 2(7):353–384.

Hooks, B. M., Lin, J. Y., Guo, C., and Svoboda, K. (2015). Dual-Channel Circuit Mapping Reveals Sensorimotor Convergence in the Primary Motor Cortex. *The Journal of Neuroscience*, 35(10):4418–4426.

Hovde, K., Gianatti, M., Witter, M. P., and Whitlock, J. R. (2019). Architecture and organization of mouse posterior parietal cortex relative to extrastriate areas. *The European Journal of Neuroscience*, 49(10):1313–1329.

Huang, Z. J. and Zeng, H. (2013). Genetic Approaches to Neural Circuits in the Mouse. *Annual Review of Neuroscience*, 36(1):183–215.

Hubel, D. H. and Wiesel, T. N. (1959). Receptive fields of single neurones in the cat's striate cortex. *The Journal of Physiology*, 148(3):574–591.

Hubel, D. H. and Wiesel, T. N. (1961). Integrative action in the cat's lateral geniculate body. *The Journal of Physiology*, 155(2):385–398.1.

Hubel, D. H. and Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of Physiology*, 160(1):106–154.2.

Hubel, D. H. and Wiesel, T. N. (1965). Receptive fields and functional architecture in two nonstriate visual areas (18 and 19) of the cat. *Journal of Neurophysiology*, 28:229–289.

Hubel, D. H. and Wiesel, T. N. (1998). Early exploration of the visual cortex. *Neuron*, 20(3):401–412.

Huffman, K. J., Molnár, Z., Dellen, A. V., Kahn, D. M., Blakemore, C., and Krubitzer, L. (1999). Formation of Cortical Fields on a Reduced Cortical Sheet. *Journal of Neuroscience*, 19(22):9939–9952.

Hutcheon, B., Miura, R. M., and Puil, E. (1996). Subthreshold membrane resonance in neocortical neurons. *Journal of Neurophysiology*, 76(2):683–697.

Hutcheon, B. and Yarom, Y. (2000). Resonance, oscillation and the intrinsic frequency preferences of neurons. *Trends in Neurosciences*, 23(5):216–222.

Ishizu, T. and Zeki, S. (2011). Toward A Brain-Based Theory of Beauty. *PLoS ONE*, 6(7).

Izhikevich, E. (2003). Simple model of spiking neurons. *IEEE Transactions on Neural Networks*, 14(6):1569–1572. Conference Name: IEEE Transactions on Neural Networks.

Jackman, S. L. and Regehr, W. G. (2017). The Mechanisms and Functions of Synaptic Facilitation. *Neuron*, 94(3):447–464.

Jacobs, B., Schall, M., Prather, M., Kapler, E., Driscoll, L., Baca, S., Jacobs, J., Ford, K., Wainwright, M., and Treml, M. (2001). Regional dendritic and spine variation in human cerebral cortex: a quantitative golgi study. *Cerebral Cortex (New York, N.Y.: 1991)*, 11(6):558–571.

Jankowski, M. M., Ronnqvist, K. C., Tsanov, M., Vann, S. D., Wright, N. F., Erichsen, J. T., Aggleton, J. P., and O'Mara, S. M. (2013). The anterior thalamus provides a subcortical circuit supporting memory and spatial navigation. *Frontiers in Systems Neuroscience*, 7:45.

Jiang, S., Guan, Y., Chen, S., Jia, X., Ni, H., Zhang, Y., Han, Y., Peng, X., Zhou, C., Li, A., Luo, Q., and Gong, H. (2020). Anatomically revealed morphological patterns of pyramidal neurons in layer 5 of the motor cortex. *Scientific Reports*, 10(1):7916. Number: 1 Publisher: Nature Publishing Group.

Jiang, X., Shen, S., Cadwell, C. R., Berens, P., Sinz, F., Ecker, A. S., Patel, S., and Tolias, A. S. (2015). Principles of connectivity among morphologically defined cell types in adult neocortex. *Science*, 350(6264):aac9462.

Johnston, D. and Wu, S. M.-S. (1995). *Foundations of Cellular Neurophysiology.* MIT Press, Cambridge, Massachusetts. Library Catalog: mitpress.mit.edu Publisher: The MIT Press.

Kaas, J. H. (2008). The evolution of the complex sensory and motor systems of the human brain. *Brain Research Bulletin*, 75(2-4):384–390.

Kalatsky, V. A. and Stryker, M. P. (2003). New Paradigm for Optical Imaging: Temporally Encoded Maps of Intrinsic Signal. *Neuron*, 38(4):529–545.

Kampa, B. M. and Stuart, G. J. (2006). Calcium Spikes in Basal Dendrites of Layer 5 Pyramidal Neurons during Action Potential Bursts. *Journal of Neuroscience*, 26(28):7424–7432.

Kasper, E. M., Larkman, A. U., Lubke, J., and Blakemore, C. (1994). Pyramidal neurons in layer 5 of the rat visual cortex. I. Correlation among cell morphology, intrinsic electrophysiological properties, and axon targets. *Journal of Comparative Neurology*, 339(4):459–474.

Kawai, R., Markman, T., Poddar, R., Ko, R., Fantana, A., Dhawale, A., Kampff, A. R., and Ölveczky, B. P. (2015). Motor cortex is required for learning but not executing a motor skill. *Neuron*, 86(3):800–812.

Keller, G., Bonhoeffer, T., and Hubener, M. (2012). Sensorimotor Mismatch Signals in Primary Visual Cortex of the Behaving Mouse. *Neuron*, 74(5):809–815.

Keller, G. B. and Mrsic-Flogel, T. D. (2018). Predictive Processing: A Canonical Cortical Computation. *Neuron*, 100(2):424–435.

Kim, E., Juavinett, A., Kyubwa, E., Jacobs, M., and Callaway, E. (2015). Three Types of Cortical Layer 5 Neurons That Differ in Brain-wide Connectivity and Function. *Neuron*, 88(6):1253–1267.

Kimura, Y., Hawkins, M. T. R., McDonough, M. M., Jacobs, L. L., and Flynn, L. J. (2015). Corrected placement of Mus - Rattus fossil calibration forces precision in the molecular tree of rodents. *Scientific Reports*, 5(1):14444. Number: 1 Publisher: Nature Publishing Group.

Klapoetke, N. C., Murata, Y., Kim, S. S., Pulver, S. R., Birdsey-Benson, A., Cho, Y. K., Morimoto, T. K., Chuong, A. S., Carpenter, E. J., Tian, Z., Wang, J., Xie, Y., Yan, Z., Zhang, Y., Chow, B. Y., Surek, B., Melkonian, M., Jayaraman, V., Constantine-Paton, M., Wong, G. K.-S., and Boyden, E. S. (2014). Independent optical excitation of distinct neural populations. *Nature Methods*, 11(3):338–346.

Klein, S., Staring, M., Murphy, K., Viergever, M. A., and Pluim, J. P. W. (2010). elastix: a toolbox for intensity-based medical image registration. *IEEE transactions on medical imaging*, 29(1):196–205.

Kolb, B. and Walkey, J. (1987). Behavioural and anatomical studies of the posterior parietal cortex in the rat. *Behavioural Brain Research*, 23(2):127–145.

Kording, K. P. and Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature*, 427(6971):244–247.

Kuffler, S. W. (1953). Discharge patterns and functional organization of mammalian retina. *Journal of Neurophysiology*, 16(1):37–68.

Kumaran, D., Hassabis, D., and McClelland, J. L. (2016). What Learning Systems do Intelligent Agents Need? Complementary Learning Systems Theory Updated. *Trends in Cognitive Sciences*, 20(7):512–534.

Labarrera, C., Deitcher, Y., Dudai, A., Weiner, B., Kaduri Amichai, A., Zylbermann, N., and London, M. (2018). Adrenergic Modulation Regulates the Dendritic Excitability of Layer 5 Pyramidal Neurons In Vivo. *Cell Reports*, 23(4):1034–1044.

Lanciego, J. L. and Wouterlood, F. G. (2011). A half century of experimental neuroanatomical tracing. *Journal of Chemical Neuroanatomy*, 42(3):157–183.

Larkman, A. U. (1991). Dendritic morphology of pyramidal neurones of the visual cortex of the rat: I. Branching patterns. *The Journal of Comparative Neurology*, 306(2):307–319.

Larkum, M. (2013). A cellular mechanism for cortical associations: an organizing principle for the cerebral cortex. *Trends in Neurosciences*, 36(3):141–151.

Larkum, M. E., Kaiser, K. M. M., and Sakmann, B. (1999a). Calcium electrogenesis in distal apical dendrites of layer 5 pyramidal cells at a critical frequency of back-propagating action potentials. *Proceedings of the National Academy of Sciences*, 96(25):14600–14604.

Larkum, M. E., Nevian, T., Sandler, M., Polsky, A., and Schiller, J. (2009). Synaptic Integration in Tuft Dendrites of Layer 5 Pyramidal Neurons: A New Unifying Principle. *Science*, 325(5941):756–760.

Larkum, M. E., Zhu, J. J., and Sakmann, B. (1999b). A new cellular mechanism for coupling inputs arriving at different cortical layers. *Nature*, 398(6725):338–341.

LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature*, 521(7553):436–444.

LeCun, Y., Boser, B., Denker, J. S., Henderson, D., Howard, R. E., Hubbard, W., and Jackel, L. D. (1989). Backpropagation Applied to Handwritten Zip Code Recognition. *Neural Computation*, 1(4):541–551.

LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2323.

Leergaard, T. B., Hilgetag, C. C., and Sporns, O. (2012). Mapping the Connectome: Multi-Level Analysis of Brain Connectivity. *Frontiers in Neuroinformatics*, 6. Publisher: Frontiers.

Li, S., Liu, N., Zhang, X., McLaughlin, D. W., Zhou, D., and Cai, D. (2019). Dendritic computations captured by an effective point neuron model. *Proceedings of the National Academy of Sciences*, 116(30):15244–15252. Publisher: National Academy of Sciences Section: PNAS Plus.

Lindsey, J., Ocko, S. A., Ganguli, S., and Deny, S. (2019). A Unified Theory Of Early Visual Representations From Retina To Cortex Through Anatomically Constrained Deep CNNs. *bioRxiv*, page 511535. Publisher: Cold Spring Harbor Laboratory Section: New Results.

Lisman, J. E. (1997). Bursts as a unit of neural information: making unreliable synapses reliable. *Trends in Neurosciences*, 20(1):38–43.

London, M. and Hausser, M. (2005). Dendritic Computation. *Annual Review of Neuroscience*, 28(1):503–532.

Lopes, J. P. and Cunha, R. A. (2019). What is the extracellular calcium concentration within brain synapses?: An Editorial for 'Ionized calcium in human cerebrospinal fluid and its influence on intrinsic and synaptic excitability of hippocampal pyramidal neurons in the rat' on page 452. *Journal of Neurochemistry*, 149(4):435–437.

Lorincz, A., Notomi, T., Tamas, G., Shigemoto, R., and Nusser, Z. (2002). Polarized and compartment-dependent distribution of HCN1 in pyramidal cell dendrites. *Nature Neuroscience*, 5(11):1185–1193.

Lui, J. H., Hansen, D. V., and Kriegstein, A. R. (2011). Development and evolution of the human neocortex. *Cell*, 146(1):18–36.

Lur, G., Vinck, M., Tang, L., Cardin, J., and Higley, M. (2016). Projection-Specific Visual Feature Encoding by Layer 5 Cortical Subnetworks. *Cell Reports*, 14(11):2538–2545.

Lyamzin, D. and Benucci, A. (2019). The mouse posterior parietal cortex: Anatomy and functions. *Neuroscience Research*, 140:14–22.

Magee, J. C. (1998). Dendritic Hyperpolarization-Activated Currents Modify the Integrative Properties of Hippocampal CA1 Pyramidal Neurons. *Journal of Neuroscience*, 18(19):7613–7624.

Mainen, Z. F. and Sejnowski, T. J. (1996). Influence of dendritic structure on firing pattern in model neocortical neurons. *Nature*, 382(6589):363–366.

Manita, S., Suzuki, T., Homma, C., Matsumoto, T., Odagawa, M., Yamada, K., Ota, K., Matsubara, C., Inutsuka, A., Sato, M., Ohkura, M., Yamanaka, A., Yanagawa, Y., Nakai, J., Hayashi, Y., Larkum, M. E., and Murayama, M. (2015). A Top-Down Cortical Circuit for Accurate Sensory Perception. *Neuron*, 86(5):1304–1316.

Manns, I. D., Sakmann, B., and Brecht, M. (2004). Sub- and suprathreshold receptive field properties of pyramidal neurones in layers 5A and 5B of rat somatosensory barrel cortex. *The Journal of Physiology*, 556(Pt 2):601–622.

Markov, N. T., Vezoli, J., Chameau, P., Falchier, A., Quilodran, R., Huissoud, C., Lamy, C., Misery, P., Giroud, P., Ullman, S., Barone, P., Dehay, C., Knoblauch, K., and Kennedy, H. (2014). Anatomy of hierarchy: feedforward and feedback pathways in macaque visual cortex. *The Journal of Comparative Neurology*, 522(1):225–259.

Markram, H., Muller, E., Ramaswamy, S., Reimann, M. W., Abdellah, M., Sanchez, C. A., Ailamaki, A., Alonso-Nanclares, L., Antille, N., Arsever, S., Kahou, G. A. A., Berger, T. K., Bilgili, A., Buncic, N., Chalimourda, A., Chindemi, G., Courcol, J.-D., Delalondre, F., Delattre, V., Druckmann, S., Dumusc, R., Dynes, J., Eilemann, S., Gal, E., Gevaert, M. E., Ghobril, J.-P., Gidon, A., Graham, J. W., Gupta, A., Haenel, V., Hay, E., Heinis, T., Hernando, J. B., Hines, M., Kanari, L., Keller, D., Kenyon, J., Khazen, G., Kim, Y., King, J. G., Kisvarday, Z., Kumbhar, P., Lasserre, S., Le Bé, J.-V., Magalhães, B. R. C., Merchán-Pérez, A., Meystre, J., Morrice, B. R., Muller, J., Muñoz-Céspedes, A., Muralidhar, S., Muthurasa, K., Nachbaur, D., Newton, T. H., Nolte, M., Ovcharenko, A., Palacios, J., Pastor, L., Perin, R., Ranjan, R., Riachi, I., Rodríguez, J.-R., Riquelme, J. L., Rössert, C., Sfyrakis, K., Shi, Y., Shillcock, J. C., Silberberg, G., Silva, R., Tauheed, F., Telefont, M., Toledo-Rodriguez, M., Tränkler, T., Van Geit, W., Díaz, J. V., Walker, R., Wang, Y., Zaninetta, S. M., DeFelipe, J., Hill, S. L., Segev, I., and Schürmann, F. (2015). Reconstruction and Simulation of Neocortical Microcircuitry. *Cell*, 163(2):456–492.

Markram, H., Toledo-Rodriguez, M., Wang, Y., Gupta, A., Silberberg, G., and Wu, C. (2004). Interneurons of the neocortical inhibitory system. *Nature Reviews Neuroscience*, 5(10):793–807.

Marshel, J. H., Garrett, M. E., Nauhaus, I., and Callaway, E. M. (2011). Functional specialization of seven mouse visual cortical areas. *Neuron*, 72(6):1040–1054.

Maunsell, J. H. and Essen, D. v. (1983). The connections of the middle temporal visual area (MT) and their relationship to a cortical hierarchy in the macaque monkey. *Journal of Neuroscience*, 3(12):2563–2586. Publisher: Society for Neuroscience Section: Articles.

McClelland, J. L. and Rogers, T. T. (2003). The parallel distributed processing approach to semantic cognition. *Nature Reviews Neuroscience*, 4(4):310–322. Number: 4 Publisher: Nature Publishing Group.

McCormick, D. A. and Pape, H. C. (1990). Properties of a hyperpolarization-activated cation current and its role in rhythmic oscillation in thalamic relay neurones. *The Journal of Physiology*, 431:291–318.

Metin, C. and Frost, D. O. (1989). Visual responses of neurons in somatosensory cortex of hamsters with experimentally induced retinal projections to somatosensory thalamus. *Proceedings of the National Academy of Sciences of the United States of America*, 86(1):357–361.

Meyer, A. (1907). The temporal lobe detour of the radiations and its importance for the diagnosis of temporal lobe lesions. *Trans. Ass. Amer. Phycns.*, 22:7–16.

Miller, K. D. (2016). Canonical computations of cerebral cortex. *Current Opinion in Neurobiology*, 37:75–84.

Mo, C., Petrof, I., Viaene, A. N., and Sherman, S. M. (2017). Synaptic properties of the lemniscal and paralemniscal pathways to the mouse somatosensory thalamus. *Proceedings of the National Academy of Sciences*, 114(30):E6212–E6221. Publisher: National Academy of Sciences Section: PNAS Plus.

Mountcastle, V. B. (1997). The columnar organization of the neocortex. *Brain*, 120(4):701–722.

Naka, A., Veit, J., Shababo, B., Chance, R. K., Risso, D., Stafford, D., Snyder, B., Egladyous, A., Chu, D., Sridharan, S., Mossing, D. P., Paninski, L., Ngai, J., and Adesnik, H. (2019). Complementary networks of cortical somatostatin interneurons enforce layer specific control. *eLife*, 8:e43696. Publisher: eLife Sciences Publications, Ltd.

Naud, R. and Sprekeler, H. (2018). Sparse bursts optimize information transmission in a multiplexed neural code. *Proceedings of the National Academy of Sciences*, 115(27):E6329–E6338.

Niedworok, C. J., Brown, A. P. Y., Jorge Cardoso, M., Osten, P., Ourselin, S., Modat, M., and Margrie, T. W. (2016). aMAP is a validated pipeline for registration and segmentation of high-resolution mouse brain data. *Nature Communications*, 7(1):1–9.

Oberlaender, M., Boudewijns, Z. S. R. M., Kleele, T., Mansvelder, H. D., Sakmann, B., and Kock, C. P. J. d. (2011). Three-dimensional axon morphologies of individual layer 5 neurons indicate cell type-specific intracortical pathways for whisker motion and touch. *Proceedings of the National Academy of Sciences*, 108(10):4188–4193.

Ohno, N., Katoh, M., Saitoh, Y., Saitoh, S., and Ohno, S. (2015). Three-dimensional volume imaging with electron microscopy toward connectome. *Microscopy (Oxford, England)*, 64(1):17–26.

Oswald, M. J., Tantirigama, M. L. S., Sonntag, I., Hughes, S. M., and Empson, R. M. (2013). Diversity of layer 5 projection neurons in the mouse motor cortex. *Frontiers in Cellular Neuroscience*, 7:174.

Pallas, S. L. (2001). Intrinsic and extrinsic factors that shape neocortical specification. *Trends in Neurosciences*, 24(7):417–423.

Paxinos, G. and Franklin, K. B. J. (2007). *The Mouse Brain in Stereotaxic Coordinates*. Elsevier Science, third edition edition. Google-Books-ID: 5Md-pQgAACAAJ.

Payeur, A., Guerguiev, J., Zenke, F., Richards, B. A., and Naud, R. (2020). Burst-dependent synaptic plasticity can coordinate learning in hierarchical circuits. *bioRxiv*, page 2020.03.30.015511. Publisher: Cold Spring Harbor Laboratory Section: New Results.

Perez-Garci, E., Gassmann, M., Bettler, B., and Larkum, M. E. (2006). The GABAB1b Isoform Mediates Long-Lasting Inhibition of Dendritic Ca2+ Spikes in Layer 5 Somatosensory Pyramidal Neurons. *Neuron*, 50(4):603–616.

Peters, A. and Feldman, M. L. (1976). The projection of the lateral geniculate nucleus to area 17 of the rat cerebral cortex. I. General description. *Journal of Neurocytology*, 5(1):63–84.

Petreanu, L., Mao, T., Sternson, S. M., and Svoboda, K. (2009). The subcellular organization of neocortical excitatory connections. *Nature*, 457(7233):1142–1145.

Phillips, W. A., Larkum, M. E., Harley, C. W., and Silverstein, S. M. (2016). The effects of arousal on apical amplification and conscious state. *Neuroscience of Consciousness*, 2016(1).

Pike, F. G., Goddard, R. S., Suckling, J. M., Ganter, P., Kasthuri, N., and Paulsen, O. (2000). Distinct frequency preferences of different types of rat hippocampal neurones in response to oscillatory input currents. *The Journal of Physiology*, 529(Pt 1):205–213.

Poirazi, P., Brannon, T., and Mel, B. W. (2003). Pyramidal Neuron as Two-Layer Neural Network. *Neuron*, 37(6):989–999.

Probst, M. (1906). Über die zentralen Sinnesbahnen und die Sinneszentren des menschlichen Gehirnes. *S. B. Akad. Wiss. Wien*, 115:103–176.

Ragan, T., Kadiri, L. R., Venkataraju, K. U., Bahlmann, K., Sutin, J., Taranda, J., Arganda-Carreras, I., Kim, Y., Seung, H. S., and Osten, P. (2012). Serial two-photon tomography for automated ex vivo mouse brain imaging. *Nature Methods*, 9(3):255–258.

Raichle, M. E., MacLeod, A. M., Snyder, A. Z., Powers, W. J., Gusnard, D. A., and Shulman, G. L. (2001). A default mode of brain function. *Proceedings of the National Academy of Sciences of the United States of America*, 98(2):676–682.

Rakic, P. (2009). Evolution of the neocortex: a perspective from developmental biology. *Nature Reviews Neuroscience*, 10(10):724–735.

Ramaswamy, S. and Markram, H. (2015). Anatomy and physiology of the thick-tufted layer 5 pyramidal neuron. *Frontiers in Cellular Neuroscience*, 9.

Ramon y Cajal, S. (1909). *Histologie du système nerveux de l'homme & des vertébrés*. Paris : Maloine.

Rancz, E. A., Moya, J., Drawitsch, F., Brichta, A. M., Canals, S., and Margrie, T. W. (2015). Widespread vestibular activation of the rodent cortex. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 35(15):5926–5934.

Rao, R. P. N. and Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1):79–87.

Reardon, T. R., Murray, A. J., Turi, G. F., Wirblich, C., Croce, K. R., Schnell, M. J., Jessell, T. M., and Losonczy, A. (2016). Rabies virus CVS-

N2cΔG strain enhances retrograde synaptic transfer and neuronal viability. *Neuron*, 89(4):711–724.

Rees, C. L., Moradi, K., and Ascoli, G. A. (2017). Weighing the Evidence in Peters' Rule: Does Neuronal Morphology Predict Connectivity? *Trends in neurosciences*, 40(2):63–71.

Resulaj, A., Ruediger, S., Olsen, S. R., and Scanziani, M. (2018). First spikes in visual cortex enable perceptual discrimination. *eLife*, 7:e34044.

Richards, B. A. and Lillicrap, T. P. (2019). Dendritic solutions to the credit assignment problem. *Current Opinion in Neurobiology*, 54:28–36.

Richards, B. A., Lillicrap, T. P., Beaudoin, P., Bengio, Y., Bogacz, R., Christensen, A., Clopath, C., Costa, R. P., de Berker, A., Ganguli, S., Gillon, C. J., Hafner, D., Kepecs, A., Kriegeskorte, N., Latham, P., Lindsay, G. W., Miller, K. D., Naud, R., Pack, C. C., Poirazi, P., Roelfsema, P., Sacramento, J., Saxe, A., Scellier, B., Schapiro, A. C., Senn, W., Wayne, G., Yamins, D., Zenke, F., Zylberberg, J., Therien, D., and Kording, K. P. (2019). A deep learning framework for neuroscience. *Nature Neuroscience*, 22(11):1761–1770. Number: 11 Publisher: Nature Publishing Group.

Rockland, K. S. and Pandya, D. N. (1979). Laminar origins and terminations of cortical connections of the occipital lobe in the rhesus monkey. *Brain Research*, 179(1):3–20.

Roe, A. W., Pallas, S. L., Hahm, J. O., and Sur, M. (1990). A map of visual space induced in primary auditory cortex. *Science (New York, N.Y.)*, 250(4982):818–820.

Roe, A. W., Pallas, S. L., Kwon, Y. H., and Sur, M. (1992). Visual projections routed to the auditory pathway in ferrets: receptive fields of visual neurons in primary auditory cortex. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 12(9):3651–3664.

Rolls, E. T. (2004). The functions of the orbitofrontal cortex. *Brain and Cognition*, 55(1):11–29.

Rothman, J. S. and Silver, R. A. (2018). NeuroMatic: An Integrated Open-Source Software Toolkit for Acquisition, Analysis and Simulation of Electrophysiological Data. *Frontiers in Neuroinformatics*, 12.

Sacramento, J., Ponte Costa, R., Bengio, Y., and Senn, W. (2018). Dendritic cortical microcircuits approximate the backpropagation algorithm. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R., editors, *Advances in Neural Information Processing Systems 31*, pages 8721–8732. Curran Associates, Inc.

Saleem, A. B., Ayaz, A., Jeffery, K. J., Harris, K. D., and Carandini, M. (2013). Integration of visual motion and locomotion in mouse visual cortex. *Nature Neuroscience*, 16(12):1864–1869. Number: 12 Publisher: Nature Publishing Group.

Saleem, A. B., Diamanti, E. M., Fournier, J., Harris, K. D., and Carandini, M. (2018). Coherent encoding of subjective spatial position in visual cortex and hippocampus. *Nature*, 562(7725):124–127. Number: 7725 Publisher: Nature Publishing Group.

Sandrini, G., Serrao, M., Rossi, P., Romaniello, A., Cruccu, G., and Willer, J. C. (2005). The lower limb flexion reflex in humans. *Progress in Neurobiology*, 77(6):353–395.

Sawamura, H., Orban, G. A., and Vogels, R. (2006). Selectivity of Neuronal Adaptation Does Not Match Response Selectivity: A Single-Cell Study of the fMRI Adaptation Paradigm. *Neuron*, 49(2):307–318.

Schaefer, A. T., Larkum, M. E., Sakmann, B., and Roth, A. (2003). Coincidence Detection in Pyramidal Neurons Is Tuned by Their Dendritic Branching Pattern. *Journal of Neurophysiology*, 89(6):3143–3154.

Schiller, J., Schiller, Y., Stuart, G., and Sakmann, B. (1997). Calcium action potentials restricted to distal apical dendrites of rat neocortical pyramidal neurons. *The Journal of Physiology*, 505(Pt 3):605–616.

Schindelin, J., Arganda-Carreras, I., Frise, E., Kaynig, V., Longair, M., Pietzsch, T., Preibisch, S., Rueden, C., Saalfeld, S., Schmid, B., Tinevez, J.-Y., White, D. J., Hartenstein, V., Eliceiri, K., Tomancak, P., and Cardona, A. (2012). Fiji: an open-source platform for biological-image analysis. *Nature Methods*, 9(7):676–682.

Schmidt, S. L., Dorsett, C. R., Iyengar, A. K., and Fröhlich, F. (2017). Interaction of Intrinsic and Synaptic Currents Mediate Network Resonance Driven by Layer V Pyramidal Cells. *Cerebral Cortex*, 27(9):4396–4410.

Schneider, D. M., Sundararajan, J., and Mooney, R. (2018). A cortical filter that learns to suppress the acoustic consequences of movement. *Nature*, 561(7723):391–395.

Serre, T., Oliva, A., and Poggio, T. (2007). A feedforward architecture accounts for rapid categorization. *Proceedings of the National Academy of Sciences of the United States of America*, 104(15):6424–6429.

Shai, A. S., Anastassiou, C. A., Larkum, M. E., and Koch, C. (2015). Physiology of Layer 5 Pyramidal Neurons in Mouse Primary Visual Cortex: Coincidence Detection through Bursting. *PLoS Computational Biology*, 11(3).

Sherman, S. M. and Guillery, R. W. (1998). On the actions that one nerve cell can have on another: distinguishing "drivers" from "modulators". *Proceedings of the National Academy of Sciences of the United States of America*, 95(12):7121–7126.

Sherman, S. M. and Guillery, R. W. (2011). Distinct functions for direct and transthalamic corticocortical connections. *Journal of Neurophysiology*, 106(3):1068–1077.

Shipp, S. (2007). Structure and function of the cerebral cortex. *Current Biology*, 17(12):R443–R449.

Silver, R. A. (2010). Neuronal arithmetic. *Nature Reviews. Neuroscience*, 11(7):474–489.

Spratling, M. W. (2002). Cortical region interactions and the functional role of apical dendrites. *Behavioral and Cognitive Neuroscience Reviews*, 1(3):219–228.

Spruston, N. (2008). Pyramidal neurons: dendritic structure and synaptic integration. *Nature Reviews. Neuroscience*, 9(3):206–221.

Spruston, N., Jaffe, D. B., Williams, S. H., and Johnston, D. (1993). Voltage- and space-clamp errors associated with the measurement of electrotonically remote synaptic events. *Journal of Neurophysiology*, 70(2):781–802.

Strich, S. J. (1968). Notes on the Marchi method for staining degenerating myelin in the peripheral and central nervous system. *Journal of Neurology, Neurosurgery, and Psychiatry*, 31(2):110–114.

Stringer, C., Pachitariu, M., Steinmetz, N., Reddy, C. B., Carandini, M., and Harris, K. D. (2019). Spontaneous behaviors drive multidimensional, brainwide activity. *Science*, 364(6437). Publisher: American Association for the Advancement of Science Section: Research Article.

Stuart, G. and Spruston, N. (1998). Determinants of Voltage Attenuation in Neocortical Pyramidal Neuron Dendrites. *Journal of Neuroscience*, 18(10):3501–3510.

Stuart, G. J. and Spruston, N. (2015). Dendritic integration: 60 years of progress. *Nature Neuroscience*, 18(12):1713–1721.

Summerfield, C., Trittschuh, E. H., Monti, J. M., Mesulam, M.-M., and Egner, T. (2008). Neural repetition suppression reflects fulfilled perceptual expectations. *Nature Neuroscience*, 11(9):1004–1006.

Sur, M., Garraghty, P. E., and Roe, A. W. (1988). Experimentally induced visual projections into auditory thalamus and cortex. *Science (New York, N.Y.)*, 242(4884):1437–1441.

Suzuki, M. and Larkum, M. E. (2020). General Anesthesia Decouples Cortical Pyramidal Neurons. *Cell*, 180(4):666–676.e13.

Südhof, T. C. (2012). Calcium Control of Neurotransmitter Release. *Cold Spring Harbor Perspectives in Biology*, 4(1).

Takahashi, N., Oertner, T. G., Hegemann, P., and Larkum, M. E. (2016). Active cortical dendrites modulate perception. *Science*, 354(6319):1587–1590.

Teeter, C., Iyer, R., Menon, V., Gouwens, N., Feng, D., Berg, J., Szafer, A., Cain, N., Zeng, H., Hawrylycz, M., Koch, C., and Mihalas, S. (2018). Generalized leaky integrate-and-fire models classify multiple neuron types. *Nature Communications*, 9(1):709. Number: 1 Publisher: Nature Publishing Group.

Thompson, J. M., Woolsey, C. N., and Talbot, S. A. (1950). Visual areas I and II of cerebral cortex of rabbit. *Journal of Neurophysiology*, 13(4):277–288.

Thomson, A. M. and Lamy, C. (2007). Functional maps of neocortical local circuitry. *Frontiers in Neuroscience*, 1(1):19–42.

Thorpe, S., Delorme, A., and Van Rullen, R. (2001). Spike-based strategies for rapid processing. *Neural Networks: The Official Journal of the International Neural Network Society*, 14(6-7):715–725.

Tigges, J., Tigges, M., Anschel, S., Cross, N. A., Letbetter, W. D., and McBride, R. L. (1981). Areal and laminar distribution of neurons interconnecting the central visual cortical areas 17, 18, 19, and MT in squirrel monkey (Saimiri). *Journal of Comparative Neurology*, 202(4):539–560. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/cne.902020407.

Tigges, J., Tigges, M., and Perachio, A. A. (1977). Complementary laminar terminations of afferents to area 17 originating in area 18 and in the lateral geniculate nucleus in squirrel monkey. *Journal of Comparative Neurology*, 176(1):87–100. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/cne.901760106.

Troyer, T. W. and Miller, K. D. (1997). Physiological gain leads to high ISI variability in a simple model of a cortical regular spiking cell. *Neural Computation*, 9(5):971–983.

Tusa, R. J., Rosenquist, A. C., and Palmer, L. A. (1979). Retinotopic organization of areas 18 and 19 in the cat. *The Journal of Comparative Neurology*, 185(4):657–678.

Ueta, Y., Otsuka, T., Morishima, M., Ushimaru, M., and Kawaguchi, Y. (2014). Multiple layer 5 pyramidal cell subtypes relay cortical feedback from secondary to primary motor areas in rats. *Cerebral Cortex (New York, N.Y.: 1991)*, 24(9):2362–2376.

van Aerde, K. I. and Feldmeyer, D. (2015). Morphological and Physiological Characterization of Pyramidal Neuron Subtypes in Rat Medial Prefrontal Cortex. *Cerebral Cortex*, 25(3):788–805. Publisher: Oxford Academic.

Van Buren, J. M. and Baldwin, M. (1958). The architecture of the optic radiation in the temporal lobe of man. *Brain: A Journal of Neurology*, 81(1):15–40.

van der Want, J. J., Klooster, J., Cardozo, B. N., de Weerd, H., and Liem, R. S. (1997). Tract-tracing in the nervous system of vertebrates using horseradish peroxidase and its conjugates: tracers, chromogens and stabilization for light and electron microscopy. *Brain Research. Brain Research Protocols*, 1(3):269–279.

Van Essen, D. C. and Maunsell, J. H. R. (1983). Hierarchical organization and functional streams in the visual cortex. *Trends in Neurosciences*, 6:370–375.

van Ooyen, A., Duijnhouwer, J., Remme, M. W. H., and van Pelt, J. (2002). The effect of dendritic topology on firing patterns in model neurons. *Network (Bristol, England)*, 13(3):311–325.

Vann, S. D., Aggleton, J. P., and Maguire, E. A. (2009). What does the retrosplenial cortex do? *Nature Reviews. Neuroscience*, 10(11):792–802.

Vetter, P., Roth, A., and Häusser, M. (2001). Propagation of action potentials in dendrites depends on dendritic morphology. *Journal of Neurophysiology*, 85(2):926–937.

Viaene, A. N., Petrof, I., and Sherman, S. M. (2011). Properties of the thalamic projection from the posterior medial nucleus to primary and secondary somatosensory cortices in the mouse. *Proceedings of the National Academy of Sciences of the United States of America*, 108(44):18156–18161.

Vogels, R. (2016). Sources of adaptation of inferior temporal cortical responses. *Cortex*, 80:185–195.

Vogels, T. P. and Abbott, L. F. (2009). Gating multiple signals through detailed balance of excitation and inhibition in spiking networks. *Nature Neuroscience*, 12(4):483–491. Number: 4 Publisher: Nature Publishing Group.

Vries, S. E. J. d., Lecoq, J., Buice, M. A., Groblewski, P. A., Ocker, G. K., Oliver, M., Feng, D., Cain, N., Ledochowitsch, P., Millman, D., Roll, K., Garrett, M., Keenan, T., Kuan, L., Mihalas, S., Olsen, S., Thompson, C., Wakeman, W., Waters, J., Williams, D., Barber, C., Berbesque, N., Blanchard, B., Bowles, N., Caldejon, S., Casal, L., Cho, A., Cross, S., Dang, C., Dolbeare, T., Edwards, M., Galbraith, J., Gaudreault, N., Griffin, F., Hargrave, P., Howard, R., Huang, L., Jewell, S., Keller, N., Knoblich, U., Larkin, J., Larsen, R., Lau, C., Lee, E., Lee, F., Leon, A., Li, L., Long, F., Luviano, J., Mace, K., Nguyen, T., Perkins, J., Robertson, M., Seid, S., Shea-Brown, E., Shi, J., Sjoquist, N., Slaughterbeck, C., Sullivan, D., Valenza, R., White, C., Williford, A., Witten, D., Zhuang, J., Zeng, H., Farrell, C., Ng, L., Bernard, A., Phillips, J. W., Reid, R. C., and Koch, C. (2018). A large-scale, standardized physiological survey reveals higher order coding throughout the mouse visual cortex. *bioRxiv*, page 359513.

Vélez-Fort, M., Rousseau, C. V., Niedworok, C. J., Wickersham, I. R., Rancz, E. A., Brown, A. P. Y., Strom, M., and Margrie, T. W. (2014). The stimulus selectivity and connectivity of layer six principal cells reveals cortical microcircuits underlying visual processing. *Neuron*, 83(6):1431–1443.

Wagor, E., Mangini, N. J., and Pearlman, A. L. (1980). Retinotopic organization of striate and extrastriate visual cortex in the mouse. *The Journal of Comparative Neurology*, 193(1):187–202.

Wahl-Schott, C. and Biel, M. (2009). HCN channels: structure, cellular regulation and physiological function. *Cellular and molecular life sciences: CMLS*, 66(3):470–494.

Wang, Q. and Burkhalter, A. (2007). Area map of mouse visual cortex. *Journal of Comparative Neurology*, 502(3):339–357. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/cne.21286.

Wang, Q., Ding, S.-L., Li, Y., Royall, J., Feng, D., Lesnar, P., Graddis, N., Naeemi, M., Facer, B., Ho, A., Dolbeare, T., Blanchard, B., Dee, N., Wakeman, W., Hirokawa, K. E., Szafer, A., Sunkin, S. M., Oh, S. W., Bernard, A., Phillips, J. W., Hawrylycz, M., Koch, C., Zeng, H., Harris, J. A., and Ng, L. (2020). The Allen Mouse Brain Common Coordinate Framework: A 3D Reference Atlas. *Cell*, 181(4):936–953.e20.

Wickersham, I. R., Lyon, D. C., Barnard, R. J. O., Mori, T., Finke, S., Conzelmann, K.-K., Young, J. A. T., and Callaway, E. M. (2007). Monosynaptic restriction of transsynaptic tracing from single, genetically targeted neurons. *Neuron*, 53(5):639–647.

Williams, S. R. and Stuart, G. J. (2000). Site independence of EPSP time course is mediated by dendritic I(h) in neocortical pyramidal neurons. *Journal of Neurophysiology*, 83(5):3177–3182.

Yamawaki, N., Li, X., Lambot, L., Ren, L. Y., Radulovic, J., and Shepherd, G. M. G. (2019). Long-range inhibitory intersection of a retrosplenial thalamocortical circuit by apical tuft-targeting CA1 neurons. *Nature Neuroscience*, 22(4):618–626.

Yamins, D. L. K. and DiCarlo, J. J. (2016). Using goal-driven deep learning models to understand sensory cortex. *Nature Neuroscience*, 19(3):356–365.

Yamins, D. L. K., Hong, H., Cadieu, C. F., Solomon, E. A., Seibert, D., and DiCarlo, J. J. (2014). Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 111(23):8619–8624.

Yilmaz, M. and Meister, M. (2013). Rapid innate defensive responses of mice to looming visual stimuli. *Current biology: CB*, 23(20):2011–2015.

Yuste, R., Gutnick, M. J., Saar, D., Delaney, K. R., and Tank, D. W. (1994). Ca2+ accumulations in dendrites of neocortical pyramidal neurons: An apical band and evidence for two functional compartments. *Neuron*, 13(1):23–43.

Zador, A. M. (2019). A critique of pure learning and what artificial neural networks can learn from animal brains. *Nature Communications*, 10(1):1–7.

Zanto, T. P., Rubens, M. T., Thangavel, A., and Gazzaley, A. (2011). Causal role of the prefrontal cortex in top-down modulation of visual processing and working memory. *Nature neuroscience*, 14(5):656–661.

Zeki, S. (1990). A century of cerebral achromatopsia. *Brain: A Journal of Neurology*, 113 (Pt 6):1721–1777.

Zeki, S. M. (1969). Representation of central visual fields in prestriate cortex of monkey. *Brain Research*, 14(2):271–291.

Zeki, S. M. (1973). Colour coding in rhesus monkey prestriate cortex. *Brain Research*, 53(2):422–427.

Zeki, S. M. (1978). The cortical projections of foveal striate cortex in the rhesus monkey. *The Journal of Physiology*, 277:227–244.

Zemankovics, R., Káli, S., Paulsen, O., Freund, T. F., and Hájos, N. (2010). Differences in subthreshold resonance of hippocampal pyramidal cells and interneurons: the role of h-current and passive membrane characteristics. *The Journal of Physiology*, 588(Pt 12):2109–2132.

Zingg, B., Chou, X.-L., Zhang, Z.-G., Mesik, L., Liang, F., Tao, H. W., and Zhang, L. I. (2017). AAV-Mediated Anterograde Transsynaptic Tagging: Mapping Corticocollicular Input-Defined Neural Pathways for Defense Behaviors. *Neuron*, 93(1):33–47.