

Investigating predictive coding as a principle of function in sensory systems



Sofia Deleniv
Balliol College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy
Michaelmas Term 2019

Abstract

According to the theory of predictive coding, the brain does not generate sensations simply by 'sponging up' bottom-up information arriving from the senses. Instead, cortical circuits are fundamentally wired to actively make predictions on the basis of experience, communicate these in a top-down fashion to areas with access to less global information, and propagate the residual errors that ensue when predictions are compared with reality (Friston, 2005). As these error signals feed forward through the processing hierarchy, they allow each successive processing stage to refine its predictions, allowing future stimulation to be better predicted and, in some ways, better 'understood' by the brain's internal model of the world. The critical predictive and error signals are thought to stem from anatomically distinct neural populations, with predictions represented both in the deep and superficial cortical layers, and errors localised to superficial sites (Kanai et al., 2015). Evidence that the superficial and deep layers are, to some degree, specialised for bottom-up and top-down communications, affiliated with activity in the γ (30-100 Hz) and α - β (8-25 Hz) frequency bands respectively (Bastos et al., 2015; Michalareas et al., 2016) has also raised the possibility that predictions and errors are encoded in different spectral channels. The goal of this thesis work was to examine the hypothesised anatomical and functional affiliation of predictive coding variables with distinct cortical layers and their oscillatory signatures. This question was pursued on a range of scales - from the behaving animal and isolated *in vitro* brain circuit, to the single cell.

In the first chapter, I piloted two behavioural paradigms in which mice made 2-Alternative Forced Choice judgements about visual stimuli while auditory cues predicting stimulus orientation were used to incentivise sensory predictions. The goal was to eventually conduct translaminar multi-electrode recordings to test whether neuronal activity patterns in the mouse visual cortex carry decodable predictive representations, and whether these are localised to any cortical layer or frequency band. I found that, while mice successfully acquired the first, stimulus location discrimination task, their behaviour showed no evidence of using auditory cues to produce orientation-specific stimulus predictions. Furthermore, mice failed to achieve above-chance performance on the second, same/different orientation judgement task. Thus, neither task was appropriate for exploring predictive coding in the visual system of the behaving mouse.

In the second chapter, I conducted two series of experiments, combining local field potential recordings with optogenetics, to explore whether the proposed oscillatory signatures of predictive coding can be found in a local circuit of mouse primary somatosensory cortex (S1) *in vitro*. Slices of S1 expressing the photosensitive channelrhodopsin under the CaMKII promoter received stimulus patterns which followed particular statistical norms. On a minority of trials, norms were violated with stimuli which either deviated from the expected stimulation amplitude (Exp 1) or locus of stimulation within the barrel column (Exp 2). Both experiments provided no evidence that power or peak frequency in the γ -band - the alleged carrier of prediction error signals - was sensitive to stimulus (un)predictability. Furthermore, I obtained no evidence that power in the α -band represented sensory predictions, on an acute or long-term temporal basis. α -power did not differentiate between conditions with diverging likely futures, and failed to scale with increasing stimulation history which goes hand-in-hand with accrual of sensory expectation (van Ede et al., 2010; van Pelt et al., 2016).

In the third chapter, I combined patch-clamp electrophysiology and computational modelling to explore how single cell computation is affected by the γ rhythm - a critical information carrier in the predictive coding scheme. Previous research has shown that γ -modulated excitation boosts the input-output slopes of cortical pyramidal cells over and above what can be achieved with arrhythmic excitation (Sohal et al., 2009). However, it is still unclear whether the same enhanced response properties can be achieved with γ -rhythmic inhibition, which is critical for inducing network γ oscillations (Mann et al., 2005; Traub et al., 1996; Whittington et al., 1995), and whether the γ -rhythm is unique in this respect. This question was explored by examining spike patterns of putative pyramidal cells in mouse S1 in response to excitatory current pulses of increasing amplitudes, while simulated rhythmic inhibitory post-synaptic potentials (IPSGs) were administered with varying degrees of γ -power, using dynamic clamp. The experiment revealed that, with increasing γ -power in the IPSG, pyramidal cells achieved higher maximal spike rates, as well as exhibiting higher input-output slopes and spike timing fidelity. A further experiment was conducted to compare the effects of varying the frequency of inhibition, finding that γ -frequency inhibition (40 Hz) appeared to be optimal in its ability to boost the input-output properties of pyramidal cells. A computational model of the single cell was able to replicate these effects by simulating conditions of low M-type K^+ channel conductance and high flux through voltage-gated Ca^{2+} channels, shedding some light on the potential underpinnings of γ 's unique effects on single-cell computation.

Acknowledgements

As my time in Oxford comes to an end, there are many people to whom I owe enormous thanks for supporting me throughout my DPhil. My first, and biggest, expression of gratitude goes out to my supervisor, Ed Mann. He constantly supported me in my pursuit of the research questions I personally cared about, and always made himself available to have a discussion or give some advice, which I appreciate immensely. I would also like to thank my second supervisor, Louise Upton, for taking the time to teach me to perform rodent surgery. I also owe thanks to my third supervisor, Mark Stokes, for welcoming me into The Stokes Lab - my second laboratory, which was a wonderful source of social and intellectual stimulation throughout my DPhil. A special thank you goes out to Pantelis Antonoudiou - a former DPhil student in our laboratory, who dedicated many days of his time to teaching me to perform *in vitro* slice preparations and LFP recordings. Pantelis' generosity, kindness, and patience is unique amongst humans - he is what people in Ukraine (where I am from) call 'a golden human being'. I would also very much like to thank Eelke Spaak, a former member the Stokes Lab, for his friendship, and for making himself so available for discussions about statistics and analysis, even after moving countries (he is now a post-doctoral researcher at the Donders Institute).

A less academic, but no less important, acknowledgement goes out to fellow DPhil student, and member of the Stokes Lab, Lev Tankelevich. He always knew how to motivate me to keep going when research was tough, and his friendship was a constant source of joy in my life throughout the years of my DPhil. I would also like to extend my thanks to Matthew Cooper and Sonia Spitzer - my office mates of several years, whose companionship and banter was irreplaceable. Thank you, Matthew for all the lunches, Game of Thrones discussions, and American road trips we had. Thank you, Sonia for your friendship.

Finally, I would like to thank Dominic Owens - a fellow scientist and my partner - for years of support. No words can describe how valuable you have been, and continue to be, in my life.

List of abbreviations and symbols

α	Alpha
β	Beta
γ	Gamma
2AFC	2-alternative forced choice
AIS	Active information storage
AMPA	α -amino-3-hydroxy-5-methyl-4-isoxazolepropionic acid
BOLD	Blood oxygenation level-dependent
CaMKII α	Ca ²⁺ /calmodulin-dependent protein kinase II, α -subunit
ChR2	Channelrhodopsin-2
cpd	Cycles per degree
DMD	Digital micromirror device
ECoG	Electrocorticography
EEG	Electroencephalography
E/IPSC	Excitatory / inhibitory post-synaptic current
E/IPSG	Excitatory / inhibitory post-synaptic conductance
E/IPSP	Excitatory / inhibitory post-synaptic potential
ERP	Event-related potential
fMRI	Functional magnetic resonance imaging
FS	Fast-spiking
GABA	γ -aminobutyric acid
LFP	Local field potential
LTP/D	Long term potentiation / depression
MEG	Magnetoencephalography
MMN	Mismatch negativity
MVPA	Multi-voxel pattern analysis
NMDA	N-methyl-D-aspartic acid
RGC	Retinal ganglion cell
SST	Somatostatin

Contents

1	CHAPTER 1 - INTRODUCTION	1
1.1	WHAT IS PREDICTIVE CODING?	1
1.2	EVIDENCE FROM COGNITIVE AND BEHAVIOURAL NEURO-SCIENCE	2
1.2.1	Perceptual priming	2
1.2.2	Mismatch negativity (MMN)	4
1.2.3	Repetition suppression	5
1.3	EVIDENCE FROM STRUCTURAL AND FUNCTIONAL BRAIN ARCHITECTURE	8
1.3.1	Structural hierarchies	8
1.3.2	Extra-classical receptive fields	13
1.4	EVIDENCE FROM OSCILLATORY ACTIVITY	15
1.4.1	What are oscillations?	15
1.4.2	Oscillations as channels for bottom-up and top-down signalling	17
1.4.3	Oscillations as carriers of predictions and errors	19
1.5	RESEARCH GOALS	22
2	CHAPTER 2	25
2.1	INTRODUCTION	25
2.1.1	Broad conceptual goals and objectives: oscillations and the laminar origins of predictions	25
2.1.2	Mouse vision as a model sensory system	26
2.1.3	Trialling a mouse visual decision-making paradigm	27

2.2	METHODS	28
2.2.1	Animals	28
2.2.2	Experimental set-up and stimuli	28
2.2.3	Behavioural training conditions	29
2.2.4	Pilot 1: Location discrimination task structure	30
2.2.5	Pilot 1: Introducing auditory cues to evoke sensory predictions	30
2.2.6	Pilot 2: Same-different judgement task	32
2.3	RESULTS - PILOT 1	33
2.3.1	Recap and objectives	33
2.3.2	Mice can learn to discriminate stimulus locations	33
2.3.3	Degrading stimulus contrast impairs discrimination accuracy	42
2.3.4	Degrading stimulus contrast slows discrimination decisions	42
2.3.5	Invalid predictive cues do not impair discrimination accuracy	43
2.3.6	Invalid predictive cues do not slow decisions	44
2.3.7	Mice may have held false beliefs about predictive cue meaning	46
2.4	RESULTS - PILOT 2	49
2.4.1	Recap and objectives	49
2.4.2	Mice cannot learn to make same-different judgements	49
2.5	DISCUSSION	53
2.5.1	Mice can make visual location discrimination decisions under head-fixed conditions	53
2.5.2	Technical caveats	55
2.5.3	Mice do not act on predictive information about task-irrelevant dimensions	56
2.5.4	Mice cannot make same-different judgements under head-fixed conditions	58
2.5.5	Did mice struggle to perceptually discriminate orthogonal orientations?	58
2.5.6	Were mice cognitively incapable of comprehending task requirements?	59
3	CHAPTER 3	61
3.1	INTRODUCTION	61
3.1.1	Broad conceptual goals and objectives: can predictive coding be local?	61
3.1.2	Experiment 1: Amplitude prediction errors	62

3.1.3	Experiment 2: Spatiotemporal prediction errors	63
3.1.4	The local field potential (LFP) as a measure of population activity	65
3.2	METHODS	68
3.2.1	Animals	68
3.2.2	Viral injections into mouse somatosensory cortex	68
3.2.3	Acute brain slice preparation	69
3.2.4	Assessing slice viability	70
3.2.5	Recording local field potentials (LFPs) in somatosensory cortex	71
3.2.6	Experiment 1 protocol: Evoking temporal sequence predictions <i>in vitro</i>	72
3.2.7	Experiment 2 protocol: Evoking spatiotemporal predictions <i>in vitro</i>	75
3.3	DATA ANALYSIS	77
3.3.1	Pre-processing LFP traces	77
3.3.2	Examining power across the frequency spectrum	77
3.3.3	Outlier removal	77
3.3.4	Statistical testing for differences	77
3.3.5	Examining cross-trial site coherence across the spectrum - Experiment 2	78
3.3.6	Examining cross-trial phase lag consistency - Experiment 2 . .	78
3.3.7	Cross-frequency coupling analysis - Experiment 2	78
3.3.8	Statistical analysis of coherence and cross-frequency coupling - Experiment 2	80
3.4	RESULTS - EXPERIMENT 1	82
3.4.1	Recap and objectives	82
3.4.2	Statistical analysis: controlling for cross-slice and cross-trial variability	82
3.4.3	Power in the γ -band is responsive to stimulus amplitude	84
3.4.4	Power in the γ -band is not responsive to stimulus probability .	85
3.4.5	Slope of temporal change in γ -band power is not responsive to stimulus probability	86
3.4.6	Peak frequency in the γ -band is not responsive to deviant probability	89
3.4.7	Slope of temporal change in γ -band peak frequency is not responsive to deviant probability	89
3.4.8	Power in the α -band is not responsive to deviant probability .	91
3.4.9	Slope of temporal change in α -band power is not responsive to deviant probability	92

3.5 RESULTS - EXPERIMENT 2	94
3.5.1 Recap and objectives	94
3.5.2 Analytical approach to exploring effects of deviant stimulation	94
3.5.3 γ -band power: no significant effects of deviation from the spatiotemporal norm	97
3.5.4 Peak frequency in the γ -band: no effect of deviation from the spatiotemporal norm	97
3.5.5 Activity in the α -band: no effect of deviation from the spatiotemporal norm	98
3.5.6 Effects of sequence reversal on γ -band power and peak frequency, and α -band power	101
3.5.7 Temporal effects of stimulus deviation on γ -band and α -band activity	103
3.5.8 Stimulus deviation does not significantly influence coherence between L2/3 and L4/5	105
3.5.9 Stimulus deviation does not influence strength of L2/3 - L4/5 phase-amplitude coupling.	108
3.6 DISCUSSION	110
3.6.1 Activity in the γ -band does not carry prediction errors in a local S1 circuit	110
3.6.2 Potential dissociation between peak frequency and power in the γ -band	112
3.6.3 α -band power bears no significant relationship to sensory predictions or errors in a local S1 circuit	113
3.6.4 Technical caveats	115
4 CHAPTER 4	117
4.1 INTRODUCTION	117
4.1.1 How does the γ rhythm affect information transfer?	117
4.1.2 Frequency-dependent scaling of principal cell responses	118
4.2 METHODS	121
4.2.1 Animals	121
4.2.2 Acute brain slice preparation	121
4.2.3 Intracellular recordings	121
4.2.4 Dynamic clamp	125
4.2.5 Experiment 1 protocol: simulating γ -rhythmic inhibition varying in power	127

4.2.6	Experiment 2 protocol: simulating γ -rhythmic inhibition varying in frequency	130
4.2.7	Modelling the single cell: model architecture	132
4.2.8	Modelling the effects of rhythmic inhibition	133
4.3	RESULTS: EXPERIMENT 1	135
4.3.1	Recap and objectives	135
4.3.2	Input resistance and stability	135
4.3.3	Inhibitory γ -power increases principal cell max. firing rates . .	136
4.3.4	Inhibitory γ -power boosts principal cell response scaling	137
4.3.5	Inhibitory γ -power increases principal cell rheobase	137
4.3.6	Inhibitory γ -rhythm increases principal cell spike fidelity . . .	138
4.4	RESULTS: EXPERIMENT 2	142
4.4.1	Recap and objectives	142
4.4.2	Input resistance and stability	142
4.4.3	γ -frequency inhibition increases principal cell max. firing rates	143
4.4.4	γ -frequency inhibition increases principal cell response scaling .	144
4.4.5	Frequency of rhythmic inhibition influences rheobase	145
4.4.6	Frequency of rhythmic inhibition affects principal cell spike fidelity	147
4.5	RESULTS: SINGLE-CELL MODEL	149
4.5.1	Recap and objectives	149
4.5.2	Roles of A-type and M-type K^+ currents in principal cell response scaling.	149
4.5.3	Role of voltage-gated Ca^+ channel conductance in principal cell response scaling.	151
4.6	DISCUSSION	156
4.6.1	Simulating perisomatic γ -band inhibition produces multiplicative increases in principal cell gain	156
4.6.2	Relative importance in rheobase and maximal spike rate in input-output slopes	158
4.6.3	Towards a better understanding of γ	159
5	GENERAL DISCUSSION	160
5.1	Investigating predictive coding in the behaving mouse	160
5.2	Investigating predictive coding <i>in vitro</i>	164
5.3	Towards a mechanistic insight of γ -oscillations and single cell computation	166

List of Figures

2.1	Location discrimination task acquisition by two females.	36
2.2	Location discrimination task acquisition by Male 1.	39
2.3	Location discrimination task acquisition by Male 2.	40
2.4	Reaction time and accuracy data under valid and invalid predictive cueing conditions.	45
2.5	Post-cue licking patterns are shaped by the false belief that auditory cues predict stimulus location.	48
2.6	Same-different discrimination task performance	51
3.1	Experimental protocol for evoking temporal expectations <i>in vitro</i> . .	74
3.2	Experimental protocol for evoking spatiotemporal expectations <i>in vitro</i> . .	76
3.3	Illustrative example of stimulus-evoked LFP in mouse S1 <i>in vitro</i> and its spectral decomposition, during standard ramp stimulus delivery. .	83
3.4	Effects of deviant amplitude stimulation, and its predictability, on power in the γ -band, <i>in vitro</i>	87
3.5	Effects of deviant amplitude stimulation, and its predictability, on temporal changes in γ -power, <i>in vitro</i>	88
3.6	Effects of deviant amplitude stimulation, and its predictability, on peak frequency (PF) in the γ -band, <i>in vitro</i>	90
3.7	Effects of deviant amplitude stimulation, and its predictability, on power in the α -band, <i>in vitro</i>	93
3.8	Raw power across the frequency spectrum for extracellular recordings acquired from L2/3 and L4/5 on the same single randomly selected trials from 2 experiments.	96
3.9	Effects of deviant forward-skips in a stimulation sequence on γ -band power and peak frequency, and α -band power, <i>in vitro</i>	99
3.10	Effects of deviant back-skips in a stimulation sequence on γ -band power and peak frequency, and α -band power, <i>in vitro</i>	100
3.11	Effects of stimulation sequence reversal on γ -band and α -band power, <i>in vitro</i>	102
3.12	Temporal effects of deviant and control spatiotemporal stimulation on γ -power, peak frequency, and α -power, <i>in vitro</i>	104
3.13	Computing coherence between L2/3 and L4/5 signals, <i>in vitro</i> , in a sample experiment.	106

3.14	Coherence and phase-locking consistency between L2/3 and L4/5 signals, <i>in vitro</i> , on trials featuring forward-skip deviations versus standard sequences.	107
3.15	Modulation index estimates for cross-frequency coupling between LFP signals in L2/3 and L4/5.	109
4.1	Electrophysiological phenotyping of a putative pyramidal cell.	123
4.2	Schematic illustration of the dynamic clamp technique.	126
4.3	Generation of γ -oscillatory synaptic inhibitory conductance trains of varying γ -power, through manipulation of rhythmicity and synchrony.	129
4.4	Experimental protocol for probing the effects IPSG γ -power on principal cell response properties.	131
4.5	Effects IPSG γ -power on principal cell response properties in L2/3 of S1.	140
4.6	Effects of IPSG γ -power on principal cell spike fidelity.	141
4.7	Effects of IPSG frequency on principal cell response properties in L2/3 of S1.	146
4.8	Effects of IPSG frequency on principal cell spike fidelity.	148
4.9	Effects of A-type and M-type K^+ currents on the relationship between frequency of inhibition and the response properties of a model principal cell.	153
4.10	Effects of Ca^{2+} current on the relationship between frequency of inhibition and the response properties of a model principal cell.	154
4.11	Effects of slow calcium-activated K^+ current on the relationship between frequency of inhibition and spike output in a model principal cell.	155

INTRODUCTION

1.1 WHAT IS PREDICTIVE CODING?

In a 1993 publication documenting some physiological properties of neurons in the monkey primary visual cortex (V1), Grosf and colleagues argued that their findings were in support of the idea that “V1 is an image-processing area, rather than simply a spatially transformed replication of the retina” (Grosf et al., 1993). I am quoting this sentence not necessarily because it is a revelation, but because it struck me that the idea of a cortical area being there to not merely reproduce the retinal image appears to be something that had to be actively advocated for as recently as the early 1990s.

What would a visual system be if the goal of each processing stage was purely to reproduce the information represented at the previous milestone of information flow? Arguably, it would be a series of cameras taking photographs of photographs taken by the cameras before. Where meaningful image processing would come in is unclear. The theory that the sensory brain cannot make sense of the world by passively processing inputs in a bottom-up fashion has been a central tenet in psychology since the mid-19th century (Gregory, 1968; Gregory, 1970; Helmholtz, 1860). According to this idea, the brain makes an ongoing effort to actively generate an internal representation of the external world in order to infer the underlying sources of patterns of external stimulation it encounters throughout life. Without such active mapping of internal models onto external inputs – something psychologists refer to as ‘perceptual hypotheses’ (Gregory, 1968; Schlagbauer et al., 2017) – the brain would arguably represent nothing more than a bitwise description of the world that is void of the meaning necessary for organisms to make decisions and take actions. Thus, the theory posits that our implicit expectations, generated through repeated exposures to diverse external inputs, fundamentally shape how we interpret each pattern of stimulation to arrive through the periphery.

A more current take on this idea - predictive coding – attempts to ground the more-or-less philosophical concept of a ‘perceptual hypothesis’ in neurophysiological reality (Friston, 2005; Friston & Kiebel, 2009; Kanai et al., 2015). It posits that cortical circuits are fundamentally wired to make predictions and communicate them to cortical regions situated lower in the perceptual processing hierarchy – regions with response properties tuned to more local-level information - in order to constrain their responses in accordance with the most likely global interpretation of an incoming pattern of sensory stimulation. This allegedly allows sensory systems to reduce informational redundancy in the neural code, which goes hand in hand with reducing the metabolic costs of generating sensations, by representing and propagating forward

primarily those inputs which were not successfully predicted (i.e. prediction errors). In turn, the error signals that are constantly being propagated through the processing hierarchy as we navigate our environment allow the brain to continuously refine its internal model of the world and thus make increasingly more accurate interpretations of stimulation patterns. According to the logic of predictive coding, the content of sensory cortical activity should therefore contain two primary pieces of information: i) internal predictions of upcoming inputs formed on the basis of experience, and ii) error signals resulting from an ongoing comparison between these predictions and the feedforward input carrying evidence from the environment.

The idea that sensory cortices actively generate predictions (or ‘perceptual hypotheses’; Gregory, 1968) and respond to their discrepancies with reality is supported by a diverse body of research examining how the brain responds to sensory expectations that are either implicitly generated as a result of sequence learning (Garrido et al., 2008; Meyer & Olson, 2011; Meyer et al., 2014), exposure to prime stimuli (Amado et al., 2018; Barbosa & Kouider, 2018) or explicit predictive cues within the task (Alink et al., 2010; Bichot et al., 2005; Gardner & Fontanini, 2014; Stokes et al., 2009). A characteristic finding of such studies is that sensory cortices actively generate patterns of anticipatory activity in neuronal populations representing upcoming stimuli (Kok et al., 2014; Stokes et al., 2009), and react to the evidence by responding to departures from the predictions engendered by these activations (Bell et al., 2016; Auksztulewicz & Friston, 2015; Garrido et al., 2008; Komatsu et al., 2015), while encoding hypothesis-confirming inputs with superior resolution (Bell et al., 2016; Summerfield et al., 2008). What follows is a summary of the main streams of research delving into the question of whether predictive coding provides a good account of perception and its neural underpinnings, obtained across a range of techniques, from behaviour, to functional magnetic resonance imaging (fMRI), magneto-/electro-encephalography (MEG/EEG), electrocorticography (ECoG), and single-unit electrophysiology.

1.2 EVIDENCE FROM COGNITIVE AND BEHAVIOURAL NEUROSCIENCE

1.2.1 PERCEPTUAL PRIMING

One behavioural phenomenon which lends some support to the predictive coding framework is priming – the tendency for stimuli to evoke quicker and less error-prone responses if they are preceded by a related stimulus (i.e. a prime), even if the prime is masked or presented briefly enough so as to evade conscious perception (Shigeno, 2017). Notably, primed stimuli have also been found to lead to weaker neural activity (termed ‘repetition suppression’) compared to those that were either not preceded by a prime or not sufficiently related to it (Barbosa & Kouider, 2018). While it is likely that some priming effects are underpinned by bottom-up neuronal

adaptation (especially in cases where the prime and stimulus have a high degree of feature similarity, and thus can be expected to recruit highly overlapping neuronal populations), a body of evidence now argues that they can be a by-product of active predictive processes taking place in the brain. Specifically, it has been proposed that priming effects might occur because of the bias produced by the prime towards a particular interpretation of the next input, which materialises in more efficient processing of and faster behaviour towards stimuli that match the implicit prediction (Gotts et al., 2012).

One finding which supports this notion is that priming can occur even when the prime and subsequent probe stimulus stem from different sensory modalities – a phenomenon termed cross-domain priming (Amado et al., 2018; Zhang et al., 2010). In one fMRI experiment showing this effect, participants heard spoken names of famous individuals followed by brief presentations of face images, on which they were asked to make recognition judgements (Amado et al., 2018). Faces which corresponded to the spoken primes were recognised significantly faster than those that did not, and evoked reduced BOLD signals in the fusiform and occipital face areas, known to harbour face-selective neurons (Kanwisher et al., 1997; Pitcher et al., 2011). This observed cross-domain priming effect argues strongly against the idea that local adaptation can account for all cases of priming, as primes and probe stimuli from different modalities will by definition recruit different neuronal populations, and thus make purely bottom-up effects implausible. Instead, it provides a convincing case that active predictive processes are at work in sensory systems, with auditory primes capable of biasing visual areas towards particular interpretations of images. Furthermore, the fMRI findings lend credence to the idea that sensory responses in a predictive brain should be weakest when the internal model of input matches reality.

One potential caveat of cross-priming effects is that they often involve a consciously perceived prime (Amado et al., 2018; Young et al., 1988), and thus arguably engage active cognitive processes. This leaves much room for speculation about the extent to which prediction is an inherent and automatic part of sensory processing, as opposed to being mediated by cognition – and thus, to some degree, intelligence. I would argue that a process proposed to constitute a fundamental coding property of sensory systems ought to be proven independent of conscious cognitive activity, and applicable to the lowest organisms in possession of a brain. There is some limited evidence that cross-domain priming is observed under conditions where the prime is subliminal owing to extremely brief duration or heavy visual masking (Faivre et al., 2014), which indicates a reasonable likelihood that the predictive processes which likely contribute to priming effects can indeed occur without cognitive mediation.

1.2.2 MISMATCH NEGATIVITY (MMN)

Another phenomenon which has been argued to reflect predictive processes in the sensory brain is the mismatch negativity (MMN) – an electrophysiological hallmark of ‘surprise’ in response to violations of sensory regularities that has been captured in EEG and ECoG recordings in humans (Garrido et al., 2008), monkeys (Komatsu et al., 2015) and rats (Harms et al., 2014). The MMN is typically studied using the so-called roving paradigm (Aukstulewicz & Friston, 2015; Garrido et al., 2008; Komatsu et al., 2015; Wacongne et al., 2011). This protocol involves presenting participants with a sensory stream repeating the same stimulus (eg. A-A-A) or stimulation pattern (eg. AB-AB-AB), peppered with occasional oddball stimuli or unexpected switches to different stimulus trains. Across the aforementioned species, such unexpected oddballs or stream switches have been found to elicit a mismatch negativity (MMN) – a component of the event-related potential which consistently emerges within 150-250 ms of the violating stimulus onset (Fishman, 2014; Garrido et al., 2008; Harms et al., 2014; Komatsu et al., 2015).

The MMN appears to be a universal feature of the brain’s response to expectancy violation – it occurs in response to discrepancies with a variety of expected stimulus features, such as duration (Jacobsen and Schroger, 2003), amplitude (Naatanen et al., 1987), and inter-stimulus interval (Ford and Hillyard, 1981). While the MMN is generally not observed during deep sleep, it is present in other lighter stages of sleep (Loewy et al., 1996; Sabri et al., 2000), and appears to be insensitive to an individual’s attentional state (Sculthorpe et al., 2009), which indicates that the phenomenon is rooted in some degree of automaticity. Furthermore, since the MMN was originally documented in the auditory modality (Hillyard et al., 1973), it has now been observed in the visual (Kremlacek et al., 2016), olfactory (Krauel et al., 1999), and somatosensory domains (Kekoni et al., 1997), implying that the MMN is a domain-general neurophysiological signature of sensory ‘surprise’.

The MMN was originally thought to reflect bottom-up change detection as a result of release from local adaptation brought on by stimulus repetition (Schroger, 1998). However, it has more recently been proposed to result from a comparative process signalling departures of inputs from top-down expectancies established by environmental regularities (Stefanics et al., 2014; Winkler, 2007). This latter hypothesis predicts that the MMN should be observed not only under conditions of a low-level sensory deviation (eg. A-A-A-B), where release from local adaptation is the most parsimonious account, but also under circumstances where the violation concerns a more complex or abstract rule. Indeed, the MMN has recently been found in response to violations of higher-level stimulus regularities (Horvath & Winkler, 2004; Paavilainen et al., 2018). In one such experiment (Paavilainen et al., 2018), the oddball stimuli were pairs of tones presented in descending frequency order, embedded in an auditory stream featuring tone pairs in ascending frequency order. The finding that such oddball pairs characterised by an abstract relationship evoked robust

MMNs presents a convincing case that this electrophysiological phenomenon is not exclusively a function of local neuronal adaptation, but takes into account higher-level statistical properties of the input stream. Furthermore, given that participants' ears were being fed two distinct auditory streams with opposing within-pair frequency relationships, it is unlikely that violations of these statistical relationships were being monitored concurrently in both streams at a cognitive level.

In the same year, researchers presented evidence that simulated sensory prediction errors accurately map onto the amplitudes of visual event-related potentials (ERPs) and MMN responses in a visual roving paradigm (Stefanics et al., 2018). In this study, participants viewed a visual stream consisting of faces of a particular colour (red or green) and emotional state (happy or fearful), which exhibited occasional unexpected stream switches. An ideal predictive observer was modelled in order to quantify the correspondence between its output and EEG recordings obtained from human participants. This model observer implemented a hierarchical belief updating system in which beliefs about external states were generated and refined both on a first-order level (eg. probability of seeing a red versus green, or happy versus fearful face in the stimulus stream) and on a second-order level, which was based on conditional statistics (eg. probability of seeing a red face given the previous stimulus). It was found that, on a trial-wise basis, visual ERPs were significantly modulated by prediction errors computed at the second-order level of the model. Furthermore, trials featuring the top and bottom 10% of model-based prediction error magnitudes mapped onto those with the strongest and weakest mismatch negativities, respectively. These results not only corroborate the evidence that the MMN responds to violations of expectancies shaped by higher-level regularities (Paavilainen et al., 2018), but draw direct parallels between neuronal responses (reflected in the ERP) and sensory prediction error.

1.2.3 REPETITION SUPPRESSION

In recent years, even phenomena considered to be proxies for non-predictive habituation have become re-examined through the lens of predictive coding. One such phenomenon is repetition suppression – the observation that stimulus-evoked neuronal responses weaken over time in response to repeated stimulus exposure (De Baene & Vogels, 2010; Krekelberg et al., 2006; Summerfield et al., 2008). Repetition suppression has been documented in a huge diversity of circumstances, from weakened single-unit responses in monkey superior temporal sulcus neurons to viewing repeated hand actions (Kuravi et al., 2016), to suppressed blood oxygenation level-dependent (BOLD) signals in human visual cortex when repeatedly viewing colours from the same linguistic category (Bird et al., 2012). Traditionally, repetition suppression was viewed as an index of neuronal fatigue – an assumption implicit in the extensive use of repetition suppression paradigms as a window into the representational tuning properties of neurons across sensory and cognitive domains (Barron et al., 2016; Boynton & Finney, 2003; Grill-Spector et al., 2006; Grill-Spector & Malach, 2001). However,

this fatigue-based interpretation may not be the whole story, according to a more recent perspective which implicates top-down predictive mechanisms (Aukstulewicz & Friston, 2016; Summerfield et al., 2008; Todorovic et al., 2011).

Discussion of the potentially predictive nature of repetition suppression was, in part, triggered by the finding that repetition suppression of BOLD activity in the human fusiform gyrus is sensitive to the relative frequency of visual stimulus repetitions across experimental blocks (Summerfield et al., 2008). That is, in blocks where stimulus repetitions were relatively frequent compared to stimulus alternations (controlling for the total number of exposures to the critical stimulus), sensory responses were suppressed to a higher degree. This led the authors to conclude that repetition suppression is a by-product of fulfilled sensory expectations - i.e. conditions of low prediction error. The finding that repetition suppression is especially pronounced when repetitions are statistically common has been replicated by others using fMRI (Kronbichler et al., 2018; Kovacs et al., 2012), EEG (Summerfield et al., 2011) and MEG recordings of human auditory cortex (Todorovic et al., 2011), although there have been some notable failures to replicate the effect on single unit and LFP adaptation in monkeys (Kaliukhovich & Vogels, 2010).

The issue of whether repetition suppression corresponds to a rapid decline in sensory prediction error is further complicated by the finding that multi-voxel fMRI patterns exhibiting such suppression co-exist within the same sensory areas with patterns exhibiting repetition enhancement (de Gardelle et al., 2013). The researchers behind this finding argued that this might reflect the existence of specialised neuronal populations representing prediction errors (which should exhibit repetition suppression) and those accumulating sensory expectations (which ought to respond more strongly as the evidence for a stimulus repeating itself grows). However, it is unclear how the course-grained snapshot of neuronal activity that is afforded by an fMRI multi-voxel pattern should map onto the much finer anatomical distinctions made by the predictive coding framework. While two co-existing fMRI patterns are alleged to provide evidence for distinct representational codes relevant to predictive coding, this framework posits that error- and prediction-encoding neurons are, to a great extent, segregated within the cortical column (Kanai et al., 2015). It is likely that the issue of dissecting the potential anatomical basis of predictive coding surpasses the resolution capabilities of neuroimaging techniques. To my knowledge, no single-unit recording studies have thus far explored potential anatomical segregation between neurons exhibiting differential responses in a stimulus repetition context.

Despite their limitations, neuroimaging studies have shed considerable light on the functional role of predictions in sensory processing. The evidence they have produced suggests that predictive modulation of neuronal activity increases the quality of expected stimulus representations, despite producing suppressed responses (Kok et al., 2012; Noudoost et al., 2017). To show this, researchers applied multivariate pattern analysis (MVPA) to fMRI data obtained from participants making orien-

tation discrimination decisions about stimuli after receiving either valid or invalid predictive auditory cues. They found that, while validly cued gratings evoked significantly weaker BOLD responses in V1 (again, supporting the notion that response suppression may represent lower prediction error), the resulting multi-voxel patterns of activity enabled superior decoding of stimulus orientation compared to when the cue had been invalid (and the stimulus thus unexpected). This improvement in decoding performance correlated with better participant performance, which suggests that the changes to representational quality picked up by the fMRI scans were of consequence for behaviour. More recently, single-unit recordings from monkey inferotemporal cortex have corroborated this evidence, finding that sensory expectations produced suppressed stimulus-evoked responses but resulted in a representational advantage, as indexed by superior stimulus decoding by neuronal firing rates (Bell et al., 2016).

An explanation put forward for this effect is that predictive signals might alter the gain of recipient neurons, producing a representational 'sharpening' effect for expected signals (Kok et al., 2012). This might be implemented by co-recruitment of different forms of inhibition which carry out both divisive (eg. fast spiking interneurons; Lee et al., 2012) and subtractive (eg. somatostatin-positive interneurons; Sturgill & Isaacson, 2015) gain control on principal cortical cells, producing an overall suppression of responses along with steeper input-output curves. However, we currently have little to no direct evidence on the possibility of predictive signals recruiting varying forms of inhibition, and to varying outcomes. A further complication is that research suggests that neurons exhibiting strong divisive effects tend to receive weaker subtractive modulation (Arandia-Romero et al., 2016), which raises questions about how these effects might co-occur during sensory predictions. Notwithstanding our lack of mechanistic understanding of enhanced representational quality for expected stimuli, these observations offer some support to the idea that predictive coding might optimise sensory systems for saving on metabolic costs, by reducing the amount of activity spent on encoding the necessary information.

Unfortunately, neither the human nor monkey study permits any conclusions about whether such representational benefits would emerge under conditions in which participants do not actively attend to the stimuli, and thus do not cognitively engage with the statistical regularities of the sensory stream. Furthermore, it is unclear, from a logical standpoint, whether it is stimuli which match expectations or those which deviate from them that should benefit from greater encoding quality. As a counterexample to the aforementioned benefits of meeting expectations, one study observed that receptive fields in cat primary auditory cortex (A1) exhibited a two-fold increase in frequency resolution when exposed to frequency deviants in a stream of auditory tones (Ulanovsky et al., 2003).

1.3 EVIDENCE FROM STRUCTURAL AND FUNCTIONAL BRAIN ARCHITECTURE

1.3.1 STRUCTURAL HIERARCHIES

Hierarchical organisation appears to be a key architectural principle of the brain. It manifests in a plethora of evidence that sensory systems are composed of processing pathways featuring regions with increasingly complex tuning properties as one proceeds further away from primary sensory sites along a functionally defined gradient (Zeki & Shipp, 1988). A classic example of this is found in the visual system, where the first cortical station, V1, comprises a retinotopic map of neurons with ‘simple’ preferences for features such as oriented edges, each tuned to a circumscribed location defined as the cell’s receptive field (Hubel & Wiesel, 1969). These cells provide inputs to regions characterised by increasingly complex tuning properties and larger receptive fields, such as the colour selective area V4 (Bartels & Zeki, 2000) and area V5 /MT, which exhibits selectivity for motion vectors (Zeki, 2015). The hierarchy culminates in areas such as the fusiform gyrus of the temporal lobe, where neuronal preferences are so specific as to encode faces in a position- and angle-invariant fashion (human - Kanwisher et al., 1997; monkey - Tsao et al., 2006), and further still – specific patterns of facial motion (Polosecki et al., 2013). Within this hierarchical framework, brain regions are conceptualised as occupying some level within the hierarchy, with higher-level areas defined on the basis of function as those containing neurons with more complex tuning properties, often also exhibiting increased spatial invariance. In this scheme, projections from lower to higher levels in the hierarchy are labelled as feedforward or bottom-up; while those projecting in the opposite direction are labelled as feedback or top-down. (See the schematic titled ‘Sensory hierarchies’ for an illustration of feedforward and feedback signalling, as well as a rough summary of the canonical connectivity motif in cortical hierarchies, discussed below.)

The brain’s hierarchical structure implies that a higher region, by virtue of receiving convergent inputs from lower regions, has access to more a global view of a sensory input than the levels below. This puts such regions in the position of being able to generate top-down signals predicting the holistic and stable objects which likely underpin the highly dynamic and variable patterns of stimulation which animate neuronal populations at earlier processing stages. In other words, a brain region higher up in the sensory hierarchy may harbour, in its tuning properties, a range of experience-driven Platonic representations of objects in the environment which may be mapped back onto a noisy and time-varying distribution of signals at its input levels.


There is a wealth of empirical evidence that both feedforward and feedback communication does take place within sensory systems. Firstly, structural anatomy supports this. A fundamental feature of cortex is the presence of a canonical pattern

of intracortical projections which direct the flow of signals into and out of the cortical column. A projection can be classified as feedforward or feedback by identifying its source and target using, for instance, unidirectional tracers (eg. [3H]leucine isotope, which harnesses anterograde transport - Morgan & Huston, 1990 - and horse radish peroxidase (HRP), which relies on retrograde axonal transport - Sherlock et al., 1975) and triangulating this with functional information about the hierarchical relationship between the connected areas. Studies using this method have yielded a wealth of knowledge about the nature of projections which link brain areas residing at different levels within processing hierarchies (summarised in the schematic 'Sensory hierarchies').

As a rule, feedforward inputs enter the cortical column at L4, from where activation spreads to the superficial L2/3, and eventually L5 (Bastos et al., 2012; Thomson & Bannister, 2003). This scheme of information transmission was first documented in the cat visual system using intracellular recordings (Douglas & Martin, 1991), and subsequently replicated in further studies exploring synaptic transmission using glutamate uncaging (Weiler et al., 2008), as well as a combination of electrophysiological recordings and histology (Thomson et al., 2002). The connectivity pattern for information arriving in a feedback direction follows a different, but similarly stereotypical scheme, with inputs bypassing L4 in favour of L2/3 and L6 (Rockland & Pandya, 1979), although it has been found that feedback signals arriving from the most high-level regions, such as the prefrontal cortex, terminate almost exclusively in L1 (Markov et al., 2013).

The manner in which signals depart from the cortical column to be transmitted to target sites also appears to have a universal blueprint, separately for feedforward and feedback communications. Here, the distance to be crossed by a signal matters for its laminar origins. Whereas feedforward signals travelling short distances along the processing hierarchy (eg. $V2 \Rightarrow V4$) leave the cortical column at L2/3 and L5, those travelling further (eg. $V2 \Rightarrow IT$ / inferotemporal cortex) originate much more heavily from L2/3 (Markov et al., 2013). Feedback signals, on the other hand, originate both in L2/3 and L6 when travelling short-distances back down the hierarchy (eg. $V4 \Rightarrow V2$) while stemming primarily from the deep layers when travelling further down to lower-level areas (eg. IT / inferotemporal cortex $\Rightarrow V1$) (Angelucci & Bressloff, 2006; Markov et al., 2013; Rockland & Pandya, 1979).

While this summary is certainly an oversimplification of the laminar distribution of interareal connections, the overarching distinctions between the laminar origins and targets of feedforward and feedback connections are consistent enough to have given rise to a widely used index for characterising whether a connection between two areas is feedforward or feedback - the SLN% (Supragranular Labelled Neurons %). This measure captures the percentage of neurons which hail from the supra- versus infragranular layers of the cortex, involved in connecting a source area to its target. Here, a high skew towards supragranular origins is taken to indicate



Sensory hierarchies. A: The macaque visual system as a model sensory hierarchy. A small number of key visual areas are marked with examples of their preferred stimuli, demonstrating a gradient of tuning complexity. This gradient begins with tuning to oriented bar stimuli in V1 (Hubel & Wiesel, 1969), angles and curves in V2 (Ito & Komatsu, 2004), simple shapes in V4 (Oleskiw, Nowack & Pasupathy, 2018), composite motion in V5 (Cui et al., 2013), and complex objects and faces in IT (inferotemporal cortex; Bondar et al., 2009). A cortical hierarchy is primarily defined functionally, with areas populated by cells with more complex tuning properties seen as occupying higher positions within the hierarchy. Connections in the low $\Rightarrow$ high direction are classified as feedforward, while those in the high $\Rightarrow$ low direction are classified as feedback. **B:** a simplified schematic of the laminar connectivity pattern linking a cortical column in one area to its higher-level (left) or lower-level (right) targets, as described by Markov et al. (2013). Feedforward projections crossing small hierarchical distances (eg. $V2 \Rightarrow V4$) originate in L2/3 and L5, terminating in L3B and L4 of the target area. Those which cross larger hierarchical distances (eg. $V2 \Rightarrow IT$) originate predominantly in L2/3. A portion of feedforward inputs are provided by pathways from L5 routed via higher-order thalamic nuclei. Feedback projections follow a different pattern. Those which cross small hierarchical distances (eg. $V4 \Rightarrow V2$) originate in L2/3 and L6, terminating in corresponding layers. Those crossing larger hierarchical distances (eg. $IT \Rightarrow V2$) draw significantly heavier contributions from L6, and terminate exclusively in L1.

a heavily feedforward mode of communication, and thus a ‘low’ position in the processing hierarchy relative to the target area. Hierarchical ‘trees’ constructed from such anatomical data agree well with those based on functional definitions of relatively lower- or higher-level areas (Barone et al., 2000; Markov et al., 2013). Furthermore, comparing what is known about the laminar origins and targets of anatomical connections in the visual (Markov et al., 2013) and somatosensory (Petersen, 2007) systems indicates that the overarching motif is shared. This provides a rationale to believe that the described motif is universal, rather than specific to any particular sensory hierarchy.

The co-existence of these microcircuits within the column, as well as the replication of this scheme across cortical sites has formed the basis for viewing them as part of a fundamental motif in cortical processing (Douglas, Martin & Whitteridge, 1989). This motif highlights the presence of an evolved system for shuttling information both in the feedforward and feedback directions, which appears to be transplantable to any particular cortical region. This raises the question: what roles do these motifs play in the grand scheme of processing? Particularly interesting are connections subserving transmission in the top-down direction, which are inherently less connected to the outside world, and yet significantly outnumber projections supporting feedforward signalling (Felleman & Van Essen, 1991; Markov et al. 2013). Proponents of the predictive coding framework argue that it is precisely this co-occurrence of bottom-up and top-down communication which logically implies that the brain is predictive, with feedforward and feedback connections playing distinct roles (Bastos et al., 2012). Indeed, evidence that the superficial and deep cortical layers give rise to distinct profiles of oscillatory activity (Spaak et al., 2012) has also given rise to the possibility that predictive coding makes use of distinct anatomical and spectral channels for communicating predictions and prediction errors - to be discussed in a future section.

It is also important to note that interareal communication does not exclusively take place via direct cortico-cortical pathways, but also recruits parallel cortico-thalamo-cortical projections, whereby signals are first relayed from cortical L5 to higher-order thalamic nuclei before being shuttled to the feedforward input layers of target cortical regions (Li, Guido & Bickford, 2003; Reichova & Sherman, 2004; Theyel et al., 2010). These transthalamic pathways are distinct from the modulatory feedback pathways which target thalamic neurons from cortical L6. In fact, not all thalamic nuclei which receive feedback take part in cortico-thalamo-cortical signalling, which is thought to subserve primarily feedforward communication (Guillery, Feig & Van Lieshout, 2001). This view is based on observations that transthalamic projections exhibit the same driving properties as observed for other bottom-up connections, including large EPSPs, the occurrence of paired-pulse depression indicative of high transmitter release probability, and the tendency to engage exclusively ionotropic receptors (Sherman & Guillery, 2011). Moreover, signals arriving via higher-order thalamic nuclei can evoke robust responses in target areas, which exceed those pro-

duced by direct cortical connections both in magnitude and duration (Zhang & Bruno, 2019). The precise function of such pathways in interareal communication is, however, still largely unknown. Calcium imaging of the mouse visual cortex, combined with layer-specific optogenetic activation of principal cells, has revealed that transthalamic pathways recruited by cortical L5 are critical for the global propagation of stimulus-evoked activity (Stroh et al., 2013). This has led to the suggestion that higher-order thalamic nuclei are essential for coordinating large-scale cooperation between brain areas, based on cognitive and attentional states (Sherman, 2016). However, further research into this topic is unfolding slowly. As such, the current discussion of feedforward and feedback connections between regions in sensory hierarchies will center on direct cortico-cortical connections with better documented effects.

Feedback and feedforward pathways are also known to differ in terms of their spatial specificity. Using reciprocal connectivity between primate visual areas V1 and V2 as a model system, researchers found that feedforward projections show a retinotopic mapping of inputs to the next processing level (Perkel et al., 1986), while feedback projections tend to be somewhat more spatially diffuse (Shmuel et al., 2005; Salin & Bullier, 1995). This observation perhaps makes sense when considered through the lens of predictive coding. Sensory representations in higher-level regions exhibit greater degrees of position-invariance than those in lower-level regions such as V1 – a tendency which becomes more pronounced as one ascends the processing hierarchy (comparison of striate and extrastriate cortex – Smith, 2001; invariant encoding of biological motion in the superior temporal sulcus – Grossman et al., 2010). This presumably allows the brain to achieve a visual and conceptual representation of an ‘apple’ that remains the same irrespective of retinal location. Thus, it is perhaps to be expected that top-down signals regarding perceptual hypotheses about the meanings of sensory inputs are somewhat less spatially specific than the signals inherent to the lower levels of the hierarchy.


The somewhat distinct postsynaptic effects of feedforward and feedback projections are also indicative of their likely distinct functions. While feedforward projections are classified as directly ‘driving’ their targets, feedback signals produce more modulatory and non-linear outcomes (Sherman & Guillery, 1998). It should be noted that this framework has more recently been found to oversimplify the physiological properties of feedforward versus feedback connections, especially for intrathalamic circuits (for summary see Bickford, 2015). However, the notion that the two classes of connection functionally differ is supported by evidence from the visual system that feedforward-driven activity is critically reliant on rapid-action AMPA glutamate receptors, while feedback activity is largely underpinned by the slower NMDA receptor-mediated transmission (implied by differential sensitivity to receptor blockage by CNQX and APV respectively; Self et al., 2012). These two differing modes of transmission set the stage for feedforward stimulation providing a default source of neuronal activation, while slower-acting feedback plays a more modulatory (or, to use anthropomorphic terms, ‘nudging’) role, in accordance with likely interpretations

of input.

1.3.2 EXTRA-CLASSICAL RECEPTIVE FIELDS

Setting aside the brain’s structural connectivity patterns as evidence for bottom-up and top-down signalling, there is clear functional evidence that such interactions take place. A testament to this is the wealth of studies finding that neurons in low-level sensory regions, such as V1, exhibit responses to information that extends beyond the scope of their classical receptive fields (Desimone & Shein, 1987; Grosf et al., 1993). These so-called extra-classical receptive fields are often found to heavily depend on feedback from higher-level regions such as V2, V4 and V5 (Hupé et al., 1998; Nassi et al., 2013; Schnabel et al., 2018). One specific extra-classical effect that is consistent with the notion of perceptual predictions being imposed by top-down modulation is the observation that neurons in V1 can respond to illusory contours that could not have possibly activated the cells in a bottom-up fashion (Grosf et al., 1993; Kok & de Lange, 2014). A stimulus that is commonly used in studying this is the so-called Kanizsa triangle (shown below, in the schematic ‘Kanizsa triangle’).

Originally, researchers discovered, using single-unit recordings in monkeys, that V1 neurons with receptive fields falling on the ‘edge’ of an illusory contour respond as if that contour was present (Grosf et al., 1993). An intuitive explanation for this is that such responses stem from feedback inputs from higher-level regions gauging the holistic extent of the illusory figure. Consistent with this, a single-unit recording study of monkey visual areas found that illusory modulation of neuronal responses in V1 occurred with a greater latency than in V2, indicating the effect was inherited from top-down (Lee & Nguyen, 2001). Interestingly, the reported delay of >100 ms has been specifically linked to the arrival of top-down signals, as indicated by the onset of other extra-classical receptive field effects in V1 (Schnabel et al., 2018; Self et al., 2012). Furthermore, application of high-field (7T) fMRI with the power to discriminate BOLD responses across cortical layers (Scheeringa et al., 2016) has indicated that illusory contours predominantly engage a different laminar profile of activity in V1 compared to contours that are physically present, and thus reach V1 in a bottom-up fashion (Kok et al., 2016). These lines of evidence suggest that illusory effects in primary visual cortex might be conferred by higher-level regions with access to more holistic information. This puts the illusory contour phenomenon within the territory of predictive coding theories.



Kanizsa triangle. A commonly used visual illusion, the Kanizsa triangle involves a trio of pacmen images oriented so as to impose the perception of a triangle. Note, however, that the contours of this shape are not physically present, and that the receptive field of a V1 neuron encompassing part of the illusory contour (red circle) would be receiving no bottom-up evidence for an edge. Any responses in V1 to such contours are thus traditionally interpreted as resulting from higher-level regions with access to a holistic representation of the illusory triangle projecting predictions about the contours' existence.

More recently, parallels between illusory contour induction and predictive coding were examined in humans in an fMRI paradigm. Here, BOLD activity was compared across the visual system during viewing of the Kanizsa triangle and a control display featuring a similar, but scrambled, array of pacmen figures that did not induce an illusory percept (Kok & de Lange, 2014). It was observed that V1 responses to illusory contours were enhanced compared to responses to real contours. The authors took this to indicate that illusory contours evoked a greater perceptual prediction error owing to the fact that, while the appropriately oriented pacmen induced the impression of an object with contours, which was presumably communicated by higher-level regions, this failed to predict their actual absence in the bottom-up input to V1. Nonetheless, the notion that illusory contour effects in V1 reflect predictive processes currently lacks causal evidence. While findings on the relative timing of such effects in lower- and higher-order areas (Lee & Nguyen, 2001), as well as the suppression of activity in V1 for correctly predicted contours (Kok & de Lange, 2014) present interesting circumstantial evidence, there have yet been no studies attempting to intervene with top-down signalling in order to probe its effects on neuronal activity and illusory perception.

1.4 EVIDENCE FROM OSCILLATORY ACTIVITY

In recent years, one particularly prominent property of brain activity has become implicated in providing the functional avenue for the top-down and bottom-up communication of sensory predictions and prediction errors in cortical circuits (Bastos et al., 2012; Kanai et al., 2015; Michalareas et al., 2016). The property in question is the widespread nature of oscillatory activity.

1.4.1 WHAT ARE OSCILLATIONS?

Oscillations are rhythmic fluctuations in the membrane potential or spiking pattern of individual neurons or neuronal ensembles. When sufficiently large numbers of neurons engage in synchronised rhythmic activity, the resulting macroscopic oscillation can be captured using a diversity of methods operating on different scales, such as extracellular electrodes directly probing neural tissue (Spaak et al., 2012), electrocorticography (ECoG) sensors placed on the cortical surface (Hill et al., 2012), or electroencephalography (EEG) sensor arrays placed non-invasively on the scalp (Barbosa & Kouider, 2018). Oscillatory activity is widely observed throughout the brain in an enormous variety of mental states. Since some of the earliest pioneers of oscillation research used EEG to document prominent oscillatory activity in the alpha-frequency band (α ; 8-12 Hz) in resting humans (Berger, 1929), a vast body of work has observed that the brain generates rhythmic activity, often in selective frequency bands, in a huge diversity of contexts, such as during attentive processing (Bichot et al., 2005; Meyer et al., 2014; Thut et al., 2006), memory retrieval (Foster

et al., 2016; Javadi et al., 2017), decision-making (Polania et al., 2014), and action selection (Brittain & Brown, 2014; Joundi et al., 2012). One would be hard-pressed to find any cognitive or perceptual task in which oscillatory activity is not observed in the brains of humans or other animals.

The ubiquity of neural oscillations has given rise to questions regarding their potential function in brain processing. One prominent theory holds that oscillatory activity constrained to a specific bandwidth provides the neural equivalent of a radio channel which neuronal populations can tune into in order to exchange information (Fries, 2005). This so-called “communication through coherence” hypothesis is premised on the idea that, when neuronal populations engage in rhythmic activity, their membrane potential fluctuations generate periodic windows of receptiveness, with inputs being most effective at evoking responses when aligned with the rhythmic depolarisation peaks of their target population. Thus, neuronal populations engaging in synchronised – or coherent - oscillatory activity maximise their potential for communication by virtue of being able to send and receive signals at the same rate.

There is a wealth of evidence in support of the notion that neuronal oscillations serve a ‘binding’ function, allowing neurons to tune in and exchange information (Bichot et al., 2005; Gray et al., 1990; Grothe et al., 2012; Rohenkohl et al., 2018; Womelsdorf et al., 2007). One of the earliest experiments to provide the basis for this theory found that visual cortical neurons activated by their preferred stimulus entered a state of rhythmic synchronization in the gamma (γ ; 30-100 Hz) frequency band (Gray et al., 1989). Another experiment bolstered the notion that γ may serve as a channel for neurons to tune into when representing a shared piece of information. Here, monkeys were instructed to look out for visual targets characterised by a particular collective property, such as colour, as they scanned across a complex multi-item array (Bichot et al., 2005). During such visual search, whenever the receptive fields of a subset of recorded neurons encountered stimuli which displayed the appropriate feature, spikes produced by the neurons in question exhibited increased coherence with the local field potential (LFP) specifically in the γ -band. These findings indicate that γ -band rhythmic activity may play the role of metaphorical glue, connecting neuronal populations encoding the same message. Moreover, there is a convincing case that γ may facilitate the transmission of this message to the appropriate recipient region(s).

In one experiment showing this, high-density ECoG recordings were obtained throughout the monkey visual system while the animal performed a task requiring it to attend to one of two stimuli, depending on the colour of a cue presented at the start of each trial (Bosman et al., 2012). The researchers observed robust synchronisation in the γ -band between the oscillatory activity of higher-level area V4 and neuronal populations in V1 encoding the location of the target stimulus. This synchrony reflected enhanced information flow between V1 and V4 in the feedforward direction, as indicated by Granger causality analysis of the time-courses of oscillatory activity in these two regions. The same synchronisation effect was not found between V4

and V1 neurons encoding the location of the distractor stimulus. This lent robust support to the idea that synchronous rhythmic activity in the γ -band might be the mechanism by which the brain flexibly enables relevant information to be shuttled between selected sender and recipient regions.

There is good reason to believe that such interareal synchronisation is behaviourally relevant, as the strength of V1-V4 γ -band synchronisation has been found to predict the speed with which monkeys perform change detection (Rohenkohl et al., 2018). Of note, it is highly plausible that enhanced V1-V4 synchrony during task performance might in fact be merely a marker, rather than the underpinning mechanism, of efficient information exchange between these sites. However, this possibility is inherently difficult to test, as it would be virtually impossible to selectively interfere with synchrony without fundamentally affecting the signals which serve as the critical information carriers.

1.4.2 OSCILLATIONS AS CHANNELS FOR BOTTOM-UP AND TOP-DOWN SIGNALLING

The potential relevance of oscillations to predictive coding came to light as it emerged that oscillatory activity in different frequency bands can be traced back to distinct cortical layers (Buffalo et al., 2011; Maier et al., 2010; Spaak et al., 2012) – layers which are, at this point, laden with expectations about encoding sensory predictions and prediction errors, owing to their anatomical connection patterns in sensory hierarchies (Rockland & Pandya, 1979; Kanai et al., 2015). This has given rise to the possibility that distinct frequency bands might provide communication channels for the top-down and bottom-up communication of sensory predictions and prediction errors, respectively (Bastos et al., 2012).

The development of this sphere of research goes back to a series of translaminar recording studies which found that activity in the deeper (or infragranular) layers of the cortex appears to have a low-frequency spectral fingerprint, peaking in the α (8-12 Hz) and β (12-25 Hz) frequency range, while the superficial (or supragranular) layers primarily oscillate at higher, γ -band (30-100 Hz) frequencies (Buffalo et al., 2011; Maier et al., 2010; Spaak et al., 2012). These differences are thought to arise from a variety of intrinsic properties of the local circuitry, such as the presence of fast-rhythmic burst neurons (Traub et al., 2005), first introduced as a distinct class of pyramidal ‘chattering’ cells (Gray & McCormick, 1996) in the superficial layers, the low-frequency pacemaker-like properties of a subpopulation of L5 pyramidal cells (Silva et al., 1991; Sun & Dan, 2009), and the nature of interactions between principal neurons and fast-spiking interneurons, which are thought to be critical regulators of γ -band activity (Salkoff et al., 2015; Sohal et al., 2009).

Given our pre-existing knowledge that feedforward connections originate predominantly in the supragranular L2/3 while feedback connections also arise from the deeper L5/6 (Angelucci & Bressloff, 2006; Rockland & Pandya, 1979), it is logical to suppose that low- and high-frequency oscillatory activity might map onto communication in opposite directions through sensory hierarchies. Several studies support this, finding that hierarchies inferred from the directional flow of oscillatory activity in distinct frequency bands between regions in a sensory hierarchy correspond quite accurately to those constructed from anatomical tracing studies with explicit knowledge of the senders and recipients of neuronal projections (Bastos et al., 2015; Michalareas et al., 2016; van Kerkoerle et al., 2014).

In the first of such studies, researchers performed Granger causality analysis of ECoG data between every possible pair of 28 sensors situated in the rhesus macaque visual system (Bastos et al., 2015). Briefly, Granger causality analysis involves examining the degree to which the past value of some variable (eg. γ -power or membrane potential) recorded at site A explains the future value of that variable at site B, and vice versa (Faes et al., 2017). Its goal is thus to infer whether site A, to some degree, causally determines the unfolding of events at site B, which would imply that information flows in the direction $A \Rightarrow B$. Knowing whether site A is above or below site B in the hierarchy, based on properties such as the relative complexity of its receptive fields and the predominance of anatomically feedforward versus feedback projections it receives from site B, helps determine whether the information flow $A \Rightarrow B$ represents feedforward or feedback communication. Applying this method to the ECoG recordings revealed that oscillatory activity in the feedforward direction showed a distinct spectral peak in the γ -band, while activity in the feedback direction was selectively affiliated with the lower-frequency β -band. Shortly thereafter, the same conclusions were reached by a study of human MEG data (Michalareas et al., 2016). Such findings are further bolstered by evidence that, while stimulating the lower-level V1 prompts it to engage in synchronous γ -band activity together with V4, stimulating the higher-level area synchronises it with V1 at a substantially lower frequency (van Kerkoerle et al., 2014).

While these spectral information channels can be artificially isolated through analysis or experimental stimulation, studies suggest that low- and high-frequency oscillations frequently interact (Dougherty et al., 2017; Fontolan et al., 2014; Popov et al., 2017; Richter et al., 2017). One study conducting cross-laminar recordings in monkey V1 reported that, during visual stimulation, firing rates in L2/3 were tightly coupled to the phase of α -oscillatory activity which was traced down to the infragranular layers (Dougherty et al., 2017). There is reason to believe that such cross-frequency interactions occur when there is a need for bottom-up information flow to be re-routed according to task needs – something which presumably requires top-down communication. Consistent with this, Richter and colleagues (2017) found that top-down β -band synchronisation within the monkey visual system in an attentional task enhanced visually-evoked γ -band synchronisation in the $V1 \Rightarrow V4$

direction. This cross-frequency coupling effect has now been documented between the human primary and association auditory cortices, using depth electrode recordings (Fontolan et al., 2014). A further study which found this effect in the human visual system using MEG, indicated that such coupling is of relevance for behaviour, as top-down modulation of stimulus-driven γ -band activity in V1 predicted performance quality on an attention task (Popov et al., 2017).

Thus, we are left with a convincing case that lower- and higher-frequency oscillatory activity may correspond to the exchange of information in the top-down and bottom-up directions respectively, and that these oscillations routinely interact to select the information that is allowed to propagate further up the processing hierarchy. A tempting question which now arises is: do we have any evidence that the information being exchanged in these two directions maps onto the distinction between sensory predictions and prediction errors?

1.4.3 OSCILLATIONS AS CARRIERS OF PREDICTIONS AND ERRORS

There is solid evidence to believe that sensory expectations pre-activate neuronal populations that are involved in encoding the expected stimulus or sensory property (Kok et al., 2014; Puri et al., 2009; Stokes et al., 2009). More recently, these neural correlates of perceptual prediction have been endowed with some spectral properties, indicating that lower-frequency oscillatory activity in the α - and β -bands might be the critical carrier of predictions.

Traditionally, power in α -band was found to avoid the hemisphere which was to process an expected input – a phenomenon which is termed α -lateralisation and thought to indicate, at least in part, increased top-down allowance for the bottom-up propagation of information from the expecting hemisphere (Myers et al., 2014; Rihs et al., 2007). More recently, various studies found that this lateralisation phenomenon is sensitive to the predictability of sensory events (Bauer et al., 2014; Sedley et al., 2016, van Pelt et al., 2016). In one such study, human participants performed a task in which they monitored a contracting visual stimulus for changes in contraction speed (Bauer et al., 2014). The delay with which the change occurred relative to stimulus onset was randomly selected from a range of values, such that the hazard rate (i.e. probability that a change was about to happen given that it hasn’t already done so) increased sharply towards the end of each trial. Under these conditions, the longer a participant waited on any given trial to see a change, the more predictable that change was going to be when it arrived. This index of predictability scaled with the strength of α/β -lateralisation in participants’ visual cortex, and was thus thought to reflect a signal of increasing stimulus expectancy. Notably, trials with more expected changes and stronger lateralisation of low-frequency activity were also associated with shorter reaction times. A similar relationship between expectation and α/β -lateralisation has also been seen in the somatosensory system (van Ede et al., 2010).

Arguably, an important piece of the puzzle missing from these findings is verifying that such low-frequency oscillatory activity actually contains decodable information about sensory predictions. A study by Broski-Guerniero and colleagues (2017) offers some more concrete evidence of this using MEG. Human participants performed a task in which they were asked to detect faces in ambiguous black-and-white images. To address the question of whether perceptual predictions about seeing a face have any particular spectral properties, the neural data were analysed using the information-theoretic measure AIS (Active Information Storage), which measures the amount of decodable information in the future of a neural process which is explained by its past. The use of this measure was rooted in the assumption that the brain may not always know when exactly a prediction needs to be measured against reality, requiring the predictive representation to persist for a certain amount of time – something which the AIS measure aims to capture.

The authors found that the brain’s face-processing areas exhibited an enhanced AIS index (indicative of a lingering face prediction), which correlated with power in the β -band on a trial-by-trial basis. Interestingly, the strength of this predictive information also correlated with the degree of top-down information flow from these face-selective areas, as indicated by Granger causality analysis. Thus, it was concluded that low-frequency oscillations may indeed provide a channel for transmitting prior knowledge and expectations. Such a conclusion, however, may require significantly more evidence. An increased AIS index in essence indicates that a pattern of neuronal activity (whether static or dynamic) returns to its state with a sufficient degree of periodicity to indicate a persistent representation. However, it is unclear how discriminative such brain activity truly is with respect to perceptual content. This is something which may perhaps be properly assessed by combining decoding analysis of neural data with a task in which participants expect particular faces, rather than faces in general. Nonetheless, this is some of the firmest evidence to date documenting a link between low-frequency oscillations, top-down communication, and pre-activation of stimulus-relevant areas during perceptual expectation.

The next relevant issue to examine is whether we have any evidence that γ -oscillatory activity indexes prediction error. Some clues that this might be the case are offered by the finding that γ -band synchrony is sensitive to the predictability of stimulus repetition. Specifically, a stimulus repeat which occurs in a block with low repetition probability evokes a significantly higher γ -band response compared to a statistically common repetition (Todorovic et al., 2011). More recently, two intracranial recording studies of human participants engaging in auditory sequence listening have shed some more light on the potential roles of γ -activity. These indicate two important properties of the γ -band: i) the ability to encode prediction errors evoked both by local anomalies in the input and more global deviations from a statistical norm (El Karoui et al., 2015), and ii) sensitivity to the precision of a sensory prediction (Sedley et al., 2016), which is argued to be an important determinant of the

degree to which sensory systems weigh current evidence against predictions in refining the internal model (Kanai et al., 2015).

One experiment dissociated lower- and higher-level prediction errors by presenting participants with local deviants (eg. A-A-A-B) which might not have been global deviants by virtue of being presented in a stimulus set where such sequences were common (El Karoui et al., 2015). Conversely, a local non-deviant (eg. A-A-A-A) presented in the same stimulus set represented a global deviant, owing to its rarity, and thus arguably generated a more complex prediction error that could be gleaned from a more integrated representation of the stimulus stream. The electrophysiological data indicated that, while local novelty triggered a rapid (> 80 ms) and transient increase in high γ -band activity (60-120 Hz) in the temporal regions, global prediction errors elicited a more sustained and widely distributed γ -band response starting with >350 ms after the critical deviation. As such, oscillations in the γ -band might play a role in responding to prediction errors rooted in varying degrees of abstraction, perhaps originating at different stages in the processing hierarchy.

In the second aforementioned experiment, researchers manipulated the precision (and thus reliability) of sensory priors by exposing participants to auditory sequences where tones varied in pitch around a mean with a particular degree of variance (Sedley et al., 2016). This was underpinned by the rationale that a sensory system engaging in a ‘rational’ mode of predictive coding would update its internal model more readily when a prediction error departs from an internal prediction rooted in a highly variable, and thus unreliable, set of perceptual experiences. Intracortical recordings revealed that sudden pitch deviations from the stream average evoked neural responses in the γ -band which correlated most strongly with a measure of prediction error scaled by precision (i.e. the inverse of the variance in the preceding stimulus stream).

It is important to consider whether the oscillatory activities corresponding to parameters such as sensory prediction (Broski-Guerniero et al. 2017) and prediction error (El Karoui et al., 2015; Sedley et al., 2016) indeed serve information-encoding functions or are merely correlative markers of encoding without having any direct bearing on the process itself. This is certainly a widely raised issue (Friston et al., 2015; Sohal, 2016). Oscillations are an intrinsic property of circuits with interconnected excitatory and inhibitory cell populations engaging in mutual activation and suppression on particular time-scales, meaning that rhythmic activity may be merely a reflection of the excitatory-inhibitory interactions that encode the necessary information (Ray, Ni & Maunsell, 2013).

In some ways, to say that oscillations are irrelevant to encoding would be to ignore the evidence of their ability to influence the input-output properties of principal neurons (Sohal et al., 2009; also explored in Chapter 3 of this thesis), which clearly points to an active role in managing communications. On the other hand, it is not clear to me whether there is anything to be truly gained from examining the

functional relevance or irrelevance of oscillations. Since it is impossible to treat them as separate from the patterned neuronal interactions which give rise to them, asking whether oscillations matter to information encoding may be equivalent to asking whether I am able to pick up an object because of the specific pattern of muscular contractions in my hand or because of the movements generated by these contractions (i.e. the movements *are* the contractions; the distinction is merely semantics).

Setting aside the debate on whether oscillations encode predictive coding variables or merely reflect their encoding, the evidence firmly supports a view of the brain as one which generates sensory predictions (Broski-Guerniero et al., 2017; Kok et al., 2014; Stokes et al., 2009), weighs them against sensory evidence (Sedley et al., 2016), and responds to resulting discrepancies (El Karoui et al., 2015; Todorovic et al., 2011). In light of this, I would like to ask: is anything still missing, and how can my thesis research contribute to the wider goal of understanding the brain as a predictive coding 'machine'?

1.5 RESEARCH GOALS

Without a question, predictive coding is powerful in its ability to explain neuronal and oscillatory phenomena. Nonetheless, there is arguably a difference between yet again verifying that the brain generates sensory predictions and responds to their violations, and striving to identify some constraints to the phenomenon which might shed light on how and where the parameters of this predictive coding scheme play out in the brain. I would argue that, while we have sufficient evidence to make the general statement that the brain engages in predictions, we still lack the detail required to understand how this is implemented.

One critical element of the predictive coding scheme that is yet to be directly tested is the supposed anatomical affiliation of sensory predictions and prediction errors with distinct neuronal populations across the cortical column (with predictions represented in the superficial and, to a greater degree, deep layers, and errors primarily in the superficial layers; Bastos et al., 2012; Kanai et al., 2015). This expectation follows logically from the non-overlapping laminar origins of feedback versus feedforward connections (Berezovsky et al., 2011; Rockland & Pandya, 1975), which has singled out the retrogradely signalling neurons of the deeper layers (and the oscillatory activity they generate) as the natural carriers of predictive signals that need to travel top-down across the sensory hierarchy. In contrast, it is the supragranular neurons giving rise to feedforward connections which are thought to be the primary source of prediction errors that propagate upwards through the hierarchy and refine future predictions.

Evidence from multi-voxel pattern analysis of fMRI data indicates that neuronal populations showing weakened responses to predictable stimulation (i.e. alleged pre-

diction error ensembles) and those exhibiting the converse properties (i.e. candidate sensory prediction carriers) co-exist within the same sensory areas (de Gardelle et al., 2013). However, the resolution of analysis afforded by fMRI paradigms has arguably been insufficient to truly illuminate any laminar segregation between those signals, with fMRI only recently acquiring the power to access layer-specific BOLD responses (Kok et al., 2016; Scheeringa et al., 2016). Often, it is common for animal research to fill in the ‘holes’ left by the global nature of neuroimaging and address questions with a greater degree of anatomical specificity. But a look at the animal research indicates that we have not yet concretely answered the question of laminar or spectral specialisation within the predictive coding scheme.

Much of recent research into neuronal and circuit dynamics in animal brains has made use of two-photon microscopy, which offers cellular-level resolution within its imaging plane, but comes with the notable caveat of having a maximal imaging depth because of issues with fluorescent signal-to-noise ratios and potential for cytotoxicity due the greater energy levels required to image at greater depths (Liu et al., 2009; Yildirim et al., 2019). Such inconvenience has confined many investigations to the superficial layers of the cortex – which have nonetheless yielded some fascinating insights into sensory predictive processes. One study imaged the responses of L2/3 V1 neurons as mice crossed a virtual environment, uncovering neurons which produced strong responses to occasional perturbations in the normal coupling between the animals’ movements and optic flow playback (Keller et al., 2012). This was strongly indicative of prediction error encoding properties in the superficial layers. Several years later, researchers imaged L2/3 of mouse V1 as the animals made their way through a virtual tunnel peppered with visual landmarks. They reported that, within hours of the animals beginning to practise their tunnel traversal, a portion of neurons (6%) produced anticipatory responses which enabled a pattern classifier to discriminate the orientation of an upcoming visual stimulus (Fiser et al., 2016). In short, we now have evidence that neurons encoding both sensory predictions and prediction errors reside within the superficial layers.

These findings leave us with some important questions about the neural signatures we might find in the deeper cortical layers. Along with L2/3, they are the primary recipients of feedback projections (Rockland & Pandya, 1979), as well as the delayed recipients of information during feedforward signal sweeps through the cortical column (Douglas & Martin, 1991; Thomson et al., 2002). This raises questions about whether these layers might house both predictive and error-based signals, perhaps on different time-scales. I would argue that, in order to assess this situation, it is worth traversing the cortical column and examining the responses of individual layers during sensory expectation and its clash with reality. It is with this rationale that, over the course of my research project, I aimed to set up a mouse behaviour paradigm that would enable me to conduct cross-laminar electrophysiological recordings of primary visual cortex during perceptual decision-making – work that will be detailed in Chapter 1.

I followed this attempt with a set of *in vitro* optogenetic stimulation experiments, detailed in Chapter 2, aiming to examine whether local intracortical circuitry is capable of acquiring sensory predictions, and whether oscillatory activities in the γ - and α/β -bands have roles to play in encoding them, as well as their departures from input. This question is interesting for several reasons. Evidence that neurons in V1 can receive contextual information which contributes to extra-classical receptive fields via lateral connections within V1 (Self et al., 2013) implies that lower-level regions may, to some degree, be equipped to generate sensory predictions using information that is not necessarily of top-down origin. Furthermore, the variety of circuit plasticity mechanisms which take place on a range of time-scales (Bliss & Lomo, 1973; Dylla et al., 2017; Fujisawa et al., 2008), and with such a degree of biochemical diversity as to implicate distinct mechanisms at work in different cortical layers (Daw et al., 2004), raises the possibility that sensory predictions – which arguably can be distilled to the learning of past inputs - are a hardwired property of all circuits. As such, the work comprising Chapter 2 of this thesis involved exposing brain slices to patterns of stimulation adhering to particular statistical regularities, which enabled me to ask – can a circuit learn to predict its input?

Finally, a goal which reflects in the work comprising Chapter 3 was to provide some more mechanistic answers regarding the role of oscillatory activity in the grand scheme of sensory processing. Given that γ -rhythmic activity has been preferentially attributed the role of propagating prediction errors through sensory hierarchies (Bastos et al., 2012), it is relevant to ask what are the effects of γ on information processing in target neurons. This is particularly interesting given some experimental evidence that γ -rhythms enhance single neuron input-output properties (Sohal et al., 2009), which raises the possibility that the oscillatory channel carrying prediction errors might be optimal for triggering target cell responses. Thus, with these goals in mind, my hope is to shed some light on the predictive brain both on a global, local-circuit and single cell level.

CHAPTER 2

2.1 INTRODUCTION

2.1.1 BROAD CONCEPTUAL GOALS AND OBJECTIVES: OSCILLATIONS AND THE LAMINAR ORIGINS OF PREDICTIONS

Over the past decade, imaging neuronal activity in the superficial layers of the mouse visual system during virtual world navigation has revealed the existence of neurons which produce stimulus-predictive responses (Fiser et al., 2016) and those which respond to violations in normal optic flow (Keller et al., 2012). These findings concur with the hypothesis that a predictive cortical system should harbour cells encoding both sensory predictions and the prediction errors which ensue when the internal model fails to predict external reality (Friston, 2005). However, the evidence still leaves much room for speculation about the implementation of predictive coding and its anatomical and spectral properties, which have been a central consideration of theories in this sphere of neuroscience.

One important, and yet largely untested, tenet of the framework is that sensory predictions and prediction errors originate in non-overlapping neural populations across the cortical column (Bastos et al., 2012; Kanai et al., 2015). This idea is premised on a combination of anatomical and functional evidence which strongly implicates the deeper (infragranular) layers of sensory cortex in top-down communication well-suited to signalling predictions, and the more superficial (supragranular) layers in feedforward signalling, ascribed the role of propagating prediction errors (Bastos et al., 2012; Berezovsky et al., 2011; Michalareas et al., 2016; Rockland & Pandya, 1979). Furthermore, evidence that oscillatory activity in the superficial and deep cortical layers is, to some degree, spectrally distinct (with a skew towards γ - (30-100 Hz) and α/β -bands (8-25 Hz) respectively; Spaak et al., 2012), has raised the hypothesis that predictions and prediction errors might be communicated through separate frequency channels (Bastos et al., 2012).

As of now, the virtually absent exploration of the deeper cortical layers leaves us with little true food for thought about the allegedly distinct encoding functions of neurons across the cortical layers, and the putative information-transmitting roles of the oscillatory patterns they produce. Thus, the objective of the research laid out in this chapter was to fill some of these empirical holes. The work was geared towards the plan of conducting translaminar recordings from the mouse V1 during a perceptual decision-making task in which animals were encouraged to internally generate stimulus predictions and were occasionally exposed to sensory violations which can

be reasonably expected to trigger prediction error signals. In examining neuronal activity during i) the period of stimulus anticipation and ii) the post-onset stimulus processing period, my goals were manifold.

Firstly, I hoped to verify that anticipatory single- or multi-unit activities in visual cortex represent stimulus-specific predictions that can be discriminated using classification analysis. Secondly, in differentiating signals on the basis of laminar origin, my hope was to produce some concrete answers about whether there is indeed a preferential distribution of predictive and error signals in distinct cortical layers. This would also allow me to characterise the properties of any potential oscillations that might be observed and shed some light on the roles of rhythmic activity in encoding the informational variables that are so central to the predictive coding framework.

2.1.2 MOUSE VISION AS A MODEL SENSORY SYSTEM

I chose the mouse visual cortex to serve as my model system for a few reasons. Firstly, there is a strong precedent for its use both in perceptual decision-making paradigms (Andermann et al., 2010; Douglas et al., 2006; Graf et al., 2018; Pinto et al., 2013) and specifically in investigations of sensory predictions (Fiser et al., 2016; Keller et al., 2012). Mice are capable of performing a variety of visual tasks, such as those requiring discrimination of stimulus location (Graf et al., 2018), coherent motion direction (Douglas et al., 2006), and grating orientation (Andermann et al., 2010; Jin et al., 2018; Jurjut et al., 2017; Pinto et al., 2013), as well as detection of orientation change (Wang et al., 2016). In such tasks, perceptual decisions can be successfully obtained using either a two-alternative forced choice (2-AFC) response paradigm, in which animals respond by choosing to lick a water spout situated to their left or right (Andermann et al., 2010; Graf et al., 2018), or a go/no-go paradigm, in which animals indicate choice by licking a water spout or actively withholding from it (Jin et al., 2019; Jurjut et al., 2017; Pinto et al., 2013). The evident breadth of behaviours that can be harnessed to interrogate visual decision-making in mice makes them an amenable model system for the purposes of my research.

A second reason why mouse vision is a sensible model for investigating broad questions about the mechanisms of sensory processing is its overarching similarity with the visual systems of other species (Bonin et al., 2011; Gao et al., 2010; Niell & Stryker, 2008). To say this is not necessarily to undermine the clear distinctions between the visual systems of mice and other animals, which certainly exist. One robust species difference concerns the degree of topographical organisation. It has long been known that the visual areas of species like monkeys (Ohki et al., 2006) and cats (Bonhoeffer & Grinvald, 1991) exhibit highly organised gradients of tuning properties, with neurons tuned to particular stimulus orientations or locations neigh-

bouring those with subtly different preferences, to form a continuous representational map. Often, these gradients are replicated across the cortical layers, giving rise to the notion of a cortical column as a unit of cortical information processing (Hubel & Wiesel, 1969; Mountcastle et al., 1955). While researchers have lamented our lack of concrete evidence as to the function of such columns (Horton & Adams, 2005), others have taken representational maps to indicate that those in possession of this organising principle enjoy higher local circuit processing efficiency, manifesting in benefits such as sharper orientation tuning (Koch et al., 2016).

Such maps are substantially weaker in the mouse visual system, where neurons with diverse functional properties are often found intermingled in a fashion which implies a loose organisation, if any at all (Sohya et al., 2007). Nonetheless, research indicates that the mapping principle is fundamentally present, with a clear degree of tuning continuity at larger scales of $>100\text{ }\mu\text{m}$ (Bonin et al., 2011; Schuett et al., 2002; although Ringach et al., 2016 identified a length constant of $>40\text{ }\mu\text{m}$). The relative lack of organisation on a small scale by no means appears to preclude neuronal selectivity, as mouse V1 neurons exhibit tuning for properties such as orientation and spatial frequency on par with the more complex visual systems of predatory species (Li et al., 2012; Niell & Stryker, 2008; Vreysen et al., 2012). Thus, despite some notable differences between the visual systems of mice and more complex animals, it is sensible to assume that the mouse visual system is capable of high quality representations of stimulus dimensions such as orientation and spatial frequency. Consequently, using a behavioural task which encourages mice to generate perceptual predictions can, in principle, be expected to produce stimulus-specific neuronal representations that would be amenable to decoding analysis.

2.1.3 TRIALLING A MOUSE VISUAL DECISION-MAKING PARADIGM

Before embarking on an electrophysiological recording project that would require chronic implantation of a multi-electrode probe into the mouse visual cortex, it was critical to pilot the task I had in mind for the animals and verify that it exhibited two main properties. Firstly, the task needed to be one that mice could learn to a high performance standard using a reasonably quick training regimen which could be replicated after probe implantation in a cohort of experimental mice. This was a critical consideration, since chronically implanted microelectrode arrays tend to suffer from a longevity issue, exhibiting a gradual decline in signal strength and making activity-based decoding analysis more difficult over time (Flint et al., 2012; Stark & Abeles, 2007). The second important property was that the sensory cues used to elicit sensory predictions within the task setting would produce verifiable effects on animal behaviour. I reasoned that this would be the most unambiguous indicator that sensory predictions were indeed being generated somewhere in the brain, and that if such predictions did play out in visual cortex, perhaps preferentially in some layer, I could expect to capture them with a translaminar probe.

2.2 METHODS

2.2.1 ANIMALS

The experiments were approved by the Committee on Animal Care and Ethical Review of the University of Oxford and licensed by the UK Home Office. I used 4 adult mice (C57BL/6, 2 males and 2 females) for the first pilot experiment, and 2 adult mice (C57BL/6, 2 males) for the second, aged 10 weeks at the start of behavioural training. In all cases, animals of the same sex were littermates and housed in the same cage. Approximately two weeks prior to commencing the training regime, I implanted stainless steel head-bars onto the skulls using dental cement, for the purpose of behavioural training under head-fixation. Animals were allowed at least one week of recovery time before starting a water restriction protocol aiming to produce a motivational state that would encourage them to perform a behavioural task in exchange for water rewards.

During the period of pre-training water restriction, mice were weighed and given a dose of water (0.7 ml) on a daily basis. Training began when the appropriate target pre-training weight was achieved ($>85\%$, as specified by the project license). Once training began, animals consumed most of their daily water intake during their week-day training sessions. On weekends, daily water doses were provided to each animal using a syringe with a metallic tip that emulated the licking spouts used during training. Animals continued to be regularly weighed to ensure that they remained within the acceptable margins of weight loss under water restriction, and were supplemented with further water in case they performed poorly on a particular session and thus did not receive sufficient amounts of water during training. A 0.7 ml daily water ration appeared to be appropriate, as it enabled the target pre-training weight to be achieved within 7-8 days of commencing water restriction.

2.2.2 EXPERIMENTAL SET-UP AND STIMULI

All training sessions were conducted in a dark sound-proof booth. Stimuli and training protocols were programmed in Python 2.7, and displayed on two 7" Raspberry Pi screens controlled by two Raspberry Pi Model 2 mini computers and situated at a 45° angle with respect to the animal. The stimuli that required lick responses were sine-wave gratings (0.04 cycles per degree (cpd), 100% contrast, black and white; 30 Hz refresh rate), drifting at 2 Hz. Throughout the course of each training session, whenever animals were not presented with gratings, both screens displayed dynamic visual noise created by randomising the visual grating image, so as to generate a noise stimulus with an equivalent number of black and white pixels.

Lick responses were detected using a design which allowed the mouse to complete an electronic circuit upon making contact with the metallic water spout (Slotnick, 2009). This was enabled by the fact that the mouse was connected to the anode of a 9V battery through constant contact with aluminium foil lining the inside of its cylinder, while the water spout was connected to the battery's cathode. Reward delivery was implemented by allowing lick detections occurring within the critical time window to trigger the opening of solenoid valves connected to a water supply, which allowed a 4 μ l portion of water to flow to the appropriate lick spout within 40 ms of the licking event.

2.2.3 BEHAVIOURAL TRAINING CONDITIONS

At the start of each training session, animals were briefly handled. Once they were head-fixed and placed into a half-cylinder coated with aluminium foil, the cylinder was secured inside the testing booth such that the animal's face was directly in front of two licking spouts and at a 45° angle with respect to the two screens used for visual stimulus presentation. A surveillance camera placed underneath the licking spouts was used to ensure that each spout was equidistant to the animal's mouth, to prevent any obvious discrepancies in effort required to make left and right responses.

Training proceeded in a number of stages during which the animals were habituated to the set-up, and trained to lick specifically during stimulus presentation before responses could be shaped into location discrimination (Pilot 1) and same-different (Pilot 2) judgements. These steps are detailed below. First, animals underwent a period of habituation to the testing apparatus, during which water was unconditionally delivered through the spouts to train the mice that licking them would result in water delivery (Long et al., 2015). The animals were then progressed to the first stage of response shaping, during which water was delivered indiscriminately from any spout provided that the animal licked it after the onset of a grating stimulus. This was done to train the animals that licking behaviour only paid off during the stimulus presentation period. Finally, animals were trained to identify the locations of drifting gratings presented on a single screen (Pilot 1) or judge whether the orientations of two simultaneously presented stimuli were same or different (Pilot 2), by receiving water rewards only in response to licking the correct spout.

The training protocol was programmed such that animals could collect up to two rewards on a given trial, separated by >1 s, provided they returned to the same spout after delivery of the first reward without engaging with the other spout in the intermediate period. This feature of the task inevitably meant that animals could only perform half the number of possible rewarded trials before reaching their daily quota of water intake; however I reasoned that it would nonetheless be conducive to training by forming strong associations between the relevant stimulus and correct licking spout.

Training occurred five days a week, with each session terminating as soon as each animal collected a total of 182 rewards (amounting to 0.73 ml of water). This was done to prevent full satiation, which would likely reduce the effort that animals would be willing to exert on the following training occasion. If animals failed to obtain all their allocated rewards within 50 minutes of session commencement, training was terminated for the day to prevent boredom (which became clear when animals would stop responding altogether). In this case, the animals were manually topped up with the required volume of water to ensure a daily water intake of 0.7 ml. Upon completing a training session, animals were rewarded with a small amount of condensed milk.

2.2.4 PILOT 1: LOCATION DISCRIMINATION TASK STRUCTURE

Water-restricted mice were trained, under head-fixation, to discriminate whether drifting sine-wave gratings were presented on the right or left screen by licking the corresponding spout. At the start of each trial, a sine grating was randomly selected to be presented either on the left or right screen for a 4 s duration. During the grating presentation window, animals were rewarded with water if they responded by licking the spout that was on the same side as the screen with the grating. Licking the incorrect spout at any point during stimulus presentation triggered a punishment waiting period of 6 s, during which both screens displayed visual noise and lick responses had no pay-off.

Following stimulus offset (on correct trials) or the punishment period (on incorrect trials), animals were given an inter-trial interval selected at random from a uniform distribution of values ranging between 6000 and 8000 ms. Licking the spouts during this period had neither beneficial nor detrimental effects on task progression. This interval was immediately followed by a conditional waiting period which extended for as long as animals continued to make lick responses. To trigger the next trial, animals needed to refrain from licking entirely for >5 s. Thus, not only did licks not pay off at this point, but they had the counterproductive effect of delaying stimulus onset for a potentially indefinite period. The purpose of this task feature was to discourage compulsive licking and encourage responses specifically during stimulus presentation.

Errors and response omissions were followed with error-correction trials. This involved presenting the grating stimulus on the same screen as the previous trial, which was repeated on subsequent trials until a correct response was given. This point triggered the task programme to resume random trial-wise selection of stimulus location.

2.2.5 PILOT 1: INTRODUCING AUDITORY CUES TO EVOKE SENSORY PREDICTIONS

Predictive auditory cues were incorporated into the protocol once animals showed consistently above-chance location discrimination for 3 training sessions. At this stage, animals were presented with one of two auditory cues (6 kHz or 12 kHz; 400 ms duration; 80 dB volume) through a loudspeaker located inside the training booth, 2 s prior to the onset of each grating stimulus. The 2 s interval was selected on the basis of two factors: i) the need for a cue-to-stimulus interval of sufficient length to maximise the chances of capturing predictive electrophysiological activity without contamination by residual auditory cue-evoked activations, and ii) evidence that mice are capable of retaining information in working memory over delays of up to 4-5 s (Liu et al., 2014).

Each of the two tones predicted, with 100% accuracy, whether the animals were going to be presented with a vertically or horizontally oriented grating. Under these conditions, animals underwent three consecutive training sessions before any further task or stimulus manipulations.

To test whether cues carrying predictions about stimulus orientation might facilitate responding under difficult conditions, stimulus contrast was progressively reduced to 60%, 40%, 20%, 10%, 7%, 5%, and 3% over the course of consecutive training sessions (at the lowest contrast value, a Gaussian mask and noise were also superimposed on the grating). The aim was to identify a level of contrast degradation that would make discriminating stimulus from noise sufficiently taxing so as to influence performance, in terms of accuracy or reaction times. To test whether the predictive auditory cues were contributing to performance at lower contrast levels (<20%), I began to interfere with the predictive value of these cues.

Testing whether animals relied on auditory cues providing predictive information about stimulus orientation involved examining whether behaviour was deleteriously affected on trials when their predictions turned out to be incorrect. Thus, each training session contained a pseudo-randomly selected 20% of trials when auditory cues were followed by a stimulus orientation which was opposite to that affiliated with a given cue (these trials are termed ‘invalid’). These were programmed to occur with equal frequency for both auditory cues, to prevent selective validity degradation effects.

2.2.6 PILOT 2: SAME-DIFFERENT JUDGEMENT TASK

Water-restricted mice were trained, under head-fixation, to discriminate whether drifting sine-wave gratings presented on two screens were of the same or different orientations, by licking the appropriate spout. Thus, one mouse was required to lick the left spout if two gratings were the same and right spout if different, while the other mouse was allocated the opposite response requirements. These response requirements remained the same for the duration of the training protocol. At the start of each trial, one sine grating was presented on top of dynamic noise on one of the two screens (location on left or right screen selected at random), for a 2 s duration, and immediately followed by a second sine grating on the opposite screen. The two gratings overlapped by 2 s. Each training session was pre-programmed to contain 50% of trials with gratings of the same orientation; and a remaining 50% of trials with different orientations. Each training session was also pre-allocated an equal number of horizontal and vertical same pairs, to avoid associating any one orientation more strongly with same status. Order of trial types was determined pseudo-randomly, the only constraint being that there could be no more than two consecutive trials featuring differently oriented grating pairs.

During the time window marking the coincident presentation of both gratings, animals were rewarded with water if they responded by licking the spout that was selected to be affiliated with a same or different judgement. Licking the incorrect spout at any point during stimulus presentation triggered a punishment waiting period of 6 s, during which the screens displayed visual noise and lick responses had no pay-off. Following stimulus offset (on correct trials) or the punishment period (on incorrect trials), animals were given an inter-trial interval selected at random from a uniform distribution of values ranging between 6000 and 8000 ms. Licking the spouts during this period had neither beneficial nor detrimental effects on task progression. This interval was immediately followed by a conditional waiting period which extended for as long as animals continued to make lick responses. To trigger the next trial, animals needed to refrain from licking entirely for >5 s. Thus, not only did licks not pay off at this point, but they had the counterproductive effect of delaying stimulus onset for a potentially indefinite period. The purpose of this task feature was to discourage compulsive licking and encourage responses specifically during stimulus presentation.

Errors and response omissions were followed with error-correction trials, in which a grating pair with the same same/different status as in the previous trial continued to be presented (orientations selected at random) until the animal provided a correct response. This meant that the relative frequencies of same versus different trials dynamically steered away from the pre-programmed 50:50 distribution in the early stages of training (assuming that animals exhibited initial biases); however equal frequencies were to be re-imposed once animals attained high levels of performance.

2.3 RESULTS - PILOT 1

2.3.1 RECAP AND OBJECTIVES

In this chapter, I piloted a behavioural task designed to evoke sensory expectations in the mouse visual cortex, with the goal of eventually recording V1 neuronal activity and exploring the anatomical and physiological properties of these representations under task conditions. To this end, I designed a A-alternative forced choicenative Forced Choice (2-AFC) paradigm which set out to combine visual decision-making demands with an implicit incentive to form internal predictions.

Head-fixed mice were trained to discriminate the locations of sine-wave gratings by licking the appropriate water spout. After criterion performance level was achieved, auditory cues were introduced during the pre-stimulus period, predicting whether the animals were about to see a horizontally or vertically oriented grating. Gradually, grating contrast was lowered over successive training sessions to make the task more challenging (for evidence that stimulus contrast influences neuronal tuning and mouse performance on the 2-AFC see Gao et al., 2010, Li et al., 2012, & Long et al., 2015). It was reasoned that making decisions more difficult might bias the animals to exploit the predictive information provided by the auditory cues, possibly by pre-activating neuronal populations with the appropriate orientation tuning to aid in stimulus detection (tuned pre-activations documented and found to boost stimulus detection in humans by Kok et al., 2014 and Stokes et al., 2009).

The raw data collected from each training session consisted of timestamps for the following events: i) individual lick detections, ii) stimulus onsets, and iii) openings of the solenoid valves which delivered water rewards. Data from each animal, whom I will refer to as Female 1 (Vanilla), Female 2 (Cardamom), Male 1 (Benji) and Male 2 (Fatty), will be examined separately.

2.3.2 MICE CAN LEARN TO DISCRIMINATE STIMULUS LOCATIONS

All animals learned to respond at high rates within the first few training sessions (Figure 2.1A for females; Figures 2.2C and 2.3C for males). Despite the evidently high degree of effort expended, neither of the females exceeded chance-level performance at any point over the course of 16 training sessions (Figure 2.1D), as shown by a series of binomial tests comparing the proportion of correct responses each animal gave during every session to a theoretical success probability of 0.5. The females' performance, in fact, fell significantly short of chance on a small number of sessions, as seen from the statistical summary in Table 2.1.

Both females exhibited a pervasive bias towards licking the right spout (Figure 2.1B). This was evidenced by the significantly skewed ratios of right to left licks which were harvested throughout each of their training sessions, as summarised in Table 2.2. Nonetheless, this bias was likely not the initial factor underpinning both animals' failure to acquire the task, but rather secondary to reinforcement histories which generated a strong preference towards one spout - on different timescales in the two females.

Thus, female 1 (Vanilla) showed a consistently higher accuracy on trials requiring right responses to the detriment of left-sided accuracy (Figure 2.1C) - however this preceded the animal's bias towards licking the right spout by five sessions (Figure 2.1B). This suggests that the animal's incidentally higher success rate on right-sided trials in the early stages of training may have caused it to evolve a persistent bias towards licking the right spout. This bias appears to have chronically disincentivised the animal from orienting towards the opposite spout with sufficient frequency to enable rule acquisition. Female 2 (Cardamom), instead, switched between higher right-sided and left-sided accuracy over the course of the training regimen (Figure 2.1C). This suggests that the animal's preference towards one spout was not chronic, but rather fluctuated on a smaller time-scale, in a manner that was driven by more recent reinforcement history. Overall, the behavioural patterns of both females indicate that performance was heavily driven by conditioning experiences which led the animals to chronically (in the case of Vanilla) or temporarily (in the case of Cardamom) treat one spout as the correct one, rather than varying their responses enough to learn that both spouts paid off, depending on stimulus location.

The females' failure to show any signs of improvement contrasted with successful task acquisition by the two males, which achieved significantly above-chance performance at earlier stages of training, with Benji exceeding chance by training session n°14 (*Accuracy*= 64%; $p= 0.003$), and Fatty doing so by training session n°12 (*Accuracy*= 73%; $p< 0.001$). Furthermore, both males showed clear evidence of a steady improvement in accuracy leading up to the first session demarcating above-chance performance, with the upward performance trend starting after session n°6 (Figures 2.2F and 2.3F). Given this disparity in performance between the females and males, training of females was terminated after session n°16 and all further analyses and training focused on the males.

These animals eventually achieved high levels of accuracy, with Benji's performance peaking at 92.3% on training session 19 and Fatty's performance peaking at 96.6% on training session 18 (Figures 2.2F and 2.3F). At this point, both animals were also markedly faster at supplying correct responses (Figures 2.2D and 2.3D), as indicated by two Mann-Whitney Wilcoxon rank sum tests (Benji: *Rank sum*= 5407, $p< 0.001$; Fatty: *Rank sum*= 3662, $p< 0.001$) for differences in median reaction times between each animal's first (Benji: *Med*= 1759.5 ms, *IQR*= 2299 ms; Fatty: *Med*=

2234 ms, $IQR= 1658$ ms) and peak accuracy session (Benji: $Med= 344$ ms, $IQR= 267.5$ ms; Fatty: $Med= 293$ ms, $IQR= 94.5$ ms).

An additional measure I used to track the animals' progress was their accuracy on trials requiring responses towards a different spout than the one which was rewarded in the previous trial (referred to as 'side switch trials' in Figures 2.2F and 2.3F). I reasoned that this was a relevant performance measure as it is theoretically possible for an animal to have a high level of overall accuracy while actually showing a significant skew towards being correct when required to select a specific spout.

Both animals showed evidence for pre-existing side biases, as evidenced by a series of χ^2 -tests comparing the proportions of successes on trials requiring left-sided versus right-sided lick responses (data shown in Figures 2.2E and 2.3E; test results in Table 2.3). These tests revealed that Benji's response accuracy was significantly higher on trials requiring left-sided licks on 12 out of 28 training sessions, while accuracy on trials requiring right-sided licks never significantly exceeded left-sided accuracy. The tests also verified that Fatty showed the converse tendency of making significantly more accurate responses on trials requiring right-sided licks on 6 out of 28 sessions, while left-sided accuracy was never better. As such, overall accuracy might conceal a selective performance advantage for one spout, and thus overestimate an animal's level of task understanding. Indeed, there was a trend for Fatty's response switch accuracy to lag behind his overall accuracy in the initial stages of training (Figure 2.3F), although a series of χ^2 -tests did not identify a single session containing a statistically significant discrepancy in accuracy. A visual inspection of Benji's overall and side switch accuracies (Figure 2.2F) shows that the two measures traced each other quite precisely over the course of training. As such, none of the two animals show evidence for their initial licking side biases significantly impairing rule acquisition on this task.




Figure 2.1: Location discrimination task acquisition by female 1 (left) and 2 (right). **A** shows response rate (% of trials on which the animal made a lick response during the post-stimulus onset response window). **B** shows mean Log_{10} ratio of Right:Left licks during the response window, as a measure of rightward bias. **C** shows mean accuracy on trials requiring a left (green) versus right (purple) lick response. **D** shows mean accuracy averaged across all trials (red) and specifically for trials requiring the animal switch spouts from the most recently rewarded one (grey).

Table 2.1: Results of binomial tests testing Vanilla and Cardamom's accuracy across sessions for a deviation from chance level performance.

<i>Animal</i>	<i>Session number</i>	<i>Test results (sig. marked with *)</i>	
Vanilla	1-2	<i>Accuracy</i> = 34.5%, <i>p</i> = 0.025	<i>Accuracy</i> = 28.6%, <i>p</i> < 0.001*
	3-4	<i>Accuracy</i> = 35.4%, <i>p</i> = 0.001*	<i>Accuracy</i> = 35.9%, <i>p</i> < 0.001*
	5-6	<i>Accuracy</i> = 42.9%, <i>p</i> = 0.064	<i>Accuracy</i> = 53.5%, <i>p</i> = 0.425
	7-8	<i>Accuracy</i> = 50.0%, <i>p</i> = 0.999	<i>Accuracy</i> = 49.7%, <i>p</i> = 0.999
	9-10	<i>Accuracy</i> = 40.4%, <i>p</i> = 0.019	<i>Accuracy</i> = 39.5%, <i>p</i> = 0.110
	11-12	<i>Accuracy</i> = 48.2%, <i>p</i> = 0.890	<i>Accuracy</i> = 47.5%, <i>p</i> = 0.580
	13-14	<i>Accuracy</i> = 38.8%, <i>p</i> = 0.003*	<i>Accuracy</i> = 50.9%, <i>p</i> = 0.879
	15-16	<i>Accuracy</i> = 47.5%, <i>p</i> = 0.580	<i>Accuracy</i> = 50.0%, <i>p</i> = 0.999
Cardamom	1-2	<i>Accuracy</i> = 42.25%, <i>p</i> = 0.235	<i>Accuracy</i> = 52.79%, <i>p</i> = 0.528
	3-4	<i>Accuracy</i> = 43.63%, <i>p</i> = 0.215	<i>Accuracy</i> = 51.40%, <i>p</i> = 0.761
	5-6	<i>Accuracy</i> = 55.92%, <i>p</i> = 0.389	<i>Accuracy</i> = 40.00%, <i>p</i> = 0.424
	7-8	<i>Accuracy</i> = 53.21%, <i>p</i> = 0.444	<i>Accuracy</i> = 46.35%, <i>p</i> = 0.348
	9-10	<i>Accuracy</i> = 59.19%, <i>p</i> = 0.029	<i>Accuracy</i> = 65.97%, <i>p</i> < 0.001*
	11-12	<i>Accuracy</i> = 55.29%, <i>p</i> = 0.192	<i>Accuracy</i> = 68.12%, <i>p</i> < 0.001*
	13-14	<i>Accuracy</i> = 23.87%, <i>p</i> = 0.001*	<i>Accuracy</i> = 40.56%, <i>p</i> = 0.002*
	15-16	<i>Accuracy</i> = 47.19%, <i>p</i> = 0.371	<i>Accuracy</i> = 52.80%, <i>p</i> = 0.230

* sig. threshold Bonferroni-corrected to $p \leq 0.003$

Table 2.2: Results of single-sample t-tests testing Vanilla and Cardamom's Log_{10} Right:Left lick ratios for a difference from 0 (i.e. no side bias).

<i>Animal</i>	<i>Session number</i>	<i>Test results (sig. marked with *)</i>	
Vanilla	1-2	$R:L= -0.31, t= -4.08, p< 0.001^*$	$R:L= -0.28, t= -5.91, p< 0.001^*$
	3-4	$R:L= -0.14, t= -3.38, p< 0.001^*$	$R:L= -0.078, t= -2.05, p= 0.042$
	5-6	$R:L= 0.16, t= 4.10, p< 0.001^*$	$R:L= 0.29, t= 7.31, p< 0.001^*$
	7-8	$R:L= 0.27, t= 6.74, p< 0.001^*$	$R:L= 0.24, t= 5.44, p< 0.001^*$
	9-10	$R:L= 0.24, t= 5.65, p< 0.001^*$	$R:L= 0.38, t= 8.27, p< 0.001^*$
	11-12	$R:L= 0.28, t= 3.89, p< 0.001^*$	$R:L= 0.26, t= 5.94, p< 0.001^*$
	13-14	$R:L= 0.23, t= 5.23, p< 0.001^*$	$R:L= 0.32, t= 7.59, p< 0.001^*$
	15-16	$R:L= 0.24, t= 5.63, p< 0.001^*$	$R:L= 0.16, t= 4.14, p< 0.001^*$
Cardamom	1-2	$R:L= -0.32, t= -5.15, p< 0.001^*$	$R:L= 0.14, t= 3.81, p< 0.001^*$
	3-4	$R:L= -0.11, t= -2.30, p= 0.215$	$R:L= 0.37, t= 9.67, p< 0.001^*$
	5-6	$R:L= 0.06, t= 1.48, p= 0.389$	$R:L= 0.34, t= 3.75, p< 0.001^*$
	7-8	$R:L= 0.18, t= 4.81, p< 0.001^*$	$R:L= 0.32, t= 8.48, p< 0.001^*$
	9-10	$R:L= 0.42, t= 9.58, p< 0.001^*$	$R:L= 0.42, t= 8.21, p< 0.001^*$
	11-12	$R:L= 0.39, t= 8.71, p< 0.001^*$	$R:L= 0.22, t= 5.79, p< 0.001^*$
	13-14	$R:L= 0.096, t= 2.79, p= 0.005$	$R:L= 0.052, t= 1.39, p= 0.291$
	15-16	$R:L= 0.13, t= 4.10, p= 0.003^*$	$R:L= 0.067, t= 1.76, p= 0.187$

* sig. threshold Bonferroni-corrected to $p \leq 0.003$




Figure 2.2: Location discrimination task acquisition by Male 1 (Benji). *A-B* show lick patterns across trials taken from one early (*A*) and later (*B*) training session. Individual licks are shown in grey; reward deliveries in response to licks following stimulus onset (marked with a sine grating) are shown in red. The animal could collect ≤ 2 rewards per trial with a >1 s interval, provided it did not lick the wrong spout in the intervening period. A 5 s no-lick period prior to stimulus onset was implemented to discourage compulsive licking. *C* shows response rate (% of trials on which the animal made a lick response during the post-stimulus onset response window). *D* shows median RTs (in s) on correct trials. *E* shows mean accuracy on trials requiring a left vs right response. *F* shows mean accuracy averaged across all trials (red) and specifically for trials requiring the animal to switch spouts from the most recently rewarded one (grey). Grey triangles demarcate sessions where stimulus contrast was reduced to 60%, 40%, 20%, 10%, 7%, 5%, and 3%.




Figure 2.3: Location discrimination task acquisition by Male 2 (Fatty). *A-B* show lick patterns across trials taken from one early (*A*) and later (*B*) training session. Individual licks are shown in grey; reward deliveries in response to licks following stimulus onset (marked with a sine grating) are shown in red. The animal could collect ≤ 2 rewards per trial with a >1 s interval, provided it did not lick the wrong spout in the intervening period. A 5 s no-lick period prior to stimulus onset was implemented to discourage compulsive licking. *C* shows response rate (% of trials on which the animal made a lick response during the post-stimulus onset response window). *D* shows median RTs (in s) on correct trials. *E* shows mean accuracy on trials requiring a left vs right response. *F* shows mean accuracy averaged across all trials (red) and specifically for trials requiring the animal to switch spouts from the most recently rewarded one (grey). Grey triangles demarcate sessions where stimulus contrast was reduced to 60%, 40%, 20%, 10%, 7%, 5%, and 3%.

Table 2.3: Results of χ^2 -tests comparing Benji's and Fatty's accuracy on trials requiring right- versus left-sided responses

<i>Animal</i>	<i>Session number</i>	<i>Test results (sig. marked with *)</i>	
Benji	1-2	$\chi^2 = 9.82, p = 0.002$	$\chi^2 = 6.54, p = 0.011$
	3-4	$\chi^2 = 1.57, p = 0.21$	$\chi^2 = 13.76, p < 0.001^*$
	5-6	$\chi^2 = 7.77, p = 0.005$	$\chi^2 = 21.038, p < 0.001^*$
	7-8	$\chi^2 = 2.94, p = 0.087$	$\chi^2 = 11.55, p < 0.001^*$
	9-10	$\chi^2 = 2.13, p = 0.14$	$\chi^2 = 9.22, p = 0.002$
	11-12	$\chi^2 = 24.51, p < 0.001^*$	$\chi^2 = 15.43, p < 0.001^*$
	13-14	$\chi^2 = 21.31, p < 0.001^*$	$\chi^2 = 26.36, p < 0.001^*$
	15-16	$\chi^2 = 16.11, p < 0.001^*$	$\chi^2 = 0.49, p = 0.48$
	17-18	$\chi^2 = 17.51, p < 0.001^*$	$\chi^2 = 6.19, p = 0.013$
	19-20	$\chi^2 = 5.69, p = 0.017$	$\chi^2 = 0.41, p = 0.52$
	21-22	$\chi^2 = 20.56, p < 0.001^*$	$\chi^2 = 1.49, p = 0.22$
	23-24	$\chi^2 = 11.79, p < 0.001^*$	$\chi^2 = 16.52, p < 0.001^*$
	25-26	$\chi^2 = 0.018, p = 0.89$	$\chi^2 = 1.23, p = 0.027$
	27-28	$\chi^2 = 3.19, p = 0.074$	$\chi^2 = 7.83, p = 0.005$
Fatty	1-2	$\chi^2 = 0.46, p = 0.49$	$\chi^2 = 10.85, p < 0.001^*$
	3-4	$\chi^2 = 0.026, p = 0.87$	$\chi^2 = 14.84, p < 0.001^*$
	5-6	$\chi^2 = 19.61, p < 0.001^*$	$\chi^2 = 22.83, p < 0.001^*$
	7-8	$\chi^2 = 0.49, p = 0.48$	$\chi^2 = 3.75, p < 0.053$
	9-10	$\chi^2 = 46.45, p < 0.001^*$	$\chi^2 = 20.48, p < 0.001^*$
	11-12	$\chi^2 = 0.054, p = 0.82$	$\chi^2 = 5.28, p = 0.021$
	13-14	$\chi^2 = 3.29, p = 0.07$	$\chi^2 = 1.39, p = 0.24$
	15-16	$\chi^2 = 7.63, p = 0.006$	$\chi^2 = 0.49, p = 0.48$
	17-18	$\chi^2 = 1.25, p = 0.26$	$\chi^2 = 0.083, p = 0.77$
	19-20	$\chi^2 = 7.60, p = 0.006$	$\chi^2 = 3.49, p = 0.062$
	21-22	$\chi^2 = 5.33, p = 0.021$	$\chi^2 = 4.57, p = 0.033$
	23-24	$\chi^2 = 0.14, p = 0.70$	$\chi^2 = 5.74, p = 0.016$
	25-26	$\chi^2 = 2.13, p = 0.14$	$\chi^2 = 0.14, p = 0.71$
	27-28	$\chi^2 = 2.21, p = 0.14$	$\chi^2 = 1.94, p = 0.16$

* sig. threshold Bonferroni-corrected to $p \leq 0.0017$

2.3.3 DEGRADING STIMULUS CONTRAST IMPAIRS DISCRIMINATION ACCURACY

To identify task conditions that might pose a sufficient challenge to visual decision-making and thus possibly increase reliance on predictive auditory cues, I gradually reduced the contrast of the grating and noise that animals were required to discriminate over the course of consecutive sessions, in the following order: 60%, 40%, 20%, 10%, 7%, 5%, and 3%. To examine whether location discrimination accuracy began significantly faltering during any of these sessions, I ran two preliminary χ^2 -tests - one per animal - comparing proportions of successes between eight training sessions, which included the final session featuring 100% contrast and the following seven sessions featuring the aforementioned contrast reductions. The tests indicated that Benji's accuracy significantly differed between these sessions ($\chi^2 = 40.32$, $p < 0.001$).

Follow-up χ^2 -tests (sig. threshold Bonferroni-corrected to $p \leq 0.006$) revealed that accuracy became significantly lower at contrasts $\leq 5\%$ compared to the maximum contrast condition (5%: $\chi^2 = 15.27$, $p < 0.001$; 3%: $\chi^2 = 14.04$, $p < 0.001$). The test also flagged up a significant difference in Fatty's accuracy amongst the seven tested sessions ($\chi^2 = 41.53$, $p < 0.001$), which follow-up tests pinned down to performance becoming significantly worse at 7% and 3% contrast compared to the maximum contrast condition (7%: $\chi^2 = 7.98$, $p = 0.005$; 5%: $\chi^2 = 1.06$, $p = 0.30$; 3%: $\chi^2 = 20.19$, $p < 0.001$).

2.3.4 DEGRADING STIMULUS CONTRAST SLOWS DISCRIMINATION DECISIONS

Next, I examined the effects of degrading stimulus contrast on reaction times (RTs), to identify conditions which proved sufficiently taxing as to cause animals to make slower perceptual decisions (shown in Figures 2.2D and 2.3D). In all cases, RTs were sampled exclusively from trials where animals produced accurate responses, and outlier RTs of < 100 ms were eliminated from analysis, as these most likely did not represent stimulus-driven responses.

To test whether RTs were significantly influenced by differing stimulus contrast conditions, I performed two Kruskal-Wallis tests - one per animal - for differences in median RTs across eight levels of the contrast variable. The tests revealed a main effect of contrast for both Benji ($K \text{ stat} = 49.09$, $p < 0.001$) and Fatty ($K \text{ stat} = 204.42$, $p < 0.001$). Given that the number of clarifying post-hoc comparisons that could be made is extremely high, median RTs in $\leq 20\%$ contrast conditions were compared specifically to those in the 100% contrast condition using a series of Mann-Whitney Wilcoxon rank sum tests (sig. threshold Bonferroni-corrected to $p \leq 0.010$), to verify

whether RTs under any of these conditions were significantly impaired compared to the condition that would reasonably support maximum performance.

Benji’s tests revealed that median RTs at a contrast level of 100% were not significantly faster than at 20% (*Rank sum* = 2258.5, $p = 0.35$) and 10% (*Rank sum* = 1152.0, $p = 0.45$), however significantly faster than at 7% (*Rank sum* = 1232.5, $p = 0.010$), 5% (*Rank sum* = 1439.5, $p = 0.006$) and 3% contrasts (*Rank sum* = 1136.0, $p = 0.009$). Fatty’s tests revealed that median RTs at 100% contrast were not significantly faster than at 20% (20% *Rank sum* = 1663.0, $p = 0.023$), however were significantly faster compared to median RTs at all contrasts $\leq 10\%$ (10% *Rank sum* = 1614.5, $p < 0.001$; 7% *Rank sum* = 985.0, $p < 0.001$; 5% *Rank sum* = 691.5, $p < 0.001$; 3% *Rank sum* = 1073.5, $p < 0.001$). Thus, the evidence suggests that both animals displayed a tendency towards significant slowing of responses with declining stimulus contrast. This took place at $\leq 7\%$ contrast in Benji, and $\leq 10\%$ contrast in Fatty. Since visual conditions during these training sessions provide a challenge that is most likely to encourage animals to use the given predictive cues to aid in stimulus detection, statistical analyses examining whether the animals did indeed do this will center around the contrast ranges found to be specifically challenging for each animal.

2.3.5 INVALID PREDICTIVE CUES DO NOT IMPAIR DISCRIMINATION ACCURACY

The animals’ accuracy and RT data indicate that task performance suffered as a result of reducing stimulus and background noise contrast. This provides some footing to explore my evidence-based (Kok et al., 2014; Stokes et al., 2009) conjecture that animals performing under visually taxing conditions might increase their reliance on predictive auditory cues to assist with stimulus detection. I reasoned that, if animals exploited the cues to pre-activate orientation-tuned V1 neuronal populations to enable faster stimulus detection and processing, the predictive cue advantage should manifest in more accurate and rapid judgements about stimulus location. This advantage should not take place where predictions made by the cues turn out to be incorrect.

To test whether animals measurably based their behaviours on predictive cues, I compared accuracy on trials featuring valid and invalid cues (80% and 20% of trials, respectively), aggregated across contrast levels associated with significantly slowed RTs in each animal ($\leq 7\%$ in Benji and $\leq 10\%$ in Fatty; Figure 2.4B). Data were submitted to a series of χ^2 -tests comparing proportions of successes between valid and invalid cueing conditions. These tests revealed no evidence that accuracy differed significantly depending on predictive cue validity, in either Benji (Valid cueing *Accuracy* = 75.2%; Invalid cueing *Accuracy* = 80.4%; $\chi^2 = 0.434$, $p = 0.51$) or Fatty (Valid cueing *Accuracy* = 78.3%; Invalid cueing *Accuracy* = 80.2%; $\chi^2 = 0.072$, $p = 0.788$).

2.3.6 INVALID PREDICTIVE CUES DO NOT SLOW DISCRIMINATION DECISIONS

I proceeded to examine whether interfering with predictive cue validity impaired the animals' performance by significantly slowing the time taken to provide correct responses. Median RTs taken from sessions aggregated across contrast levels associated with significantly slowed reaction times in each animal ($\leq 7\%$ in Benji and $\leq 10\%$ in Fatty; Figure 2.4B) were entered into a pair of Mann-Whitney Wilcoxon tests to compare median RTs on validly and invalidly cued trials. The tests revealed that median RTs did not significantly differ between valid and invalid cueing conditions in both Benji (Valid cueing $Med= 741.5$ ms, $IQR= 1399$ ms, Invalid cueing $Med= 634.5$ ms, $IQR= 1389.75$ ms; $Rank\ sum= 1133.5$, $p= 0.24$) and Fatty (Valid cueing $Med= 420$ ms, $IQR= 242.5$ ms, Invalid cueing $Med= 411$ ms, $IQR= 234$ ms; $Rank\ sum= 829.9$, $p= 0.44$).

It is possible that predictive cue quality influenced RTs without shifting the central tendency of the distribution but rather by affecting the distribution profile itself. In an attempt to more closely understand these possibilities, I fit each animal's RT distribution, obtained under valid and invalid cueing conditions, to an ex-Gaussian probability density function composed of a Gaussian convolved with an exponential curve (using the open-source Ex-Gauss Matlab code library; Zandbelt, 2014). The parameters estimated by the model-fitting algorithm consist of: μ and σ , which represent the mean and spread of the normally distributed component of the function, as well as τ , representing the exponentially distributed element that is rooted in the positive skew of the RT distribution and provides a measure of the prevalence of slow responses. The ex-Gaussian distribution has previously been found to provide a good fit to empirical RT distributions obtained in human experiments (Balota & Spieler, 1999; Tarantino et al., 2014), which provided a rationale to use the model to understand the influence of predictive cue validity on mouse RTs.

Figure 2.4C shows the RT distributions extracted from responses given by each animal on valid and invalid cueing trials, collapsed across the selected contrast conditions, as well as the distributions predicted by the ex-Gaussian curve fitted to each observed distribution. The resulting parameter estimates of the ex-Gaussian RT distributions for Benji (Valid cueing: $\mu= 233.21$, $\sigma= 106.36$, $\tau= 871.37$; Invalid cueing: $\mu= 258.49$, $\sigma= 102.27$, $\tau= 948.66$) and Fatty (Valid cueing: $\mu= 215.32$, $\sigma= 84.29$, $\tau= 332.84$; Invalid cueing: $\mu= 224.44$, $\sigma= 53.86$, $\tau= 351.48$) indicate that providing animals with incorrect auditory predictions increased their tendency to make slow responses, as reflected in the trend that the exponential component, τ , increases under invalid cueing conditions in both animals. Notably however, a series of χ^2 tests comparing observed RT distributions with those predicted by the ex-Gaussian point to a significant caveat.




Figure 2.4: Reaction time and accuracy data under valid and invalid predictive cueing conditions. **A** shows the mean accuracy and **B** shows the median reaction times of Mouse 1 (left) and 2 (right) on trials where auditory cues made valid (blue) versus invalid (red) predictions about stimulus orientation. **C** shows the RT histograms of Mouse 1 (top) and 2 (bottom) aggregated across all sessions featuring a mixture of valid and invalid cues, segregated by predictive cue validity (blue: valid; red: invalid). The overlying striated lines show the histogram profile predicted by the ex-Gaussian distribution fitted to the data.

The obtained χ^2 statistics (Benji: Valid cueing - $\chi^2=92.79$, Invalid cueing - $\chi^2=17.19$; Fatty: Valid cueing - $\chi^2=24.49$, Invalid cueing - $\chi^2=7.98$) show that for only one of the four distributions does the test statistic fall under the critical value of 11.1 given 5 degrees of freedom (constrained by the number of probability bins used to compare the empirical and predicted distributions). As such, the ex-Gaussian model does not provide sufficiently good fit to the animals' RT data as to allow reasonable inferences about potential effects of cueing validity on response speed.

2.3.7 MICE MAY HAVE HELD FALSE BELIEFS ABOUT PREDICTIVE CUE MEANING

The animals' accuracy and RT data provide no reason to believe that they were using predictions offered by the auditory cues to aid in stimulus detection and boost location discrimination performance. This was evidenced by the finding that location judgements towards stimuli whose orientations violated cued predictions (i.e. invalidly cued trials) were not less accurate or slower compared to when stimulus orientations complied with prediction. This leads to two primary possible explanations for why the anticipated cue validity effect was not found.

One possibility is that mice treated the auditory cues as irrelevant and uninformative features of their environment, and thus performed the location discrimination task without any input from the auditory cues. This is plausible, given that the auditory cues predicted a stimulus feature (i.e. orientation) which was not explicitly linked to that on which the animals based their judgements (i.e. location). However, one observation convinced me that mice were not, in fact, ignoring the predictive cues. During my surveillance of the animals during training, I noticed a tendency to begin vigorously licking the spouts immediately following auditory cue onset. This was verified by two one-sample t-tests which found that the mean number of spout licks obtained during the post-cue period significantly deviated from zero (Benji: $t=9.79$, $p<0.001$; Fatty: $t=12.47$, $p<0.001$) - the number of licks obtained prior to cue onset due to the mandatory 'silent period' required to trigger the next trial. Clearly, the auditory cues did not pass by entirely unnoticed.

Importantly, the animals were capable of rapidly switching the spout they were licking upon stimulus onset. This is supported by two one-sample t-tests finding that the proportion of trials on which both Benji ($M=0.33$, $SD=0.12$, $t=8.58$, $p<0.001$) and Fatty ($M=0.29$, $SD=0.08$, $t=11.98$, $p<0.001$) switched from the spout they last licked in the post-cue period upon stimulus onset significantly differed from zero. This suggests that the auditory cues prompted the animals to prepare responses that they perhaps ultimately anticipated to pay off, but did not commit to, given their ability to switch licking spouts following stimulus onset. The implication of this is that animals may have regarded the auditory cues as providing them with information that formed the basis for a plausibly beneficial anticipatory response. One question

that arises from this is - what information did the mice think they were being given?

I suspected that the animals operated on the false assumption that the cues actually provided them with predictive information about the likely location of an upcoming stimulus. While the cues were programmed to have no global statistical relationship with stimulus location, it is possible that animals based their impressions on memories of a small number of recent trials, which would increase the risk of perceived statistical dependency between cue type and stimulus location. To explore this possibility, I examined whether the animals' choice of spout during the post-cue period on any given trial was significantly biased by the location of the stimulus (and thus, the pay-off spout) in the most recently rewarded trial featuring the same auditory cue (illustrated in Figure 2.5A). A look at one animal's licking patterns across individual trials from several sessions indicates that there may be a relevant relationship (Figure 2.5B-C), as when an auditory cue had previously been followed by a right-sided stimulus, the animal tended to lick the right spout in anticipation of the next stimulus.

To statistically assess whether animals treated the auditory cues as supplying them with location-predicting information, I extracted their average Log_{10} -ratio of Right:Left licks from the post-cue period, as an indicator of right side bias in cue-evoked licking. This measure was obtained across all training sessions featuring auditory cues and collapsed across cues predicting horizontally and vertically oriented gratings, since the 'false belief' hypothesis assumes that orientation information was considered irrelevant, and that both types of auditory cues would be equally affected by the false belief if it did take place. Two independent-samples t-tests - one per animal - were then used to test for significant differences between mean post-cue Log_{10} R:L lick ratios on trials when an auditory cue was last associated with a right- versus left-sided stimulus. The tests revealed that, on trials with a cue that was last followed by a right-sided stimulus, post-cue right licking bias was significantly stronger both in Benji (Previous left: $M = -0.17$, $SD = 1.00$; Previous right: $M = 0.19$, $SD = 0.88$, $t = 7.34$, $p = 0.01$) and Fatty (Previous left: $M = 0.11$, $SD = 1.20$; Previous right: $M = 0.58$, $SD = 1.19$, $t = 6.95$, $p = 0.01$), compared to when that cue had previously been associated with a left-sided stimulus.

Thus, examining the animals' anticipatory licking patterns leads to an important conclusion. The finding that mice were not significantly influenced by the correctness with which auditory cues predicted stimulus orientation can be explained by the likelihood that they held a false belief that cues told them about stimulus location rather than orientation. In this case, a cue providing invalid orientation information could not be disruptive of performance, as this was not the stimulus dimension the animals attended. Instead, animals acted upon their false belief about the informative content of predictive cues by orienting towards the spout they considered likely to pay off upon stimulus onset.




Figure 2.5: Post-cue licking patterns are shaped by the false belief that auditory cues predict stimulus location. **A** demonstrates the procedure used to determine whether mice acted on false beliefs about auditory cue meaning. For each occurrence of a particular auditory cue (6 kHz or 12 Hz; here denoted as red or green), the analysis involved identifying the most recent occurrence of that particular cue and registering the location of the stimulus which appeared on that trial. Thus, the post-cue licking pattern obtained during each trial (demarcated with the mouse) can be related to (pink arrow) stimulus location upon the most recent encounter with the same cue (demarcated in pink). **B** and **C** show left and right licks made by one animal (Benji) during the post-cue period (0-2 s) on trials segregated on the basis of the previous cue encounter involving a right-side stimulus (**B**) or left-side stimulus (**C**). Here, left and right licks are shown in red and blue respectively. **D** shows each animal's mean Log_{10} ratio of Right:Left licks for cues which, on previous encounter, were associated with a left or right stimulus location.

2.3 RESULTS - PILOT 2

2.4.1 RECAP AND OBJECTIVES

The goal of this experiment was to pick up from the failure of the previously piloted task (see Results section 2.3) and pilot a further task with the stated goal of evoking sensory expectations in mouse visual cortex. In this 2-alternative forced choice (2-AFC) paradigm, head-fixed mice were required to judge whether the orientations of two sine-wave gratings were same or different, by licking the appropriate water spout. Upon the animals' achievement of a high level of performance on this task, my plan was to introduce pre-stimulus auditory cues that were predictive of stimulus orientation, and combine these with manipulations of the relative frequencies of same versus different stimulus pairings in various experimental blocks. It was my reasoning that this would offer a window into the effects of environmental probabilities on perceptual hypotheses. However, as discussed below, the animals failed to attain above-chance accuracy on this task after several weeks of training, and the pilot was terminated.

The raw data collected from each training session consisted of timestamps for the following events: i) individual lick detections, ii) stimulus onsets, and iii) openings of the solenoid valves which delivered water rewards. Data from each male, whom I will refer to as Male 1 (Saruman) and Male 2 (Frodo), will be examined separately.

2.4.1 MICE CANNOT LEARN TO MAKE SAME-DIFFERENT JUDGEMENTS

Despite the fact that both animals showed high response rates (Figure 2.6A), neither of them exceeded chance-level performance over the course of >20 training sessions (Figure 2.6C), as shown by a series of binomial tests comparing proportions of correct responses given on each session to a theoretical 0.5 success probability (results shown in Table 2.4). The finding that performance was, in fact, significantly below chance rate on a proportion of trials implied that animals were not merely guessing, but rather being driven by counterproductive licking biases which overrode potential understanding of task rules. In the following explorations of this possibility, all results are interpreted with respect to a Bonferroni-corrected significance threshold of $p \leq 0.002$, to account for sequential testing across >20 training sessions for each animal.

Both animals indeed showed a significant bias towards licking one spout in the early stages of training (Figure 2.6B). In Saruman, this is evidenced by a series of single-sample t-tests finding that the cross-trial mean Log_{10} ratio of Right:Left licks

harvested during the response window was significantly lower than zero during the first five training sessions, (Sessions (S): S1 $t = -4.56$, $p < 0.001$; S2 $t = -4.95$, $p < 0.001$; S3 $t = -5.76$, $p < 0.001$; S4: $t = -10.33$, $p < 0.001$; S5: $t = -5.96$, $p < 0.001$). This bias manifested in a significant accuracy advantage for trials requiring left versus right lick responses (Figure 2.6C), as shown by a series of χ^2 -tests comparing proportions of successes between the two trial types (Sessions (S): S1 $\chi^2 = 9.083$, $p = 0.003$; S2 $\chi^2 = 17.02$, $p < 0.001$; S3 $\chi^2 = 29.27$, $p < 0.001$; S4 $\chi^2 = 31.72$, $p < 0.001$; S5 $\chi^2 = 15.20$, $p < 0.001$). At the same time, a series of binomial tests showed that right-sided accuracy consistently fell below chance during most sessions in this time period (Sessions (S): S1 *Accuracy* = 27.8%, $p < 0.001$; S2 *Accuracy* = 36.5%, $p = 0.007$; S3 *Accuracy* = 33.3%, $p < 0.001$; S4 *Accuracy* = 31.1%, $p < 0.001$; S5 *Accuracy* = 37.8%, $p = 0.023$), implying an active and detrimental licking bias was a more likely explanation for poor overall performance than guessing.

A bias towards licking predominantly one spout was even more pronounced in Frodo. A series of single-sample t-tests found that Frodo's cross-trial mean Log_{10} ratio of Right:Left licks during the response window was significantly higher than zero during 8 of the first 10 training sessions, (Sessions (S): S1 $t = 1.90$, $p < 0.05$; S2 $t = 3.83$, $p < 0.001$; S3 $t = 7.26$, $p < 0.001$; S4 $t = 7.50$, $p < 0.001$; S5 $t = 1.55$, $p = 0.122$; S6 $t = 5.65$, $p < 0.001$; S7 $t = 2.66$, $p = 0.009$; S8 $t = 5.46$, $p < 0.001$; S9 $t = 10.78$, $p < 0.001$; S10 $t = 7.59$, $p < 0.001$), implying a robust bias towards the right spout. A series of χ^2 -tests confirmed that the time-course of his rightward bias was associated with a significantly greater accuracy for right-sided compared to left-sided trials (Sessions (S): S1 $\chi^2 = 15.52$, $p < 0.001$; S2 $\chi^2 = 27.45$, $p < 0.001$; S3 $\chi^2 = 45.55$, $p < 0.001$; S4 $\chi^2 = 50.19$, $p < 0.001$; S5 $\chi^2 = 17.15$, $p < 0.001$; S6 $\chi^2 = 38.14$, $p < 0.001$; S7 $\chi^2 = 10.15$, $p = 0.001$; S8 $\chi^2 = 41.28$, $p < 0.001$; S9 $\chi^2 = 31.06$, $p < 0.001$; S10 $\chi^2 = 56.95$, $p < 0.001$). A series of binomial tests also mirrored Saruman's results, showing that Frodo's accuracy on trials requiring licks of the non-preferred left spout fell significantly short of chance across most sessions during which the animal exhibited the right-side licking bias (Sessions (S): S1 *Accuracy* = 26.6%, $p < 0.001$; S2 *Accuracy* = 28.3%, $p < 0.001$; S3 *Accuracy* = 17.8%, $p < 0.001$; S4 *Accuracy* = 22.3%, $p < 0.001$; S5 *Accuracy* = 37.5%, $p = 0.014$; S6 *Accuracy* = 24.6%, $p < 0.001$; S7 *Accuracy* = 41.1%, $p = 0.113$; S8 *Accuracy* = 24.5%, $p < 0.001$; S9 *Accuracy* = 11.7%, $p < 0.001$; S10 *Accuracy* = 16.8%, $p < 0.001$).

Observations across the two animals show that both mice had a pre-existing preference for licking one spout - a bias that actively interfered with their ability to orient towards the non-preferred spout and thus obtain sufficient reward experience for task rule acquisition. This bias required a non-trivial amount of training to neutralise. Ultimately, animal performance plateaued at chance level for trials requiring both left and right judgements, providing no indication that accuracy would exceed chance within a reasonable time frame.




Figure 2.6: Same-different discrimination task performance by Male 1 (Saruman; left) and Male 2 (Frodo; right) **A** shows the response rate (defined as % of trials on which the animal made a lick response during the response window). **B** shows mean Log_{10} ratio of Right:Left licks during the response window, as a measure of rightward bias. **C** shows mean accuracy on trials requiring a left (green) versus right (purple) lick response. **D** shows mean accuracy averaged across all trials (red) and specifically for trials requiring the animal switch spouts from the most recently rewarded one (grey).

Table 2.4: Results of binomial tests testing Saruman and Frodo’s accuracy across sessions for a deviation from chance level performance.

<i>Animal</i>	<i>Session number</i>	<i>Test results (sig. marked with *)</i>	
Saruman	1-2	<i>Accuracy= 34.8%, $p < 0.001^*$</i>	<i>Accuracy= 48.4%, $p = 0.075$</i>
	3-4	<i>Accuracy= 45.3%, $p = 0.31$</i>	<i>Accuracy= 46.8%, $p = 0.43$</i>
	5-6	<i>Accuracy= 51.2%, $p = 0.82$</i>	<i>Accuracy= 45.9%, $p = 0.25$</i>
	7-8	<i>Accuracy= 39.1%, $p = 0.001^*$</i>	<i>Accuracy= 39.4%, $p = 0.007$</i>
	9-10	<i>Accuracy= 57.1%, $p = 0.11$</i>	<i>Accuracy= 48.4%, $p = 0.75$</i>
	11-12	<i>Accuracy= 49.3%, $p = 0.94$</i>	<i>Accuracy= 45.9%, $p = 0.25$</i>
	13-14	<i>Accuracy= 46.4%, $p = 0.42$</i>	<i>Accuracy= 49.2%, $p = 0.82$</i>
	15-16	<i>Accuracy= 44.6%, $p = 0.186$</i>	<i>Accuracy= 37.8%, $p = 0.001^*$</i>
	17-18	<i>Accuracy= 44.4%, $p = 0.13$</i>	<i>Accuracy= 35.5%, $p < 0.001^*$</i>
	19-20	<i>Accuracy= 45.4%, $p = 0.26$</i>	<i>Accuracy= 30.4%, $p < 0.001^*$</i>
	21-22	<i>Accuracy= 47.4%, $p = 0.52$</i>	<i>Accuracy= 40.9%, $p = 0.045$</i>
	23-24	<i>Accuracy= 46.4%, $p = 0.42$</i>	<i>Accuracy= 48.0%, $p = 0.68$</i>
	25-26	<i>Accuracy= 47.7%, $p = 0.63$</i>	<i>Accuracy= 43.5%, $p = 0.082$</i>
	27	<i>Accuracy= 39.8%, $p = 0.007$</i>	
Frodo	1-2	<i>Accuracy= 37.3%, $p = 0.009$</i>	<i>Accuracy= 39.8%, $p = 0.023$</i>
	3-4	<i>Accuracy= 30.4%, $p < 0.001^*$</i>	<i>Accuracy= 34.9%, $p < 0.001^*$</i>
	5-6	<i>Accuracy= 43.5%, $p = 0.13$</i>	<i>Accuracy= 37.4%, $p = 0.003$</i>
	7-8	<i>Accuracy= 51.4%, $p = 0.80$</i>	<i>Accuracy= 35.0%, $p = 0.001^*$</i>
	9-10	<i>Accuracy= 27.2%, $p < 0.001^*$</i>	<i>Accuracy= 27.0%, $p < 0.001^*$</i>
	11-12	<i>Accuracy= 29.3%, $p < 0.001^*$</i>	<i>Accuracy= 32.8%, $p = 0.008$</i>
	13-14	<i>Accuracy= 34.2%, $p < 0.001^*$</i>	<i>Accuracy= 34.5%, $p < 0.001^*$</i>
	15-16	<i>Accuracy= 36.8%, $p = 0.003$</i>	<i>Accuracy= 38.2%, $p = 0.004$</i>
	17-18	<i>Accuracy= 49.1%, $p = 0.88$</i>	<i>Accuracy= 50.8%, $p = 0.93$</i>
	19-20	<i>Accuracy= 43.6%, $p = 0.042$</i>	<i>Accuracy= 47.9%, $p = 0.98$</i>
	21-22	<i>Accuracy= 39.2%, $p = 0.009$</i>	<i>Accuracy= 42.8%, $p = 0.13$</i>
	23	<i>Accuracy= 51.0%, $p = 0.87$</i>	

* sig. threshold Bonferroni-corrected to $p \leq 0.002$

2.5 DISCUSSION

This series of pilot experiments highlighted several strengths and limitations of the head-fixed mouse as a behavioural model for studying perceptual decision-making. In the first pilot experiment, water-restricted and head-fixed mice were trained to perform a 2-alternative forced choice (2-AFC) task, in which they were required to judge the locations of visual sine-wave gratings by licking a water port to the left or right side of their mouth. This experiment showed that mice were capable of acquiring the task within a short training period (~ 10 sessions), and to a high standard of performance ($>90\%$ accuracy; Figures 2.2F and 2.3F). However, the hypothesised possibility that their behaviour could be significantly influenced by cues providing predictive information about a behaviourally irrelevant stimulus attribute (i.e. orientation) did not find its roots in the evidence. Mice were not significantly less accurate or slowed in indicating stimulus location after being given invalid predictive cues (Figure 2.4). Instead, the animals' pre-stimulus licking patterns suggested that they operated under the false impression that cues were, in fact, predicting stimulus location, rather than orientation.

This was evidenced by their tendency, during the anticipatory post-cue period, to lick the spout which was rewarded on the most recent last trial featuring a particular cue (Figure 2.5). Thus, if an animal's most recent encounter with a 12 kHz auditory cue was immediately followed by a left-sided grating stimulus, then the animal exhibited a significant bias towards licking the left spout upon its next encounter with that same auditory cue. The finding that mice may have misinterpreted orientation-predictive cues as providing spatial information makes the behavioural paradigm inappropriate for the stated goal of exploring orientation-specific expectations and anticipatory activation patterns in the mouse visual cortex. Thus, the 2-AFC location discrimination task did not form the basis for proceeding with training and chronic electrophysiological recordings in an experimental cohort of mice. However, it does extend our current knowledge of mouse capabilities within the realm of 2-AFC tasks. This is explored in greater detail below.

2.5.1 MICE CAN MAKE VISUAL LOCATION DISCRIMINATION DECISIONS UNDER HEAD-FIXED CONDITIONS

To the extent of my knowledge, this is the first demonstration of mice providing visual stimulus location judgements using a 2-AFC task design, under head-fixation. This contributes to a growing body of decision-making research exploring the behavioural repertoire of the mouse that can be harnessed to probe perceptual processing using a variety of paradigms, and in particular the 2-AFC task design which, as discussed below, enjoys a range of benefits over other tasks.

Until quite recently, perceptual decision-making was largely explored using freely moving paradigms. In these tasks, the limits of animals' visual discriminative abilities are probed by requiring them to swim towards target stimuli (Prusky et al., 2000; Wang et al., 2016) or make their way towards the correct touch-screen to provide a nose poke response (Busse et al., 2011; Graf et al., 2018; Turner et al., 2017), and thus indicate stimulus detection and location decisions. In tasks such as these, mice usually learn at a relatively high pace, showing above-chance discriminative responses within 10-15 training sessions (Turner et al., 2017).

While freely moving paradigms are immensely useful where the primary measure of interest is behaviour, they have several noteworthy limitations. These include the fact such paradigms often preclude precise stimulus delivery, making it difficult to determine parameters such as the spatial frequency of a visual display or odour concentration. Any degree of parameter certainty that is afforded by taking the animal's motion into account using tracking equipment will not be equivalent to controlled stimulus presentation, which limits the scope of perceptual investigation that can be carried out using freely moving animals. A further consideration which disadvantages freely moving paradigms is that they do not provide sufficient mechanical stability to enable neuronal signals to be visualised using methods such as calcium or voltage-sensitive dye imaging. Consequently, head-fixed behaviour has been gaining a foothold in the field as a means of enabling both sensory processing, cognition, and their underlying neuronal activities to be characterised using a variety of techniques, including the 2-AFC and go/no-go paradigms (Dombeck et al., 2007; Schwarz et al., 2010; Guo et al., 2014).

Importantly, while the go/no-go paradigm is a popular and effective method of extracting perceptual judgements (Abraham et al., 2012; Lee et al., 2012; Pinto et al., 2013), it comes with a number of caveats which cast the alternative – 2-AFC – design into a more favourable light. Firstly, mice are innately predisposed to lick the spout. The implication is that withholding a lick to provide a no-go response involves a host of impulse-regulating processes which likely contaminate the behaviour and resulting neuronal activity with cognitive and premotor components that are absent during a go response. More recently, a group of authors argued that the demands on balancing impulsiveness with instrumental learning lead the go/no-go task to underestimate mouse capabilities (Berditchevskaia et al., 2016). Furthermore, mice and rats tend towards opposing response policies in go/no-go paradigms, which highlights the likelihood that animals' behaviour on the task is suffused with their evolved roles in predator-prey niches (Jones et al., 2017). Such inherently different response policies imply that the go/no-go paradigm is not a just another neutral method for interrogating sensory discrimination capabilities across species. Arguably, all of these limitations can be redressed using a response harvesting system which places symmetrical demands on both judgements – something embodied by the 2-AFC paradigm. It is thus arguably important to establish the breadth of contexts and sensory modalities in which the 2-AFC task can be effectively used to study mouse perceptual decision-

making.

Evidence suggests that this paradigm can, in principle, support rapid task acquisition. In a 2-AFC somatosensory object localisation task (Guo et al., 2014), mice were required to lick a left or right water spout to judge whether a bar was touching their anterior or posterior whisker pad, which they learned to do at above-chance rate within 5-10 training sessions (corroborating previous findings in rats by Krupa et al., 2001). Introducing a brief delay between stimulus offset and response window to challenge working memory caused the animals' performance to instantly decline; however recovery was achieved after 5-10 further training sessions. This, combined with the time-course of learning the initial delay-free version of the 2-AFC task, suggests that 5-10 training sessions might be a standard time-course for mice to acquire a 2-AFC task featuring two water spouts.

The behavioural pilot described in this chapter provided the first demonstration of the 2-AFC paradigm being used to collect location discrimination responses from head-fixed mice in the visual modality. My observation that both mice I trained showed above-chance accuracy between the 10th and 15th training sessions, however, indicates that the visual 2-AFC may require longer training times than the previously explored somatosensory version (Guo et al., 2014). This might be attributable to rodents' fundamental reliance on the whisker system, and the likelihood that any behaviour would be acquired more rapidly within this modality.

2.5.2 TECHNICAL CAVEATS

It is undeniably important to note that the task was successfully learnt only by the male mice which underwent training. Two females which were trained concurrently with the males failed to show any signs of task acquisition (Figure 2.1C-D) at a time when the males reached accuracy levels exceeding 70% (Figures 2.2F and 2.3F), causing their training to be terminated after session n°16. The females' behavioural patterns suggested that both animals may have treated reinforcements as signalling that a particular spout was the exclusive source of reward, as opposed to providing rewards when stimuli were presented in specific locations. This likely prevented them from sampling each water spout at sufficient rates during each training session to achieve effective rule comprehension. The likelihood of this is reflected in one female's chronic tendency to respond accurately on trials requiring right-sided responses, while left-sided accuracy consistently fell below 50% (Figure 2.1C;left), and in the other female's tendency to vacillate between higher right-sided and left-sided accuracy on different training sessions (Figure 2.1C; right). The females' failure to learn the task leaves some uncertainty about the general utility of 2-AFC paradigms. Females have successfully been included in previous studies using the 2-AFC response harvesting system, however they are notably fewer in number than males in one study (Guo et al., 2014), while another study does not specify relative numbers of males and females undergoing behavioural training (Burgess et al., 2017). One further difficulty

with contextualising the failure of the two females to learn the task outlined in the first part of this chapter is that studies often focus on post-acquisition task performance and any interventions which may affect it, rather than exploring the speed and nature of progress in task acquisition, as well as any sex differences which might arise. This points to the likelihood that females are capable of learning the same 2-AFC tasks as males, albeit on different time-scales - something which appears to be under-explored.

It is also likely that laboratory mice, as a population, exhibit considerable variability in cognitive ability, such that training both males and females in larger numbers would ultimately reveal no significant difference in performance between the sexes. The studies described in this chapter certainly suffer from the shortcoming of training small numbers of animals, which prevents this work from clarifying the issue of potential sex differences in task performance. The decision to pilot tasks with small numbers was taken for two principal reasons. Firstly, it was my rationale that a paradigm in which behaviour is to be combined with chronic electrophysiological recordings should be learnable by the vast majority of laboratory mice, to prevent implantation of chronic recording devices into a group of animals of which only a small fraction would be able to fulfil task requirements. Thus, I reasoned that piloting a few mice should be sufficient to indicate how attainable a task is to the typical member of the C57BL/6 strain. Secondly, the decision to trial small numbers of animals was driven by practical constraints, as I was training animals in a single behavioural booth offered by another laboratory. Concurrently training large numbers of animals was not viable in this case - something that can certainly be done in future studies. Despite mixed results, however, successful task acquisition by the males was sufficient to enable a test of the 2-AFC location discrimination task's suitability for the purpose of exploring predictive coding.

2.5.3 MICE DO NOT ACT ON PREDICTIVE INFORMATION ABOUT TASK-IRRELEVANT DIMENSIONS

One of the main goals of this pilot was to test whether mice judging stimulus location would make use of auditory cues which predicted the orientations of upcoming grating stimuli to aid performance under challenging conditions. Such behavioural effects were to serve as the basis for exploring predictive neuronal activity in the animals' visual cortex. While dissociating the stimulus dimension determining behaviour (i.e. location) and the dimension of sensory prediction (i.e. orientation) seems counterintuitive, I opted to do this for several reasons.

One rationale was that, if task-irrelevant cues evoked any predictive cortical representations, these would likely be relatively free from explicit cognitive control, which might otherwise come into play if the cues were directly task-relevant. As a consequence, anticipatory activity in V1 might provide a window into automatic sensory predictions generated as a result of exposure to the statistical properties of

the environment, rather than cognitively imposed expectations. From my perspective, this would provide a robust indicator of the automaticity of predictive coding in sensory systems. A further rationale behind my choice to provide task-irrelevant auditory cues was that it would minimise the risk of visual input becoming irrelevant. Animals given task-relevant cues (i.e. those carrying direct predictions of upcoming stimulus location) would arguably lose any incentive to wait for the visual stimulus, as the auditory cue itself would enable them to automatically prepare a left versus right lick response. This would effectively minimise the need for the brain to represent any visual predictions, since the auditory cue would become the primary driver of choice behaviour while vision takes a functional back-seat. Finally, evidence suggests that dissociating the behaviourally relevant and predicted stimulus dimensions should not be problematic for the likelihood of evoking anticipatory cortical representations, and by extension any behaviours which may be rooted in them. Researchers have observed predictive stimulus-specific activations in mouse V1 under conditions where absolutely no requirements were placed on the animals to engage in perceptual decision-making and behaviour (i.e. their objective was purely to traverse a virtual tunnel and reach the end-goal of a water reward; Fiser et al., 2016). This indicates that sensory systems are capable of registering regularities and recruiting predictive neural codes even when they have little direct applicability to behaviour.

To test whether mice exhibited signs of acting upon task-irrelevant predictive information, I introduced auditory cues (400 ms duration; 6 kHz or 12 kHz) which preceded by 2 s the grating stimuli whose locations the mice were required to judge, and predicted whether these would be horizontally or vertically oriented. I reasoned that, if the animals were sufficiently challenged by the task – brought about by conditions of degraded stimulus contrast - they might, to some degree rely on the orientation-specific predictive information to pre-activate neuronal ensembles in V1 (found to boost stimulus detection in humans by Kok et al., 2014 & Stokes et al., 2009) to aid grating detection, and consequently localisation.

To explore this possibility, I examined the accuracy and speed of animals' location judgements on trials where the auditory cues provided valid versus invalid predictions, at stimulus contrasts that had been lowered to a sufficient degree so as to produce significantly slowed performance. Behavioural analysis provided no evidence to believe that mice relied on the orientation-specific predictions provided by the auditory cues, which was expected to manifest in less accurate or slowed responses on invalidly cued trials (Figure 2.4). Instead, the animals' post-cue licking patterns (examined in Results section 2.3.6) suggested that mice behaved as if the auditory cues predicted stimulus location. This manifested in vigorous licking which was significantly biased towards the spout associated with the location of the stimulus during the most recent trial featuring the same auditory cue. It is important to note that stimulus location was selected at random on a trial-wise basis, and had no global statistical relationship with the auditory cues. Thus, mice must have pooled evidence to support their false belief from the recent past, inferring that a cue which most

recently co-occurred with a right-sided stimulus must predict that a right response should be made upon the next encounter with that cue.

This finding is somewhat unexpected, and without much experimental precedent to which one could refer. The evidence that mice operate on the basis of what appears to be a false belief about predictive cue meaning suggests that sensory systems may not register those statistical properties of the environment that are not directly relevant to behavioural demands, where those demands are clearly outlined. Instead, animals may be motivated to attribute meaning to sensory inputs in accordance with perceived behavioural context.

2.5.4 MICE CANNOT MAKE SAME-DIFFERENT JUDGEMENTS UNDER HEAD-FIXED CONDITIONS

In the second pilot study detailed in this chapter, I attempted to train head-fixed and water-restricted mice to judge whether two sine-wave gratings presented on the left and right screens were of the same or different orientations, using two licking spouts. Associations between the left versus right water spout and ‘same’ versus ‘different’ stimulus pairing status were counterbalanced across the two mice. After >20 sessions of training, neither animal exhibited above-chance accuracy, which led to termination of the training protocol (Figure 2.6C-D). There are two main possible explanations for the animals’ failure to acquire the same-different task, rooted in low-level perceptual and cognitive factors respectively.

2.5.5 DID MICE STRUGGLE TO PERCEPTUALLY DISCRIMINATE ORTHOGONAL ORIENTATIONS?

One possible underpinning of failed task acquisition is that mice were incapable of visually discriminating the two orthogonally oriented gratings which would form the internal basis for judging if there was a difference. The limited visual capacities of rodents have indeed been an oft-cited reason for their relatively delayed rise in popularity in visual research, as rodents, and mice in particular, are known to have lower visual acuity than a range of species such as human and non-human primates (Merigan & Katz, 1990), ferrets (Price & Morgan, 1987) and cats (Timney et al., 1978). This inferior acuity is, to a great extent, rooted in species differences in the densities of retinal ganglion cells (RGC) and strength of foveation, which constrains spatial resolution. While primates have a RGC density of over 100,000 cells/mm² (McGregor et al., 2018; Wassle et al., 1990), rodents exhibit a substantially lower density of 2000-3000 cells/mm² (mice: Drager & Olsen, 1981; rats: Salinas-Navarro et al., 2009). Furthermore, while RGC density varies 2-3 fold between the rodent retinal periphery and the superior retinal region considered to be an analogue of the

fovea (Drager & Olsen, 1981; Salinas-Navarro et al., 2009), this difference factor is in the region of 2000-3000 in the primate retina (Wassle et al., 1990). Together, these parameters point to fundamental physiological limitations which severely limit the resolving capabilities of rodent visual systems.

However, several lines of evidence indicate that mice trained in this pilot experiment should not have struggled to discriminate between the vertically and horizontally oriented gratings they were presented. Firstly, mouse visual cortex appears to show an affinity for representing spatial frequencies in the region of 0.04 cycles per degree (cpd) - the frequency used in the pilot. Maximally preferred spatial frequencies of mouse V1 neurons range between 0.02 and 0.08 cpd, with a median single unit preference of 0.036 cpd (Niell & Stryker, 2008), as well as a high sensitivity to inputs at and around 0.06 cpd, as reflected in LFP analysis (Porciatti et al., 1999). Perhaps more importantly, mouse behaviour points to 0.04 cpd being an appropriate frequency for enabling the animals to discriminate orthogonal orientations. Experiments measuring mouse visual acuity using the optomotor (Abdeljalil et al., 2005; Prusky et al., 2004; Schmucker et al., 2005) and optokinetic (Sinex et al., 1979) reflexes place their estimates of maximal contrast sensitivity and acuity at 0.06-0.1 cpd. Crucially, a frequency of 0.04 cpd was successfully used to extract orientation discrimination judgements from head-fixed mice in a go/no-go paradigm (Pinto et al., 2013). This argues against the idea that limited perceptual discrimination was the primary bottleneck in the animals' ability to acquire the task featured in this pilot experiment.

2.5.6 WERE MICE COGNITIVELY INCAPABLE OF COMPREHENDING TASK REQUIREMENTS?

Perhaps a more plausible explanation for the animals' failure to learn the same-different discrimination task is that the task rules exceed this species' cognitive capabilities. While a number of similar tasks have been successfully used in the past (Glickfeld et al., 2013; Wang et al., 2018), there is no published evidence of mice performing the particular task used in this pilot. This leaves open the possibility that, much like Thomas Edison, I have "just found [one of] 10,000 ways that won't work" (Oxford Essential Quotations) - though in asking a mouse to make a perceptual judgement, rather than finding the right lightbulb filament.

One published task that depended on the animals' ability to discriminate orientations - with success - required mice to indicate their detection of orientation changes to on-screen grating stimuli (0.1 cpd) by releasing a lever which they otherwise held over the course of a trial (Glickfeld et al., 2013). Animals tended to acquire this task rapidly, and were able to discriminate changes as small as 8° with near 100% accuracy. In another successful task posing apparently similar demands to those of the current pilot experiment (Wang et al., 2018), mice were presented with an initial grating in one hemifield, which was followed by a second grating of identical

orientation in the opposite hemifield. After a brief interval, the first grating either underwent an orientation change or not – something that mice were asked to signal using a go/no-go response system.

The animals’ success on these two tasks, coupled with their failure on the same-different judgement task, points to a potential strategy that mice use when indicating change detection – one that would be of no help on the same-different task piloted here. In both of the described studies, mice could succeed by performing a simple temporal comparison between patterns of stimulation within the same hemifield. Thus, a vertical grating undergoing a subtle tilt – if resolved by the visual system – could evoke something akin to a brief arousal signal, allowing the animal to indicate that it can see a difference between the past and present orientations.

In contrast, the same-different task piloted in this chapter required mice to make explicit spatial comparisons between the two stimuli in front of them. Deciding whether two distinct objects exhibit the potentially arbitrary property of matching or not matching might require an entirely distinct cognitive process than the one involved in deciding if the current version of an input matches its past version (i.e. change detection). It is entirely possible that mice simply do not possess the cognitive ability to carry out this process. Of further importance, mice which received training in both aforementioned published experiments indicated their detection of change through action or inaction (i.e. releasing / continuing to hold a lever or licking / not licking a water spout). The go/no-go response system, despite its aforementioned shortcomings, might in fact represent a more intuitive output modality than the 2-AFC system for yes/no decisions (i.e. change/no change; same/not same). In this respect, it is unclear whether the animals’ failure to indicate that two stimuli were same or different was rooted primarily in their inability to comprehend these concepts or the difficulty of mapping ‘same’ and ‘different’ judgements onto arbitrarily selected spout sides. Regardless of the precise underpinning reasons, however, both the location discrimination paradigm and the same-different task examined in this chapter proved inappropriate for examining sensory predictions and errors in the mouse visual system. I now turn to the *in vitro* slice work to ask a different question about the possible implementations of predictive coding.

CHAPTER 3

3.1 INTRODUCTION

3.1.1 BROAD CONCEPTUAL GOALS AND OBJECTIVES: CAN PREDICTIVE CODING BE LOCAL?

The high prevalence of top-down projections in the brain (Markov et al., 2013), as well as the evidence that these enable higher-level regions to influence lower processing levels in a manner which suggests predictive modulation (eg. encoding of illusory contours in V1; Grosf et al., 1993; Kok et al., 2016), has given rise to the theory that a predictive mode of operation is virtually inevitable where hierarchical processing takes place (Friston, 2005; Kanai et al., 2015). The objective of the current chapter was to ask a complementary question: does a predictive system *need* to be embedded in a hierarchy?

The classical receptive field - widely regarded as a property of bottom-up analysis combined with local lateral interactions - makes a convincing case for perceptual predictions being ubiquitous even where hierarchies do not intervene. After all, the processes of lateral inhibition which enable retinal bipolar and ganglion cells to be more responsive to stimuli which border a contrasting stimulus property in the receptive field (Cook & McReynolds, 1998; Kramer & Davenport, 2015) can be viewed as indicative of a wiring scheme preparing the visual system to detect breaks from continuity – a spatial prediction error at the lowest level of analysis, if you will. The existence of a plethora of circuit plasticity mechanisms (Bliss & Lomo, 1973; Dylla et al., 2017; Fujisawa et al., 2008) also indicates that a circuit is, to some degree, equipped to learn from its stimulation history on its own. A range of plasticity mechanisms make plausible candidates for the basis of local circuit prediction, allowing predictable inputs to become less effective drivers of activity. The possibilities include short-lasting homosynaptic depression (Gover & Abrams, 2009), long-term depression (Glanzman, 2009; Kandel et al., 2013), input-selective potentiation of inhibitory synapses on target neurons (Das et al., 2011; Fischer & Carew, 1993; Ramaswami, 2014), and anti-Hebbian plasticity, whereby presynaptic inputs repeatedly preceding a target’s spike become depressed (Grant et al., 1996; Koch et al., 2013; Letzkus, Kampa & Stuart, 2006; Roberts & Leen, 2010).

Thus, the goal of two series of experiments described in this chapter was to test the possibility that neural signatures attributed to predictive coding can be observed on a local-circuit level, using recordings of local field potentials (LFP) from slices of

the mouse barrel cortex as a model *in vitro* circuit. The main question of interest here was: can a local circuit exposed to a particular pattern of stimulation ‘learn’ its statistical regularities and respond when such regularities are violated by deviant stimuli? Does this response to violation have a specific affiliation to oscillatory activity in the γ -band, as posited by predictive coding theories (Bastos et al., 2012; Kanai et al., 2015)? In this chapter, the question was addressed by driving excitatory neurons of the mouse primary somatosensory cortex (S1) to express the algal photosensitive cation channel channelrhodopsin-2 (ChR2; Boyden et al., 2005), which allowed these cell populations to be activated with spatiotemporal patterns of light intended to exhibit particular regularities and occasional violations.

An important consideration in attempting to find the physiological signatures of predictive coding in a brain slice is the stimulation protocol. What types of conditions might be reasonably expected to generate sensory predictions and responses to violations of these predictions in a slice?

3.1.2 EXPERIMENT 1: AMPLITUDE PREDICTION ERRORS

An ECoG study which inspired the stimulation protocol used in the first set of experiments in this chapter dissociated ‘local’ and ‘global’ prediction errors (El Karoui et al., 2015). Here, sequence elements that were local deviants (eg. A-A-A-B), were not necessarily global deviants, if in a stimulus set where such sequences were frequent. Such a local-global distinction provides a basis to explore the extent to which sensory regions are capable of detecting low-level versus complex violations of stimulus stream regularities. Thus, the stimulation protocol featured in the first part of this chapter referenced this prediction error dichotomy to probe *in vitro* responses of the somatosensory circuit exposed to stimulation intensity prediction errors. Slices of S1 were exposed to sine-wave photo-stimulus trains which ramped up in amplitude over the course of a 4.1 s period. On 20% of trials, the ramp featured one of four different amplitude anomalies halfway through the trial - a small or large amplitude deviation in the positive or negative direction (Figure 3.1B-C). The key manipulation was the frequency with which one selected anomaly (in this case, the large positive amplitude deviant) occurred in the stimulus set, with one set of slices being exposed to it with 5% frequency, and another set receiving it on 80% of stimulation trials. This allowed me to ask whether the slice’s response to a large amplitude deviant in its stimulus train is modulated by how frequent, and thus predictable, that deviant is in the global context of the entire stimulus stream? In other words – does a slice behave differently when a local deviant is also a global deviant, versus when it is not?

I selected stimulation amplitude as a parameter of interest because of empirical evidence suggesting that activation strength in mouse S1 neuronal populations is directly relevant to encoding information. In one study, optogenetic stimulus trains modulated at a rate of 10 Hz were delivered to layer 4 of the mouse primary so-

matosensory cortex while the animals navigated a virtual tunnel in which walls were not physically present, but rather implied by scaling LED power with proximity to the 'wall' (Sofroniew et al., 2015). Using this stimulation scheme, animals successfully navigated the curves of the illusory tunnel with an equal level of accuracy as when navigating a set-up in which walls were physically present. Thus, mice would rapidly turn left if the amplitude of optogenetic stimulation in the left hemisphere encoding proximity signals from the right wall was increased in order to signal wall approach. The finding that mice could use stimulation amplitude variation within their barrel cortex to derive the required spatial information implies that this stimulation mode can serve as a distance code, although it does not necessarily imply that it does so during natural behaviour. As such, I reasoned that an optogenetic stimulation protocol featuring amplitude ramps with deviations would constitute an *in vitro* simulation of a navigation scenario in which a signal encoding wall approach might be rapidly disrupted with a locally deviant signal which may or may not be expected given its frequency in the stimulus set. Furthermore, while the authors of the aforementioned study did not explore the effects of delivering stimulation at a range of frequencies, their successful use of 10 Hz stimulation indicates that this frequency might permit meaningful information encoding within S1. With this in mind, both experiments featured in this chapter presented optogenetic stimulation at a rate of 10 Hz.

3.1.3 EXPERIMENT 2: SPATIOTEMPORAL PREDICTION ERRORS

In the second series of experiments, prediction error concerned the site of photostimulation – a parameter choice based on the fundamental relationship between neuronal location within its designated cortical site and informational content.

One of the most universal and prominent features of sensory cortices across modalities and species is the presence of a place code, whereby neurons situated at different sites have distinct representational properties. Here, change along directions parallel to the cortical surface often co-varies with representations of space, such as retinotopic (Hubel & Wiesel, 1969; Schuett et al., 2002) and somatotopic locations (Sumser et al., 2017; although tonotopy in auditory cortex, eg. Humphries et al., 2010, is not strictly spatial, as it corresponds to the cochlea's, rather than environment's, spatial representation of sound). The barrel field of the rodent primary somatosensory cortex in particular has become a prototype for cortical functional mapping, due to its robust spatial representation of the animal's whisker pad. Here, L4 features cylindrical structures, termed barrels, populated with neurons whose responses are primarily driven by one principal whisker, with precise correspondence between a barrel's cortical location and the site of its principal whisker on the pad (Woolsey & Van der Loos, 1970). On the scale of a single barrel, neurons in the layers above and below are considered to be constituents of a barrel column – a proposed computational unit of the whisker system (Petersen, 2007).

More recently, a mapping principle has also been uncovered on the subcolumnar scale. The receptive fields of barrel neurons are known to be multi-dimensional, exhibiting sensitivity to a range of features aside from whisker identity, such as velocity and angular direction of whisker deflection (Andermann & Moore, 2006; Bale & Maravall, 2018). We now know, due to electrophysiological recordings (Andermann & Moore, 2006) and voltage-sensitive dye imaging (Tsytsarev et al., 2010) of the rat S1 superficial layers, that these neurons form a graded directional selectivity map within individual barrel columns. Here, principal whisker deflections in the direction of a particular neighbouring whisker generate neuronal activations whose centre of gravity is biased towards the barrel representing that nearer whisker.

Evidence for such a tuning gradient is weaker in the mouse, where researchers have observed what can be better described as a salt-and-pepper arrangement. Here, similarly tuned units cluster at certain sites within the barrel column, but without a graded progression of tuning properties across the column’s extent (Peron et al., 2015). For the purposes of the current experiment however, the question of whether the mouse barrel cortex replicates the rat’s organised directional tuning map is not critical. The essential property exhibited by both rats and mice is the tendency for neuronal populations at different sites within their individual barrel columns to have different informational properties. This means that exposing tissue within a single barrel to a particular spatiotemporal pattern of stimulation might, in a rough way, provide an *in vitro* equivalent to a context in which a particular whisker experiences a rapid sequence of deflections in various directions. This holds true for both species, even if the underlying tuning map is organised in a less graded fashion in the mouse.

With this equivalence in mind, the stimulation protocol used in the second series of experiments of this chapter aimed to probe individual barrel columns for signs of predictive coding by exposing slices of ChR2-expressing barrel cortex to temporal sequences of differently localised light patches. These were delivered using a Digital Micro-mirror Device (DMD) at a rate of 10 Hz (i.e. 100 ms per patch), in accordance with previous findings that stimulating mouse S1 at this frequency provided functional sensory signals during navigation (Sofroniew et al., 2015). Light patches were targetted specifically to L2/3, which receives direct input from L4 upon the arrival of feedforward thalamocortical signals (Thomson & Bannister, 2003), and which exhibits the sub-columnar angular tuning gradient first discovered in the rat (Tsytsarev et al., 2010) and subsequently found in a significantly looser format in the mouse (Peron et al., 2015). Furthermore, an inter-trial interval of 10 s was chosen based on a similar rationale to that in experiment 1 of this chapter - that is, to replicate the intervals that behaving animals would encounter when performing a task such as virtual navigation. On a majority of trials (85%), these patches proceeded in a particular order – an attempt to roughly simulate a specific sequence of whisker deflections. However, a small number of trials featured deviations from the standard sequence, such as premature stimulation at a particular location, or total reversal of

sequence direction. The goal of this experiment was simple: that is, to ask whether the slice responded to stimulation events differently under conditions when these were expected or not, given their position in the sequence.

3.1.4 THE LOCAL FIELD POTENTIAL (LFP) AS A MEASURE OF POPULATION ACTIVITY

In both series of *in vitro* experiments, microelectrodes were used to conduct extracellular recordings of the local field potential (LFP). An LFP signal can be generated by currents flowing across any excitable membrane, with contributions potentially stemming from various cellular compartments of both neurons and glial cells in the vicinity of the recording electrode (Buzsaki et al., 2012). These currents, measured at a given location in the extracellular medium, produce a specific potential with respect to a reference value obtained from outside the tissue, and are used as a measure of population activity. It is important to note that LFP signals are thought to be underpinned predominantly by currents emanating from pyramidal cells (Murakami & Okada, 2006), as the consistent alignment of their dendrites along one direction increases the distance over which electrical fields travel and gives these cells an open-field configuration (Lorente de No, 1947; Ness et al., 2018). This, in turn, allows such fields to persist for long enough to be picked up by extracellular current sensors.

The precise nature of the component signals of the LFP is not entirely clear, given the breadth of potential contributing factors. However, it is currently thought that the LFP primarily represents the summed synaptic activity of neuronal ensembles contained within the small volume of tissue surrounding a recording electrode (Katzner et al., 2009). The reasons for this are several. Firstly, the posited importance of synaptic activity is based on the notion that extracellular currents which give rise to a measurable LFP signal must have a considerable degree of temporal overlap (Buzsaki et al., 2012). From a probabilistic standpoint, the relatively slow nature of synaptic currents makes them a primary candidate (Elul, 1971). Relevant to this point is also a plethora of modelling work showing that the LFP can be relatively well-explained by assuming that the extracellular medium has low-pass filtering properties which preferentially attenuate higher-frequency events (>100 Hz) such as action potentials, while preserving the slower signals produced by synaptic events (Bedard et al., 2004; Bedard et al., 2006; Linden et al., 2011), although this is disputed (Dehghani et al., 2010; Logothetis et al., 2007). One of the first studies to crystallise the link between LFP and synaptic activity examined concurrent extracellular and intracellular patch-clamp recordings in the mouse cortex, finding that the LFP could predict subthreshold membrane potential fluctuations of individual neurons as accurately as the recorded synaptic currents (Haider et al., 2016). This strongly suggests that the LFP and synaptic activity closely correspond. As such, the LFP provides a good indicator of cortical activity.

The LFP is a popular measure of ensemble activity which can be harvested using microelectrodes directly inserted into tissue (Esghaei et al., 2017; Xing et al., 2010). Importantly, for the purposes this chapter, the LFP signal can be spectrally decomposed using methods such as Fourier (Burns et al., 2010) or wavelet (Manning et al., 2009) analysis, to estimate whether neuronal activity exhibits oscillatory components at frequencies of interest. Thus, cortical LFP recordings during a ‘predictive coding’ stimulation protocol would lend themselves well to addressing some current questions about whether parameters of the predictive coding scheme, such as sensory predictions and prediction errors, are preferentially encoded by lower- and higher-frequency α - β (8-25 Hz) and γ (30-150 Hz) oscillations respectively (Bastos et al., 2012; Kanai et al., 2015).

While activities in the two frequency bands are predominantly generated in the deep and superficial cortical layers respectively (Buffalo et al., 2011; Spaak et al., 2012), there is some reason to believe that probing cortical LFPs using a single well-positioned extracellular recording electrode should allow oscillatory signals from various layers to be sampled. Researchers have combined analysis of cross-correlations in spatially separated LFP recordings with direct cortical cuts at specific intervals to show that the LFP is volume-conducted to a distance of 400-500 μm (Dubey & Ray, 2016; Engel et al., 1990). Given the 900 μm thickness of C57BL/6 mouse cortex (Lerch et al., 2008), this spatial spread can be exploited to capture signals from across the layers using an electrode positioned in the vicinity of L4. This was the strategy in the first series of experiments, where the LFP was acquired and analysed from a single extracellular electrode placed in the centre of the cortical column. In the second series of experiments, recordings were conducted simultaneously using two electrodes placed in the deep and superficial layers, for better spatial resolution of the signal.

The ultimate goal of acquiring LFP recordings during the aforementioned stimulation protocols was to ask questions such as: Is there an observable response to violation of expectancy that is consistent with predictive coding occurring on a local-circuit level? Is activity in the γ -band particularly sensitive to prediction errors? Are prediction errors accompanied by changes in oscillatory interactions, such as those that can be captured using coherence analysis?

This last question is interesting given evidence that dynamic fluctuations in oscillatory coherence between neuronal populations enable flexible exchange of information according to task demands (Gutierrez et al., 2010; Richter et al., 2017; Roberts et al., 2013). Translaminar coherence is particularly relevant from a predictive coding perspective, given its proposed distinction between the top-down communication of predictions via low-frequency oscillations, and the feedforward transmission of prediction errors by higher-frequency rhythms (Bastos et al., 2012; Kanai et al., 2015). Indeed, localised interactions across the layers come in a variety of flavours (Senzai et al., 2019; Spaak et al., 2012), and can take the form of i) coupling oscillatory activities

stemming from different layers in one or several frequency bands (termed amplitude-amplitude coupling; Shirvalkar et al., 2010) or ii) modulation of oscillatory activity in one frequency band by the phase of an ongoing oscillation in a different, usually lower-frequency band (termed phase-amplitude coupling; Seymour et al., 2017; Spaak et al., 2012). As such, spectrally decomposed LFP signals simultaneously acquired from the deep and superficial layers of the barrel column in the second series of experiments were subjected to cross-frequency coupling analysis to explore the possibility that prediction error is associated with some form of perturbation in translaminar coherence.

3.2 METHODS

3.2.1 ANIMALS

Experiments were approved by the Committee on Animal Care and Ethical Review of the University of Oxford and licensed by the UK Home Office. Sixteen adult mice (C57BL/6, 9 males and 7 females, aged 6-8 weeks) underwent surgery to receive intracranial injections of Chr2 for the *in vitro* optogenetic work in Experiment 1; fifteen adult mice (C57BL/6, 7 males and 8 females, aged 6-8 weeks) did so using the same surgical procedures for the work in Experiment 2.

3.2.2 VIRAL INJECTIONS INTO MOUSE SOMATOSENSORY CORTEX

Mice received primary somatosensory (S1) cortical injections of adeno-associated virus serotype 5 (AAV5) expressing Chr2-mCherry under the CaMKII α promoter (AAV5-Chr2-CaMKII α -mCherry), with a concentration of 10^{12} genome copies/ ml, for the purpose of enabling optogenetic activation of excitatory cell populations during stimulation and physiological recordings from brain slices *in vitro*. The viral construct was acquired from the Vector Core Facilities at the North Carolina Gene Therapy Centre (courtesy of Dr. Karl Deisseroth).

At the start of the procedure, mice were anaesthetised with isoflurane (4% in 2 L O₂/min) and received a quick head shave before being placed in a stereotaxic frame (Stoelting Co.). The scalp was then disinfected with a chlorhexidine digluconate solution prior to administering a 40 μ l injection of marcaine under the scalp area, as well as a subcutaneous injection of a mix of buprenorphine (0.12 mg/kg) and meloxicam (0.1 mg/kg) for systemic post-operative pain relief. Throughout the procedure, anaesthesia was maintained with 1-1.5% isoflurane in 2 L O₂/min; the occasional adjustment was made with the aim of ensuring a constant breathing rate of 1 Hz, indicative of a stable anaesthetic state. Body temperature was maintained at 36° C using a homoeothermic blanket and a heat probe placed under the animal's abdomen (Harvard Apparatus).

Following a scalp incision, a dental drill was used to make a drill tip-wide craniotomy at two sites overlying S1 (site 1: 2 mm posterior and 3.2 mm lateral to Bregma; site 2: 1.5 mm posterior and 3 mm lateral to Bregma). Each craniotomy site was injected with AAV5-Chr2-CaMKII α -mCherry in 3 x 100 nl /min doses at two separate depths (-350 μ m and -750 μ m), for a total of 1.2 μ l virus per animal, using a Hamilton Microliter Syringe. The edges of the scalp were then coated with

a small amount of marcadine before suturing with braided vicryl sutures (Ethicon 5-0 W/ 16 mm tapered needle). Animals were allowed to recover for 30-60 minutes in a warmed (33° C) chamber prior to being moved to their cages and housed with the other animals that underwent surgery on the same day (usually 4-5 mice which had been littermates). Mice were kept on a 12h light / 12 dark cycle, and monitored daily for the first week post-surgery to ensure satisfactory recovery.

It should be noted that, prior to the aforementioned surgeries taking place, an earlier batch of 20 mice (C57BL/6, 10 males and 10 females, aged 6-8 weeks) had received S1 injections of another construct - AAV5-ChR2-CaMKII α -eYFP - for experimental purposes. However, brain slices from these animals were not included in the analysis reported in this chapter, owing to them exhibiting poor fluorescent marker gene expression in the target cortical region, weak photostimulation-evoked responses, and no verifiable activity in the γ frequency range. Discussions in our laboratory concluded that this was likely due to expiry of the viral construct, which had been purchased several years before commencing this work.

3.2.3 ACUTE BRAIN SLICE PREPARATION

For both series of experiments, work was carried out over a period of 6-16 weeks after animals received viral injections into S1, to allow the virus to reach high levels of expression. To start the brain extraction and slicing procedure, mice were anaesthetised in a chamber filled with isoflurane (4% in 2 L O₂/min) until they exhibited a breathing rate indicative of deep anaesthesia (0.5-1 Hz) and a total loss of the pedal reflex. At this point, they were rapidly decapitated and the brain was dissected in warmed (32-35° C) sucrose solution (150 mM sucrose, 40 mM NaCl, 25 mM NaHCO₃, 15 mM glucose, 7.4 mM MgSO₄, 3 mM KCl, 1.25 mM NaH₂PO₄ and 1 mM CaCl₂) - a method which has previously been used in our laboratory, with a good degree of success, to cut slices for experiments involving LFP recordings.

Brains were sliced in coronal sections for Experiment 1; and thalamocortical sections for Experiment 2, with the aim of increasing the preservation of barrels neighbouring each other along the row dimension, (see Agmon & Connors, 1991). Briefly, acquiring thalamocortical sections involves cutting the front of the intact brain at a 45° angle and subsequently placing it cut side down for slicing. Slices were produced at a thickness of 350 μ m, using a VT 1200S vibratome (Leica) and incubated in an interface storage chamber filled with room-temperature aCSF (containing 126 mM NaCl, 24 mM NaHCO₃, 10 mM glucose, 3.5 mM KCl, 2 mM CaCl₂, 1.25 mM NaH₂PO₄, and 1 mM MgSO₄) for at least 1 hour before commencing recordings. All solutions used during the dissection, slicing, and incubation procedures were saturated with 95% O₂/ 5% CO₂.

3.2.4 ASSESSING SLICE VIABILITY

Prior to data acquisition, it was important to establish that potential recording sites in S1 were suitable for the experimental paradigm. That is, it was critical to ascertain that a candidate recording site exhibited: i) robust fluorescence indicative of successful expression of the ChR2-CaMKII α -mCherry construct, ii) light-evoked activation, a pre-requisite for an optogenetic stimulation protocol, and iii) activity in the γ -frequency, loosely defined as 30-150 Hz. Prior to testing for the latter two properties, fluorescence was assessed in a submerged recording chamber under a microscope (Olympus BX5I WI microscope). Suitability assessment then proceeded differently in Experiments 1 and 2.

Experiment 1: After ascertaining fluorescent expression of the reporter mCherry gene, the brain slice was taken to the interface chamber where recordings were to be conducted, and an extracellular electrode was inserted into the vicinity of L4 in the putative S1 region, identified on the basis of its location relative to prominent anatomical landmarks such as the hippocampus or lateral ventricles. A Thor Labs blue LED (M470F3, 470 nm emission) was placed directly on top of the tissue of interest, and a 2 s train of sine-modulated light pulses (10 Hz) was applied to the tissue to test whether the underlying cell population produced activations that fulfilled the criteria of exhibiting light-evoked activity with significant oscillatory components in the γ -band. If a visual inspection of the LFP trace confirmed this was the case, then the site was used to conduct recordings under experimental conditions. On the other hand, if the target site failed to produce any γ -oscillatory activity in the context of clear light-evoked activations, then nearby sites in the slice would continue to be probed with the electrode until an appropriate site was located. Figure 3.1D provides a representative example of an LFP trace that exhibited γ -oscillatory activity and was thus considered suitable for experimental recordings.

Slices inevitably differ in terms of absolute ChR2 expression levels, and consequently in terms of photosensitivity. Such variability may play a particularly important role in this experiment, where stimulation amplitude is the critical manipulated variable. While using trains of light pulses to probe candidate recording sites for light-evoked and γ -band activity, I took the opportunity to gauge the tissue's sensitivity to light stimulation. To prevent photosensitive slices from exhibiting response ceiling effects during the amplitude progression of the stimuli used in the experiment, I scaled absolute LED output by a factor which yielded a satisfactory increase in response strength with increasing power. Thus, absolute levels of light being delivered to the tissue differed across slices, while leaving the rate of change in light amplitude which occurred during stimulus presentation unaltered.

Experiment 2: After confirming fluorescent expression of the reporter mCherry gene, an extracellular electrode was inserted into a putative recording site in L2/3 of S1 and a recording was obtained while a single standard light sequence (see Figure 3.2 and Section 3.2.7) was delivered onto a patch of tissue below the electrode. A Morlet wavelet transform was immediately carried out on the test recording (using Igor Pro software), to estimate power across the spectrum and test whether γ -band activity was present. If γ was confirmed, a second recording electrode was inserted into L4/5 within the same barrel column as the L2/3 electrode, and the experimental protocol began. Targeting of specific layers and barrel columns was visually guided by the pattern of fluorescence expression across the cortex, which was abundant above and below L4 and clearly outlined the lateral extent of each barrel column.

3.2.5 RECORDING LOCAL FIELD POTENTIALS (LFPs) IN SOMATOSENSORY CORTEX

Electrophysiological data were acquired at a rate of 10 kHz. Execution of the experimental protocol, data acquisition, and pre-processing of the electrophysiological recordings were all controlled by custom-written scripts run in Igor Pro software (Version 6.3.7).

Experiment 1: Extracellular recordings were conducted in an interface recording chamber comprising two compartments: i) a bottom compartment filled with distilled H₂O that was continuously bubbled with 95% O₂/ 5% CO₂ and maintained at 33-34° C) and ii) an overlying compartment perfused with aCSF, which contained the brain slice. Slice visualisation and targeting of the recording electrode were performed with the aid of a Wild Heerbrugg dissection microscope. Recordings were acquired using a borosilicate glass electrode (1-5 M Ω resistance), filled with aCSF.

Experiment 2: Extracellular recordings were conducted in a submerged recording chamber that was continuously perfused with warmed aCSF (30° C). Slice visualisation and targeting of the recording electrodes were performed with the aid of an Olympus BX5I WI microscope. Dual recordings were acquired using two borosilicate glass electrodes (1-5 M Ω resistance), filled with aCSF.

3.2.6 EXPERIMENT 1 PROTOCOL: EVOKING TEMPORAL SEQUENCE PREDICTIONS *IN VITRO*

Each stimulus consisted of a train of sine-modulated light pulses (4.1 s duration; 10 Hz) which progressively increased in amplitude (max. LED power 30 mW/mm²) as a function of time, such that the peak amplitude of the sine at time-point, Amp_t occurring at 100 ms intervals, was determined by the following formula:

$$Amp_t = Amp_{t-1} + (Amp_{t-1} \times 0.10 \times (t-1))$$

Thus the amplitude of any given light pulse was 10% higher than that of the immediately preceding pulse. There were two main types of stimulus: i) a standard ramp, in which the described amplitude progression occurred in an uninterrupted fashion, and ii) an oddball ramp in which the progression was interrupted at 1.9 s by the occurrence of a single oddball pulse, which deviated from the projected amplitude by one of four amounts (+35%, +70%, -35% or -70%). An illustration of the standard ramp and the four possible oddball ramps is provided in Figure 3.1C.

To test whether local S1 circuitry engages in a predictive mode of information processing, the goal was to record LFPs during patterns of photostimulation that were either ‘predictable’ or ‘unpredictable’ by virtue of having high or low levels of prevalence in the stimulus set. Importantly, these stimulation patterns needed to be physically identical so as to ensure that any potential effects on the neural response were not attributable to trivial stimulus differences. Given this rationale, the experiment was based on a between-group design which set out to compare LFP recordings obtained from two groups of slices in which the frequencies of two types of photostimulus were reversed.

One group of slices (22 slices; 12 slices from males and 10 slices from females) received a stimulus set in which the standard ramp occurred on 80% of trials, while each of the oddball ramps (-35%, -70%, +35%, and +70%) occurred with a 5% frequency. A second group of slices (15 slices; 8 slices from males and 7 slices from females) received a stimulus set in which one of the oddball ramps (+ 70%) was selected to occur on 80% of trials, while each the remaining stimulus types (-35%, -70%, +35%, and standard ramp) occurred with a 5% frequency. This frequency reversal across slice groups enabled the same stimulus (i.e. ramp featuring a +70% amplitude deviant) to assume the role of the ‘predictable’ and ‘unpredictable’ stimulus on the basis of having a prevalence rate of 80% versus 5%, respectively, in different slices.

The experimental protocol proceeded as follows. Slices were initially exposed to 20 repetitions of the stimulus assigned to have a high frequency, based on the rationale that this would begin habituating the neural circuitry and engage potential predictive

coding mechanisms. This was followed by 100 trials in which the five stimulus types (standard ramp and four deviant ramps) occurred with their assigned frequencies, as described above. Stimuli were presented in pseudorandom order, with two significant constraints: 1) two low-frequency stimuli may never occur on consecutive trials, and 2) a given low-frequency stimulus may not be identical to the previously occurring low-frequency stimulus. These constraints were primarily intended to prevent the potential development of local ‘pockets’ of probabilistic information, such as those which might be caused by the occurrence of the same low-frequency stimulus on multiple consecutive trials. In all cases, stimuli were separated by 16 s inter-trial intervals. Each slice was used for a maximum of one experiment, yielding a total of 120 discrete trial recordings: 100 recordings obtained during the delivery of a high-frequency ‘expected’ stimulus, and 20 recordings obtained while applying the four lower-frequency ‘unexpected’ stimulus types.

The stimulus frequency of 10 Hz was selected on the basis of the finding that optogenetic stimulation of mouse S1 with sinusoidally modulated light pulses delivered at 10 Hz effectively replaced natural neuronal activation patterns triggered from the periphery by direct contact with walls, and enabled animals to navigate a virtual corridor with a high degree of accuracy (Sofroniew et al., 2015). The relatively long inter-trial interval of 16 s was also selected to mimic the intervals that behaving animals would encounter in a task context, in order to explore whether such stimulation conditions could yield evidence for predictive coding on a local circuit level.

Electrophysiological recordings were initially amplified 10-fold by Axoclamp 2A (Molecular Devices), before undergoing further 100-fold amplification and low-pass filtering (1 kHz) by LPBF-48DG, NPI Electronic.




Figure 3.1: Experimental protocol for evoking temporal expectations *in vitro*. **A** illustrates the experimental set-up. Slices of mouse somatosensory cortex expressing ChR2 under the CaMKII α promoter were held in an interface chamber, and received direct exposure to blue light modulated in a sinusoidal fashion at 10 Hz, while a single extracellular electrode was used to acquire LFP recordings from the vicinity of L4. **B** demonstrates the standard ramp stimulus used in the experiment, featuring a sinusoidal ramp steadily ascending in amplitude at a rate of 10% per cycle. **C** zooms in on a 300 ms time-bin which contains one of the five possible amplitude variants used in the experiment – the standard amplitude pulse (grey), +35% (pink), +70% (red), -35% (light blue), and -70% (dark blue) amplitude deviants. The frequencies of the +70% amplitude deviant and standard amplitude pulse were reversed in two groups of slices. **D** provides a representative 750 ms LFP trace obtained during the trialling and assessment stage of the protocol, during which slices were probed for the presence of light-evoked γ -band responses. The higher-frequency fluctuations nested in the troughs of the stimulus-evoked 10 Hz LFP oscillation were taken as reflective of γ -band activity. Slices exhibiting such a light-driven response characteristic were used for experimental recordings.

3.2.7 EXPERIMENT 2 PROTOCOL: EVOKING SPATIOTEMPORAL PREDICTIONS *IN VITRO*

Stimuli were patterns of light projected directly on the brain slice surface using a DMD modulating the spatiotemporal output of a LED (470 nm emission; max. power 6 mW/mm²), both from Vialux. Patterns were projected through an x40 magnification lens, onto a tissue surface square with a width which did not exceed the span of a single barrel column.

A single stimulation trial consisted of a sequence of 200x200 pixel squares of light which were presented, one at a time, at 8 evenly spaced consecutive locations along a circle of 768-pixel diameter, at a 10 Hz frame rate. Thus, each trial lasted a total of 800 ms. The experiment consisted of four trial types (illustrated in Figure 3.2), collectively intended to enable an exploration of the potential effects on the LFP response of physiological stimulation which deviates from statistical norm. These four trial types were as follows:

- ***Standard sequence:*** squares were presented sequentially along a pre-determined direction around the circle (whether this was clock-wise or counter-clockwise was counterbalanced across experiments).
- ***Reverse deviant sequence:*** squares were presented sequentially along a direction opposite to the standard.
- ***Forward skip deviant sequence:*** traversal along the standard direction was interrupted by a skipping forward of two locations, with a randomly selected start of any frame between the 3rd and 5th elements of the sequence.
- ***Backward skip deviant sequence:*** a forward-skip was immediately followed by a backward skip of two locations.

Each experiment consisted of 215 trials, separated by 10 s inter-trial intervals. The stimulation protocol started with 15 consecutive presentations of the standard sequence followed by a pseudorandomly ordered stimulus set interleaving 170 standard sequences (85% frequency), 10 reverse deviant sequences (5% frequency), 10 forward-skip deviant sequences (5% frequency), and 10 backward-skip deviant sequences (5% frequency). The only constraint placed on the otherwise random order of stimuli was that deviant sequences were separated from each other by at least one standard sequence. Each experiment involved the use of one slice; no slice was used for more than one experiment.




Figure 3.2: Experimental protocol for evoking spatiotemporal expectations *in vitro*. **A** demonstrates the experimental set-up. Slices of mouse somatosensory cortex expressing ChR2 under the CaMKII α promoter were held in a submerged chamber. Two extracellular microelectrodes simultaneously recorded LFPs from L2/3 and L4/5, while a DMD was used to deliver spatiotemporal patterns of light stimulation. **B** shows the expression pattern of fluorescent marker gene mCherry in a representative slice, demonstrating preferential expression in the superficial and deep layers. White triangles highlight two septa which mark the boundaries of a single barrel column. The white square outlines the extent of L2/3 where optogenetic stimulation was targeted. A microelectrode is located in L4, for reference. **C-F** illustrate the four trial types used in the experiment. **C**: Standard sequence trial, **D**: Forward-skip deviant sequence trial, **E**: Reverse deviant sequence trial, **F**: Backward-skip deviant sequence trial.

3.3 DATA ANALYSIS

3.3.1 PRE-PROCESSING LFP TRACES

Raw traces were subjected to the following zero-phase filters: i) a notch filter centred on 50 Hz used to eliminate noise generated by mains electricity, ii) a 0.2 Hz high-pass filter, and iii) 200 Hz low-pass filter.

3.3.2 EXAMINING POWER ACROSS THE FREQUENCY SPECTRUM

To estimate the distribution of power across the frequency spectrum contained within the LFP signal, pre-processed traces were subjected to a continuous Morlet wavelet transform (with a width of 6 cycles), using Igor Pro software. The resulting Igor binary wave files were imported to MATLAB (Version 2016a), using an open-source toolbox of functions (courtesy of Jakub Bialek). The bulk of data analysis took place in MATLAB (R2016a).

Prior to examining any stimulus-evoked activity across the spectrum, estimates of power at each time-point and frequency were normalised by mean estimates at the same frequency from a pre-stimulus baseline period of 500 ms.

3.3.3 OUTLIER REMOVAL

When examining estimates of parameters of interest (eg. γ -power or α -power), outliers were removed in order to make the data more amenable to parametric analysis (as tested using a Shapiro-Wilk test). For each experiment, any trials featuring estimates that deviated from the lower and upper quantiles of the experimental estimates of that parameter by ± 2 IQRs were removed prior to computing the experimental mean. Furthermore, any individual experiments yielding estimates departing from the lower and upper quantiles of the cross-experiment dataset by ± 2 IQRs were removed prior to statistical analysis. This procedure eliminated no more than 2 data-points each time it was applied.

3.3.4 STATISTICAL TESTING FOR DIFFERENCES

Where more than 2 levels of a variable were concerned, and/or statistical interactions between variables were plausible, the test of choice was an ANOVA. Where data failed to meet the normal distribution requirement for an ANOVA and

a statistical interaction was not necessary to explore, the test of choice was the non-parametric equivalent of a one-way ANOVA – a Friedman’s analysis of variance test. Comparisons of data where ≤ 2 levels of a variable were concerned involved either t-tests or Wilcoxon rank-sum tests, depending on whether data violated the normality requirement of parametric analysis.

3.3.5 EXAMINING CROSS-TRIAL SITE COHERENCE ACROSS THE SPECTRUM - EXPERIMENT 2

Owing to the fact that Experiment 2 involved recording LFPs simultaneously from L2/3 and L4/5, I was able to estimate coherence between these two channels to examine whether stimulation contexts associated with a prediction error give rise to changes in coherence between signals obtained from the superficial and deep cortical layers. To test this, I first conducted a continuous Morlet wavelet transform for each trial, using pre-processed LFPs from the two laminar recordings (see Section 3.3.2). These channel-specific estimates of power were then used to obtain a cross-spectrum estimating signal correlation between the two channels across the trials for each individual experiment and condition, as a function of frequency f ($S_f^{1,2}$), as follows:

$$S_f^{1,2} = \sum_{t=1}^{n^{o}trials} W_f^1 \times \sum_{t=1}^{n^{o}trials} W_f^{*2}$$

Where W_f^1 is the wavelet transform of the signal obtained from channel 1 (L2/3), and W_f^{*2} is the complex conjugate of the wavelet transform of the signal obtained from channel 2 (L4/5). Coherency between the signals was then computed by normalising the squared magnitude of the cross-spectrum by the product of the magnitudes obtained from the wavelet transforms of individual channels, as follows:

$$\frac{|S_f^{1,2}|}{(\sum_{t=1}^{n^{o}trials} W_f^1 \times W_f^{*1}) \times (\sum_{t=1}^{n^{o}trials} W_f^2 \times W_f^{*2})}$$

3.3.6 EXAMINING CROSS-TRIAL PHASE LAG CONSISTENCY - EXPERIMENT 2

Coherency is influenced both by covariance in power across the spectrum between two channels, as well as consistency in the phase relations of the two signals (Roach & Mathalon, 2008; Torrence & Compo, 1998). Thus, in addition to examining coherence, I isolated a measure of cross-trial phase lag consistency between the two signals as a function of condition. Phase lag consistency was computed separately for each experiment and condition, by obtaining the magnitude of the imaginary components of the complex Morlet wavelet description at each frequency and time-point, summed over trials.

3.3.7 CROSS-FREQUENCY COUPLING ANALYSIS - EXPERIMENT 2

To examine i) whether low-frequency activity recorded from the deeper L4/5 channel significantly modulated higher-frequency activity concurrently recorded from the more superficial L2/3, and ii) whether such cross-frequency coupling was sensitive to deviant versus standard stimulation conditions, I performed exploratory cross-frequency coupling analysis, using methods originally described in Canolty et al. (2006). The analysis was carried out using the openly available cross-frequency coupling MATLAB toolbox (courtesy of Andrew Quinn, University of Oxford), and proceeded as follows.

First, pre-processed extracellular recordings from both channels (encompassing ± 400 ms of stimulus duration; 1.6 s in total) were split into 30 windows of 100 ms lengths, with 50 ms overlap. Such windowing was necessary to ensure that cross-frequency coupling over the course of a trial could be estimated with some degree of temporal resolution (by default, the analysis would yield a single time-blind estimate). Within each window, recordings were band-pass filtered to produce two sets of real-valued signals centred on frequencies that could plausibly provide oscillatory channels for the interplay between modulating low-frequency activity in L4/5 and high-frequency activity in L2/3. Recordings from L2/3 were subjected to parallel filters with centre frequencies ranging from 50 to 90 Hz (a broadly defined γ -band), separated by 20 Hz steps and comprising 20 Hz bandwidths; filters applied to recordings from L4/5 were of centre frequencies 8 and 10 Hz (covering the traditional α -frequency band), in 2 Hz steps and 3 Hz bandwidths. This was done separately for each individual trial of every experiment.

Next, Hilbert transforms were applied to each band-pass filtered signal to obtain a collection of complex-valued analytic representations of the signals at each of the channels and frequency bands of interest. In the next step, phase information was extracted from the set of L4/5 complex representations to create a set of analytic phase time series (θ_α) focused on each of the two aforementioned α -range frequency bands where possible modulating activity might be taking place. Conversely, the amplitudes of the set of L2/3 complex representations were used to create an equivalent analytic time series of magnitude values ($\text{Amp}_\gamma(t)$) at each of the γ -frequency bandwidths of interest. Using these two types of time series, the procedure constructs a collection of composite complex-valued signals combining each of the two α -phase time series with each of the γ -amplitude time series (yielding a total of 8 composite phase-amplitude series, Z), as follows:

$$Z(t) = \text{Amp}_\gamma(t) \times e^{i \times \theta_\alpha(t)}$$

The length of this vector provides a measure of coupling between the phase of ongoing low-frequency activity in L4/5 and the amplitude of higher-frequency activity in L2/3, while the average of this length across trials captures the degree of consistency in this coupling. One caveat of this value is its sensitivity to baseline differences in γ -amplitude – i.e. given experiments with identical cross-frequency dependencies, those with higher γ -amplitudes would yield larger average vector lengths.

A normalised modulation index (MI_{norm}) was thus derived to ensure that cross-frequency coupling estimates were independent of amplitude. This was produced by normalising each experiment’s average vector length by a metric of distance between it and the average vector length taken from 250 surrogate composite phase-amplitude time series. In each surrogate series, the phase time-series was circularly shifted by an integer randomly selected from a uniform distribution ranging from 0 to the number of samples in the time-series, thus preserving the amplitude information while disrupting any potential phase-amplitude relations. The distance metric was computed as follows:

$$D = (\text{mean}(V_{TRUE}) - \text{mean}(V_{SURR})) / \text{std}(V_{SURR})$$

where V_{TRUE} is the vector of the true composite phase-amplitude time-series and V_{SURR} is the vector of the surrogate time-series. Next, the distance metric was used to scale the average vector length and obtain MI_{norm} , as follows:

$$MI_{norm} = (D \times e^{i \times \text{angle}(\text{mean}(V_{TRUE}))})$$

3.3.8 STATISTICAL ANALYSIS OF COHERENCE AND CROSS-FREQUENCY COUPLING - EXPERIMENT 2

I approached the question of whether deviant and standard stimulation produce differences in i) coherence or ii) cross-frequency coupling strength between the superficial and deep layers without an a priori hypothesis of the spectral locus of the expected effect. The time when this effect was expected was also not constrained, as a change in power coherence, phase relations, or phase-amplitude coupling could occur immediately in response to a deviation, or manifest in a response to future stimulation. This open-ended approach to examining a time- and frequency-indexed measure comes with a high false positive risk, due to the extremely high number of frequency-time points available for testing. To address this issue, I used a non-parametric statistical test recommended by Maris, Schoffelen & Fries (2007), which takes into account prior ignorance of the spectral or temporal locus of an effect.

For coherence and phase lag consistency data, the test proceeds by transforming the frequency- and time-indexed coherence values (see Section 3.3.5) for the two conditions of interest (Dev vs Ctrl) into a matrix of frequency- and time-indexed z-scores ($Z_{f,t}$; Brillinger, 1981), separately for each experiment, as follows:

$$\frac{(\tanh^{-1}(\text{Coh}_{f,t}^{1,2(Dev)}) - 1/d.f.^{(Dev)} - 2) - (\tanh^{-1}(\text{Coh}_{f,t}^{1,2(Ctrl)}) - 1/d.f.^{(Ctrl)} - 2)}{\sqrt{(1/d.f.^{(Dev)} - 2) + (1/d.f.^{(Ctrl)} - 2)}}$$

For cross-frequency coupling data, z-scores are obtained by subtracting and consequently dividing each experiment's mean MI_{norm} across deviant trials by the cross-experiment mean and standard deviation of MI_{norm} indices obtained in the control condition, separately for each corresponding time-window and modulating/carrier frequency combination. For each experiment, this yields a three-dimensional matrix of z-scores (indexed by modulating frequency, carrier frequency, and time-window) which serves as the foundation for statistical analysis.

Analysis proceeds by identifying each experiment's maximal absolute value z-score across the range of specified dimensions (i.e. time-point and frequency for coherence and phase lag consistency data; time-window, modulating frequency, and carrier frequency for cross-frequency coupling data) and obtaining an average maximal z-score across experiments. This z-score is then assigned a statistical significance level by a randomisation test which compares this z-score to a reference distribution of average maximal z-scores constructed on the basis of 10,000 random permutations of data point-label associations. A resulting Monte-Carlo p-value (p_{MC}) of < 0.05 would indicate that the deviant and control conditions differ significantly somewhere in the time-frequency plane (for coherence and phase-lag consistency data) or in some time-window (for cross-frequency coupling data); however it does not specify the spectral and temporal location of the effect. Nonetheless, it can be used to guide plot interpretation.

3.4 RESULTS - EXPERIMENT 1

3.4.1 RECAP AND OBJECTIVES

In this series of experiments, extracellular recordings of LFPs were obtained from a total of 37 slices of mouse S1 *in vitro*, (n=4 from two different mice, n=3 from five different mice, n=2 from five different mice, and the remainder of single slices taken from four different mice), while cortical tissue was subjected to optogenetic stimulus ramps. Figure 3.3 provides an illustrative example of the stimulus-evoked LFP and its spectral decomposition. The main objective was to test whether stimulus-evoked responses in barrel cortex are modifiable by the statistical regularities of the stimulus stream. Specifically, the question was whether responses to a local amplitude deviant (referred to as LPD, or Large Positive Deviant throughout this section) in the stimulus ramp differ depending on whether that deviant is statistically unpredictable or not, owing to low (5%; n=22 slices) or high (80% n=15 slices) frequency of occurrence in the stimulus set. Throughout this results section, these experimental conditions will be referred to as Condition 1 (LPD rare) and Condition 2 (LPD common).

3.4.2 STATISTICAL ANALYSIS: CONTROLLING FOR CROSS-SLICE AND CROSS-TRIAL VARIABILITY

Individual slices differ in terms of photosensitivity and baseline magnitude of γ -activity. Furthermore, differences can also emerge over time within a slice, over the course of an experiment. To facilitate statistical comparisons of conditions in which a large amplitude deviant constituted the ‘expected’ versus ‘unexpected’ stimulus, each trial featuring the deviant was matched with the immediately preceding trial in the same experiment featuring a non-deviant pulse, to serve as the control. I reasoned that a paired-samples approach using trial pairs with maximal temporal proximity would enable a comparison between the large and standard stimulation amplitudes which maximally controlled for parameters that might change over the course of an experiment (eg. channelrhodopsin saturation, habituation), and enable optimal comparison of conditions under which a local amplitude deviant was a global deviant by virtue of low frequency in the stimulus set versus when it was common, and thus expected.




Figure 3.3: Illustrative example of stimulus-evoked LFP in mouse S1 *in vitro* and its spectral decomposition, during standard ramp stimulus delivery. **A** illustrates the standard sinusoidal ramp stimulus used to modulate the delivery of blue light onto the tissue surface. **B** demonstrates a representative LFP recording obtained from L4/5, aligned to the stimulus. **C** demonstrates log-power as a function of time and frequency in the LFP trace, showing clear spectral peaks at the stimulus frequency (10 Hz) and its harmonic frequency (20 Hz), as well as distinct modulations of power at the higher frequencies of the γ -spectrum (50-150z). Here, power is modulated by the 10 Hz rhythm, reflecting the presence of nested γ -oscillations.

3.4.3 POWER IN THE γ -BAND IS RESPONSIVE TO STIMULUS AMPLITUDE

The γ -band can be loosely defined as encompassing frequencies within the bounds of 30-150 Hz. However, given the likelihood that oscillatory processes taking place at lower and higher γ -frequencies represent physiologically and functionally distinct processes (Ray & Maunsell, 2011; Ray et al., 2008; Van der Meer & Redish, 2009), the γ -band was split into a series of narrower 15 Hz-wide frequency bands, ranging from 30 to 150 Hz, for the purposes of analysis.

To ascertain whether γ -power was sensitive to stimulus amplitude, time-frequency descriptions of the LFP (see Sections 3.3.1 and 3.3.2) were used to estimate maximal power at each of the defined γ -bandwidths during the 100 ms following the onset of each large amplitude deviant. This estimate was normalised by the equivalent estimate obtained from the preceding 100 ms (illustrated in Figure 3.4A), to control for slice differences in absolute levels of γ -power and capture an estimate of stimulus-driven change. The same normalised maximal γ -power estimate was obtained for the same 100 ms period marking the standard pulse taken from each matched control trial.

Normalised estimates of power within the discretised γ -bandwidths in response to deviant and standard stimuli were aggregated across the two experimental conditions (Figure 3.4B), and submitted to a repeated measures two-way ANOVA with Frequency Band (60-75 Hz, 75-90 Hz, 90-110 Hz, 110-130 Hz, and 130-150 Hz) and Stimulation Amplitude (Large Deviant and Standard Pulse) as repeated measures variables, to test whether power in any of the frequency bands was sensitive to stimulus amplitude.

The ANOVA revealed a significant main effect of Stimulus Amplitude ($F(1, 35) = 83.28, p < 0.001, \eta_p^2 = 0.704$), owing to stimulus-evoked power across the γ -bands being, on average, significantly higher during the large amplitude deviant ($M = 0.49, SD = 0.25$) compared to the standard pulse ($M = 0.087, SD = 0.009$). A significant main effect of Frequency Band ($F(1,7) = 4.22, p = 0.007, \eta_p^2 = 0.11$) indicated that power was, on average, significantly higher in some frequency bands; however this main effect was not pursued further, due to the vast number of post-hoc comparisons this would involve combined with the relatively small informational gain of doing so. Finally, the ANOVA found no significant Frequency Band X Stimulus Amplitude interaction ($F(1, 35) = 1.15, p = 0.33, \eta_p^2 = 0.032$). Thus, γ -power in the LFP was found to be significantly higher when slices were exposed to larger stimulus amplitudes, delivered through increased LED power. The next, and crucial question, is: is this amplitude effect modulated by whether that level of stimulation was expected, by virtue of being common (Condition 2) versus rare (Condition 1)?

3.4.4 POWER IN THE γ -BAND IS NOT RESPONSIVE TO STIMULUS PROBABILITY

To test the critical hypothesis that stimulus-driven γ -band power is sensitive to stimulus probability, I extracted the differences between normalised maximal γ -power estimates obtained during the large amplitude deviants and their matched control trials. These γ -power contrast measures (Figure 3.4C), averaged across trials within each experiment, were submitted to a repeated-measures two-way ANOVA with Frequency Band (60-75 Hz, 75-90 Hz, 90-110 Hz, 110-130 Hz, and 130-150 Hz) as the repeated-measures variable and Large Deviant Probability (Low (5%) and High (80%)) as the between-subjects' variable. The ANOVA reported no significant main effects of Large Deviant Probability ($F(1, 34) = 2.21$, $p = 0.147$, $\eta_p^2 = 0.061$), and Frequency Band ($F(1, 7) = 0.844$, $p = 0.551$, $\eta_p^2 = 0.024$), as well as no significant Large Deviant Probability X Frequency Band interaction ($F(1, 7) = 1.092$, $p = 0.369$, $\eta_p^2 = 0.031$). These results contradict what would be expected if the probability of the amplitude deviant significantly modulated stimulus-driven γ -band power.

Nonetheless, a close-up visual inspection of the stimulus-evoked γ -activity taking place in the slice (Figure 3.4D) suggests that splitting the range of γ -frequencies into five equal bandwidths may not be entirely representative of the oscillatory activity taking place. Instead, the time-frequency transform suggests that the γ -activity is more accurately described as occupying two discrete bandwidths - a lower one spanning 30 to 50 Hz, and a higher one extending from 75 to 150 Hz. Given these observations, a further analysis was undertaken on the basis of a more data-driven outline of the γ bands within which an effect of stimulus predictability might be found. As done in Section 3.4.3 (above), time-frequency descriptions of the LFP were used to estimate maximal power in the two observed γ -bandwidths (30-50Hz and 75-150Hz) during the 100 ms following the onset of each large amplitude deviant. This estimate was normalised by the equivalent estimate obtained from the preceding 100 ms (illustrated in Figure 3.4A). The same normalised maximal γ -power estimate was obtained for the same 100 ms period marking the standard pulse taken from each matched control trial.

The resulting normalised estimates of γ -power in response to deviant and standard stimulation, aggregated across the two experimental conditions, are shown Figure 3.4E. As done in the analysis of five equally spaced γ -bands featured in this section, the differences between normalised maximal γ -power estimates obtained during the large amplitude deviants and their matched control trials were extracted for both the 30-50 Hz and 75-150 Hz bands, as a function of whether a deviant was rare or common in the stimulus set (Figure 3.4F). The obtained γ -power contrast measures were averaged across trials within each experiment and submitted to two independent-samples t-tests (one per frequency band), comparing power contrast es-

timates between conditions of low versus high deviant probability. The tests revealed no significant difference in γ -power contrast between condition 1 (deviant rare; $M=0.32$, $SD=0.33$) and condition 2 (deviant common; $M=0.28$, $SD=0.22$) in the 30-50 Hz band ($t=-0.45$, $p=0.655$). The difference between condition 1 (deviant rare; $M=0.48$, $SD=0.39$) and condition 2 (deviant common; $M=0.35$, $SD=0.18$) was also not significant in the 75-150 Hz band ($t=-1.37$, $p=0.178$). Thus, analysis of γ -power using a more data-driven delineation of γ -bands supported the results of the previous analysis suggesting that the probability of deviant stimulation did not have a significant impact on stimulus-driven γ -power.

3.4.5 SLOPE OF TEMPORAL CHANGE IN γ -BAND POWER IS NOT RESPONSIVE TO STIMULUS PROBABILITY

Next, I examined whether γ -power estimates exhibit a particular relationship with time on the time-scale of an experiment, and whether the slope of this relationship differs between conditions where the large amplitude deviant is rare (Condition 1) or common (Condition 2). A set of two linear regressions was performed per experiment – separately for large amplitude deviant trials and their matched control trials – to estimate the slope of the relationship between raw maximum γ -power obtained during the critical 100 ms period and trial number (Figure 3.5A). Slope estimates (Figure 3.5B), segregated by trial type and condition, were then submitted to a repeated-measures two-way ANOVA with Stimulus Amplitude (Large Amplitude Deviant and Standard Ramp) as the repeated-measures variable, and Large Deviant Probability (Low (5%) and High (80%)) as the between-subjects' variable. The ANOVA found no significant main effects of Stimulus Amplitude ($F(1, 34)=0.157$, $p=0.69$, $\eta_p^2=0.005$), and Large Deviant Probability ($F(1, 34)=1.21$, $p=0.28$, $\eta_p^2=0.035$), as well as no significant Stimulus Amplitude x Large Deviant Probability interaction ($F(1, 34)=2.69$, $p=0.11$, $\eta_p^2=0.075$).

Next, I examined whether γ -power changes at all over the course of an experiment. The derived slope estimates were aggregated across slices and trial types, and submitted to a one-sample Wilcoxon rank sum test, which reported that the median slope of temporal change in stimulus-evoked γ -power ($Med=-2.59$, $IQR=292.61$) did not significantly differ from zero ($Rank\ sum=345$, $p=0.85$).




Figure 3.4: Effects of deviant amplitude stimulation, and its predictability, on power in the γ -band, *in vitro* (n=37). **A** illustrates the procedure used to obtain normalised estimates of γ -power in response to large amplitude deviants and standard pulses, shown in **B**. Max. γ -power was estimated within a defined frequency band during 100 ms marking the critical pulse stimulus, and subsequently normalised by an equivalent power estimate from the immediately preceding 100 ms. Normalisation was used to gauge the degree of change taking place from one stimulus pulse to the next; the same method was used in analyses of γ -band PF. **B** shows normalised max. power, for a range of γ -frequencies, during photostimulation with large amplitude deviants (red) and standard pulses (blue), averaged across slices. **C** shows mean subtractive contrasts between the normalised max. γ -power estimates shown in **B**, for Conditions 1 (deviant rare; green; n=22) and 2 (deviant common; grey; n=15). **D** shows a 100 ms snapshot of a time-frequency transform data of a representative trial, highlighting γ -activity in two discrete bandwidths: 30-50 Hz and 75-150 Hz, used to conduct a further data-driven analysis of γ -power, shown in **E-F**. **E** shows normalised max. power, for each of the two empirically defined γ -bands, during photostimulation with large amplitude deviants (red) and standard pulses (blue), averaged across slices. **F** shows mean subtractive contrasts between the normalised max. γ -power estimates shown in **D**, for Conditions 1 (deviant rare; green; n=22) and 2 (deviant common; grey; n=15).




Figure 3.5: Effects of deviant amplitude stimulation, and its predictability, on temporal changes in γ -power, *in vitro* (n=37). **A** shows trial-wise estimates of stimulus-evoked γ -power (within the broadly defined γ -band of 30-150 Hz), from a single experiment. Note that power is shown in arbitrary units resulting from normalising raw time-wise power estimates by values obtained from a 500 ms baseline period preceding stimulus delivery. Changes in γ -power over the time-scale of an experiment were quantified by taking slope estimates from linear regressions (plotted as dashed lines) fitted to the relationship between single-trial γ -power estimates and trial number, separately for trials featuring standard pulse stimuli and large amplitude deviants (dots). **B** shows a violin plot of the resulting γ -power slope estimates, obtained as described in A. Black boxes delineate the lower and upper quartiles of the data distribution; the median is marked in white.

3.4.6 PEAK FREQUENCY IN THE γ -BAND IS NOT RESPONSIVE TO DEVIANT PROBABILITY

I turned to examining whether peak frequency - defined as the frequency exhibiting the highest power - within the broadly defined γ -band was sensitive to the predictability of large amplitude deviants. Time-frequency descriptions of the LFP (see Sections 3.3.1 and 3.3.2) were used to extract peak frequency estimates within the range of 30-150 Hz, as a function of time (Figure 3.6A).

Peak frequency estimates were averaged within each 100 ms period featuring a large amplitude deviant, and normalised by the equivalent estimate obtained in immediately preceding 100 ms of the same trial. Normalised peak frequency estimates were obtained using the same procedure for matched control trials. The resulting estimates (Figure 3.6B) were submitted to a repeated-measures two-way ANOVA with Stimulus Amplitude (Large Amplitude Deviant and Standard Pulse) as the repeated-measures variable, and Large Deviant Probability (Low (5%) and High (80%)) as the between-subjects variable. The ANOVA reported no significant main effects of Stimulus Amplitude ($F(1, 34) = 1.65$, $p = 0.21$, $\eta_p^2 = 0.046$), and Large Deviant Probability ($F(1, 34) = 1.27$, $p = 0.27$, $\eta_p^2 = 0.036$), as well as no significant Stimulus Amplitude X Large Deviant Probability interaction ($F(1, 34) = 0.064$, $p = 0.80$, $\eta_p^2 = 0.002$). This indicates that deviant predictability was not a significant modulator of peak frequency in the γ -band during stimulus-evoked activation.

3.4.7 SLOPE OF TEMPORAL CHANGE IN γ -BAND PEAK FREQUENCY IS NOT RESPONSIVE TO DEVIANT PROBABILITY

Next, I examined whether peak frequency exhibits a particular relationship with time on the time-scale of an experiment, and whether the slope of this relationship differs between conditions where the large amplitude deviant is rare (Condition 1) or common (Condition 2).

A set of two linear regressions was performed per experiment - separately for large amplitude deviant trials and their matched control trials - to estimate the slope of the relationship between raw peak frequency and trial number (Figure 3.6C). Slope estimates (Figure 3.6D), segregated by trial type and condition, were submitted to a repeated-measures two-way ANOVA with Stimulus Amplitude (Large Amplitude Deviant and Standard Ramp) as the repeated-measures variable, and Large Deviant Probability (Low (5%) and High (80%)) as the between-subjects' variable. The ANOVA found no significant main effects of Stimulus Amplitude ($F(1, 34) = 0.47$, $p = 0.49$, $\eta_p^2 = 0.14$), and Large Deviant Probability ($F(1, 34) = 1.54$, $p = 0.22$, $\eta_p^2 = 0.043$), as well as no significant Stimulus Amplitude X Large Deviant Probability interaction ($F(1, 34) = 0.026$, $p = 0.87$, $\eta_p^2 = 0.001$).




Figure 3.6: Effects of deviant amplitude stimulation, and its predictability, on peak frequency (PF) in the γ -band, *in vitro* (n=37). **A** shows raw γ -band PF as a function of time, averaged across trials featuring standard amplitude pulses (blue) and large amplitude deviants (red). Results are shown separately for slices in Condition 1 (top; deviant rare; n=22) and Condition 2 (bottom; deviant common; n=15). **B** shows γ -band PF estimates averaged over the 100 ms duration of each standard pulse and large amplitude deviant, separately for slices in Condition 1 (left; deviant rare) and Condition 2 (right; deviant common). Estimates were normalised by values obtained from the previous 100 ms within the same trial. **C** shows trial-wise stimulus-evoked γ -band PF estimates from a single experiment. Changes in γ -band PF over the time-scale of an experiment were quantified by taking slope estimates from linear regressions (plotted as dashed lines) fitted to the relationship between single-trial estimates of peak frequency and trial number, separately for trials featuring standard pulse stimuli and large amplitude deviants (dots). **D** shows a violin plot of the resulting γ -band PF slope estimates, obtained as described in C. Black boxes delineate the lower and upper quartiles of the data distribution; the median is marked in white.

3.4.8 POWER IN THE α -BAND IS NOT RESPONSIVE TO DEVIANT PROBABILITY

Next, I examined whether power in the α -band immediately preceding, or during, large amplitude deviants and control standard pulses discriminated between conditions where such deviations were rare (Condition 1) versus common (Condition 2). Time-frequency descriptions of the LFP (see Sections 3.3.1 and 3.3.2) were used to estimate maximal power in the α -band (5-15 Hz) during the 100 ms period which i) preceded and ii) followed, the onset of each large amplitude deviant. Equivalent α -power estimates were obtained for matched control trials. It should be noted that assessing stimulus-evoked activity in the α -band is highly problematic given that stimulus trains were delivered at 10 Hz, making it likely that any observed 10 Hz rhythmic activity *in vitro* is, to a large degree, directly entrained by photostimulation. However, the 10 Hz bandwidth used to operationalise α in the following analyses creates room for examining rhythmic activity at other, nearby frequencies, and raises the possibility that a meaningful response discriminating different stimulus conditions might be observed.

First, I tested whether α -power discriminated between stimulation conditions which were to be followed by two different pulses, which were either expected or unexpected. Maximal α -power estimates from the 100 ms period preceding the critical pulse on deviant and matched control trials (Figure 3.7A) were submitted to a repeated-measures two-way ANOVA with Stimulus Amplitude (Large Amplitude Deviant and Standard Pulse) as the repeated-measures variable, and Large Deviant Probability (Low (5%) and High (80%)) as the between-subjects' variable. The ANOVA found no significant main effects of Stimulus Amplitude ($F(1, 34) = 3.44$, $p = 0.072$, $\eta_p^2 = 0.092$), and Large Deviant Probability ($F(1, 34) = 0.084$, $p = 0.77$, $\eta_p^2 = 0.002$), as well as no significant Stimulus Amplitude X Large Deviant Probability interaction ($F(1, 34) = 0.425$, $p = 0.52$, $\eta_p^2 = 0.012$). This indicated that α -power did not distinguish between stimulation conditions associated with different predictable futures, as would be reflected in a significant main effect of Large Deviant Probability.

Next, estimates of maximal α -power during the deviants and equivalent time periods in the matched control trials (Figure 3.7B) were submitted to a two-way ANOVA of the same format. The ANOVA found no significant main effects of Stimulus Amplitude ($F(1, 34) = 2.79$, $p = 0.10$, $\eta_p^2 = 0.076$), and Large Deviant Probability ($F(1, 34) = 0.058$, $p = 0.811$, $\eta_p^2 = 0.002$), as well as no significant Stimulus Amplitude X Large Deviant Probability interaction ($F(1, 34) = 0.29$, $p = 0.59$, $\eta_p^2 = 0.008$). Thus, stimulus-driven α -power was found to not be sensitive to stimulus amplitude or probability.

3.4.9 SLOPE OF TEMPORAL CHANGE IN α -BAND POWER IS NOT RESPONSIVE TO DEVIANT PROBABILITY

Next, I explored whether α -power exhibits a particular relationship with time on the time-scale of an experiment, and whether the slope of this relationship differs between conditions where the large amplitude deviant is rare (Condition 1) or common (Condition 2). A set of two linear regressions was performed per experiment – separately for large amplitude deviant trials and their matched control trials – to estimate the slope of the relationship between raw maximal stimulus-evoked α -power captured during the 100 ms pulse duration, and trial number (Figure 3.7C). Slope estimates (Figure 3.7D), segregated by trial type and condition, were submitted to a repeated-measures two-way ANOVA with Stimulus Amplitude (Large Amplitude Deviant and Standard Ramp) as the repeated-measures variable, and Large Deviant Probability (Low (5%) and High (80%)) as the between-subjects' variable. The ANOVA found no significant main effect of Stimulus Amplitude ($F(1, 34) = 0.33$, $p = 0.57$, $\eta_p^2 = 0.01$) and a moderately significant main effect of Large Deviant Probability ($F(1, 34) = 4.18$, $p = 0.049$, $\eta_p^2 = 0.110$), reflecting the observation that slopes of temporal α -power change were somewhat higher in Condition 1 (large deviant probability low; $M = 314.41$, $SD = 795.16$) than in Condition 2 (large deviant probability high; $M = -241.87$, $SD = 795.21$). The Stimulus Amplitude x Large Deviant Probability was not significant ($F(1, 34) = 0.045$, $p = 0.833$, $\eta_p^2 = 0.001$).

Given prior evidence of an association between ramp-up in low-frequency power and the build-up of sensory expectation (van Ede et al., 2010; van Pelt et al., 2016), it is relevant to test whether there is a significant net change in α -power over experimental time. The derived slope estimates were aggregated across slices and stimulus types, and submitted to a one-sample Wilcoxon rank sum test, which found that the median α -power slope ($Med = -12.24$, $IQR = 695.06$) did not significantly differ from zero ($Rank\ sum = 331$, $p = 0.97$).




Figure 3.7: Effects of deviant amplitude stimulation, and its predictability, on power in the α -band, *in vitro* (n=37). Note that power is shown in arbitrary units resulting from normalising raw time-wise power estimates by values obtained from a 500 ms baseline period preceding stimulus delivery. **A-B** demonstrate the mean max. α -power estimates obtained during the 100 ms preceding (A) and following (B) the onset of large amplitude deviants (red) and standard pulses taken from matched control trials (blue). Results are shown separately for slices in Condition 1 (deviant rare; left; n=22) and Condition 2 (deviant common; right; n=15). **C** shows trial-wise stimulus-evoked α -power estimates from a single experiment. Changes in α -power over the time-scale of an experiment were quantified by taking slope estimates from linear regressions (plotted as dashed lines) fitted to the relationship between raw single-trial estimates of α -power and trial number, separately for trials featuring standard pulse stimuli and large amplitude deviants (dots). **D** shows a violin plot of the resulting α -power slope estimates, obtained as described in C. Black boxes delineate the lower and upper quartiles of the data distribution; the median is marked in white.

3.5 RESULTS - EXPERIMENT 2

3.5.1 RECAP AND OBJECTIVES

In this series of experiments, extracellular recordings were obtained simultaneously from the superficial (L2/3) and mid-to-deep (L4/5) layers of 24 slices of mouse somatosensory cortex *in vitro*, (includes n=5 from one mouse, n=3 from one mouse, and n=2 from one mouse, with the remainder being single slices from 5 different mice.) during exposure to spatiotemporal patterns of optogenetic stimulation. Each stimulation event was a sequence of light squares, appearing one at a time for 100 ms at different sites of cortical tissue contained within a single cortical barrel column. While the standard, and most frequent (85%), stimulus involved a sequence of eight squares traversing the radius of a circle, the stimulus stream was peppered with a minority of trials in which the spatiotemporal pattern deviated from the statistical norm, in the form of skipping sequence elements or a total reversal of sequence direction. The main objective of the experiment was to probe whether electrophysiological activity, *in vitro*, could distinguish predictable forms of stimulation from those delivered outside their normal temporal context.

3.5.2 ANALYTICAL APPROACH TO EXPLORING EFFECTS OF DEVIANT STIMULATION

Figure 3.8 shows the raw magnitude squared of Morlet time-frequency decompositions of LFP recordings obtained from randomly selected single trials of several experiments, in L2/3 and L4/5 respectively (prior to baselining). Given that most clear γ -band modulation appeared to take place <50 Hz, γ was restricted to a 30-45 Hz bandwidth for purposes of analysis. The three parameters of interest in this results section are: i) power in the γ -band (30-45 Hz), ii) peak frequency in the γ -band (γ -band PF), and iii) power in the α -band (8-12 Hz). A critical question, given the three deviant trial types used in the experiment, was whether any of these parameters responded significantly to premature stimulation of a patch of tissue, as exemplified by the forward-skip. Furthermore, it was interesting to explore the slice's response during photostimulus delivery to a patch of tissue following its failure to receive the temporally appropriate stimulation as a result of a forward-skip. As such, the back-skip, in which stimulation was delivered to one of the previously 'skipped' locations in a back-skip deviant trial type, would prove informative.

When examining the neural response to forward-skips and back-skips, estimates of maximal γ -band and α -band power, as well as γ -band PF, were averaged within the 100 ms period featuring each appropriate deviant event, yielding 20 individual trial estimates for forward-skips and 10 for back-skips. Given likely variation in photosensitivity and stimulus-evoked LFP power between slices and different patches

of tissue within the same slice, I reasoned that the most robust way to examine sensitivity to deviation involved setting up a series of appropriate comparisons. As such, each deviant trial was paired with a standard sequence trial which immediately preceded it in the same experiment (to serve as a matched control trial) and three different measures were extracted, as follows:

1. *Deviant stimulus measure:* The relevant estimate obtained from the deviant trial was normalised by the corresponding estimate obtained from the 100 ms period marking photostimulation of the same patch of tissue in the matched control trial (trial-1). This measure aimed to compare signals between conditions where stimulus delivery to the same patch of tissue occurred within the appropriate temporal context (i.e. standard sequence trial) or not (i.e. during a forward-skip or back-skip). Note that a tissue patch receiving stimulation during a forward-skip at time t receives stimulation at $t+200$ ms in a standard sequence trial. During a back-skip event, the patch of tissue receiving stimulation at time t receives stimulation at $t-100$ ms in a standard sequence trial. As such, the deviant stimulus measure compares signals which differ, not only in terms of deviance, but also time-point during trial. Controlling for this time-shift requires a specific control, outlined below.

2. *Time control measure:* For every trial featuring a forward-skip or back-skip, the two immediately preceding standard sequence trials were selected (trial-1 and trial-2). The relevant estimate was obtained in trial-1 from a 100 ms period following time-point t , which corresponded to deviant event onset in the deviant trial to which the two standard sequence trials were yoked. Note that in this case, the tissue was not responding to deviant stimulation, but to non-deviant stimulation taking place at the same time within the trial as the nearest future deviant event. This estimate was normalised by the equivalent estimate obtained from trial-2 during a 100 ms period starting with $t+200$ ms for forward-skips and $t-100$ ms for back-skips. The result is a comparison of signals separated by the same amount of time as those contributing to the deviant stimulus measure; however, without incorporating any deviance.

3. *Deviance control measure:* The relevant estimate obtained from a 100 ms period preceding the forward-skip deviation ($t-100$ ms) was normalised by the estimate obtained from the same time period in the matched control trial. This measure aimed to capture the contrast between standard and deviant trials in a manner which could not have possibly been influenced by departure from statistical expectancy.

If deviant photostimulus delivery indeed influenced some facet of the LFP signal, one would expect the deviant stimulus measure to significantly differ both from the time and deviant control measures, which control both for the timing and stimulus deviance differences that are inherent in the signals being compared for the deviant stimulus measure. These three measures are illustrated in Figure 3.9A and Figure 3.10A, for forward-skip and back-skip trials respectively.




Figure 3.8: Raw power across the frequency spectrum for extracellular recordings acquired from L2/3 (left) and L4/5 (right) on the same single randomly selected trials from 2 sample experiments. Raw extracellular recordings are shown above their corresponding time-frequency transforms.

3.5.3 γ -BAND POWER: NO SIGNIFICANT EFFECTS OF DEVIATION FROM THE SPATIOTEMPORAL NORM

To examine whether γ -power in either L2/3 or L4/5 was sensitive to deviation from the spatiotemporal norm, as triggered by photostimulation outside its normal sequence context, time-frequency descriptions of the LFP (see Sections 3.3.1 and 3.3.2) were used to estimate power in the γ -band across time (30-45 Hz) and obtain deviant stimulus, and stimulus and time control measures of maximum γ -band power, as per description in Section 3.5.2, separately for forward-skips (illustrated in Figure 3.9A) and back-skips (illustrated in Figure 3.10A). To test whether the critical deviant stimulus measure of γ -power differed from the control measures for either laminar site and deviant type, measures were averaged across trials for each slice and submitted to a series of Friedman's analysis of variance tests for difference of score distributions, separately for recordings from L2/3 and L4/5.

A Friedman's test looking at forward-skips (Figure 3.9B) revealed no significant difference between the distributions of deviant stimulus measures and the two control measures in L2/3 ($\chi^2 = 5.25$, $p = 0.072$) and L4/5 ($\chi^2 = 0.54$, $p = 0.76$). In the back-skip condition (Figure 3.10B), the test similarly flagged up no significant differences both in L2/3 ($\chi^2 = 4.20$, $p = 0.12$) and L4/5 ($\chi^2 = 3.25$, $p = 0.19$). Thus, there is no evidence that power in the γ -band responded to violations of statistical norms.

3.5.4 PEAK FREQUENCY IN THE γ -BAND: NO EFFECT OF DEVIATION FROM THE SPATIOTEMPORAL NORM

To examine whether γ -band PF in either L2/3 or L4/5 was sensitive to deviation from the spatiotemporal norm, time-frequency descriptions of the LFP (see Sections 3.3.1 and 3.3.2) were used to estimate peak frequency within the defined γ -band (30-45 Hz), and obtain deviant stimulus, and stimulus and time control measures of γ -band PF, as per description in Section 3.5.2. These were submitted to a series of Friedman's tests, separately for L2/3 and L4/5. In all cases, follow-up post-hoc tests involved Bonferroni-corrected significance thresholds of $p = 0.025$, to account for the necessary double comparison of the deviant stimulus measure with two control measures.

In the forward-skip condition, (Figure 3.9B), the test showed a significant difference between the three distributions of γ -band PF estimates in L2/3 ($\chi^2 = 9.45$, $p = 0.009$). This difference was explored with a series of post-hoc Dunn's pair-wise tests, which showed, taking threshold correction into account, that the distribution of deviant stimulus measures failed to significantly differ both from those of the deviant ($t = -0.36$, $p = 0.39$) and time control measures ($t = 0.91$, $p = 0.033$). The Friedman's test also flagged a significant difference between the three distributions in L4/5 ($\chi^2 =$

11.45, $p=0.003$), which was explored with a further pair of post-hoc Dunn’s pair-wise tests. These revealed that the distribution of deviant stimulus measures differed significantly from that of the deviance control measure ($t=1.09$, $p=0.011$), but not of the time control measure ($t=-0.27$, $p=0.52$).

In the back-skip condition (Figure 3.10B), the test indicated a significant difference between the three distributions in L2/3 ($\chi^2=6.17$, $p=0.046$). Dunn’s pair-wise post-hoc tests indicated that this was due to the distribution of deviant stimulus estimates differing significantly from that of the deviance ($t=0.92$, $p=0.025$) but not time control ($t=0.833$, $p=0.041$) measures. In L4/5, Friedman’s test also reported a significant difference between the three γ -band PF measures ($\chi^2=9.50$, $p=0.009$). Dunn’s pair-wise tests pointed to the distribution of deviant stimulus measures being significantly different from that of the deviance ($t=1.25$, $p=0.002$), but not time control measures ($t=1.84$, $p=0.06$). Thus, testing failed to provide consistent evidence that peak frequency responded to violations of statistical norms.

3.5.5 ACTIVITY IN THE α -BAND: NO EFFECT OF DEVIATION FROM THE SPATIOTEMPORAL NORM

Finally, a series of Friedman’s tests looked for differences between the distributions of maximal α -power estimates captured by the deviant stimulus measure, and the two control measures, separately for L2/3 and L4/5. In the forward-skip condition (Figure 3.9B), the test revealed that the three distributions did not significantly differ in L2/3 ($\chi^2=4.00$, $p=0.135$), but did so in L4/5 ($\chi^2=12.20$, $p=0.002$). This was explored with post-hoc Dunn’s tests, which revealed a significant difference between the distribution of deviant stimulus measures and that of the time ($t=1.30$, $p=0.004$), but not deviance control measures ($t=-0.10$, $p=0.82$).

In the back-skip condition (Figure 3.10B), Friedman’s test found no significant difference between the distributions of deviant stimulus measures and the two controls in L2/3 ($\chi^2=4.67$, $p=0.097$), however did report a significant difference in L4/5 ($\chi^2=9.80$, $p=0.007$). Post-hoc Dunn’s tests suggested that this was underpinned by the distribution of deviant stimulus measures being significantly different from that of time ($t=1.10$, $p=0.014$), but not the deviance control measures ($t=-0.20$, $p=0.655$). As such, there is no robust evidence that α -power responds to deviant stimulus delivery at either laminar site.




Figure 3.9: Effects of deviant forward-skips in a stimulation sequence on γ -power and peak frequency, and α -power, during L2/3 stimulation, *in vitro* (n=24). **A** illustrates the three measures used to test if the LFP distinguishes tissue stimulation during a forward-skip from non-deviant stimulation. Top: Deviant stimulus measure. An estimate of interest (eg. γ -power) obtained during stimulation of a tissue patch targeted by a forward-skip (i.e. a deviant event) is normalised by the equivalent estimate obtained when stimulating the same tissue patch in the preceding standard sequence trial (control). Middle: Time control issue. The estimate obtained from the control trial during 100 ms yoked to the forward-skip in the matched deviant trial is normalised by the estimate obtained from the previous control trial during 100 ms, delayed by 200 ms. Thus, two stimulus events are compared with a 200 ms time-shift as featured in the deviant stimulus measure, without differing degrees of event deviance. Bottom: Deviance control measure. The estimate from 100 ms preceding the forward-skip in a deviant trial is normalised by the estimate from the same period in the control trial. This measure captures the difference between the deviant and control trials before any deviation occurs. To support local-circuit predictive coding, the deviant stimulus measure is expected to differ both from the time and deviance control measures. **B** shows violin plots of deviant stimulus, time control, and deviance control estimates of γ -power (top), α -power (middle), and γ -band PF (bottom), separately for L2/3 and L4/5. Black boxes delineate the lower and upper quartiles of the data distribution; the median is marked in white.




Figure 3.10: Effects of deviant back-skips in a stimulation sequence on γ -band power and peak frequency, and α -band power, during L2/3 stimulation, *in vitro* (n=24). **A** illustrates the three measures used to test if the LFP distinguishes tissue stimulation during a back-skip from non-deviant stimulation. Top: Deviant stimulus measure. An estimate of interest (eg. γ -power) obtained during stimulation of a tissue patch targeted by a back-skip (i.e. a deviant event) is normalised by the equivalent estimate obtained when stimulating the same tissue patch in the preceding standard sequence trial (control). Middle: Time control issue. The estimate obtained from the control trial during 100 ms yoked to the back-skip in the matched deviant trial is normalised by the estimate obtained from the previous control trial during 100 ms, leading by 100 ms. Thus, two stimulus events are compared with a 100 ms time-shift as featured in the deviant stimulus measure, without differing degrees of event deviance. Bottom: Deviance control measure. The estimate from 100 ms preceding the forward-skip in a deviant trial is normalised by the estimate from the same period in the control trial. This measure captures the difference between the deviant and control trials before any deviation occurs. To support local-circuit predictive coding, the deviant stimulus measure is expected to differ both from the time and deviance control measures. **B** shows violin plots of deviant stimulus, time control, and deviance control estimates of γ -power (top), α -power (middle), and γ -band PF (bottom), separately for L2/3 and L4/5. Black boxes delineate the lower and upper quartiles of the data distribution; the median is marked in white.

3.5.6 EFFECTS OF SEQUENCE REVERSAL ON γ -BAND POWER AND PEAK FREQUENCY, AND α -BAND POWER

In one deviant trial type (5% frequency), the order in which a sequence of tissue patches received stimulation was reversed with respect to the standard sequence (Figure 3.11A). This trial type can shed some light on whether exposure to a particular stimulation sequence can recruit plasticity mechanisms (eg. LTP; Bliss & Lomo 1973) which enable the next locus of stimulation to be 'anticipated' based on predictable recent stimulus delivery to the neighbouring tissue patch. Such an 'anticipatory' effect can perhaps be visualised as one in which stimulation of a tissue patch shortly followed by activation of another results in the neurons currently undergoing stimulation exerting a long- or short-term depressing effect on neurons about to receive stimulation, via mechanisms which may involve GABAergic intervention (Das et al., 2011) or metabotropic glutamate receptors (Cho et al., 2000). In this case, running a reverse stimulation sequence might offer a read-out of the normal interactions between neighbouring clusters of neurons which may amount to some form of circuit-level predictive coding.

Comparing LFP signals in response to stimulation sequences running in opposing directions is not straightforward. The electrode may consistently obtain a stronger signal during stimulation of one patch of tissue versus another, in which case aligning LFPs recorded during reverse sequences may generate a false impression of a difference. A visual inspection of the average γ - and α -power trace recorded during standard and reverse sequences (Figure 3.11B) reveals that the precise locus of stimulation appeared not to be the main determinant of spectral power, but rather sequence onset and trial time. That is, both γ - and α -activity exhibited an initial onset-driven spike in power (unlikely to be spectrally specific, as LFPs are often characterised by a prominent onset-driven broadband component; Ray & Maunsell, 2010), followed by an exponential decay in power over the time-course of a trial, with minor variations which can likely be attributed to stimulation of distinct tissue patches. In order to cancel out some of the temporal decay effects and make signals extracted from the two sequence types more comparable, an exponential curve was fit to the time-course of γ/α -power, separately for each reverse sequence trial and its immediately preceding control trial, with the aim of capturing trial-specific estimates of temporal decay (Figure 3.11D). Each time-course was then scaled by its own time-reversed exponential curve to yield a time-course with an effectively reduced decay (Figure 3.11E).

A series of cluster-based permutation tests (one for each parameter and laminar site; see fieldtriptoolbox.org/tutorial/cluster_permutation_timelock/ for test details) were provided each experiment's average reverse and control time-series, to explore whether the two trial types significantly differed at any temporal locus. These tests revealed that in neither L2/3 (γ : $p_{MC}=0.90$; α : $p_{MC}=0.87$) nor L4/5 (γ : $p_{MC}=0.92$; α : $p_{MC}=0.75$), did the two variables significantly differ between trial types.




Figure 3.11: Effects of stimulation sequence reversal on γ -band and α -band power, during L2/3 stimulation, *in vitro* (n=24). **A** illustrates the standard sequences (blue) and reverse sequences (red) used for this analysis. **B** show trial time-lines, averaged across experiments, comparing γ -power between standard and reverse sequences, separately for extracellular recordings obtained in L2/3 (left) and L4/5 (right). **C** shows the same for α -power. **D** shows a sample γ -power trace from a single trial recorded in L2/3 (blue), starting with the time-point associated with the average peak in onset-driven power at 150 ms post stimulus onset. An exponential curve was fit to the data, to capture the degree of temporal decay in power (pink). **E** shows the result of scaling the γ -power trace shown in D by the time-reversed exponential curve. The intended effect is to cancel out the decay while preserving the smaller-scale fluctuations.

3.5.7 TEMPORAL EFFECTS OF STIMULUS DEVIATION ON γ -BAND AND α -BAND ACTIVITY

Next, I tested whether the LFP response to deviant stimulation distinguished itself on the temporal scale of an experiment. It is possible, for example, that stimulus-evoked responses in the γ -band weaken over time to photo-stimuli delivered within their appropriate temporal context, while doing so at a slower rate for unexpected stimulation events. This might be reflective of a stimulus history recruiting plasticity mechanisms which contribute to response suppression (Garrido et al., 2009; Randlett et al., 2019), without providing the same readout of circuit habituation during stimulation associated with prediction error.

To examine this question, I focused on the forward-skip deviant type. Trial-wise estimates of maximal γ -power, α -power, and γ -band PF were extracted from a 100 ms period spanning each of the following events: i) each forward-skip event (deviant event), ii) each non-deviant event where the same patch of tissue received stimulation in the immediately preceding standard sequence (stimulus control event), and iii) each event immediately preceding the forward-skip deviation in every sequence featuring a forward-skip (deviance control event). These three events are illustrated in Figure 3.12A.

Linear regressions were fitted to the relationship between data points taken from each of the three events and trial number within each experiment (illustrated in Figure 3.12B for a sample L2/3 recording), to compute the slope of change over time, separately for α -power, γ -power, and γ -band PF, as well as each laminar recording site. Slope estimates (Figure 3.12C-E) were submitted to a series of Friedman's tests, which reported no significant differences between the distributions of γ -power ($\chi^2 = 0.20$, $p = 0.90$), γ -PF ($\chi^2 = 0.80$, $p = 0.67$) and α -power ($\chi^2 = 2.36$, $p = 0.31$) slope estimates for the three event types in L2/3. The same was found for γ -power ($\chi^2 = 2.91$, $p = 0.24$), γ -PF ($\chi^2 = 5.09$, $p = 0.078$) and α -power ($\chi^2 = 0.18$, $p = 0.91$) slopes in L4/5. Thus, the three measures did not significantly differ in terms of their relationship with experimental time, in either recorded layer.

Finally, I set out to verify if any of the three variables exhibited significant change over experimental time. Slope estimates were collapsed across the three measures for each parameter of interest, and submitted to a series of one-sample Wilcoxon rank sum tests for significant differences from a median of zero, separately for each laminar site. In L2/3, the tests revealed that γ -power ($Med = 4.0 \times 10^{-4}$, $IQR = 0.05$, $p = 0.81$), α -power ($Med = 3.0 \times 10^{-4}$, $IQR = 0.10$, $p = 0.38$), and γ -PF ($Med = 4.13 \times 10^{-2}$, $IQR = 0.16$, $p = 0.81$) slopes all showed no significant differences from zero. Similarly, all slopes failed to significantly differ from zero in L4/5 (γ -power: $Med = 3.8 \times 10^{-3}$, $IQR = 0.10$, $p = 0.79$; α -power: $Med = 1.4 \times 10^{-3}$, $IQR = 0.10$, $p = 0.59$; γ -PF: $Med = 0.12$, $IQR = 0.26$, $p = 0.33$).




Figure 3.12: Temporal effects of deviant and control spatiotemporal stimulation on γ -power, peak frequency, and α -power, during L2/3 stimulation, *in vitro* (n=24). **A** illustrates the three event types during which the parameters of interest were estimated for purposes of linear regression. For all trials featuring a forward-skip, the 100 ms period of light delivery to the patch of tissue targeted by the skip was used to estimate the deviant signal (red). The 100 ms period during which that same patch of tissue was stimulated in a standard trial (i.e. non-deviant stimulation) was used to extract a stimulus control signal (blue). Finally, the 100 ms period preceding any deviation in each trial featuring a forward-skip was used to estimate a signal controlling for deviance (green). **B** shows trial-wise stimulus-evoked γ -power estimates from a single experiment, for each of the three event types described in **A**. Changes in γ -power over the time-scale of an experiment were quantified by taking slope estimates from linear regressions (plotted as dashed lines) fitted to the relationship between single-trial estimates of γ -power and trial number, separately for the three event types (shown in red, blue and green). **C-E** show violin plots of the resulting slope estimates, found separately for changes in γ -power (**C**), peak frequency (**D**), and α -power (**E**), as described in **B**. Black boxes delineate the lower and upper quartiles of the data distribution; the median is marked in white.

3.5.8 STIMULUS DEVIATION DOES NOT SIGNIFICANTLY INFLUENCE COHERENCE BETWEEN L2/3 AND L4/5

Next, I tested whether deviant stimulation of a tissue patch during forward-skip events influences the level of: i) coherence, and ii) phase lag consistency between L2/3 and L4/5 signals captured by the two extracellular electrodes, at any frequency. Prior to computing coherence or phase lag consistency across trials, all deviant trials featuring a forward-skip and the immediately preceding standard sequence trials serving as their controls were aligned to capture ± 500 ms relative to forward-skip onset in the pair's deviant trial. Then, cross-trial coherence was computed separately for deviant and control trials, as described in Section 3.3.5 and illustrated in Figure 3.13.

Figure 3.14 demonstrates resulting estimates of magnitude squared same-frequency coherence between L2/3 and L4/5 recordings across the frequency spectrum. Statistical testing (see Section 3.3.8) for the significance of z-scores associated with comparing deviant and control trials across time and frequencies (bottom of Figure 3.14A) revealed that the mean difference between deviant and control trials did not exceed chance rate at any time-point or spectral locus ($z = 3.97$, $p_{MC} = 0.52$). Thus, there is no evidence that coherence between superficial and deep layers was significantly influenced by deviant stimulation.

Next, the consistency of phase lags (PLC) between same-frequency activities of L2/3 and L4/5 was computed for each frequency and experiment (see Section 3.3.5; Figure 3.14B). Testing for differences (see Section 3.3.8) between the PLC z-scores associated with comparing deviant and matched control trials (bottom of Figure 3.14B) revealed no significant difference between deviant and control trials at any time or frequency ($z = 5.63$, $p_{MC} = 0.21$).




Figure 3.13: Computing coherence between L2/3 and L4/5 signals, *in vitro*, in a sample experiment. Computing coherence for a single experiment involves obtaining a squared product of the cross-trial sum of wavelet transforms of LFP recordings obtained from L2/3 and that of the complex conjugates of wavelet transforms of LFP recordings from L4/5, normalised by the product of summed magnitudes obtained from the two layers, as detailed in Section 3.3.5. The left panel features LFP recordings obtained from L2/3 during three sample trials taken from a representative experiment, alongside plots of the magnitude of the corresponding complex wavelet signal description, as a function of time and frequency (denoted as $W_{L2/3}$). The right panel features LFP recordings concurrently obtained from L4/5 during the same three trials, alongside plots of the magnitude of the complex conjugate of the corresponding complex wavelet signal description, as a function of time and frequency (denoted as $W^*_{L4/5}$). The sums of these parameters computed across all trials within an experiment ($\sum_{t=1}^{n^{o} trials} W_{L2/3}$ for L2/3 and $\sum_{t=1}^{n^{o} trials} W^*_{L4/5}$ for L4/5) are used to compute the degree of coherence between L2/3 and L4/5 signals, shown in the lowest panel.




Figure 3.14: Coherence and phase-locking consistency between L2/3 and L4/5 signals, during L2/3 stimulation, *in vitro*, on trials featuring forward-skip deviations versus standard sequences (n=24; 10 trial comparisons per slice). **A** shows mean magnitude-squared coherence across the spectrum and time, for trials featuring forward-skip deviations (top) and control trials featuring standard sequences (middle). Differences in coherence values between the two trials types were transformed into z-scores (bottom); and submitted to a permutation test for significant differences. No loci were found to significantly differ between deviant and standard trial types. **B** shows mean phase-locking consistency (PLC) across the spectrum and time, for trials featuring forward-skip deviations (top) and control trials featuring standard sequences (middle). Differences in coherence values between the two trials types were transformed into z-scores (bottom), and submitted to a cluster-based permutation test for significant differences. No loci were found to significantly differ between deviant and standard trial types. In both cases, z-scoring was performed as described by Brillinger (1981).

3.5.9 STIMULUS DEVIATION DOES NOT INFLUENCE STRENGTH OF L2/3 - L4/5 PHASE-AMPLITUDE COUPLING

Exploratory cross-frequency coupling analysis was performed to examine the possibility that any cross-frequency coupling between signals in L2/3 and L4/5 (Seymour et al., 2017; Spaak et al., 2012) may be significantly disrupted during an unpredictable stimulus event. Each trial featuring a forward-skip and its matched control trial was split into 30 x 100 ms time windows with 50% overlap and normalised phase-amplitude modulation indices were obtained for each window (Canolty et al., 2006) and each combination of candidate modulating (7-9 Hz, 9-11 Hz) and carrier (40-60 Hz, 60-80 Hz, 80-100 Hz, 100-120 Hz) frequency bands, as described in Section 3.3.7.

Figure 3.15 demonstrates the resulting modulation indices obtained across the range of described carrier frequency bands and time windows; data from each panel are centred on a different modulating frequency band. Statistical testing (see Section 3.3.8) revealed that the mean difference between deviant and control trials did not exceed chance rate for any time window or modulating / carrier frequency combination ($z = 2.72$, $p_{MC} = 0.99$). Thus, there is no evidence that cross-frequency coupling between superficial and deep layers, *in vitro*, was significantly influenced by deviant stimulation.




Figure 3.15: Modulation index estimates for cross-frequency coupling between LFP signals in L2/3 and L4/5, during L2/3 stimulation, *in vitro* (n=24; 10 trials per slice and condition). **A** demonstrates modulation indices across the range of tested carrier frequency bands (bandwidth 20 Hz) and time windows (100 ms duration; 50 ms overlap) for a modulating frequency band of 7-9 Hz, segregated by trials featuring forward-skip deviations (top) and control trials featuring standard sequences (bottom). **B** demonstrates equivalent values for a modulating frequency band of 9-11 Hz.

3.6 DISCUSSION

The two sets of experiments examined in this chapter explored whether the purported physiological hallmarks of predictive coding can be observed on the level of a local neuronal circuit isolated from the interactive context of a large-scale hierarchical network. To this end, LFP recordings were obtained from slices of mouse S1 serving as an *in vitro* model circuit, while the tissue was being exposed to temporal (Exp 1) or spatiotemporal (Exp 2) patterns of optogenetic stimulation exhibiting particular regularities. Occasionally, stimuli delivered to the tissue violated these regularities, for instance by deviating from the expected stimulus amplitude (Exp 1) or the location of photostimulation (Exp 2).

I obtained no robust evidence that slices of somatosensory cortex generated predictions based on their optogenetic stimulation histories. This manifested in a range of observations, such as the finding that spectral power and peak frequency in the γ -band were not sensitive to stimulus (un)predictability. Furthermore, I found no evidence to believe that power in the α -band corresponded with the build-up of sensory predictions either acutely (i.e. in the immediate lead-up to a stimulation event) or on a large temporal scale (i.e. by scaling up with growing stimulation history).

3.6.1 ACTIVITY IN THE γ -BAND DOES NOT CARRY PREDICTION ERRORS IN A LOCAL S1 CIRCUIT

In the first experiments detailed in this chapter, slices of S1 tissue were exposed to trains of sine-modulated light pulses which ramped up in amplitude over the course of the stimulus train. I explored the potential link between γ -activity and sensory prediction error by examining power and peak frequency in the γ -band during delivery of a large amplitude deviant pulse which was associated with distinct degrees of predictability in two sets of slices, where it occurred with either high (80%) or low (5%) frequency. Examining stimulus-evoked γ -band power in response to amplitude deviants and non-deviant pulses positioned at the same time-points in control trials showed that γ -power was sensitive to the amplitude of optogenetic stimulation in both conditions, increasing to a greater extent during the amplitude deviant compared to control pulse (Figure 3.4B). However, γ -power differences between trials delivering amplitude deviants and non-deviant pulses were not sensitive to deviant predictability (Figure 3.4C). Thus, while γ -power was stronger in response to tissue stimulation at higher LED power, this was not affected by whether higher power at that particular time-point in a trial was a statistical commonality.

In a further exploration of the link between γ -power and stimulus predictability, a series of linear regressions determined the rate at which stimulus-evoked γ -power captured during amplitude deviants and control pulses changed over the time-course of an experiment, as a function of deviant predictability (Figure 3.4D-E). This was underpinned by the speculation that γ -power might exhibit faster decline over a growing number of trials in response to a statistically common versus rare stimulus (as the common stimulus is better predicted, and thus evokes a weaker prediction error). This may be underpinned by the recruitment of experience-dependent plasticity (Cho et al., 2000; Das et al., 2011; Bozorgmehr et al., 2013; Randlett et al., 2018). However, testing provided no evidence that stimulus-evoked γ -power changed over a slice's stimulation history in a manner which was sensitive to stimulus amplitude or its predictability.

The second series of experiments detailed in this chapter provides corroborating evidence. In these experiments, extracellular recordings were obtained from L2/3 and L4/5 in S1 slices exposed to photo-stimuli targeting distinct patches of tissue in a particular sequence, at the level of a single barrel column, where a minority of trials (15%) featured deviant stimulus delivery. Comparing γ -power measured during forward- and back-skip deviant events with two control measures revealed no significant differences that would point to γ -power being sensitive to deviation from the norm, at either laminar site. Furthermore, exploring slopes of change in γ -power across the time-course of an entire experiment revealed no significant differences between slopes computed using estimates taken from the deviant event and those taken from control events.

In fact, one-sample Wilcoxon rank sum tests revealed that γ -power did not significantly change over the course of the experimental protocol, in either series of experiments. This finding does not sit well with the supposed association between γ -oscillatory activity and prediction error, which predicts that γ -power should decrease with an increasingly long stimulation history, as experience accrues and sensory prediction error declines. The observed absence of significant change in γ -power over time in this series of experiments does not necessarily undermine the link between γ -oscillations and prediction errors, but rather indicates a high likelihood that an isolated cortical circuit is not sensitive to sensory prediction errors (and consequently, does not engage in predictive coding).

3.6.2 POTENTIAL DISSOCIATION BETWEEN PEAK FREQUENCY AND POWER IN THE γ -BAND

Analysis of peak frequency in the γ -band (γ -band PF) yielded somewhat different results from those obtained with γ -band power. Turning to the first series of experiments, peak frequency estimates taken during amplitude deviants and non-deviant pulses in control trials were compared as a function of amplitude deviant predictability. Testing revealed that, neither stimulus amplitude nor deviant pulse predictability had any significant main effect on γ -band PF (Figure 3.6A-B). Moreover, no relationship was found between stimulus amplitude or predictability and the slope with which peak frequency in the stimulus-evoked γ -band response changed over the time-course of an experiment (Figure 3.6D).

These findings are, to a large degree, corroborated by the second series of experiments. Here, I examined whether deviant photostimulation events, in the form of forward-skip or back-skip departures from the standard sequence, yielded different peak frequency estimates than two control estimates (Figures 3.9B and 3.10B). In all cases, distributions of γ -band PF across experiments failed to differ from both control measures. Where the distribution of deviant γ -band PF measures differed significantly from only one of the two controls, it did so compared to the deviance control measure which captured PF during a time-window which took place 100 ms before the onset of any deviant stimulation. In contrast, the distribution of PF estimates taken during deviant stimulation consistently failed to differ from that of the time control PF measures, for both laminar sites and deviant event types. This strongly suggests that any differences between the deviant and deviance control measures can be explained by the temporal difference between the time-points when the PF estimates were obtained. A further important result was that slopes of γ -band PF change over the time-course of an experiment also did not differ significantly between deviant and control stimulation events (Figure 3.12B). Thus, the notion that peak frequency in the γ -band might be meaningfully sensitive to deviations from the norm, at least on the level of a local barrel circuit, is lacking evidence.

Interestingly, however, the first series of experiments showed that, unlike γ -power, which was significantly greater during tissue stimulation at higher LED power, peak frequency in the γ -band was not necessarily determined by optogenetic drive. This finding is consistent with a body of evidence suggesting that γ -band power and peak frequency are not entirely related variables (Jia et al., 2013; Lima et al., 2010). Traditionally, the two were treated as correlated facets of γ -oscillatory activity, with similar modulation by network drive (Whittington et al., 1995; Traub et al., 1996), and similar contributions to behavioural performance on perceptual and cognitive tasks (Edden et al., 2009; Scheeringa et al., 2016; Womelsdorf et al., 2007). In V1, the two variables respond similarly to manipulations such as stimulus contrast increases (Ray & Maunsell, 2010). However, other studies failed to replicate this relationship,

finding instead that higher-contrast stimuli evoke higher γ -power without significant shifts in the spectral locus of the peak (Bartolo et al., 2011; Henrie & Shapley, 2005). One experiment further dissecting such differences in monkey V1 found that some manipulations – such as increases in visual stimulus size – modulated γ -power and peak frequency of the LFP in opposing directions (Jia et al., 2013).

This dissociation has led to the hypothesis that the two facets of γ -oscillatory activity may, to some extent, involve dissociable mechanisms (Jia et al., 2013; Lima et al., 2010). Whereas models indicate that the genesis of γ -rhythms is largely a function of a compromise between fast inhibition and slower feedback mediated by principal cell-interneuronal connections (Brunel & Wang, 2003), peak frequency may be partially determined by functional involvement from feedback or lateral connections, which may change the time constant of communications (Jia et al., 2013; Kang et al., 2011). In light of this conjecture, it is interesting that the peak frequency of γ -activity in sensory (Lima et al., 2010) and motor (Heinrichs-Graham et al., 2018) cortices shifts under conditions where sensory stimuli and motor plans, respectively, compete for selection – conditions which are likely to recruit lateral inhibition or top-down control (Coultrip et al., 1992; Desimone & Duncan, 1995).

The possible association between peak frequency and top-down or lateral signalling raises some interesting questions about its role in the supposed γ -frequency channels encoding prediction errors under conditions where sensory predictions provide differing degrees of constraint. For instance, deviations from low-precision sensory predictions, generated by less consistent stimulation histories (Friston, 2005; Stefanics et al., 2018) might be encoded by γ -oscillations centred at different frequencies compared to errors marking departures from deeply ingrained predictions, as the former are likely preceded by weaker predictive constraints on bottom-up information flow. Whether this is the case remains to be tested. However, for the time being it suffices to say that my experiments provided no evidence that peak frequency in the γ -band was significantly influenced by stimulus predictability in a local barrel cortical circuit.

3.6.3 α -BAND POWER BEARS NO SIGNIFICANT RELATIONSHIP TO SENSORY PREDICTIONS OR ERRORS IN A LOCAL S1 CIRCUIT

One interesting question that can be pursued, based on the experiments detailed in this chapter, is whether α -activity could be plausibly linked to the representation of sensory expectations (or responses to deviations from them; although this is not a widely held view) in a local circuit.

Turning to the first series of experiments, I found no evidence that stimulus-evoked α -power obtained during the 100 ms duration of a deviant differed between

conditions with different degrees of deviant predictability (nor was α -power significantly sensitive to stimulation magnitude). The results of the second experiment series support this finding, as the distributions of α -power estimates obtained during forward-skip and back-skip deviant events did not differ significantly from the distributions of two control estimates, in either L2/3 or L4/5 (Figures 3.9B and 3.10B). Thus, α -power was not found to be sensitive to whether stimulation is (un)expected. This is perhaps more peripheral compared to the finding that α -power does not significantly distinguish between conditions likely to contain different sensory predictions.

This was explored using data from the first experiment series, where I tested for differences between maximal α -power estimates obtained during the 100 ms period preceding each amplitude deviant and non-amplitude pulse, under conditions of low and high deviant probability (Figure 3.7A-B). While LED power during this period was identical for both pulse types and conditions, the most likely upcoming stimulus was not. It is plausible that a circuit generating some form of sensory expectation of the immediate future which reflects in α -rhythmic activity would produce levels of α -power which discriminate between events most likely followed by large amplitude deviants versus non-deviant pulses. An analysis of variance revealed, however, that this was not the case.

It is plausible however, that α -power might represent the build-up of sensory expectation *in vitro* without necessarily discriminating between different contents of expectation (content-specific signals might be encoded by a different variable, such as phase-locked spike patterns). To explore this possibility, I tested whether stimulus-evoked α -power exhibited a gradual ramping up in power over the course of an experiment, as slices were exposed to more stimulus exemplars which ought have contributed stronger sensory predictions. Both in the first (Figure 3.7C-D) and second (Figure 3.12E) series of experiments, I failed to find that α -power significantly increased over experimental time. This does not fully line up with human studies which have found links between the build-up of sensory expectation and lateralisation of power in the α/β -bands in the sensory cortices (Bauer et al., 2014; Sedley et al., 2016, van Ede et al., 2010; van Pelt et al., 2016). A plausible conclusion is that the local barrel cortex circuit is unlikely to have the capacity to generate representations of statistical properties in its input stream and produce sensory predictions. While researchers have certainly observed α -band oscillatory activity on the scale of a single cortical column, where it frequently entrains higher-frequency oscillations and spikes across the layers (Sakata & Harris, 2009; Spaak et al., 2012), these signals – the purported carriers of predictive information – perhaps need to be seeded by predictive inputs from beyond the local circuit itself.

3.6.4 TECHNICAL CAVEATS

It is critical to address the technical limitations of attempting to find the physiological signatures of predictive coding *in vitro*. One important question is whether the parameters of optogenetic stimulation - i.e. duration of stimulus events, inter-trial intervals and number of stimulus exposures - is compatible with what we know about stimulation conditions which result in synaptic plasticity such as long-term potentiation (LTP; Bliss & Lomo, 1973). The use of long inter-trial intervals and a 10 Hz stimulus delivery rate in both experiments described in this chapter was based specifically on evidence of the functionality of such stimulation in supplanting natural stimulus-evoked activity in mouse S1 during behaviour (Sofroniew et al., 2015), rather than any consideration of stimulus parameters conducive to plasticity. Nonetheless, evidence suggests that our stimulation protocol - involving a 10 Hz stimulus frequency, and inter-trial interval of 16 s (Experiment 1) and 10 s (Experiment 2) - is plausible for plasticity induction *in vitro*.

Traditionally, long-term potentiation - a model for studying the synaptic underpinnings of learning and memory - has been studied in the hippocampus, where it is induced using prolonged high-frequency (typically, 100 Hz) stimulation of afferent fibres (Bliss & Collingridge, 1993; Bliss & Lomo, 1973). However, a variety of other induction methods have been found effective, including theta-burst stimulation, used both in the hippocampus (Nguyen & Kandel, 1997) and sensory cortex (Sarihi et al., 2008). This involves periodically presenting short 100 Hz bursts of 4-5 stimuli at the lower frequency of 5 Hz, thereby emulating potential modulatory effects on spiking behaviour of the hippocampal theta rhythm (Buzsaki, 2002). In a further departure from the classic high-frequency stimulation induction protocol, researchers have found that repeated low-frequency stimulation at 1 Hz is also capable of inducing a longer-latency, but lasting, long-term potentiation of synapses in the CA1 area of the hippocampus, in a manner which does not depend on NMDA receptors (Gonzalez et al., 2014; Lante et al., 2006), unlike the vast proportion of observed long-term potentiation events (Park et al., 2014). This suggests that a high frequency of stimulus presentation may not be necessary for inducing plasticity in a slice.

Furthermore, an experiment comparing the effectiveness of burst stimulation delivered at inter-burst intervals ranging between 100 ms and 20 s found that LTP at Schaffer collateral-CA1 synapses could be obtained under most stimulation conditions, with the longest inter-burst intervals of 10 and 20 s evoking a shorter-lasting (5-10 min), but nonetheless robust, synaptic potentiation (Grover et al., 2009). These observations were interpreted by the authors as consistent with the prevalence of a range of low-frequency rhythms observed in the hippocampus, such as the 0.01-2 Hz sharp wave ripples seen during slow wave sleep (Buzsaki, 1986) which logic would suggest to have the capability of modulating neuronal activity and plasticity. Importantly, these observations have implications for the work outlined in this chapter,

as they indicate that the featured inter-trial intervals might plausibly support the induction of short-lasting plasticity, although they are not suitable for inducing long-term potentiation. A noteworthy aspect of the observed synaptic potentiation effects is that they emerged after a total of 80 stimulation trials (Grover et al., 2009), indicating that the numbers of stimulus exposures in the work described in this chapter (120 trials in Experiment 1; 215 trials in Experiment 2) were plausibly sufficient for inducing synaptic plasticity.

Finally, it is important to acknowledge that there is a considerable degree of uncertainty about the discreteness of neuronal populations receiving stimulation by the travelling square stimuli employed in Experiment 2 of this chapter, which has implications for the plausibility of plasticity induction through communication between nearby cell populations. As covered in Section 3.1.3, the decision to expose L2/3 of mouse S1 to spatiotemporal sequences of differently located squares of photostimulation was based on evidence obtained both from calcium and voltage-sensitive dye imaging that neuronal populations scattered across the lateral extent of the rodent barrel column are tuned to different angles of whisker deflection (Andermann & Moore, 2006; Tsytsarev et al., 2010). It was thus assumed that photostimulus delivery to discrete patches of tissue within L2/3 would engage neurons with distinct coding preferences, and thus roughly simulate a scenario in which a mouse experiences a particular sequence of whisker deflections with a certain degree of regularity (thus allowing stimulus expectancies to arise, and be subsequently broken). Given the prevalence of recurrent excitatory circuits within L2/3 (Feldmeyer, Lübke & Sakmann, 2006; Lefort et al., 2009), such a stimulation protocol might be expected to recruit some form of synaptic plasticity through repetitive signalling between connected populations of stimulated neurons.

Importantly, however, this assumption may not hold given considerable uncertainty about the degree to which targeting discrete patches of tissue recruits discrete and non-overlapping neuronal populations. Morphological reconstructions of excitatory neurons within mouse barrel cortex reveal that the dendritic arbours of cells within L2/3 extend up to 200 μm from the soma, with some sparser extensions seen at distances of up to 300 μm (Lefort et al., 2009). Such a dendritic arbour would, in fact, permeate the lateral extent of the average mouse barrel column (Aronof & Petersen, 2008). This makes it probable that the stimulation protocol outlined in this chapter may not, in fact, have targetted discrete cell populations with each differently localised square patch of photostimulation, but may instead have activated heavily overlapping groups of neurons with dendritic arbours permeating multiple stimulation zones. As such, it is difficult to draw firm conclusions from the experiments described - something that future studies can address by evoking and monitoring the neuronal activity with single-cell resolution, using methods such as two-photon targeted optogenetic stimulation (Zhang et al., 2018).

CHAPTER 4

4.1 INTRODUCTION

4.1.1 HOW DOES THE γ RHYTHM AFFECT INFORMATION TRANSFER?

The question of whether γ -oscillatory activity plays a role in predictive coding has thus far been largely examined by correlating power in the γ -band to relevant sensory parameters such as prediction error (El Karoui et al., 2015; Todorovic et al., 2011). Aside from this, our understanding of predictive processing in the brain might greatly benefit from asking the question on a smaller scale. Namely – how does γ -rhythmic activity, the hypothesised channel for feeding prediction errors through processing hierarchies (Bastos et al., 2012; Kanai et al., 2015), influence the response properties of neurons which take part in the rhythm? Exploring this might offer some insights into how prediction errors are propagated through sensory systems, and to what effect on their target populations.

One notable study indicated that rhythmicity at 40 Hz selectively enhanced the input-output curves of principal neocortical neurons (Sohal et al., 2009). In this experiment, pyramidal cell spikes were measured in response to excitatory postsynaptic current trains (EPSCs) that were either: i) asynchronous, by virtue of random sampling of input timings from a uniform distribution, ii) rhythmic at a theta (θ) frequency, with inputs being delivered on each half of a 125 ms cycle, or iii) rhythmic at a γ -frequency, with inputs delivered on each half of a 25 ms cycle. Target cell responses exhibited increased response gain specifically for γ -modulated EPSC trains, as well as enhanced mutual information between spike frequency and input strength. Modulation in the θ -band, on the other hand, produced only a marginally higher input-output slope compared to the non-rhythmic input. These findings formed the basis for concluding that γ might play a special role in boosting information transfer.

Nonetheless, the study leaves some uncertainties. For instance, there is the question of whether gain enhancement is specific to rhythmic excitation, or whether inhibition can be recruited to the same effect. This is an important consideration, since evidence suggests that inhibition is absolutely critical for generating the γ -rhythm (Bartos et al., 2007; White et al., 2000). This is evidenced, for instance, by the fact that intracellular γ -frequency postsynaptic potentials in principal cells reverse around the equilibrium potential for the inhibitory γ -aminobutyric acid type A (GABA_A) receptors (Penttonen et al., 1998). Both soma-targeting fast-spiking

(FS) and dendrite-targeting somatostatin (SST) interneurons are implicated in the γ -rhythm (Sohal et al., 2009; Veit et al., 2017), although FS interneurons are thought to be particularly relevant. Their autocorrelograms often peak in the γ -frequency band (Buzsáki et al., 1983) and they phase-lock to ongoing γ -frequency activity in the surrounding LFP (Penttonen et al., 1998; Hasenstaub et al., 2005). Furthermore, the causal role of perisomatic inhibition delivered by FS interneurons during network γ -oscillatory activity is supported by current source density analysis (Mann et al., 2005).

The time-course of inhibitory post-synaptic currents (IPSCs) is thought to be critical for the frequency of an oscillating network, with *in vivo* animal work finding that pharmacologically prolonging IPSCs results in slower pyramidal neuron rhythmicity (Traub et al., 1996; Whittington et al., 1995; Whittington et al., 2000). Such findings have led to the proposal that interneuronal connections onto pyramidal cells are pivotal in the induction of synchronised activity in the entire network, while the time constants of interneuron-to-interneuron, as well as interneuron-to-pyramidal cell connections within circuits set the frequency of these interactions (Bartos et al., 2007).

In this chapter, I set out to test the effects of γ -rhythmic inhibitory post-synaptic conductances (IPSGs) on the response properties of cortical pyramidal cells, and examine whether differing degrees of inhibitory γ -power produce varying degrees of gain enhancement. To this end, I used patch-clamp electrophysiology to record the spike outputs of individual pyramidal neurons in mouse S1 in response to increasing levels of excitation (delivered using square current pulses), during administration of simulated IPSGs of varying degrees of γ -power (delivered using dynamic clamp). My goal was to test how IPSG trains with differing γ -power affected response gain in principal neurons and thereby extend the work in our laboratory which had begun answering the question of whether γ -rhythmic inhibition plays a special role in enhancing the information transmission properties of excitatory neurons.

4.1.2 FREQUENCY-DEPENDENT SCALING OF PRINCIPAL CELL RESPONSES

In the second series of experiments detailed in this chapter, I examined whether effects exerted by rhythmic inhibition on principal cell input-output curves critically depend on frequency. In other words: is the γ -band special in matters of information transmission, as indicated by previous findings (Sohal et al., 2009), or can inhibition delivered at other frequencies in principle generate the same response scaling effects?

Exploring this is relevant to the greater scientific endeavour of dissecting the role of oscillatory dynamics in circuit processing. Thus far, the γ -rhythm in particular has

been implicated in facilitating the transmission and ‘binding’ of information across neuronal populations (Bichot et al., 2005; Fries, 2005; Grothe et al., 2012; Womelsdorf et al., 2007). As proposed by Pascal Fries - one of the progenitors of this hypothesis – this might, in part, be due to gamma’s capacity to modulate excitation at a sufficiently rapid rate so as to largely evade slow-onset feedback inhibition and thereby make it a particularly effective driver of postsynaptic targets (Fries, 2016). The aforementioned study examining the effects of γ - versus θ -band modulation of EPSCs on principal cell output certainly supports the view of gamma as an enhancer of signal transmission, finding that γ -modulated excitatory inputs were associated with particularly steep input-output response scaling (Sohal et al., 2009).

Oscillatory activity at lower frequencies has, in contrast, been traditionally ascribed an inhibitory role in circuit processing (Klimesch et al., 2007). In support of this view is a range of studies documenting an association between reduced power in the α/β frequency bands and increased stimulus-evoked neuronal activity and facilitated sensory detection. One human EEG study, for instance, found that α -band desynchronisation at posterior scalp sites during the pre-stimulus period predicted higher-amplitude event-related potentials (ERPs) and more accurate stimulus recall, on a trial-by-trial basis (Myers et al., 2014). In other studies, reduced α -band power immediately before stimulus onset was associated with a higher probability of participants detecting a visual stimulus (Busch & VanRullen, 2010), as well as misinterpreting visual noise as a signal (Iemi et al., 2017). This link between lower-frequency activity and inhibition is not necessarily surprising, as the inherently longer periods of reduced network excitability that it implies provide greater opportunity for briefly incoming stimuli to be suppressed. Indeed, stimuli arriving in the sensory cortex during one half of the ongoing EEG α cycle evoke significantly weaker ERPs (Barry et al., 2004) and are much more likely to go undetected than stimuli arriving in synchrony with other parts of the cycle (Busch, Dubois & VanRullen, 2009). The finding that the topography of low frequency activity can be consistently and voluntarily redistributed to areas where incoming sensory stimuli need to be filtered out of awareness to meet task demands (Ikkai et al., 2016; Rihs et al., 2007; Worden et al., 2000, further indicates that such rhythms have the capacity to suppress excitation and signal transmission. Thus, rhythms in the α/β -bands can plausibly be expected to influence scaling of principal cell responses in a manner which opposes that of the γ -rhythm.

Finding that this is the case would by no means be showing the obvious, as inhibition does not appear to be the end of the story for α -oscillatory activity. Crucially, the complications are especially interesting from the perspective of predictive coding. In one study, for example, researchers performed multi-electrode array recordings in the macaque visual system during a decision-making task featuring saccades (Neupane et al., 2017), focusing on the pre-saccadic period, during which representations of visual stimulus locations undergo a remapping to account for the inevitable migration of a fixated object along the retina (Duhamel et al., 1992; Neupane et al., 2016;

Wang et al., 2016). Analysis of LFPs across multiple sites in V4 revealed that the pre-saccade period was characterised by a spike in α -band coherence between neuronal populations encoding a stimulus' pre- and post-saccadic spatial locations. The relative phases of α -activity at these sites indicated that information flowed from the future to the present receptive field, leading the authors to conclude that α played a binding role – something traditionally ascribed to γ -band activity (Bichot et al., 2005). More specifically, the researchers suggested that α facilitates so-called 'predictive forward mapping' – that is, the binding of information pertaining to predictions as well as sensory reality (Neupane et al., 2017). As such, it may well be the case that α - and γ -oscillatory activity play equivalent roles in boosting information transmission, albeit perhaps with fundamental differences in terms of the direction along which facilitated flow occurs. This is an especially tantalising possibility given the links drawn between the two frequency bands and the feedforward and feedback communication of sensory prediction errors and predictions respectively (Bastos et al., 2012; Buffalo et al., 2011).

Given our current lack of knowledge about the outcomes of rhythmic inhibition at various frequencies, I set out to examine the effects of simulated rhythmic IPSPs delivered at a range of frequencies on the input-output properties of principal cells in mouse S1 *in vitro*. I used patch-clamp electrophysiology to record the spike outputs of individual cells receiving increasingly strong excitatory inputs (in the form of square current pulses), while dynamic clamp was used to administer IPSPs at a frequency of either 10, 20, 40, 60, 80 or 100 Hz. My goal was to gain a better understanding of rhythmic inhibition, and the potential relevance of different frequencies to information transmission in cortical circuits.

4.2 METHODS

4.2.1 ANIMALS

The experiments were approved by the Committee on Animal Care and Ethical Review of the University of Oxford and licensed by the UK Home Office. 16 adult mice (C57BL/6, aged 6-12 weeks; 8 males and 8 females) were used for the first set of experiments and 16 adult mice (C57BL/6, aged 6-12 weeks; 8 males and 8 females) for the second set of experiments with somatosensory cortex *in vitro*.

4.2.2 ACUTE BRAIN SLICE PREPARATION

To start the brain extraction and slicing procedure, mice were anaesthetised in a chamber filled with isoflurane (4% in 2 L/min O₂) until they exhibited a breathing rate indicative of deep anaesthesia (0.5-1 Hz) and a total loss of the pedal reflex. At this point, they were rapidly decapitated and the brain was dissected in NMDG-HEPES recovery solution (93 mM NMDG, 2.5 mM KCl, 1.2 mM NaH₂PO₄, 30 mM NaHCO₃, 20 mM HEPES, 25 mM glucose, 10 mM MgSO₄, 0.5 mM CaCl₂. Osmolarity 300-310 mOsm/L; pH adjusted to 7.3-7.4 with HCl) that had been cooled on ice over a period of 30-60 mins. Coronal slices of 350 μ m thickness were obtained using a VT 1200S vibratome (Leica). They were given 10-15 minutes to recover in an interface storage chamber filled with warmed (28-33° C) NMDG-HEPES solution prior to being washed in aCSF (126 mM NaCl, 24 mM NaHCO₃, 10 mM glucose, 3.5 mM KCl, 2 mM CaCl₂, 1.25 mM NaH₂PO₄, and 1 mM MgSO₄) and being transferred to an interface storage chamber filled with room-temperature aCSF for at least 1 hour of recovery before commencing recordings. The described slicing and recovery protocol was selected due to observations previously made in our laboratory that it yielded good quality slices and patch clamp recordings. All solutions used during the dissection, slicing, and incubation procedures were saturated with 95% O₂ / 5% CO₂.

4.2.3 INTRACELLULAR RECORDINGS

Experiments were conducted in a single submerged chamber perfused with warmed aCSF (28-33° C) continuously bubbling with 95% O₂ / 5% CO₂. Intracellular recordings were acquired using borosilicate glass pipettes (5-12 M Ω) filled with an internal solution (110 mM KGluconate, 40 mM HEPES, 2 mM ATP-Mg, 0.3 mM GTP-NaCl, 4 mM NaCl; pH 7.2-7.3; Osmolarity 270-290 mOsm/L). Signals were ac-

quired through the MultiClamp 700B amplifier (Molecular Devices), digitised at 8000 Hz by a data acquisition board (ITC-18, InstuTECH), and recorded using Igor Pro software. Slice and cell visualisation, as well as targeting of the recording electrode, were performed with the aid of oblique illumination and an Olympus BX51W1 microscope.

Following breakthrough of the cell membrane, bridge balance was adjusted to compensate for series resistance. Before commencing any experiments on a given cell, a series of hyperpolarising and depolarising square current pulses of varying amplitudes was administered to rule the cell out as an interneuron and classify it as a putative pyramidal cell on the basis of its electrophysiological phenotype. For this purpose, membrane potential responses only to the depolarising current pulses were visually examined, as they provided sufficient information to make high-confidence judgements identifying pyramidal neurons. The classification took into account several electrophysiological features, such as the half-width of the action potential (<1 ms in interneurons, and significantly longer in pyramidal cells; Bean, 2007), the relative velocities of the rise and fall phases of the action potential (substantially slower fall velocity in pyramidal cells compared to interneurons; McCormick et al., 1985), and the intrinsic firing pattern upon stimulation with a suprathreshold depolarising current (with regular-spiking cells exhibiting spike frequency accommodation, and fast-spiking interneurons exhibiting a non-adapting high-frequency firing pattern; Connors & Gutnick, 1990).

It is important to note that a regular spiking pattern does not unambiguously discriminate between a pyramidal cell and a regular-spiking interneuron. However, a further characteristic differentiates them - namely the observation that action potentials of regular-spiking interneurons are followed by robust short-term after-hyperpolarisation potentials not seen in pyramidal cells (Malkin et al., 2013). Thus, a visual judgement was made on the basis of all the aforementioned characteristics to identify each encountered cell as a putative pyramidal neuron. All cells which conformed to the pyramidal phenotype were included in analysis. Figure 4.1 illustrates the electrophysiological phenotyping procedure for a representative cell, highlighting the features of interest. Tables 4.1 and 4.2 summarise the electrical properties of putative pyramidal cells recorded from L2/3 of S1 for the first and second experiments in this chapter, respectively.




Figure 4.1: Electrophysiological phenotyping of a putative pyramidal cell. **A** shows the membrane potential of a representative neuron (bottom) in response to a progression of square excitatory current pulses of increasing magnitude (top) administered as part of the neuronal classification procedure. **B** provides a close-up of the initial 250 ms window of one of the membrane potential traces from **A** (marked in red), which shows spike frequency accommodation, characteristic of regular spiking cells in cortex (Connors & Gutnick, 1990). **C** provides a further yet close-up of one of the spikes from **B** (marked in red), and demonstrates two important features of the action potential which strongly indicate a pyramidal, rather than interneuronal, cell origin. Namely, it can be observed that the half-width of the action potential is longer than 1 ms, and that the falling phase of the action potential is significantly less steep than the rising phase (Bean, 2007; McCormick et al., 1985).

Table 4.1: Electrical properties of putative pyramidal cells recorded in L2/3 of mouse S1 for Experiment 1 (n=16)

	<i>C (pF)</i>	<i>Sag ratio</i>	<i>Thresh. (mV)</i>	<i>AP half-width (ms)</i>	<i>AHP (mV)</i>
M ± SD	183.27 ±113.99	0.97 ±0.047	-46.41 ±5.64	1.54 ±0.42	1.4 ± 1.5
Min.	50.80	0.84	-54.28	1.20	0.00
Max.	454.00	1.006	-33.48	2.90	4.8

Table 4.2: Electrical properties of putative pyramidal cells recorded in L2/3 of mouse S1 for Experiment 2 (n=12)

	<i>C (pF)</i>	<i>Sag ratio</i>	<i>Thresh. (mV)</i>	<i>AP half-width (ms)</i>	<i>AHP (mV)</i>
M ± SD	145.73 ±138.50	0.99 ± 0.011	-46.27 ±11.09	1.61 ±1.06	0.86 ± 0.77
Min.	281.00	0.965	-69.31 mV	1.023	0.028
Max.	513.00	1.002	-21.53 mV	4.92 ms	2.40

C - Capacitance; **Sag ratio** - defined as the ratio between the steady-state and the peak voltage response to a hyperpolarising current of -140 pA; **Thresh.** - Threshold for firing an action potential; **AHP** - Afterhyperpolarisation amplitude, defined as the difference between firing threshold and the membrane potential trough following an action potential, and determined using the first action potential in a spike train.

4.2.4 DYNAMIC CLAMP

Perisomatic inhibition – the sort delivered via GABA_A receptors by fast-spiking interneurons (Freund & Katona, 2007) – is highly effective at synchronising pyramidal cells (Lytton & Sejnowski, 1991; Mann et al., 2005), and considered a ‘key ingredient of gamma oscillations’ (Buzsaki & Wang, 2012). Dynamic clamp is the optimal way to mimic the effects on membrane potential exerted by perisomatic shunting inhibition during γ -oscillatory activity. It uses a high-resolution feedback loop (50 kHz update rate) to output a current that would flow through a theoretical receptor (in this case, GABA_A) in accordance with the ongoing interaction between conductance and membrane potential changes which unfold over time during receptor stimulation. In these experiments, dynamic clamp was performed using IgorPro software, which continuously updated the injected inhibitory current given a command conductance (which simulates GABA_A receptor opening) according to the following equation:

$$I_{inh}(t+1) = G_{inh}(t) \times (V_m(t) - E_{GABA})$$

Here, $I_{inh}(t)$ is the inhibitory current injected at time-point t , G_{inh} is the commanded inhibitory conductance, E_{GABA} is the reversal potential for GABAergic synapses and V_m is the membrane potential of the cell. The dynamic clamp procedure is illustrated in Figure 4.2.

The amplitude of the IPSP train was tailored to each individual cell before commencing the experimental protocol. It was selected on the basis of a high γ -power IPSP at the chosen amplitude in combination with a square current pulse independently sufficient to induce an action potential being able to produce a 5 mV fluctuation in membrane potential. The experiments were performed in whole cell dynamic clamp mode.




Figure 4.2: Schematic illustration of the dynamic clamp technique. Using dynamic clamp to simulate inhibitory synaptic currents involves injecting a current (I_{inh}) that is determined using a high-resolution feedback loop which samples the membrane potential (V_m), computes its distance from the reversal potential for GABA (E_{GABA}), and scales the discrepancy by the command inhibitory conductance (G_{inh}). Dynamic clamp was performed by the ITC18.

4.2.5 EXPERIMENT 1 PROTOCOL: SIMULATING γ -RHYTHMIC INHIBITION VARYING IN POWER

My goal was to study how simulated changes in γ -power in interneuronal output influence principal cell output over a gradient of excitatory input strengths. As such, the experiment needed to make use of two classes of stimulus: i) conductance trains at simulated inhibitory synapses, varying in γ -oscillatory power, and ii) excitatory current inputs of varying magnitudes. These inputs would be applied concurrently with the simulated inhibitory inputs to provide a readout of pyramidal cell response scaling during varying degrees of excitation under conditions of varying γ -power of inhibition. The stimuli used in this experiment were originally generated by Pantelis Antonoudiou - a former graduate student at our laboratory.

Stimuli used to reproduce GABA_A-receptor mediated inhibitory rhythms were 2 s synaptic conductance trains generated at a 40 Hz frequency. This was selected because rhythmic inputs at this frequency have previously been found to produce a multiplicative gain enhancement of pyramidal cell output (Sohal et al., 2009), and because Pantelis had, in previous years, set a precedent by examining the effects of various properties of inhibitory conductance trains at this frequency.

The synaptic conductance trains which simulated GABA_A receptor opening were produced by first generating a binary synaptic input train in which 1000 synchronous synaptic events occurred rhythmically at 40 Hz. These inputs were then convolved with the temporal profile of a typical inhibitory post-synaptic conductance event (sIPSG), as follows:

$$\alpha^3 x \times e^{-\alpha x} \times G_{inh} \times e/\alpha$$

where $\alpha = 200$, and $G_{inh} = 5 \times 10^{-18}$ S. This generated an inhibitory post-synaptic conductance (IPSG) train with a peak amplitude of 1.04 pS, illustrated in Figure 4.3.

Varying the power of the resulting γ -inhibition involved modifying the structure of the underlying input train before convolving it with the IPSG. The way this was done was based on the rationale that variations in γ -power obtained in an LFP might reflect variations in one or more parameters. One important factor might be the consistency of the intervals which separate the co-occurrence of multiple synaptic inhibitory events (termed rhythmicity); while another might be the precision with which these inhibitory events align with the peak of the ongoing γ -rhythm (termed synchronicity). Work previously done by Pantelis in our laboratory confirmed that, while lower rhythmicity and synchronicity both have the effect of lowering γ -power, they have strongly opposing effects on the temporal variability of the γ -oscillation (with reduced synchronicity driving γ -variability to lower values). Thus, to produce inhibitory post-synaptic conductances with varying levels of γ -power while controlling for temporal variability, stimuli were based on different levels of power-modulating rhythmicity whose effects on stimulus variability were offset by specific synchronicity values.

To produce various levels of rhythmicity of the 40 Hz IPSG trains, the interval be-

tween synchronous inhibitory event bursts (inter-burst interval; IBI) was sampled from a γ distribution ($M=ab$, $\sigma^2 = ab^2$), with a mean of 25 ms, and a standard deviation which corresponded to $CV \times \mu$, where CV was a coefficient of variation varying between 0 (perfect rhythmicity) and 1 (complete arrhythmia). Conversely, modifying the synchronicity the inhibitory γ -rhythm involved introducing a certain degree of temporal jitter to the onset of each individual synaptic event with respect to the onset of the peak to which all events would simultaneously align in a perfectly synchronous setup. This jitter parameter was sampled from a uniform distribution with the range $[\pm \text{jitter}]$ and scaled by the 25 ms interval which separates event bursts in a 40 Hz synaptic event train. Figure 4.3 provides an illustration of the independent effects of varying rhythmicity and synchronicity indices on the IPSP train. The stimuli generated for the experiment comprised the following categories of γ -power (and the following values of rhythmicity and synchronicity): Low γ -power (CV = 0.67, Jitter = 0.94), Mid γ -power (CV = 0.3, Jitter = 0.24), and High γ -power (CV = 0, Jitter = 0). In addition, a flat square pulse of inhibitory conductance matched for amplitude with the rhythmic inhibitory stimuli was also generated (which I will refer to as ‘flat’ throughout the analysis).

To produce the excitatory inputs that would be administered to read out cells’ response properties, a series of 12 square pulses of excitatory currents (each of a 2 s duration) were generated, separated by fixed 20 pA or 40 pA current steps, depending on cell sensitivity, with the goal of preventing a firing rate plateau in the early stages of the experiment. The first amplitude of the series was individually tailored to each cell, such that it was just below the amplitude of the pulse which was required for the cell to fire an action potential during the examination of its intrinsic electrophysiological phenotype.

To study the effect that different levels of inhibitory γ -power have on principal cell output, the experimental protocol consisted of applying each of the four inhibitory conductances (flat, low- γ , mid- γ and high- γ) in randomly selected order for a total of 5 repetitions (order was determined separately for each repetition cycle), in pairing with each of the 12 excitatory current pulses presented in ascending magnitude. The procedure is illustrated in Figure 4.4. Trials were separated by 10 s, for a total protocol duration of 50 minutes. Upon completion of the stimulation protocol for each level of excitation, a pair of 1 s pulses (-40 pA and -80 pA) were applied with a 1 s interval, to track the input resistance of the membrane as the experiment progressed.




Figure 4.3: Generation of γ -oscillatory synaptic inhibitory conductance trains of varying γ -power, through manipulation of rhythmicity and synchrony. **A** demonstrates the manipulation of rhythmicity, using sample 500 ms trains of 40 Hz synaptic events (top) with descending degrees of rhythmicity (left to right). These are produced by setting the interval between 1000 synchronous synaptic events to an amount drawn from a γ -distribution with a μ of 25 ms and a standard deviation of $CV \times \mu$, where CV is the coefficient of variation (0, 0.2, and 0.5 are shown here). The timelines of synaptic inhibitory trains shown on top are convolved with a kernel representing the temporal envelope of a typical single inhibitory post-synaptic conductance event (generated by the formula: $\alpha^3 x \times e^{-\alpha x} \times G_{inh} \times e/\alpha$, where $\alpha = 200$, and $G_{inh} = 5 \times 10^{-18}$ S), yielding inhibitory post-synaptic conductance (IPSP) trains with varying levels of γ -rhythmicity (bottom). **B** demonstrates the manipulation of synchronicity, using sample 500 ms trains of perfectly rhythmic 40 Hz synaptic events (top) with descending degrees of synchronicity (left to right). These are achieved by jittering the timing of each of the 1000 individual synaptic events (by $J \times (IBI/2)$), where J represents a number drawn at random from a uniform distribution with the range $[\pm \text{jitter}]$, and IBI represents the 25 ms interburst-interval which separates each burst of 1000 events in a perfectly synchronous 40 Hz spike train. Larger jitter values introduce a growing degree of timing noise to the individual events, which degrades the amplitude of the resulting IPSP trains (bottom).

4.2.6 EXPERIMENT 2 PROTOCOL: SIMULATING RHYTHMIC INHIBITION VARYING IN FREQUENCY

The goal of the second experimental protocol was to test whether any multiplicative scaling effects that γ -rhythmic inhibition has on principal cell responses can be obtained using other frequencies. In other words: is the γ -rhythm special in terms of its modulatory properties? Addressing this question required two classes of stimulus: i) inhibitory post-synaptic conductance trains oscillating at various frequencies while being matched for power (defined as area under the curve), and ii) excitatory square current pulses, to be applied concurrently with each inhibitory conductance train to trigger a readout of principal cell responses to excitation under conditions of varying rhythmic inhibitory conductance frequencies.

Rhythmic inhibitory conductance (IPSG) trains were generated as sinusoidal waves at the following frequencies: 10 Hz, 20 Hz, 40 Hz, 60 Hz, 80 Hz, and 100 Hz, each of a 2 s duration. These IPSGs were administered to principal cells in cortical L2/3 using dynamic clamp mode (see Figure 4.2 for illustration), while being paired with 2 s excitatory square current pulses which increased in magnitude in 12 steps over the course of the experiment. The excitatory current steps were separated by 10 pA, 20 pA or 40 pA, in a manner which took into account each cell's response sensitivity, so as to avoid reaching a firing rate plateau in the early stages of the experiment. Furthermore, the first amplitude of the series was individually tailored to each cell, to be just below threshold for firing an action potential (which was assessed during the electrophysiological phenotyping procedure which preceded any experimental recordings).

The experimental protocol consisted of applying each of the six IPSG trains (10 Hz, 20 Hz, 40 Hz, 60 Hz, 80 Hz, and 100 Hz) in randomly selected order for a total of 5 repetitions (order was determined separately for each repetition cycle), for each of the 12 excitatory current pulses presented in ascending magnitude. Trials were separated by 10 s, for a total protocol duration of 60 minutes. Upon completion of the stimulation protocol for each level of excitation, a pair of 1 s pulses (-40 pA and -80pA) were applied with a 1 s interval, to track the input resistance of the membrane as the experiment progressed.




Figure 4.4: Experimental protocol for probing the effects IPSG γ -power on principal cell response properties. A cell's membrane potential was being recorded while receiving 2 s square pulses of excitatory current of increasingly large amplitudes (I_{exc} Amp) starting with an initial amplitude, defined as just below the minimum required for an action potential, for a total of 12 steps, each 20 pA or 40 pA higher than the last (step size adjustments were made to prevent spike frequency ceilings during the protocol). Each excitatory input was co-administered with a range of different inhibitory post-synaptic conductance trains (IPSGs) in Experiment 1 and 2. In Experiment 1, the excitatory current pulses were paired with IPSGs defined by each of four levels of γ -power (high, medium, low, and flat control stimulus; shown above), presented in shuffled order. In Experiment 2, the excitatory current pulses were paired with IPSGs defined by each of six frequencies (10 Hz, 20 Hz, 40 Hz, 60 Hz, 80 Hz, and 100 Hz), presented in shuffled order. This procedure was repeated 5 times for each amplitude of excitatory input before moving on to the next amplitude.

4.2.7 MODELLING THE SINGLE CELL: MODEL ARCHITECTURE

To better understand the physiological underpinnings of the effects of rhythmic inhibition on pyramidal cell response properties, I attempted to compare my observations to the output of a model pyramidal neuron published by Golomb et al. (2006), which simulates the cell as a single compartment comprising the main ionic currents known to take place in the soma and proximal dendrites. To outline the primary features of the model cell, its change in membrane potential over time was represented by a series of coupled differential equations, each pertaining to a particular ionic flux. The currents included in the model were:

- I_{Na} : Transient Na^+ current responsible for the depolarising phase of the action potential.
- I_{K-dr} : Delayed rectifier K^+ current responsible for the falling phase of the action potential.
- I_{Ca} : High-threshold Ca^{2+} current through voltage-gated channels, required for vesicular release.
- I_{K-fast} : Fast calcium-activated K^+ current, which contributes to rapid repolarization (Storm, 1987; Yue and Yaari, 2004).
- I_{K-slow} : Slow calcium-activated K^+ current, which mediates slow after-hyperpolarization and spike frequency adaptation (Madison & Nicoll, 1984).
- I_{K-M} : M-type K^+ current, an inactivating muscarinic-sensitive receptor known to regulate membrane potential over long time-scales, regulating transitions between regular spiking and bursting modes of firing (Bertram et al., 1995).
- I_{K-A} : A-type K^+ current, a rapidly activating / inactivating voltage-sensitive current known to play a role in regulating spike frequency (Hagiwara et al., 1962; Tierney & Harris- Warrick, 1992).

These ionic currents were defined in accordance with factors such as ion-specific membrane conductances, membrane potential, distance from reversal potential for the relevant ion, and scaled by dynamical variables based on the activation curves of each current, as follows: $I_{Na}(V, h) = G_{Na}m^3(V)h(V - V_{Na})$, $I_{K-dr}(V, n) = G_{K-dr}n^4(V - V_K)$, $I_{K-A}(V, b, a) = G_{K-A}a^3(V)b(V - V_K)$, $I_{K-M}(V, z) = G_{K-M}z(V - V_K)$, $I_{Ca}(V, r) = G_{Ca}r^2(V - V_{Ca})$, $I_{K-fast}(V, c, d) = G_{K-fast}d([Ca^{2+}]_i)c(V - V_K)$, $I_{K-slow}(V, q) = G_{K-slow}q(V - V_K)$. The dynamics of the activation functions used to define these currents, as well as the parameters and time constants involved, are detailed in Golomb et al. (2006), and shown in Table 4.3. The conductances and reversal potentials were as follows: $G_{Na} = 35 \text{ mS/cm}^2$; $G_{K-dr} = 6 \text{ mS/cm}^2$, $G_{K-fast} = 10 \text{ mS/cm}^2$, $G_{K-slow} = 5 \text{ mS/cm}^2$, $G_{K-A} = 0.14 \text{ mS/cm}^2$, $G_{K-M} = 0.1 \text{ mS/cm}^2$, $G_{Ca} = 0.3 \text{ mS/cm}^2$ (the last three were explicitly manipulated in the simulations). $E_{Na} = 55 \text{ mV}$, $E_K = -90 \text{ mV}$, $E_{Ca} = 120 \text{ mV}$.

The kinetics of intracellular calcium concentration, $[Ca^{2+}]_i$, were determined as follows:

$$\frac{d[Ca^{2+}]_i}{dt} = -u \times I_{Ca} - \frac{[Ca^{2+}]_i}{\tau_{Ca}}$$

where $u = 0.13 \text{ cm}^2/\text{ms} \times \mu\text{A}$, and $\tau_{Ca} = 13 \text{ ms}$.

All simulations were done in MATLAB (R2017b), using the fourth-order Runge-Kutta method with a Δt of 0.05 ms, as described in the original publication of this model (Golomb et al., 2006).

4.2.8 MODELLING THE EFFECTS OF RHYTHMIC INHIBITION

The goal was to examine the firing patterns of the model cell in response to a series of stimuli which mimicked the experimental set-up described in Section 4.2.6. The cell was injected with a series of 75 excitatory square current pulses (2 s duration), starting with 700 pA/cm², and increasing in 40 pA/cm² steps. Each excitatory current pulse was paired with simulated rhythmic inhibition oscillating in a sinusoidal fashion at one of six different frequencies: 10 Hz, 20 Hz, 40 Hz, 60 Hz, 80 Hz, and 100 Hz.

Rhythmic perisomatic inhibition was simulated by generating a sinusoidal conductance at the required frequency which, for each time-step in the simulation, was converted into the current I_{GABA} through scaling by $(V - E_{GABA})$, where E_{GABA} represented the reversal potential for GABA of -70 mV. Thus, the full equation computing the change in the model cell's membrane potential over time was:

$$C \frac{dV}{dt} = -G_{Leakage}(V - E_{Leakage}) - I_{Na} - I_{K-dr} - I_{K-A} - I_{K-M} - I_{Ca} - I_{K-fast} - I_{K-slow} - I_{GABA} + I_{app}$$

where $C = 1 \text{ } \mu\text{F}/\text{cm}^2$, $G_{Leakage} = 0.05 \text{ mS}/\text{cm}^2$, $E_{Leakage} = -70\text{mV}$, and I_{app} was the applied excitatory current.

To identify a set of conditions which replicated the relationship between IPSP frequency and pyramidal cell input-output properties observed in the experimental section of this chapter, I simulated the cell featuring manipulations in contributions stemming from the A- versus M-type K^+ current, as well as the voltage-gated Ca^{2+} current, which in turn regulates the fast and slow calcium-activated K^+ currents (I_{K-fast} and I_{K-slow}).

Table 4.3: Kinetics equations and parameters for the single pyramidal cell model

<i>Current, variable</i>	<i>Kinetics</i>	<i>Parameters</i>
I_{Na}, m	$\frac{dm}{dt} = (1 + e^{\frac{-(V-\theta_m)}{\sigma_m}})^{-1}/dt$	$\theta_m = -35$ mV $\sigma_m = 10$ mV
I_{Na}, h	$\frac{dh}{dt} = \phi[(1 + e^{\frac{-(V-\theta_h)}{\sigma_h}})^{-1} - h]/\tau_h(V)$ $\tau_h(V) = 0.1 + 0.75(1 + e^{\frac{-(V-\theta_{ht})}{\sigma_{ht}}})^{-1}$	$\theta_h = -45$ mV $\sigma_h = -7$ mV $\theta_{ht} = -40.5$ mV $\sigma_{ht} = -6$ mV $\phi = 1$
I_{Kdr}, n	$\frac{dn}{dt} = \phi[(1 + e^{\frac{-(V-\theta_n)}{\sigma_n}})^{-1} - n]/\tau_n(V)$ $\tau_n(V) = 0.1 + 0.5(1 + e^{\frac{-(V-\theta_{nt})}{\sigma_{nt}}})^{-1}$	$\theta_n = -35$ mV $\sigma_n = 10$ mV $\theta_{nt} = -27$ mV $\sigma_{nt} = -15$ mV $\phi = 1$
I_{K-A}, a	$\frac{da}{dt} = [(1 + e^{\frac{-(V-\theta_a)}{\sigma_a}})^{-1} - a]/dt$	$\theta_a = -50$ mV $\sigma_a = 20$ mV
I_{K-A}, b	$\frac{db}{dt} = [(1 + e^{\frac{-(V-\theta_b)}{\sigma_b}})^{-1} - b]/\tau_b,$	$\theta_b = -80$ mV $\sigma_b = -6$ mV $\tau_b = 15$ ms
I_{K-M}, z	$\frac{dz}{dt} = [(1 + e^{\frac{-(V-\theta_z)}{\sigma_z}})^{-1} - z]/\tau_z,$	$\theta_b = -39$ mV $\sigma_b = 5$ mV $\tau_b = 75$ ms
I_{Ca}, r	$\frac{dr}{dt} = [(1 + e^{\frac{-(V-\theta_r)}{\sigma_r}})^{-1} - r]/\tau_r,$	$\theta_r = -20$ mV $\sigma_r = 10$ mV $\tau_r = 1$ ms
I_{K-fast}, c, d	$\frac{dc}{dt} = [(1 + e^{\frac{-(V-\theta_c)}{\sigma_c}})^{-1} - c]/\tau_c,$ $\frac{dd}{dt} = 1 + \frac{a_c}{[Ca^{2+}]_i}^{-1}$	$\theta_c = -30$ mV $\sigma_c = 7$ mV $\tau_c = 2$ ms $a_c = 6$
I_{K-slow}, q	$\frac{dq}{dt} = [(1 + \frac{a_q^4}{[Ca^{2+}]_i^4})^{-1} - q]/\tau_q,$	$a_q = 2$ $\tau_q = 450$ ms

4.3 RESULTS: EXPERIMENT 1

4.3.1 RECAP AND OBJECTIVES

The main goal of this experiment was to verify a series of previous observations in the laboratory that i) γ -frequency (40 Hz) inhibitory synaptic conductances produce a multiplicative gain in principal cell responses to excitatory stimulation, and ii) that the gain effect is sensitive to variations in the γ -power of the inhibitory rhythm, as should be reflected in a steeper relationship between excitatory drive and firing rate during IPSGs of high γ -power.

To investigate this question, 16 putative pyramidal cells (profiled on the basis of electrophysiological characteristics; see Figure 4.1) were patched in L2/3 of S1. Their response properties were examined during stimulation with excitatory currents of increasing amplitudes, while dynamic clamp was used to co-administer simulated inhibitory postsynaptic currents (IPSCs) to mimic the effects of γ -rhythmic perisomatic inhibition. These IPSCs were based on γ -rhythmic inhibitory conductances (IPSGs; 40Hz), which differed in terms of γ -power (low, mid, and high), while conserving temporal variability, due to complementary manipulations of rhythmicity and synchronicity (see Section 4.1.5 and Figure 4.3). Figure 4.5A demonstrates a representative cell's spiking behaviour in response to a 260 pA excitatory current injection during application of IPSGs with high, medium, and low γ -power. Statistical analysis of variables extracted from the cells' spiking behaviours primarily involved repeated-measures ANOVAs; Greenhouse-Geisser corrected values are reported where a Mauchly's test indicated a violation of the sphericity assumption. Where results were followed-up by post-hoc t-tests, the significance threshold was corrected to $p \leq 0.017$ (unless otherwise specified), to account for the number of comparisons.

4.3.2 INPUT RESISTANCE AND STABILITY

Whole-cell recordings were acquired from S1 using whole-cell dynamic clamp mode. To quantify cell stability and recording quality over the course of an experiment, I measured input resistance (R_i) and resting membrane potential (V_m) at regular intervals throughout the protocol. Resistance was sampled by examining the membrane potential response to a square current pulse of -80 pA, which was applied at the end of each of the 12 blocks featuring increasingly higher excitatory current amplitudes (see Methods section 4.1.5). Membrane potential was sampled during a 500 ms baseline period preceding the excitatory current injection of each trial, and subsequently segregated by block number to track average membrane potential as the experiment unfolded. To evaluate cell stability over the duration of an experiment, average resistance and membrane potential were compared between the first and final blocks of each experiment (shown Figure 4.5C).

Two paired-samples t-tests revealed that input resistance did not significantly differ

between the first and final blocks of each experiment (Ri first block: $M = 373.69 \text{ M}\Omega$, $SD = 223.06$; Ri final block: $M = 314.02 \text{ M}\Omega$, $SD = 191.01$; $t = 1.73$, $p = 0.105$); and neither did resting membrane potential (Vm first block: $M = -70.33 \text{ mV}$, $SD = 0.93$; Vm final block: $M = -70.07 \text{ mV}$, $SD = 1.34$; $t = 0.69$, $p = 0.49$). A number of cells did not survive to the end of the stimulation protocol, however. Cells were only included in the analysis if they provided data for >8 out of 12 levels of excitation examined in the protocol.

4.3.3 INHIBITORY γ -POWER INCREASES PRINCIPAL CELL MAX. FIRING RATES

First, I set out to examine whether the spike rate ceiling, on average, differs between inhibitory γ -power conditions. Maximal spike rates achieved in each condition were identified for every cell and submitted to a repeated-measures one-way ANOVA with IPSP γ -power (Flat, Low, Medium, and High) as the repeated-measures variable. Figure 4.5D demonstrates these firing rate maxima for each individual cell (grey) as well as the cohort average (red).

The ANOVA revealed a significant main effect of IPSP γ -power ($F(1.37, 45) = 54.57$, $p < 0.001$, $\eta_p^2 = 0.78$). This was followed up by a series of post-hoc t-tests comparing neighbouring γ -power conditions. These found that the average maximal firing rate evoked during the flat non-rhythmic control IPSP ($M = 6.16$, $SD = 10.62$) was significantly lower compared to the low γ -power condition ($M = 17.02$, $SD = 7.08$, $t = -8.34$, $p < 0.001$), which was in turn significantly lower than the medium γ -power condition ($M = 23.52$, $SD = 10.61$, $t = -5.91$, $p < 0.001$). Finally, maximal firing rates in the high γ -power condition were, on average, significantly higher than those in the medium γ -power condition ($M = 32.00$, $SD = 17.66$, $t = 4.34$, $p = 0.001$). Thus, the evidence suggests that increasing γ -power in inhibitory conductance went hand-in-hand with a significant growth in maximal firing rate.

4.3.4 INHIBITORY γ -POWER BOOSTS PRINCIPAL CELL RESPONSE SCALING

In addition to observing general effects on spike frequency, it is useful to ask whether differing levels of inhibitory γ -power influence principal cells by modulating the relationship between excitatory drive and cell output. Thus, I set out to examine whether putative pyramidal cells' input-output curves (i.e. firing rates as a function of injected excitatory current amplitude) were sensitive to manipulations in inhibitory γ -power. Figure 4.5B demonstrates the input-output relationship in each γ -power condition for one representative cell from the cohort.

The nature of this relationship was quantified for each cell by fitting a linear regression to its input-output curve, starting with the first non-zero spike rate and ending with the maximal spike frequency, separately for each γ -power condition. The obtained slope estimates are shown in Figure 4.5F, individually for each cell (grey) and as a cohort average (green). These were submitted to a repeated-measures one-way ANOVA with IPSPG γ -power (Flat, Low, Medium, and High) as the repeated-measures variable, which revealed a significant main effect of IPSPG γ -power ($F(1.12, 45) = 18.70, p < 0.001, \eta_p^2 = 0.56$).

This finding was explored with a series of post-hoc t-tests, which revealed that the input-output slopes in the flat non-rhythmic control IPSPG were, on average, significantly lower compared to all three IPSPG γ -power conditions (Flat: $M = 0.82, SD = 1.23$; t-test results - High: $M = 3.90, SD = 2.84, t = 4.52, p < 0.001$; Medium: $M = 2.13, SD = 1.22, t = 6.22, p < 0.001$; Low: $M = 1.65, SD = 0.97, t = 4.36, p < 0.001$). Furthermore, mean input-output slopes were significantly higher in the high γ -power condition compared to the medium ($t = 3.39, p = 0.004$) and low ($t = 3.88, p = 0.001$) γ -power conditions. Medium γ -power, in turn, produced significantly higher mean input-output slopes than low γ -power ($t = 2.73, p = 0.015$). Thus, the evidence indicates that γ -rhythmic inhibitory inputs produce a multiplicative scaling of principal cell responses that is proportional to the γ -power in the inhibitory conductance.

4.3.5 INHIBITORY γ -POWER INCREASES PRINCIPAL CELL RHEOBASE

To further examine the dimensions of the principal cell response that might be influenced by various degrees of IPSPG γ -power, I assessed the effect of γ -power on the rheobase, defined as the minimum excitatory current required for the cell to generate at least one action potential over the course of a single trial. Rheobase was extracted separately for each cell and γ -power condition, and expressed as a % of the excitatory current which produced the maximal spike rate in each experiment. Figure 4.5F shows the resulting rheobase values, individually for each cell (grey) and as a cohort average (blue).

A repeated-measures one-way ANOVA with IPSPG γ -power (Flat, Low, Medium, and High) as the repeated-measures variable revealed a significant main effect of γ -power ($F(1.48, 45) = 78.73, p < 0.001, \eta_p^2 = 0.84$). A series of post-hoc t-tests produced the fol-

lowing findings: rheobase was significantly higher in the high compared to medium γ -power condition (High: $M = 50\%$, $SD = 17$; Medium: $M = 36\%$, $SD = 12$; $t = 5.94$, $p < 0.001$), which was in turn significantly higher than rheobase in the low γ -power condition ($M = 33\%$, $SD = 10$; $t = 4.56$, $p < 0.001$). All three γ -power conditions were associated with significantly lower mean rheobase values compared to those in the flat non-rhythmic control IPSPG condition (Flat: $M = 78\%$, $SD = 21$; t-test results - High: $t = 8.003$, $p < 0.001$; Medium: $t = 10.31$, $p < 0.001$; Low: $t = 9.69$, $p < 0.001$). These findings show that γ -power significantly influences the cells' baseline sensitivity, as reflected in higher thresholds for eliciting spiking when excitatory stimulation is accompanied by IPSPGs with greater γ -power.

4.3.6 INHIBITORY γ -RHYTHM INCREASES PRINCIPAL CELL SPIKE FIDELITY

It is interesting to examine whether γ -power variability in rhythmic inhibitory conductances can modulate the reliability of spike timing in pyramidal cells across repetitions of the same excitatory input. This issue is relevant to the broader question of how γ -rhythmic inhibition regulates information processing in principal cells, given the plethora of evidence that spike timing, both absolute and relative, is a critical parameter of the neural code (Meister & Berry, 1999) which allows inputs to be decoded with substantially higher accuracy compared to an exclusively rate-based code (Harvey et al., 2013; Lundstrom & Fairhall, 2006; Reich et al., 2000; Ruyter van Steveninck & Bialek, 1988).

The experimental paradigm was well-suited to pursuing this question, as each level of excitatory current pulse was repeated for a total of 5 times along with each of the four inhibitory conductance types (low, mid, and high γ -power, and flat control IPSPG). This enabled me to examine the replicability of principal cells' spike patterns across stimulus repetitions, by calculating a cross-trial correlation index. This index was obtained by averaging spike timing cross-correlations, obtained at a 0 ms time-lag, between every possible pairing of trials at each level of excitation (normalised by spike rate). Figure 4.6A demonstrates the mean cross-correlation index (CCI) observed under each IPSPG γ -power condition, aggregated across excitatory input strengths.

Indices from individual experiments were submitted to a repeated-measures one-way ANOVA with IPSPG γ -power (Low, Medium, and High) as the repeated-measures variable (the flat non-rhythmic control IPSPG condition was excluded due to the generally low firing rates which allowed cross-correlation estimates to be made in only a minority of cells). The test revealed a significant main effect of IPSPG γ -power ($F(1.07, 24) = 344.732$, $p < 0.001$, $\eta_p^2 = 0.97$). This was followed up by two post-hoc t-tests (sig. threshold Bonferroni-corrected to 0.025), which showed that the mean cross-correlation index in the high γ -power condition (High: $M = 0.69$, $SD = 0.19$) was significantly higher than that in the low γ -power condition (Low: $M = 0.44$, $SD = 0.17$; $t = 18.94$, $p < 0.001$), and not significantly different compared to the medium γ -power condition (Medium: $M = 0.71$, $SD = 0.16$; $t = 1.24$, $p = 0.23$).

One interesting question that arises from looking at cross-correlation index (CCI) as

a function of excitatory drive and γ -power in an example cell (Figure 4.6B) is: is the slope of the apparent growth in spike fidelity over increasing excitatory stimulation significantly modulated by inhibitory γ -power? To explore this question, I fit a series of linear regressions to the relationship between CCI and excitatory current step number, and extracted the I-CCI slope for each level of γ -power, starting with the minimal current step required to evoke at least one spike in each condition.

Given the generally low firing rates in the flat non-rhythmic control IPSP condition (see Figure 4.5D) and the tendency for any spiking to occur at only a few of the highest excitation steps (see Figure 4.5B for example), slopes were not estimated for this condition. Cells for which the slope estimates explained $<50\%$ of variance in the cross-correlation values in any of the three γ -power conditions were discarded from this analysis, yielding a total of 12 cells for which cross-correlation slope estimates were available for high, medium, and low inhibitory γ -power conditions. The obtained I-CCI slope values are shown in Figure 4.6C, individually for each cell (grey) and as a cohort average (purple).

Running a repeated-measures one-way ANOVA with IPSP γ -power (Flat, Low, Medium, and High) as the repeated-measures variable revealed that IPSP γ -power had a significant main effect on mean I-CCI slope ($F(1.23, 22) = 6.70, p = 0.017, \eta_p^2 = 0.38$). A pair of follow-up post-hoc t-tests (sig. threshold Bonferroni-corrected to 0.025) revealed that the average I-CCI slope in the high γ -power condition (High: $M = 0.072, SD = 0.041$) was not significantly different from that in the low γ -power condition (Low: $M = 0.051, SD = 0.028; t = -2.12, p = 0.058$), however significantly higher than that in the medium γ -power condition (Medium: $M = 0.044, SD = 0.030; t = 3.25, p = 0.008$). These findings indicate that spike timing reliability increases at a faster rate as a function of input strength during inhibitory inputs of high γ -power.




Figure 4.5: Effects IPSG γ -power on principal cell response properties in L2/3 of S1 (n=16). **A** shows representative recordings from a putative pyramidal cell (top) being administered IPSGs (bottom) of high (left), medium (middle), and low (right) γ -power, while concurrently receiving excitatory stimulation with a 260 pA square current pulse. **B** shows the input-output curves of the same cell, for conditions of high, medium and low γ -power, and the flat IPSG control. **C** shows average input resistance (R_i ; top) and resting membrane potential (V_m ; bottom) across recorded cells in the first and final blocks of the experiment. **D** shows maximal firing rates exhibited at each γ -power condition. **E** shows input-output (I-O) slopes for each γ -power condition. **F** shows rheobase for each γ -power condition, defined as the minimum excitatory current required for a single AP, and expressed as a % of the input yielding the maximal spike rate in each experiment. ** *D-F show individual cells' results in grey, and cohort averages in colour.*




Figure 4.6: Effects of IPSG γ -power on principal cell spike fidelity (n=16). **A** shows the mean cross-correlation index (CCI) in spike timing, aggregated across excitatory levels for conditions of high, medium and low γ -power, and the flat IPSG control. **B** shows the relationship between excitatory drive and CCI in an example putative pyramidal cell, segregated according to conditions of high, medium and low γ -power (flat IPSG control values unavailable due to low firing rates in this condition). **C** shows input versus CCI (I-CCI) slopes for each γ -power condition. ***C shows individual cells' results in grey, and cohort averages in colour.*

4.4 RESULTS: EXPERIMENT 2

4.4.1 RECAP AND OBJECTIVES

The main goal of this experiment was to examine whether the effects on principal cell response properties by co-administration of excitatory current and inhibitory conductance (IPSG) oscillating at 40 Hz can be replicated using IPSGs of different frequencies. In other words, is the γ -rhythm special in its ability to induce multiplicative scaling of cell responses and enhancing spike fidelity, as was observed in the previous series of experiments detailed in this chapter?

To investigate this question, 12 putative pyramidal cells (profiled on the basis of electrophysiological characteristics; see Figure 4.1) were patched in L2/3 of S1. Their responses to excitatory currents of increasing amplitudes were examined while dynamic clamp was used to administer rhythmic sine-wave modulated inhibitory postsynaptic conductances (IPSGs) at the following frequencies: 10 Hz, 20 Hz, 40 Hz, 60 Hz, 80 Hz, and 100 Hz. Such broad coverage of frequencies normally ascribed contrasting roles in circuit processing (Bichot et al., 2005; Fries, 2005; Kanai et al., 2015; Myers et al., 2014) allowed me to explore their effects on single cell-level processing. Figure 4.7A demonstrates a representative cell’s spiking behaviour in response to a 180 pA excitatory current injection during application of IPSGs at each examined frequency. Statistical analysis of variables extracted from the cells’ spiking behaviours primarily involved repeated-measures ANOVAs; Greenhouse-Geisser corrected values are reported where a Mauchly’s test indicated a violation of the sphericity assumption. Where results were followed-up by post-hoc t-tests, the significance threshold was corrected to $p \leq 0.01$ (unless otherwise specified), to account for the number of comparisons.

4.4.2 INPUT RESISTANCE AND STABILITY

Whole-cell recordings were acquired from S1 using whole-cell dynamic clamp mode. To quantify cell stability and recording quality over the course of an experiment, I measured input resistance (R_i) and resting membrane potential (V_m) at regular intervals throughout the protocol. Resistance was sampled by examining the membrane potential response to a square current pulse of -80 pA, which was applied at the end of each of the 12 blocks featuring increasingly higher excitatory current amplitudes (see Methods section 4.1.5). Membrane potential was sampled during a 500 ms baseline period preceding the excitatory current injection of each trial, and subsequently segregated by block number to track average membrane potential as the experiment unfolded. To evaluate cell stability over the duration of an experiment, average resistance and membrane potential were compared between the first and final blocks of each experiment (shown Figure 4.7C).

Two t-tests revealed that input resistance did not significantly differ between the first and final blocks of each experiment (R_i first block: $M = 302.30 \text{ M}\Omega$, $SD = 297.15$; R_i final

block: $M = 392.69 \text{ M}\Omega$, $SD = 286.37$; $t = 0.605$, $p = 0.56$); and neither did resting membrane potential (V_m first block: $M = -70.74 \text{ mV}$, $SD = 0.61$; V_m final block: $M = -69.94 \text{ mV}$, $SD = 1.51$; $t = 1.62$, $p = 0.13$). A number of cells did not survive to the end of the stimulation protocol. Cells were only included in the analysis if they provided data for >8 out of 12 levels of excitation examined in the protocol.

4.4.3 γ -FREQUENCY INHIBITION INCREASES PRINCIPAL CELL MAX. FIRING RATES

First, I set out to examine whether the maximum achievable spike rate, on average, differs between conditions of varying inhibitory postsynaptic conductance (IPSG) frequencies, with reference to the 40 Hz rhythm. Maximum spike rate was extracted for each individual cell and IPSG frequency. Figure 4.7D demonstrates these for each individual cell (grey) as well as the cohort average (red). Firing rate maxima were submitted to a one-way repeated-measures ANOVA with IPSG Frequency (10 Hz, 20 Hz, 40 Hz, 60 Hz, 80 Hz, and 100 Hz) as the repeated-measures variable.

The ANOVA revealed a significant main effect of IPSG Frequency ($F(2.48, 50) = 46.51$, $p < 0.001$, $\eta_p^2 = 0.82$). A series of post-hoc t-tests comparing max. firing rates for all IPSG frequencies to the 40 Hz condition ($M = 25.57$, $SD = 2.11$) revealed that maximal firing rates were significantly higher in the 40 Hz condition compared to the 10 Hz ($M = 13.46$, $SD = 4.35$, $t = 9.79$, $p < 0.001$) and 20 Hz ($M = 16.30$, $SD = 4.62$, $t = 6.96$, $p < 0.001$) conditions, with further comparisons failing to survive the corrected significance threshold, although showing a trend for higher values compared to 40 Hz (40 Hz vs 60 Hz: $t = 2.98$, $p = 0.014$; 40 Hz vs 80 Hz: $t = 2.88$, $p = 0.016$; 40 Hz vs 100 Hz: $t = 1.38$, $p = 0.19$). As such, the evidence suggests that increasing IPSG frequency up to 40 Hz yields significant firing rate increases during excitation, with no further significant changes taking place as IPSG frequency is increased further.

4.4.4 γ -FREQUENCY INHIBITION INCREASES PRINCIPAL CELL RESPONSE SCALING

Next, I examined whether differing frequencies of inhibitory conductance influence principal cells by modulating the relationship between excitatory drive and cell output. Thus, I explored whether putative pyramidal cells' input-output curves (i.e. firing rates as a function of injected excitatory current amplitude) are sensitive to manipulations in IPSP frequency. Figure 4.7C demonstrates the input-output relationship obtained at each inhibitory conductance frequency, for one representative cell from the cohort.

The nature of this relationship was quantified for each cell by fitting a linear regression to the input-output curve, starting with the first non-zero spike rate and ending with maximal spike frequency, separately for each IPSP frequency condition. The obtained input-output slope estimates are shown in Figure 4.7F, individually for each cell (grey) and as a cohort average (green).

Individual cells' slopes were submitted to a one-way repeated-measures ANOVA with IPSP Frequency (10 Hz, 20 Hz, 40 Hz, 60 Hz, 80 Hz, and 100 Hz) as the repeated-measures variable. The ANOVA revealed a significant main effect of IPSP Frequency ($F(2.52, 55) = 14.67$, $p < 0.001$, $\eta_p^2 = 0.57$), which was explored by a series of post-hoc t-tests comparing slopes obtained in all IPSP frequency conditions to those of the 40 Hz condition ($M = 0.21$, $SD = 0.12$). The tests revealed that input-output slopes were significantly higher in the 40 Hz condition compared to the 10 Hz ($M = 0.084$, $SD = 0.06$, $t = 4.19$, $p = 0.001$) and 20 Hz ($M = 0.11$, $SD = 0.086$, $t = 3.22$, $p = 0.008$) conditions, while failing to produce any significant differences in slope compared to other, higher frequencies (40 Hz vs 60 Hz: $t = 1.07$, $p = 0.31$; 40 Hz vs 80 Hz: $t = 1.76$, $p = 0.11$; 40 Hz vs 100 Hz: $t = 2.69$, $p = 0.21$). This evidence indicates that increasing IPSP frequency produces a statistically meaningful multiplicative scaling of principal cell responses up to 40 Hz, with further marginal increases in scaling (as reflected in growing input-output slope values as frequencies go up to 100 Hz; Figure 4.7E) failing to be statistically significant.

4.4.5 FREQUENCY OF RHYTHMIC INHIBITION INFLUENCES RHEOBASE

Next, I assessed the effect of IPSP frequency on rheobase, defined as the minimum excitatory current required for the cell to generate at least one action potential over the course of a single trial. Rheobase was extracted separately for each cell and IPSP frequency condition, and expressed as a % of the excitatory current which produced the maximal spike rate in each experiment. In some cases, the lowest excitatory current pulse delivered was sufficient to evoke >1 AP at a particular IPSP frequency, making it impossible to empirically determine the rheobase for that specific frequency. In such situations, rheobase was analytically determined by running a linear regression on the relationship between firing rate and excitatory current amplitude for the particular cell and IPSP frequency in question. Rheobase was then defined as the current amplitude at which the projected firing rate would be just above 0, and, similarly to empirically determined rheobase values, it was converted into a % of the current which produced the maximal spike rate. This extrapolation procedure occasionally yielded negative values, as seen in Figure 4.7F, which shows individual cells' rheobases (grey) and the cohort average in each IPSP frequency condition (blue).

Individual cells' rheobase values were submitted to a one-way repeated-measures ANOVA with IPSP Frequency (10 Hz, 20 Hz, 40 Hz, 60 Hz, 80 Hz, and 100 Hz) as the repeated-measures variable. The ANOVA revealed a significant main effect of IPSP Frequency ($F(2.57, 55) = 6.99, p = 0.008, \eta_p^2 = 0.39$), which was explored by a series of post-hoc t-tests comparing rheobases in all IPSP frequency conditions to those of the 40 Hz condition ($M = 35.69, SD = 28.88$). These revealed that none of the mean rheobase estimates obtained under differing IPSP frequency conditions differed significantly from those in the 40 Hz condition (10 Hz vs 40 Hz: $t = 2.065, p = 0.063$; 20 Hz vs 40 Hz: $t = 1.96, p = 0.075$; 60 Hz vs 40 Hz: $t = 1.43, p = 0.18$; 40 Hz vs 80 Hz: $t = 2.05, p = 0.065$; 40 Hz vs 100 Hz: $t = 3.04, p = 0.011$). As such, the results indicate that IPSP frequency has a significant bearing on cells' baseline sensitivity (as evidenced by the main frequency effect in the ANOVA), however the follow-up t-tests suggest that this cell cohort may lack the power to point out specific differences to a significance level which survives threshold correction.




Figure 4.7: Effects of IPSG frequency on principal cell response properties in L2/3 of S1 (n=12). **A** shows representative recordings from a putative pyramidal cell being administered IPSGs of the following frequencies, from left to right: 10 Hz, 20 Hz, and 40 Hz (top), 60 Hz, 80 Hz, and 100 Hz (bottom), while concurrently receiving excitatory stimulation with a 180 pA square current pulse. **B** shows average input resistance (R_i ; top) and resting membrane potential (V_m ; bottom) across recorded cells in the first and final blocks of the experiment. **C** shows the input-output curve of the same cell as shown in **A**, segregated according to IPSG frequency condition. **D** shows maximal firing rates exhibited at each IPSG frequency condition. **E** shows input-output (I-O) slopes for each IPSG frequency condition. **F** shows the rheobase for each IPSG frequency condition, defined as the minimum excitatory current required for a single AP, and expressed as a % of the input yielding the maximal spike rate in each experiment. Where rheobase could not be empirically determined (if min. administered excitation exceeded AP threshold for a given IPSG frequency) rheobase was determined by a linear regression on the relationship between firing rate and excitatory current magnitude, and defined as the current yielding >0 spikes. ** *D-F* show individual cells' results in grey, and cohort averages in colour.

4.4.6 FREQUENCY OF RHYTHMIC INHIBITION AFFECTS PRINCIPAL CELL SPIKE FIDELITY

Next, I examined whether variability in the frequency at which rhythmic inhibitory conductance is delivered modulates the reliability of spike timing in pyramidal cells across repetitions of the same excitatory input. The experimental paradigm was well-suited to pursuing this question, as each level of excitatory current pulse was repeated for a total of 5 times along with each of the six inhibitory conductance frequencies. This enabled me to examine the replicability of the principal cells' spike patterns across stimulus repetitions, by calculating a cross-trial correlation index. This index was obtained by averaging spike timing cross-correlations, obtained at a 0 ms time-lag, between every possible pairing of trials at each level of excitation (normalised by spike rate). Figure 4.8A demonstrates the mean cross-correlation index (CCI) observed under each IPSPG frequency condition, aggregated across excitatory input strengths.

Whether frequency made any difference to spike fidelity was tested by submitting individual cells' mean CCIs to a one-way repeated-measures ANOVA with IPSPG Frequency (10 Hz, 20 Hz, 40 Hz, 60 Hz, 80 Hz, and 100 Hz) as the repeated-measures variable. The ANOVA revealed a significant main effect of IPSPG Frequency ($F(2.51, 55) = 17.66$, $p < 0.001$, $\eta_p^2 = 0.62$). This effect was explored by running a series of post-hoc t-tests comparing mean CCI at an IPSPG of 40 Hz with those of all other frequencies. The tests revealed that the mean cross-correlation index at 40 Hz ($M = 0.34$, $SD = 0.20$) was significantly higher compared to those at all other frequencies (10 Hz: $M = 0.18$, $SD = 0.11$, $t = 4.42$, $p = 0.001$; 60 Hz: $M = 0.26$, $SD = 0.14$, $t = 2.96$, $p = 0.01$; 80 Hz: $M = 0.20$, $SD = 0.096$, $t = 3.43$, $p = 0.006$; 100 Hz: $M = 0.12$, $SD = 0.06$, $t = 4.66$, $p = 0.001$), apart from 20 Hz (20 Hz: $M = 0.29$, $SD = 0.15$, $t = 1.66$, $p = 0.13$).

I decided to supplement these findings by examining whether the apparent increase in cross-correlation as a function of increasing excitatory drive (as illustrated Figure 4.6B) is significantly affected by the frequency of concurrently administered IPSPG. To explore this question, I fit a series of linear regressions to the relationship between CCI and excitatory current step number, for each IPSPG frequency condition, starting with the minimal current step required to evoke at least a single spike in each condition. Frequencies for which the resulting slope estimates explained $< 50\%$ of variance in the cross-correlation values were discarded from analysis. The obtained I-CCI slopes are shown in Figure 4.6C, individually for each cell (grey) and as a cohort average (purple).

Individual cells' I-CCI slope values were submitted to a one-way repeated-measures ANOVA with IPSPG Frequency (10 Hz, 20 Hz, 40 Hz, 60 Hz, 80 Hz, and 100 Hz) as the repeated-measures variable. The ANOVA revealed no significant main effect of IPSPG Frequency ($F(5, 55) = 2.26$, $p = 0.07$, $\eta_p^2 = 0.24$). Overall, the findings indicate that spike timing reliability is generally highest when excitation coincides with rhythmic inhibitory conductances at frequencies of 20-40 Hz, with a significant decline in reliability as one moves further away to lower and higher frequencies. The evidence that the increase in reliability with stronger excitatory stimulation occurs at different rates depending on inhibitory conductance frequency, however, did not receive robust support.




Figure 4.8: Effects of IPSG frequency on principal cell spike fidelity (n=12). **A** shows the mean cross-correlation index (CCI) in spike timing, aggregated across levels of excitation delivered simultaneously with IPSGs of each examined frequency. **B** shows the relationship between excitatory drive and CCI in an example putative pyramidal cell, segregated according to IPSG frequency. **C** shows input versus CCI (I-CCI) slopes for each IPSG frequency. ** *C* shows individual cells' results in grey, and cohort averages in colour.

4.5 RESULTS: SINGLE CELL MODEL

4.5.1 RECAP AND OBJECTIVES

In an attempt to understand the physiological mechanisms underlying the observation that increasing the frequency of rhythmic inhibition enhances principal cell response scaling (with statistically significant changes taking place $\leq 40\text{Hz}$; see Figure 4.7E), a pyramidal neuron was simulated according to methods published by Golomb et al. (2006). In short, the model cell consisted of a single compartment at the confluence of the most predominant ionic currents known to flow across the soma and proximal dendrites, detailed in Section 4.1.7. The cell was subjected to stimulation conditions which mimicked the experimental protocol (see Section 4.1.6). This simulated protocol involved pairing each of a sequence of 75 excitatory square current pulses of 2 s duration (each pulse increasing in amplitude by 40 pA/cm^2) with simulated rhythmic inhibition delivered at a range of frequencies (10 Hz, 20 Hz, 40 Hz, 60 Hz, 80 Hz, and 100 Hz). The goal was to begin understanding which variables make the heaviest contribution to the aforementioned growth in response scaling with increasing frequency of inhibition, as well as the observation that increases up to 40 Hz are, statistically, the most relevant, with further frequency increases making smaller contributions.

In the results sections below, I explore the effects of modulating the degrees of i) A-type K^+ channel conductance (G_{K-A}), ii) M-type K^+ channel conductance (G_{K-M}) and iii) voltage-gated Ca^{2+} channel conductance (G_{Ca}), on the relationship between frequency of inhibition and the model cell's input-output properties. At each step, any excitatory current injected into the cell was compounded with a noise term drawn from a normal distribution of $\mu=0$ and $\sigma=1$, based on one of five different random number generators. The data presented below treats the outputs produced under each number generator as a distinct data-point, to better gauge the robustness of any trend in the face of noise.

4.5.2 ROLES OF A-TYPE AND M-TYPE K^+ CURRENTS IN PRINCIPAL CELL RESPONSE SCALING

The main goal of the simulations was to explore the relative importance of the A-type and M-type K^+ currents in the relationship between frequency of inhibition received by the model cell during excitation and the cell's output properties. While it is unclear whether the M-current is present in all principal cells, the existing literature indicates that this current plays an active role in regulating cell output across a range of cell types and brain regions - including hippocampal pyramidal cells (Golomb et al., 1996; Hu, Vervaeke & Storm, 2002; Gu, Vervaeke & Storm, 2005; Madison & Nicoll, 1984), cortical pyramidal (Peng et al., 2017) and stellate (Nigro et al., 2014) cells, multiple types of interneuron (Grigorov et al., 2014; Lawrence et al., 2006; Li et al., 2014), as well as cells outside the forebrain (Bordas, Kovacs & Pal, 2015). The A-current is, similarly, widely observed in the cortex, (Guan et al., 2011), as well as hippocampus (Hoffman et al., 1997), where its

activity in the dendrites regulates spike thresholds and EPSP summation (Biro et al., 2018). As such, both types of current are likely relevant parameters to explore in attempting to understand the mechanistic underpinnings of rhythmic inhibitory influences on cell output.

The stimulation protocol pairing rhythmic inhibition with increasing excitation was simulated for each of two values of G_{K-A} and G_{K-M} , as follows: $G_{K-A}=0$ mS/cm² or 1 mS/cm²; $G_{K-M}=0$ mS/cm² or 1.4 mS/cm². While these conductances can naturally assume a vast range of values between these limits under physiological conditions, I explored the effects of switching these two conductances on or off in a binary fashion, with the aim of working towards a relatively global understanding of the importance of these variables (and in the interests of minimising the number of time-consuming simulations to be run). The minimum and maximum conductance values used in the simulations were not subjected to any fine-tuning, but rather based directly on the lower and upper limits featured in the original Golomb et al. (2006) publication. Sample membrane potential traces of the model cell are shown in Figure 4.9A, featuring conditions of $G_{K-A}=1.4$ mS/cm², $G_{K-M}=0$ mS/cm², $G_{Ca}=0.20$ mS/cm², during co-administration of excitation (step n°38) and rhythmic inhibition at each of the aforementioned frequencies.

For all simulations featured in this section, G_{Ca} was fixed at a value of 0.020 mS/cm². To quantify the input-output properties of the model cell under each of these conditions, I extracted spike rates obtained during each of the 75 excitatory current pulses, separately for each examined frequency of inhibition. The slope of the input-output relationship for each curve was determined by running a linear regression to the subset of data-points on the curve which started with the first non-zero value and ended with the point at which the maximal firing rate was achieved (as in the experimental sections of this chapter). Figure 4.9B-C provides example input-output curves taken from two conditions: $G_{K-A}=1.4$ mS/cm², $G_{K-M}=0$ mS/cm² and $G_{K-A}=1.4$ mS/cm², $G_{K-M}=1$ mS/cm², for each inhibitory frequency.

The resulting input-output slope estimates for each of the four combinations of G_{K-A} and G_{K-M} ($G_{K-A}=0$ mS/cm², $G_{K-M}=0$ mS/cm²; $G_{K-A}=0$ mS/cm², $G_{K-M}=1$ mS/cm²; $G_{K-A}=1.4$ mS/cm², $G_{K-M}=0$ mS/cm²; $G_{K-A}=1.4$ mS/cm², $G_{K-M}=1$ mS/cm²) are shown Figure 4.9D-E. Estimates resulting from each of the five distinct random number generators are shown in grey; the group average is shown in green. As the figure shows, input-output slopes exhibited a positive relationship with frequency of inhibition only under conditions in which the M-type K⁺ current was switched off by setting channel conductance to 0 mS/cm². In cases where G_{K-M} was at its maximal value of 1 mS/cm², the relationship was inverted, and inconsistent with empirical observations from patch-clamp recordings (Figure 4.7E).

It is important to note that input-output curves under conditions of maximal G_{K-M} (Figure 4.9C) initially seem at odds with the result that such conditions yield a negative relationship between inhibitory frequency and I-O slope. Here, firing rates appear to scale more steeply with excitation at higher inhibitory frequencies, similarly to conditions of minimal G_{K-M} . Note, however, that these conditions resulted in firing rates hovering at low values for a substantial number of steps before beginning to scale with excitation. This delay - which appears longer at higher frequencies of inhibition, and is not observed under

conditions of minimal G_{K-M} (as well as experimental conditions) - produced the gradual lowering of I-O slopes with frequency that is summarised in Figure 4.9D-E.

In contrast to the effects of G_{K-M} , the simulations showed that the A-type K^+ current was less important in producing a positive relationship between frequency of inhibition and input-output slope. A positive relationship could be obtained both when $G_{K-A}=0$ mS / cm² and $G_{K-A}= 1.4$ mS/cm², as long as $G_{K-M}= 0$ ms/cm².

4.5.3 ROLE OF VOLTAGE-GATED Ca^{2+} CHANNEL CONDUCTANCE IN PRINCIPAL CELL RESPONSE SCALING

In the previous section, I observed the importance of low M-type current levels in generating a positive relationship between frequency of inhibition and response scaling in a model cell, approximating experimental findings detailed in Results section 4.4.4. Having accordingly fixed G_{K-M} at 0 mS / cm² and G_{K-A} at 1.4 mS/cm² (the latter an arbitrary selection), this section explores the effects of varying Ca^{2+} channel conductance (G_{Ca}), which regulates the rapid flux of calcium ions through voltage-gated calcium channels (I_{Ca}), which in turn triggers fast and slow calcium-activated K^+ currents (I_{K-fast} and I_{K-slow}) which contribute to rapid repolarisation, and slow after-hyperpolarization and spike frequency adaptation respectively (Madison & Nicoll, 1984). It was my reasoning that this was a relevant parameter to explore as I suspected that the same mechanism which contributes to after-hyperpolarization is one which also contributes to the decreasing return on input-output slopes of increasing frequency of inhibition.

Figure 4.10A-B shows sample input-output curves taken under low and high G_{Ca} conditions (A: $G_{Ca} = 0.05$ mS / cm²; B: $G_{Ca}= 0.30$ mS/cm²), while Figure 4.10C-E shows input-output slopes estimated for three values of G_{Ca} (0.05 mS/cm², 0.25 mS/cm², 0.30 mS/cm²), separately for each frequency of inhibition. As can be seen from the I-O slopes as a function of frequency of inhibition (Figure 4.10C-E), the slope-frequency profiles at the higher G_{Ca} values of 0.25 and 0.30 mS/cm² approximate the empirically observed increase in slope as a result of frequency increases ≤ 40 Hz, followed by a diminishing return on further increases in frequency of inhibition (see Figure 4.7E). In contrast, the slope-frequency profile observed at the lowest G_{Ca} setting of 0.05 mS/cm² shows a more dramatic increase in slope as frequency of inhibition exceeds 40 Hz, which does not agree with empirical observations.

Looking at individual input-output curves (Figure 4.10A-B) offers insights into the potential underpinning of this effect. The maximal firing rate achieved within 75 excitatory current pulses is notably higher at $G_{Ca}= 0.05$ mS/cm² compared to $G_{Ca}= 0.30$ mS/cm², with the model cell reaching a max. of 200 Hz and 80 Hz under these two conditions, respectively. This cap on possible firing rate imposed by increased conductance through voltage-gated Ca^{2+} channels may be a significant factor responsible for diminishing the return on increases in frequency of inhibition beyond 40 Hz. Model results shown in Figure 4.11 indicate that this effect is largely underpinned by increased activation of the slow calcium-activated K^+ current at higher frequencies of inhibition. This is likely due to faster inhibition producing rapid membrane potential modulations which enable more spikes to be

produced in short periods of time and thus lead to a rapid increase in conductance through the slow calcium-activated K^+ channel (known to be sensitive to number of preceding spikes; Madison & Nicoll, 1984). As a consequence, the cell periodically enters brief stretches of silence, effectively lowering the number of stimulus-evoked spikes that can be achieved under conditions of rapid inhibition (Figure 4.11B). Simulations run under conditions in which the slow calcium-activated K^+ current was switched off by clamping conductance through the channel (G_{K-slow}) to 0 mS/cm² (Figure 4.11 C-D) demonstrate this current's causal role in generating a diminishing return on input-output slopes as a function of frequency of inhibition beyond 40 Hz. In the absence of the K_{slow} current, input-output slopes scale more steeply with inhibitory frequency, with no apparent cap at 40 Hz, as seen in experimental data.




Figure 4.9: Effects of A-type and M-type K^+ currents on the relationship between frequency of inhibition and the response properties of a model principal cell. **A** shows sample traces from the model pyramidal cell being administered rhythmic inhibition of the following frequencies, from left to right: 10 Hz, 20 Hz, and 40 Hz (top), 60 Hz, 80 Hz, and 100 Hz (bottom), while concurrently receiving excitatory stimulation with square current pulse $n^{\circ}38$. **B-C** show the input-output curves of the model cell for each tested frequency of inhibition and two levels of G_{K-M} : 0 mS/cm^2 and 1 mS/cm^2 while G_{K-A} remained fixed at 1.4 mS/cm^2 . **D-G** show the input-output slopes for each frequency of inhibition and each combination of G_{K-A} and G_{K-M} as follows, in order: D - $G_{K-A} = 0 \text{ mS/cm}^2$ and $G_{K-M} = 1 \text{ mS/cm}^2$; E - $G_{K-A} = 1.4 \text{ mS/cm}^2$ and $G_{K-M} = 1 \text{ mS/cm}^2$; F - $G_{K-A} = 1.4 \text{ mS/cm}^2$ and $G_{K-M} = 0 \text{ mS/cm}^2$; G - $G_{K-A} = 0 \text{ mS/cm}^2$ and $G_{K-M} = 0 \text{ mS/cm}^2$. G_{Ca} was, in all cases, fixed at 0.20 mS/cm^2 . **D-G show model cell results under conditions of five different noise kernels in grey, and the cohort average in green.




Figure 4.10: Effects of Ca^{2+} current on the relationship between frequency of inhibition and the response properties of a model principal cell. **A-B** shows the input-output curves of the model cell for each tested frequency of inhibition, under conditions of low ($G_{Ca} = 0.05 \text{ mS/cm}^2$) and high ($G_{Ca} = 0.30 \text{ mS/cm}^2$) calcium channel conductance. **C-E** show the input-output slopes at each frequency of inhibition for each of three G_{Ca} values, as follows, in order: C - $G_{Ca} = 0.05 \text{ mS/cm}^2$; D - $G_{Ca} = 0.25 \text{ mS/cm}^2$; E - $G_{Ca} = 0.30 \text{ mS/cm}^2$. In these simulations, G_{K-A} and G_{K-M} were fixed at 1.4 mS/cm^2 and 0 mS/cm^2 respectively. ***C-F show model cell results under conditions of five different noise kernels in grey, and the cohort average in green.*




Figure 4.11: Effects of slow calcium-activated K^+ current (K_{slow}) on the relationship between frequency of inhibition and spike output in a model principal cell. **A** shows a membrane potential trace (top) from the model pyramidal cell being administered rhythmic inhibition of 40 Hz frequency while concurrently receiving stimulation with excitatory current step $n^{\circ}45$. The trace below demonstrates the unfolding of slow calcium-activated K^+ conductance ($G_{K\text{-slow}}$) at the same time. **B** shows the same variables for 100 Hz rhythmic inhibition. The light blue areas mark a temporary spike-free period in the membrane potential trace (top) which is shortly preceded by an increase in $G_{K\text{-slow}}$ (bottom) and continues until $G_{K\text{-slow}}$ drops to a level which appears to permit the resumption of spiking. **C** shows a membrane potential trace from the model pyramidal cell being administered rhythmic inhibition of 100 Hz frequency while concurrently receiving stimulation with the same excitatory current as in A and B, after switching off the slow calcium-activated K^+ current by setting conductance through the channel ($G_{K\text{-slow}}$) to 0 mS/cm². **D** shows the model cell's input-output slopes as a function of frequency of inhibition, under conditions of switched off K_{slow} current. ** In these simulations, the following conductances were fixed, as follows: $G_{K\text{-A}} = 1.4 \text{ mS/cm}^2$, $G_{K\text{-M}} = 0 \text{ mS/cm}^2$, and $G_{Ca} = 0.30 \text{ mS/cm}^2$.

4.6 DISCUSSION

In the series of experiments detailed in this chapter, I tested how variations in γ -power of inhibitory conductance trains (IPSGs) delivered at 40 Hz affected the output properties of cortical pyramidal cells being concurrently injected with excitatory current. Here, variations in γ -power were achieved by manipulating the synchrony and rhythmicity of the input trains underlying the IPSG, so as to produce trains with low, medium and high levels of γ -power. Importantly, owing to the fact that perturbing synchronicity and rhythmicity both compromise γ -power while having opposite effects on conductance train variability over time, specific values of these two variables could be combined to generate trains which differed in γ -power, but exhibited similar degrees of temporal variability.

The results complement previous reports that rhythmic excitation at 40 Hz is capable of eliciting stronger multiplicative gain in principal cells compared to excitation at a lower frequency (Sohal et al., 2009), and elaborate on previous findings from our own laboratory. I found that, much like rhythmic excitation at 40 Hz, simulated rhythmic perisomatic inhibition at this frequency is capable of producing a multiplicative gain effect on principal cell responses. This gain effect reflected in steeper input-output slopes, which were significantly greater during application of rhythmic IPSGs compared to flat control conductances, but also significantly scaled with the degree of γ -power in the IPSG. In addition, I found that increasing γ -power subtractively decreased gain, as rheobase significantly increased in tandem with the level of γ -power in the IPSG. Lastly, I also examined the effects of γ -power on principal cell spike fidelity, finding that γ -power has a significant effect.

I extended these findings through a series of experiments comparing the effects of rhythmic inhibitory conductances modulated at 40 Hz and a range of other frequencies. My results suggest that the 40 Hz rhythm might play a special role in boosting the output properties of cortical principal cells, with interesting implications for predictive coding accounts which associate γ -oscillatory activity with the propagation of sensory prediction errors (Bastos et al., 2012).

4.6.1 SIMULATING PERISOMATIC γ -BAND INHIBITION PRODUCES MULTIPLICATIVE INCREASES IN PRINCIPAL CELL GAIN

Examining the slope of change in principal cell firing rates as a function of increasing excitatory current amplitude revealed that input-output slopes were, on average, significantly higher when excitation was co-administered with a rhythmic 40 Hz inhibitory conductance compared to a flat inhibitory control conductance (Figure 4.5E). Moreover, slopes became significantly higher for every step made from low, to medium, and high IPSG γ -power conditions. This suggests that perisomatic inhibition oscillating at 40 Hz is effective at boosting the input-output slopes of principal cells in a manner that is directly sensitive to the integrity of γ .

Experiments comparing the effects of rhythmic IPSGs delivered at 40 Hz to various other frequencies suggested that 40 Hz may, in fact, be the optimal frequency for boosting principal cell input-output slopes. I found that increasing IPSP frequency increased slopes, however only up to 40 Hz, with further increases no longer generating significant return on slope (Figure 4.7E). This indicates that 40 Hz may be the highest modulatory frequency necessary to produce the maximal boost in response properties that can be achieved in principal cells.

The reasons for this effect are likely manifold. One potentially important consideration concerns the inactivation dynamics of voltage-gated sodium channels. Generally, sustained or slowly oscillating depolarization decreases sodium conductance across the membrane due to lingering sodium channel inactivation, making rapid modulation of membrane potential, such as that achieved by 40 Hz inhibition, more effective at consistently triggering spikes during excitation (Bean, 2007; Hodgkin & Huxley, 1952; Noble et al., 1996). This is supported by evidence that faster modulations of membrane potential boost gain (Azouz & Gray, 2000; Henze & Buzsaki, 2001). Thus, the higher rate of membrane modulations produced by rhythmic 40 Hz compared to flat (Exp 1) and lower frequency (Exp 2) inhibition might be responsible for the principal cells' abilities to eventually reach higher maximal firing rates, and thus result in higher input-output slopes. This is exemplified by the finding that increasing γ -power enabled cells to achieve higher maximal firing rates (Figure 4.5C-D), as well as the observation that maximal firing rates associated with modulation frequencies below 40 Hz were significantly lower than those at 40 Hz (Figure 4.7D).

Importantly, the relationship between frequency of inhibition and maximal firing rate – as well as input-output slope – was not linear. Namely, both parameters significantly increased with IPSP frequency only up to 40 Hz. The model of a single principal cell offers some insights into the conditions which yield such a positive frequency-scaling relationship, as well as the plateau at frequencies of inhibition exceeding 40 Hz. Namely, I found that setting conductance of M-type K^+ channels to 0 mS/cm² recapitulated the positive frequency-scaling relationship observed in experiments, with a high conductance value of 1 mS/cm² reversing the trend.

Given the evidence that an active M-type current plays a role in spike frequency accommodation, effectively putting a break on firing rates (Madison & Nicoll, 1984), the importance of closing this channel may lie in this enabling higher firing rates. Indeed, a visual inspection of model cell input-output slopes (Figure 4.9 B-C), where G_{K-M} values are clamped at 0 and 1 mS/cm² (for a fixed value of G_{K-A}) makes clear that the model cell exhibits substantially higher maximal firing rates when the M-type current is shut off; notably conditions of higher inhibitory frequencies resulted in higher absolute firing rates compared to lower frequencies.

Furthermore, and perhaps more importantly, a permissive G_{K-M} was associated with a phenomenon whereby the model cell, upon firing its first AP for a given frequency of inhibition, hovered at low spike rates for a substantial number of excitatory current steps before taking off to higher values (unlike when conductance was set to 0 mS/cm², which

allowed firing rates to instantly scale with stimulation). This contributed to an overall lowering of input-output slopes, most dramatically for higher frequencies of inhibition. Such a delayed scaling of firing rate with excitation may, in part, be due to the M-type current's known effect of increasing AP threshold (Brown & Adams, 1980).

While switching off the M-type K^+ current was important in ensuring a positive relationship between frequency of inhibition and input-output slope, simulations showed that higher voltage-gated calcium channel conductances were important for the plateauing of this relationship at frequencies above 40 Hz. Indeed, Figure 4.10D shows that, at a low G_{Ca} level of 0.05 mS/cm², I-O slopes rise even more dramatically at frequencies above compared to below 40 Hz. In contrast, G_{Ca} values ≥ 0.20 mS/cm² produced a relative slowing of the rise in slope with frequency of inhibition. It is my conjecture that this is due to the resulting calcium flux activating the slow K^+ current (I_{K-slow}) which is known to produce slow after-hyperpolarisation (Madison & Nicoll, 1984), which likely generates a firing rate ceiling. This means that, at a certain point, increases in membrane potential modulation frequency yield no net change in parameters such as maximal firing rate and input-output slope. Indeed, Figure 4.10B-C shows that firing rates are generally substantially lower when G_{Ca} is set is 0.30 mS/cm² compared to 0.05 mS/cm².

4.6.2 RELATIVE IMPORTANCE IN RHEOBASE AND MAXIMAL SPIKE RATE IN INPUT-OUTPUT SLOPES

Evidence suggests that membrane potential noise is quite effective at eliciting action potentials with high temporal fidelity compared to constant excitatory inputs (Mainen & Sejnowski, 1995). This connection between noise and spike fidelity has been previously observed in our laboratory, where 40 Hz IPSP trains with compromised rhythmicity (and consequently, higher temporal variability) were associated with higher degrees of spike fidelity, as indexed by spike timing cross-correlations across presentations of the same excitatory stimulus. My experiments, however, examined the effects of IPSP trains which varied specifically in terms of γ -power, while retaining similar levels of temporal variability. I found that such trains resulted in differing levels of spike fidelity, with mid-to-high γ -power IPSP trains in particular being associated with significantly higher spike fidelity levels compared to IPSPs with low γ -power and arrhythmic flat IPSPs (Figure 4.6B).

Importantly, the second experimental series showed that spike fidelity was highest for IPSPs of 20 Hz and 40 Hz, with significant drop-off with either positive or negative departures in frequency (Figure 4.8B). Combined with the observation that increasing IPSP frequency up to 40 Hz provides the most statistically significant boost to maximal achievable firing rates (Figure 4.7D) and input-output slopes (Figure 4.7E), this evidence reinforces the notion that the γ -frequency of 40 Hz may have special properties in terms of facilitating principal cell processing.

Whereas previous work implicating noisy membrane potential fluctuations in greater spike reliability may have suggested that γ rhythms boost input-output slopes at the cost of degraded temporal precision of information transmission, my results imply that this

may not be the case. Rather, rhythmic perisomatic inhibition in the γ -band may enhance multiplicative gain on principal cell responses as well as the cross-trial reliability of signals.

4.6.3 TOWARDS A BETTER UNDERSTANDING OF γ

My experiments provide solid evidence that rhythmic perisomatic inhibition is capable of boosting the response properties of pyramidal neurons in a manner specific to the γ -band, as multiplicative scaling effects and benefits to spike fidelity were most robust when cells were subjected to IPSPs of 40 Hz frequency. However, while dynamic clamp can effectively simulate the perisomatic inhibitory influence of fast-spiking PV interneurons on target cells, these findings would benefit from experimental validation by studies which directly and selectively recruit these interneurons to produce entirely naturalistic rhythmic inhibition. It may also be of interest to examine the functional outcomes and biological relevance of manipulating interneuronal firing patterns and frequencies. Furthermore, whilst the simulated inhibitory inputs used in these experiments focused on mimicking the outputs of fast-spiking interneurons, known to be critical for network γ -oscillations, I did not attempt to simulate rhythmic dendritic inhibition, such as that produced by somatostatin-positive interneurons (Scheytlejens & Arckens, 2016). These have been found to be highly effective at influencing γ -activity (Veit et al., 2017) and modulating the input-output properties of pyramidal cells (Lovett-Baron et al., 2012; Mehaffey et al., 2005). Thus, future explorations of links between γ -oscillatory activity and single-cell function would benefit from examining the effects of dendritic, as well as somatic, dynamic clamping as well.

Finally, it is important to bridge such research with the conceptual question of the place of γ -oscillations, and oscillations in general, in the grand scheme of predictive coding in sensory systems. Researchers have already successfully combined calcium imaging with animal behaviour to discover single cells responding to mismatch between optic flow and motion-generated visual prediction, residing within the mouse V1 (Keller et al., 2012). Our understanding of such cells and the possible oscillatory dynamics implicated in prediction error coding would certainly be enhanced if we performed intracellular recordings from such cells during sensory experiences such as those involving visuo-motor mismatch. Exploring the effects of disrupting interneuronal dynamics might be a further step in the direction of better understanding how prediction errors and oscillations fit in the puzzle of sensory perception.

GENERAL DISCUSSION

The question of whether the sensory brain generates predictions about its future inputs is, at this point, not highly contested. There is a wealth of evidence from across neuroscientific disciplines to support the idea that predictions take place in sensory cortices, manifesting in pre-activations of neuronal populations tuned to the anticipated stimulus property, which can be gleaned using a range of techniques from human fMRI (Stokes et al., 2009) to calcium imaging in the rodent (Fiser et al., 2016). Moreover, we have robust evidence that brains respond to departures from statistical norms and learned stimulus expectations (El Karoui et al., 2015; Keller et al., 2012; Komatsu et al., 2015; Harms et al., 2014; Sedley et al., 2016). This lends credence to the notion that the two informational variables expected in a brain engaging in predictive coding – namely, sensory predictions and prediction errors – are a biological reality.

However, little is currently known about the possible implementations of predictive coding, and the constraints within which this scheme operates. One recent proposal aiming to give predictive coding some anatomical constraints suggested that sensory predictions and prediction errors ought to be represented by distinct neuronal populations residing in the deep and superficial cortical layers (Kanai et al., 2015), in part owing to their differential propensity to give rise to top-down versus feedforward projections (Barone et al., 2000; Markov et al., 2013). These distinct anatomical sites have, in turn, been associated with oscillatory activities in different frequency bands (Michalareas et al., 2016; Spaak et al., 2012), giving rise to the suggestion that the two parameters of predictive coding might have both distinct laminar origins and spectral channels (Bastos et al., 2012; Kanai et al., 2015). The work described in this thesis aimed to make its contribution to the field by examining this, thus far unexplored, hypothesis.

5.1 INVESTIGATING PREDICTIVE CODING IN THE BEHAVING MOUSE

In the first chapter of this thesis, I aimed to test whether sensory predictions and prediction errors are encoded in the visual cortex of awake behaving mice during stimulus anticipation and processing, and whether these variables are selectively affiliated with distinct cortical layers and spectral properties. To this end, I piloted two 2-Alternative Forced Choice (2-AFC) paradigms in head-fixed water-restricted behaving mice, with the intention of eventually implanting chronic multi-electrode arrays and combining sensory decision-making behaviour with long-term recordings in the primary visual cortex. My goal was to develop a task which incentivised animals to generate sensory expectations whose underlying neuronal activity patterns might be captured in multi-unit and LFP recordings, and classified using pattern analysis.

In the first piloted task, mice were required to use lick the left or right water spout to indicate the locations of oriented visual gratings, which were preceded by auditory cues predicting grating orientation. Upon the animals' achievement of criterion performance (Figures 2.1 and 2.2F), stimulus contrast was increasingly degraded on successive training

sessions to introduce a stimulus detection challenge which might incentivise reliance on the predictive cues. This possibility was suggested by prior evidence that humans pre-activate tuned neuronal populations in their visual cortex when anticipating hard-to-detect stimuli (Stokes et al., 2009). Despite the mice being significantly challenged at low stimulus contrasts, as reflected in slower reaction times and reduced accuracies (Figures 2.1 and 2.2D-F), I obtained no evidence that the animals had, to any degree, relied on the orientation-predictive cues. This manifested in the finding that performance was not significantly affected by whether the cues had correctly predicted stimulus orientation (as was the case on 80% of trials) or not (Figure 2.3).

This outcome counters the expectation that, if animals used predictive information to boost stimulus detection, performance should be impaired on trials featuring incorrect predictions. Instead, I obtained evidence that mice likely based their behaviours on the false belief that cues predicted stimulus location, rather than orientation. This reflected in their significant propensity to lick the same spout that provided a reward during the most recent previous encounter with a particular auditory cue (Figure 2.4).

The location discrimination task offers some new insights into the behavioural capabilities of head-fixed mice. While researchers have previously shown mice to be capable of acquiring the 2-AFC task with water spouts for the purpose of making location judgements in the somatosensory and olfactory domains (Guo et al., 2014), this is, to the best of my knowledge, the first demonstration of mice successfully learning to perform this task with visual stimuli. As the mouse visual system becomes an increasingly popular platform for genetic and physiological dissections (Niell & Stryker, 2008; Pinto et al., 2013; Seabrook et al., 2017), it is arguably important to broaden our knowledge of the species' capabilities within this sensory modality.

It is relevant to note that, while this task was concurrently piloted with two male and two female mice of the same age, only the males successfully acquired the task rules within several weeks of training. As the females exhibited no signs of improvement at a time when the males were exceeding 70% accuracy after a steady trend of increasing performance, their training protocol was terminated after sixteen sessions. The possibility still remains that females might be capable of acquiring the task rules on a different time-scale to males - something which is not entirely clear from the existing behavioural literature (Burgess et al., 2017), and may need to be explored with further long-term studies taking animals' sex into account. Furthermore, the small animal numbers recruited in the task pilot mean that the observed sex differences in learning may have been spurious, due to under-sampling of individuals from a mouse line which exhibits considerable variability of non-trivial parameters such as gene expression patterns (Vedell, Svensson & Churchill, 2011) and emotional traits (Jakovcevski, Schachner & Morellini, 2008), although less is known about specifically cognitive variability. The current study cannot address this possibility, which is a significant shortfall. However, this may be dealt with in the future by piloting tasks with animal numbers large enough for individual variability in performance to recede into the background of behavioural trends as statistical noise rather than giving rise to the appearance of potentially spurious sex differences.

Considering sample size issues, along with the observed lack of predictive cueing effects on animals' performance, the 2-AFC location discrimination task proved inadequate as a basis for exploring predictive coding in the mouse, on a behavioural or neurophysiological level. Taking into account the mice which successfully learnt the task, my failure to find evidence that animals generated orientation-specific sensory predictions stands in stark contrast to earlier findings of stimulus-specific predictive activations in V1, as mice repeatedly traversed a virtual tunnel studded with predictable visual landmarks (Fiser et al., 2016). Importantly, the animals in this study were only required to run along a spherical treadmill to cross the virtual environment, and were given no task demands to explicitly incentivise sensory prediction-making. This evidence for predictive neuronal activity strongly suggests that sensory systems are capable of automatically generating predictions on the basis of an environment's statistical properties, even when such predictions are not behaviourally relevant. The fact that my observations did not agree with this, and instead pointed to mice operating on false beliefs about the meanings of predictive cues, indicates that there are limitations to harnessing active decision-making processes in exploring predictive coding.

It is likely that cognitive requirements - and neuronal activation patterns underlying behaviour - differ fundamentally between a passive task such as virtual tunnel traversal and an instrumental conditioning task in which animals are rewarded for actively making stimulus judgements. Performing a reward-based task inevitably requires the recruitment of reinforcement circuitry such as the dopaminergic ventral tegmental area (Arsenault et al., 2014) and the nucleus accumbens - a prominent target of dopaminergic projections and critical player in behavioural reinforcement (Corbit, Muir & Balleine, 2001; Eaton, Libey & Fetz, 2017). It is important to note dopamine's role in modulating neuronal activity, including its enhancement of thalamocortical afferent excitability (Govindaiah, Wang & Cox, 2010), as well influence on spike rates and signal-to-noise ratios in cortical neurons encoding sensory information (Jacob, Ott & Nieder, 2013). It is therefore possible that the requirement for animals to work for rewards in the tasks explored in this thesis recruited distinct circuit mechanisms to those which produced observable predictive and error-based neuronal responses in the primary visual cortex of mice engaging in passive traversal paradigms (Fiser et al., 2016; Keller et al., 2012). Here, predictive representations in V1 might have been generated by the continuous transmission of spatiotemporal information encoded during past encounters with the virtual tunnel's sensory milieu - a role potentially carried out by the entorhinal cortex (Bellmund, Deuker & Doeller, 2019) and hippocampus (Aqrabawi & Kim, 2018; Howard, 2018; interactions with cerebellum during visuo-motor predictions reported by Onuki et al., 2013).

Other factors likely differentiate active decision-making and passive traversal paradigms, which point to passive tasks being more appropriate for exploring predictive coding. Namely, when mice are asked to make decisions, they appear to be biased towards misinterpreting behaviourally irrelevant regularities in their environment as serving their current behavioural needs. This raises the possibility that the use of decision-making tasks might make it difficult, if not impossible, to evoke and investigate sensory predictions that are free from cognitive intervention. Furthermore, the boundaries of mouse cognitive skills are still, to some degree, unclear, as evidenced by the animals' complete failure to acquire the second behavioural task which I piloted in my first chapter. In this task, mice were trained to judge whether two visual gratings were of the same or different orientations, with the intention of

eventually combining manipulations of the relative frequencies of ‘same’ versus ‘different’ trial types and auditory cues with the goal of generating hemisphere- and orientation-specific predictive representations. However, the animals failed to perform above chance level after four weeks of continuous training, leading me to abort the training protocol.

It is likely that future research into sensory predictive processes in behaving mice would benefit strongly from using tasks which give mice no explicit decision-making demands – such as the virtual tunnel traversal paradigm - and instead rely on the evidently inherent tendency for the visual cortex to predict future stimuli (Fiser et al., 2016). Thus far, traversal paradigms have primarily been combined with calcium imaging, offering evidence for both predictive (Fiser et al., 2016) and error-based (Keller et al., 2012) responses in L2/3 of the mouse visual cortex. While insights gained from calcium imaging are largely limited to the superficial layers, we can expand our understanding of the anatomical and physiological basis of predictive coding by combining traversal paradigms with translaminar probe recordings.

There is, nonetheless, no need to discard explicit decision-making tasks from the list of experimental methods that can be used to probe the neural underpinnings of predictive coding. One of the main lessons that can be derived from this thesis is that decision-making tasks can be problematic primarily when animals are supplied with predictive cues acting on a stimulus dimension that differs from that used to guide behaviour. This is evidenced by my observation that mice given auditory cues predicting stimulus orientation are not affected in their judgements of stimulus location, and moreover appear to erroneously treat the cues as predictive of the behaviourally relevant dimension (i.e. location). If such a dimensional clash is avoided, decision-making tasks could provide a fruitful avenue for researching predictive and error representations in sensory cortices.

It might, for instance, be interesting to examine mouse behaviour and sensory cortical activity in a detection task with threshold grating stimuli, in which cues delivered to other senses (eg. auditory or whisker-based) predict specific stimulus orientations. Many interesting questions can be asked in this context. For instance, what neuronal activity patterns can be observed on ‘catch’ trials where no stimulus is presented but animals report a detection event? Might predictive neuronal representations generated by the cues lead to the false treatment of visual noise as a stimulus of a particular orientation? How do such representations differ from those observed when that stimulus is physically present? Such a task could also offer insights into how orientation prediction errors are encoded in the primary visual cortex as a function of angular deviation from prediction. Is there a universal signal which scales with angle of error, irrespective of its direction? Or does the signal reflect in differential activations of neuronal ensembles tuned to distinct orientations? A stimulus detection task such as this would likely require the use of the go/no-go response harvesting system which, despite its shortcomings (covered in Section 2.5.1), appears to intuitively map onto the cognitive processes involved in signalling a binary event such as stimulus detection (Glickfeld et al., 2013; Wang et al., 2018). By sampling signals from across the cortical layers as mice go about detecting stimuli and responding to expectancy violations, we can begin to gain an understanding of the extent to which predictive and error-based codes are anatomically segregated in the sensory cortex, and begin answering questions about the potential oscillatory mechanisms involved.

5.2 INVESTIGATING PREDICTIVE CODING *IN VITRO*

In the second chapter of this thesis, I examined the possibility of two hypothesised oscillatory channels for predictive coding parameters (Bastos et al., 2012) using two optogenetic stimulation paradigms, *in vitro*. Slices of mouse somatosensory cortex expressing ChR2 under the CaMKII promoter were exposed to streams of stimuli consisting of temporal (Experiment 1) and spatiotemporal (Experiment 2) patterns of photostimulation, peppered with infrequent deviations, in the form of a fluctuation in LED power (Experiment 1) or premature stimulation of a patch of tissue which violated standard sequence dynamics (Experiment 2). My goal was to examine whether oscillatory activity in the α - or γ -bands discriminated between expected and unexpected photostimuli, in a way that would indicate that either or both of these frequencies could map onto the critical prediction error variable.

My experiments provided no evidence that power (Figures 3.4, 3.8 and 3.9) or peak frequency (Figures 3.5, 3.8 and 3.9) in the γ -band differed significantly between stimuli that were expected or unexpected, given stimulus stream statistics. Moreover, I found no evidence that the rate at which either power or peak frequency in the γ -band changed over the time-course of an experiment (~ 1 hour) significantly differed between expected and unexpected stimuli. This is a somewhat surprising outcome if we assume that γ -oscillatory activity reflects prediction error and that statistically common and uncommon stimuli can be expected to suppress prediction error at vastly different rates with increasing stimulation history. That is, a predictable stimulus would rapidly reduce prediction error signals over the course of an experiment due to the number of stimulus encounters that would have accrued to establish it as the dominant object of prediction, while an unexpected stimulus would not.

Indeed, a phenomenon widely observed in the auditory cortex – known as stimulus-specific adaptation (SSA; Angeloni & Geffen, 2018; Ulanovsky et al. 2003) attests to the plausibility of this supposition. Here, neuronal responses to stimuli presented more frequently in an input stream tend to undergo rapid adaptation, in contrast to responses to oddball stimuli, which either adapt more slowly or even exhibit temporal facilitation (Natan et al., 2015; Nelken, 2014). This effect has even been observed in subcortical nuclei such as the medial geniculate body, although inactivation studies imply that these effects are, at least partially, inherited top-down from A1 (Bauerle et al., 2011; Antunes & Malmierca, 2011). In A1 itself, the onset of SSA effects can be as rapid as tens of milliseconds after tone onset (Ulanovsky et al., 2003; Taaseh et al., 2011), which indicates that the effect may be generated in the local circuitry, without necessarily requiring top-down inputs from across the processing hierarchy. Nonetheless, my work suggests that an entirely isolated primary sensory circuit may, after all, not be capable of differential adaptation to expected and unexpected stimuli.

Further undermining the possibility of a connection between γ -oscillatory activity and prediction error in local somatosensory circuitry was the finding that γ -power did not, on average, show significant change over the course of an entire experiment. This leads me

to conclude that a local barrel cortex circuit, devoid of global connectivity, is incapable of responding to violations of an expectancy that might theoretically be established if the recipient of a stimulation stream were capable of tracking its long-term statistics.

Several other findings indicate that local barrel cortex circuitry does not respond to prediction errors simply because it does not represent stimulus predictions in the first place. Aside from the finding that power in the α -band did not significantly discriminate between expected and unexpected stimuli (Figures 3.6, 3.8 and 3.9), Experiment 1 showed that it also did not discriminate between conditions in which the next most likely stimulus was an amplitude deviant versus non-deviant pulse (Figure 3.6). Furthermore, I found that stimulus-driven α -power showed no significant change over the time-course of an experiment, both in the first and second experiment series. These findings suggest that α -band activity, *in vitro*, fails both to distinguish between conditions associated with acutely different sensory predictions, and track the accumulation of sensory prediction as an experiment progresses, and stimulus statistics, and thus ‘expectations’, become clearer.

My evidence for a lacking association between α - and γ -oscillatory activity and predictive coding variables can, of course, be interpreted as indicative of the fact that such an association simply does not exist. However, to conclude this would be to ignore the solid body of evidence documenting links between γ -band synchrony in the intact (and awake) brain and incidences of violated sensory regularities (El Karoui et al., 2015; Sedley et al., 2016; Todorovic et al., 2011). Thus, I am inclined to conclude that the data are more indicative of the likelihood that local barrel circuitry simply does not engage in predictive coding.

It is possible that directly applying spatiotemporal patterns of optogenetic stimulation onto cortical tissue represents a protocol that is too artificial to conclusively answer the question of whether local sensory cortical circuits can represent predictions and errors. Experimental evidence suggests that this is not necessarily the case, as sine-modulated optogenetic stimulus trains delivered directly to mouse S1 have been found capable of functionally supplanting natural whisker-evoked activity patterns during navigation (Sofroniew et al., 2015). Nonetheless, a more naturalistic stimulation protocol may be required to provide a more conclusive test of the predictive properties of local circuits. As an example, some insights may be gained from examining oscillatory dynamics in the primary sensory cortices of animals exposed to stimulus patterns with occasional violations. In A1, for instance, researchers have previously found stronger neuronal responses (Ulanovsky et al., 2003), and increases in LFP γ -band power (El Karoui et al., 2015) during deviant stimulus processing, strongly hinting at prediction error encoding at the earliest cortical stages of sound analysis. Examining whether the same properties are observed if A1 were functionally or structurally disconnected from its higher-level targets might shed some light on the locality of this effect. Thus, isolating a local circuit while retaining the pathways enabling a naturalistic flow of input that cannot be provided *in vitro* might allow us to better understand the anatomical constraints, and oscillatory properties, of predictive coding.

A more naturalistic stimulation protocol - involving stimulus presentation to the awake or anaesthetised animal - would also bypass the issue of cell targeting uncertainty, which was a significant shortcoming of the work described in this thesis. Due to the broad

spread of L2/3 pyramidal cells' dendritic arbours across the barrel column (Lefort et al., 2009), presenting spatially varying optogenetic stimulation within a single column leaves open the possibility that discrete stimulus patches may nonetheless recruit heavily overlapping cell populations as a result of depolarising their dendrites at multiple sites. As such, the question of whether a computational unit of a sensory area such as S1 (i.e. a barrel column) can learn to respond to stimulus regularities and their violations might be best addressed by targetting that unit from the periphery, rather than attempting to directly pattern the activity of its tissue.

5.3 TOWARDS A MECHANISTIC INSIGHT OF γ -OSCILLATIONS AND SINGLE CELL COMPUTATION

If the γ -band is responsible for propagating the primary informational currency of sensation – i.e. prediction error - through cortical hierarchies, it is arguably critical to establish a better understanding of how activity within this frequency band influences the processing properties of its target neurons, on a single-cell level. This was the task of work outlined in Chapter 3. I conducted patch-clamp recordings of putative L2/3 pyramidal cells of somatosensory cortex to examine their stimulus-driven responses to excitatory currents while dynamic clamp was used to concurrently administer inhibitory conductance trains (IPSGs) oscillating at a γ -frequency of 40 Hz. This, in effect, simulated the rhythmic perisomatic inhibition provided by fast-spiking PV interneurons (Freund & Katona, 2007; Lytton & Sejnowski, 1991), known to play a key role in synchronising pyramidal cells and generating network γ -oscillations (Buzsaki & Wang, 2012; Mann et al., 2005).

In the first series of experiments, excitatory square current pulses of increasing amplitudes were paired with 40 Hz IPSGs differing in terms of γ -power (low, mid, and high), while conserving temporal variability, owing to complementary manipulations of rhythmicity and synchronicity (see Section 4.1.5 and Figure 4.3). I observed that rhythmic inhibition produced higher maximal spike rates (Figure 4.5D) and input-output slopes (Figure 4.5F) in principal cells compared to arrhythmic flat inhibition - an increase which scaled in proportion with γ -power in the IPSG. Furthermore, I found that rheobase was significantly lower in all γ -power conditions compared to the arrhythmic IPSG, indicative of γ -frequency inhibition having a permissive effect on firing threshold. Thus, my findings demonstrate that mimicking rhythmic perisomatic inhibition at 40 Hz increases the baseline sensitivity of principal neurons, as well as effectively having a multiplicative gain effect on output that scales with degree of γ -power. At the same time, mid-to-high levels of γ -power in the IPSG were associated with significantly greater cross-trial spike timing reliabilities compared to low γ -power levels (Figure 4.6A). This is indicative of higher levels of inhibitory γ -power providing a boost to response reliability. High IPSG γ -power also produced significantly steeper slopes of the positive relationship between excitatory current amplitude and spike fidelity (Figure 4.6C). In other words, reliability of principal cell firing patterns scaled more rapidly with input strength when cells were subject to rhythmic inhibitory conductances with high γ -power.

In a further series of experiments, I explored whether the precise frequency at which principal cells receive rhythmic perisomatic inhibition matters for their response properties. In other words, is γ -frequency inhibition particularly well-suited to enhancing gain and spike timing reliability in principal cells, or can rhythmic inhibition at other frequencies replicate the effect? I conducted patch-clamp recordings to examine stimulus-driven activity in putative pyramidal neurons in S1 during exposure to inhibitory conductances oscillating at frequencies covering the spectrum between the α - and high γ -band. Results strongly indicated that γ may indeed have special properties.

I found that increasing IPSP frequency up to 40 Hz resulted in pyramidal cells producing higher firing rates, with no significant increases compared to 40 Hz up to, and including, 100 Hz, the highest frequency examined (Figure 4.7D). Furthermore, increasing IPSP frequency - only up to 40 Hz - resulted in higher mean input-output slopes, with further increases in frequency having no significant further impact (Figure 4.7E). Examining spike timing reliability also revealed that the highest reliability was generally obtained when excitation was paired with 20 Hz and 40 Hz IPSPs, with a gradual fall-off in fidelity with increasing distance from those frequencies (Figure 4.8A). The evidence leads me to conclude that inhibitory conductance rhythmicity of 40 Hz may represent the optimal frequency for boosting principal cell processing.

The special multiplicative gain and spike fidelity effects exerted by 40 Hz inhibitory oscillations appear to take place without strong frequency-driven effects on firing threshold. Simulations of the experimental protocol in a model pyramidal cell indicated that the influence of inhibitory frequency on neuronal input-output properties can be observed under conditions of low conductance through M-type K^+ channels, combined with heightened Ca^{2+} ion flux (Figures 4.9 and 4.10). These two factors appeared to tug on firing rates in opposite manners; however the result of the combination of low G_{K-M} and high G_{Ca} was that the model neuron was able to reach higher spike rates with faster rhythmic inhibition, however (importantly) with diminishing returns with continuing increases in frequency. Dissecting the currents taking place in the model cell under varying frequencies of rhythmic inhibition indicated that the effect of high G_{Ca} was underpinned by its activation of the slow Ca^{2+} -activated K^+ current, K_{slow} (Figure 4.11). Specifically, as higher rates of inhibition produced faster membrane potential modulation, and consequently higher instantaneous spike rates, the occurrence of a number of spikes in rapid succession eventually caused K_{slow} to accumulate to a level which temporarily shut off further neuronal activity. As a result, the model cell alternated between modes of high activity and silence, contributing to an overall lowering of firing rate.

It is relevant to consider the functional implications of the evidence that γ -oscillations are particularly well-suited to enhancing the gain of cortical principal cells. Multiplicative changes in gain generally sensitise neurons to small variations in their encoding parameter, such that smaller changes in the parameter value in the input can have a larger impact on cell output. An example of such an effect in sensory cortex is found during attentional states. Depending on the exact stimulus feature animals are instructed to look out for, researchers have documented enhanced gain of orientation-tuned neurons in macaque visual cortex (McAdams & Maunsell, 1999) or directionally-tuned cells in motion-sensitive area MT (Treue and Martinez-Trujillo, 1999). Thus, multiplicative gain control is likely a

domain-general mechanism for priming the relevant neuronal populations to be particularly perceptive of subtle shifts in their variable of interest. This is interesting to contemplate from the perspective of predictive coding, which endows specifically γ -oscillatory activity with the ability to encode sensory prediction errors, proposed to be the main informational variable which propagates through a processing hierarchy. The observed effects of γ -rhythmic inhibition on principal cell response properties point to the alleged prediction error channel having the capacity to induce target cells to be particularly attuned to their input. This perhaps makes sense; unpredicted inputs are reasonably the ones that need accurate registering for the purposes of increasing the quality of future predictions.

In some ways, this conclusion is at odds with our current knowledge of the encoding of predicted versus unpredicted inputs in sensory cortices. While neuronal responses and γ -power are generally observed to be higher during the processing of unpredicted stimuli, the relatively weaker stimulus-driven responses to predicted inputs are nonetheless associated with greater encoding quality, as measured by superior stimulus decoding based on multidimensional variables such as fMRI multi-voxel patterns (Kok et al., 2012; Noudoost et al., 2017) and multi-unit responses (Bell et al., 2016). This evidence, which points to predicted stimuli enjoying a particular encoding quality advantage, does not entirely fit my finding that the γ -rhythm increases principal cell gain (consequently making cells capable of more fine-grained representations) and spike fidelity across stimulus repetitions, if γ is indeed affiliated with encoding unpredicted inputs. Increased spike fidelity is indeed a relevant dimension, as multivariate analysis of voxel or population activity patterns inherently relies on those patterns which are consistently retained over stimulus repetitions, implicating them in encoding as opposed to noise.

Importantly, multivariate analysis often does not consider the temporal precision of signal reliability over stimulus presentations (this information would not be available with multi-voxel patterns in the first place). Rather, it tends to focus on cross-trial consistency of the spatial distribution of graded responses across the swathe of tissue being examined, occasionally with some temporal windowing. Thus, the population-level index of activation reliability captured by multi-voxel pattern analysis differs from the single cell index of reliability. The question of which one is more critical for encoding information – the spatial distribution of activity across neurons versus their individual spike patterns – perhaps goes to the heart of a more general quest to understand the mechanisms by which neurons perform their encoding functions. We appear to be left with a combined uncertainty about the encoding quality of unexpected versus expected stimuli, the potential role played by γ -oscillatory dynamics, as well as the roles of population versus single-unit activities in any potential encoding quality differences. We can begin resolving these uncertainties experimentally. It might be possible to gain some insights from experiments in which multi- and single-unit recordings from the sensory cortices of animals exposed to various stimuli are paired with manipulations of network γ -oscillatory activity, using pharmacological agents such as carbachol (known to reliably evoke γ -oscillations both in cortex: Kilb and Luhman, 2003 and hippocampus: Fellous & Sejnowski, 2000). Exploring the quality of decoding using both population- and single-cell level dynamics under varying oscillatory conditions can begin providing some more specific answers about neuronal encoding of stimuli associated with different prediction errors. Hopefully these insights might trickle down to offer some food for thought on broader questions about how the cortex processes information.

REFERENCES

1. Abdeljalil, J., Hamid, M., Abdel-Mouttalib, O. ... & Serge, P. (2005). The optomotor response: a robust first-line visual screening method for mice. *Vision Research*, 45. DOI: 10.1016/j.visres.2004.12.015.
2. Abraham, N. M., Guerin, D., Bhaukaurally, K. & Carleton, A. (2012). Similar odor discrimination behaviour in head-restrained and freely moving mice. *PLoS ONE*, 7: 12: e51789. DOI: 10.1371/journal.pone.0051789.
3. Agmon, A. & Connors, B. W. (1991). Thalamocortical responses of mouse somatosensory (barrel) cortex in vitro. *Neuroscience*, 41. DOI: 10.1016/0306-4522(91)90333-j.
4. Alink, A., Schwiedrzik, C. M., Kohler, A. ... & Muckli, L. (2010). Stimulus predictability reduces responses in primary visual cortex. *Journal of Neuroscience*, 30. DOI: 10.1523/JNEUROSCI.3730-10.2010.
5. Amado, C., Kovacs, P., Mayer, R. . . & Kovacs, G. (2018). Neuroimaging results suggest the role of prediction in cross-domain priming. *Scientific Reports*, 8. DOI: 10.1038/s41598-018-28696-0.
6. Andermann, M. L., Kerlin, A. M. & Reid, R. C. (2010). Chronic cellular imaging of mouse visual cortex during operant behaviour and passive viewing. *Frontiers in Cellular Neuroscience*, 4. DOI: 10.3389/fncel.2010.00003.
7. Andermann, M. L. & Moore, C. I. (2006). A somatotopic map of vibrissa motion direction within a barrel column. *Nature Neuroscience*, 4. DOI: 10.1038/nn1671.
8. Angeloni, C. & Geffen, M. (2018). Contextual modulation of sound processing in the auditory cortex. *Current Opinion in Neurobiology*, 49. DOI: 10.1016/J.CONB.2017.10.012.
9. Angelucci, A. & Bressloff, P. C. (2006). Contribution of feedforward, lateral and feedback connections to the classical receptive field center and extra-classical receptive field surround of primate V1 neurons. *Progress in Brain Research*, 154. DOI: 10.1016/S0079-6123(06)54005-1.
10. Angelucci, A., Levitt, J. B., Walton, E. J. S. . . . & Lund, J. S. (2002). Circuits for local and global signal integration in primary visual cortex. *Journal of Neuroscience*, 22. DOI: 10.1523/JNEUROSCI.22-19-08633.2002.
11. Antunes, F. M. & Malmierca, M. S. (2011). Effect of auditory cortex deactivation on

- stimulus-specific adaptation in the medial geniculate body. *Journal of Neuroscience*, 31. DOI: 10.1523/JNEUROSCI.1915-11.2011.
12. Agrabawi, A. J. & Kim, J. C. (2018). Hippocampal projections to the anterior olfactory nucleus differentially convey spatiotemporal information during episodic odour memory. *Nature Communications*, 9. DOI:10.1038/s41467-018-05131-6.
 13. Aronoff, R. & Petersen, C. C. H. (2008). Layer, column and cell-type specific genetic manipulation in mouse barrel cortex. *Frontiers in Neuroscience*, 2. DOI: 10.3389/neuro.01.001.2008.
 14. Arsenault, J. T., Rima, S., Stemmann, H. & Vanduffel, W. (2014). Role of the primate ventral tegmental area in reinforcement and motivation. *Current Biology*, 24. DOI: 10.1016/j.cub.2014.04.044.
 15. Auksztulewicz R. & Friston K. (2015) Attentional enhancement of auditory mismatch responses: a DCM/MEG study. *Cerebral Cortex*, 25. DOI: 10.1093/cercor/bhu323.
 16. Auksztulewicz, R. & Friston, K. (2016). Repetition suppression and its contextual determinants in predictive coding. *Cortex*, 80. DOI: 10.1016/j.cortex.2015.11.024.
 17. Azouz, R. & Gray, C. M. (2000). Dynamic spike threshold reveals a mechanism for synaptic coincidence detection in cortical neurons in vivo. *PNAS*, 97. DOI: 10.1073/pnas.130200797.
 18. Baillet, S. (2014). Forward and inverse problems of MEG/EEG. In *Encyclopedia of Computational Neuroscience*. Springer, New York.
 19. Bale, M. R. & Maravall, M. (2018). Organization of sensory feature selectivity in the whisker system. *Neuroscience*, 368. DOI: 10.1016/j.neuroscience.2017.09.014.
 20. Barbosa, L. S. & Kouider, S. (2018). Prior expectation modulates repetition suppression without perceptual awareness. *Scientific Reports*, 8. DOI: 10.1038/s41598-018-23467-3.
 21. Barone, P., Batardiere, A., Knoblauch, K. & Kennedy, H. (2000). Laminar distribution of neurons in extrastriate areas projecting to visual areas V1 and V4 correlates with the hierarchical rank and indicates the operation of a distance rule. *Journal of Neuroscience*, 20. DOI: 10.1523/JNEUROSCI.20-09-03263.2000.
 22. Barron, H. C., Garvert, M. M. & Behrens, T. E. J. (2016). Repetition suppression: a means to index neural representations using BOLD? *Philosophical Transactions of the Royal Society B*. 371, DOI: 10.1098/rstb.2015.0355.
 23. Barry, R. J., Rushby, J. A., Johnstone, S. J. ... & Lawrence, C. A. (2004). Event related potentials in the auditory oddball as a function of EEG alpha phase at stimulus onset. *Clinical Neurophysiology*, 115. DOI: 10.1016/j.clinph.2004.06.004.

24. Bartels, A., & Zeki, S. (2000). The architecture of the colour centre in the human visual brain: new results and a review. *The European Journal of Neuroscience*, 12. DOI: 10.1046/j.1460-9568.2000.00905.x.
25. Bartolo, M. J., Gieselmann, M. A., Vuksanovic, V. ... & Thiele, A. (2011). Stimulus-induced dissociation of neuronal firing rate and local field potential gamma power and its relationship to the resonance blood-oxygenation level-dependent signal in macaque primary visual cortex. *European Journal of Neuroscience*, 34. DOI: 10.1111/j.1460-9568.2011.07877.x.
26. Bartos, M., Vida, I. & Jonas, P. (2007). Synaptic mechanisms of synchronized gamma oscillations in inhibitory interneuron networks. *Nature Reviews Neuroscience*, 8. DOI: 10.1038/nrn2044.
27. Bastos, A. M., Usrey, W. M., Adams, R. A. ...& Friston, K. J. (2012). Canonical microcircuits for predictive coding. *Neuron*, 76. DOI: 10.1016/j.neuron.2012.10.038.
28. Bastos, A. M., Vezoli, J., Bosman, C. A. ...& Fries, P. (2015). Visual areas exert feedforward and feedback influences through distinct frequency channels. *Neuron*, 85. DOI: 10.1016/j.neuron.2014.12.018.
29. Bauerle, P., von der Behrens, W., Kossel, M. & Gaese, B. H. (2011). Stimulus-specific adaptation in the gerbil primary auditory thalamus is the result of a fast frequency-specific habituation and is regulated by the corticofugal system. *Journal of Neuroscience*, 31. DOI: 10.1523/JNEUROSCI.5814-10.2011.
30. Bauer, M., Stenner, M. P., Friston, K. J. & Dolan, R. J. (2014). Attentional modulation of alpha/beta and gamma oscillations reflect functionally distinct processes. *Journal of Neuroscience*, 34. DOI: 10.1523/JNEUROSCI.3474-13.2014.
31. Bean, B. P. (2007). The action potential in mammalian central neurons. *Nature Reviews Neuroscience*, 8. DOI: 10.1038/nrn2148.
32. Bedard, C., Kroger, H. & Destexhe, A. (2004). Modeling extracellular field potentials and the frequency-filtering properties of extracellular space. *Biophysical Journal*, 86. DOI: 10.1016/S0006-3495(04)74250-2.
33. Bedard, C., Kroger, H. & Destexhe, A. (2006). Does the 1/f frequency scaling of brain signals reflect self-organized critical states. *Physical Review Letters*, 97: 118102. DOI: 10.1103/PhysRevLett.97.118102.
34. Bell, A. H., Summerfield, C., Morin, E. L. ...& Ungerleider, L. G. (2016). Encoding of stimulus probability in macaque inferior temporal cortex. *Current Biology*, 26. DOI: 10.1016/j.cub.2016.07.007.

35. Bellmund, J. L. S., Dueker, L. & Doeller, C. F. (2019). Mapping sequence structure in the human lateral entorhinal cortex. *eLIFE*: e45333 DOI: 10.7554/eLife.45333.
36. Berditchevskaia, A., Caze, R. D. & Schultz, S. R. (2016). Performance in a GO/NOGO perceptual task reflects a balance between impulsive and instrumental components of behaviour. *Scientific Reports*, 6: 27389. DOI: 10.1038/srep27389.
37. Berezovsky, V. K., Nassi, J. J. & Born, R. T. (2011). Segregation of feedforward and feedback projections in mouse visual cortex. *Journal of Comparative Neurology*, 519. DOI: 10.1002/cne.22675.
38. Berger, H. (1929). *Über das Elektrenkephalogramm des Menschen* (On the human electroencephalogram). *Archiv f. Psychiatrie u. Nervenkrankheiten* 87.
39. Bertram, R., Butte, M. J., Kiemel, T. & Sherman, A. (1995). Topological and phenomenological classification of bursting oscillations. *Bulletin of mathematical Biology*, 57. DOI: 10.1007/BF02460633.
40. Bichot, N. P., Rossi, A. F. & Desimone, R. (2005). Parallel and serial neural mechanisms for visual search in macaque area V4. *Science*, 308. DOI: 10.1126/science.1109676.
41. Bickford, M. E. (2015). Thalamic circuit diversity: modulation of the driver / modulator framework. *Frontiers in Neural Circuits*, 9. DOI: 10.3389/fncir.2015.00086.
42. Bird, C. M., Berens, S. C., Horner, A. J. & Franklin, A. (2014). Categorical encoding of color in the brain. *PNAS*, 111. DOI: 10.1073/pnas.1315275111.
43. Biro, A. A., Bremaud, A., Falck, J. & Ruiz, A. J. (2018). A-type K⁺ channels impede supralinear summation of clustered glutamatergic inputs in layer 3 neocortical pyramidal neurons. *Neuropharmacology*, 140. DOI: 10.1016/j.neuropharm.2018.07.005.
44. Bliss, T. V. & Collingridge, G. L. (1993). A synaptic model of memory: Long-term potentiation in the hippocampus. *Nature*, 361. DOI: 10.1038/361031a0.
45. Bliss, T. V. & Lomo, T. (1973). Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *Journal of Physiology*, 232. DOI: 10.1113/jphysiol.1973.sp010273.
46. Bondar, I. V., Leopold, D. A., Richmond, B. J. ...& Logothetis, N. (2009). Long-term stability of visual pattern selective responses in monkey temporal lobe neurons. *PLoS ONE*: e8222. DOI:10.1371/journal.pone.0008222.
47. Bonhoeffer, T. & Grinvald, A. (1991). Iso-orientation domains in cat visual cortex are arranged in pinwheel-like patterns. *Nature*, 353. DOI: 10.1038/353429a0.

48. Bonin, V., Histed, M. H., Yurgenson, S. & Reid, R. C. (2011). Local diversity and fine-scale organization of receptive fields in mouse visual cortex. *Journal of Neuroscience*, 31. DOI: 10.1523/JNEUROSCI.2974-11.2011.
49. Bordas, C., Kovacs, A. & Pal, B. (2015) The M-current contributes to high threshold membrane potential oscillations in a cell type-specific way in the pedunculopontine nucleus. *Frontiers in Cellular Neuroscience*, 9. DOI: 10.3389/fncel.2015.00121.
50. Bosman, C. A., Schoffelen, J-M., Brunet, N. ... & Fries, P. (2012). Attentional stimulus selection through selective synchronization between monkey visual areas. *Neuron*, 75. DOI: 10.1016/j.neuron.2012.06.037.
51. Boyden, E. S., Zhang, F., Bamberg, E., Nagel, G. & Deisseroth, K. (2005). Millisecond-timescale, genetically targeted optical control of neural activity. *Nature Neuroscience*, 8. DOI: 10.1038/nn1525.
52. Boynton, G. M. & Finney, E. M. (2003). Orientation-specific adaptation in human visual cortex. *Journal of Neuroscience*, 23. DOI: 10.1167/4.11.42.
53. Bozorghmehr, T., Ardiel, E. L., McEwan, A. H. & Rankin, C. H. (2013). Mechanisms of plasticity in a *Caenorhabditis elegans* mechanosensory circuit. *Frontiers in Physiology*, 4. DOI: 10.3389/fphys.2013.00088.
54. Brillinger, D. R. (1981). *Time series – data analysis and theory*. San Francisco: Holden Day.
55. Brittain, J-S. & Brown, P. (2014). Oscillations and the basal ganglia: motor control and beyond. *Neuroimage*, 85. DOI: 10.1016/j.neuroimage.2013.05.084.
56. Broski-Guerniero, A., Paasch, G-F., Wollstadt, P. ... & Wibral, M. (2017). Information-theoretic evidence for predictive coding in the face-processing system. *Journal of Neuroscience*, 37. DOI: :10.1523/JNEUROSCI.0614-17.2017.
57. Brown, D. A. & Adams, P. R. (1980). Muscarinic suppression of a novel voltage-sensitive K⁺ current in a vertebrate neurone. *Nature*, 283. DOI: 10.1038/283673a0.
58. Brunel, N. & Wang, X. J. (2003). What determines the frequency of fast network oscillations with irregular neural discharges? 1. Synaptic dynamics and excitation-inhibition balance. *Journal of Neurophysiology*, 90. DOI: 10.1152/jn.01095.2002.
59. Brunet, N. M., Bosman, C. A., Vinck, M. ... & Fries, P. (2014). Stimulus repetition modulates gamma-band synchronization in primate visual cortex. *PNAS*, 111. DOI: 10.1073/pnas.1309714111.

60. Buffalo, E. A., Fries, P., Landman, R. . . . & Desimone, R. (2011). Laminar differences in gamma and alpha coherence in the ventral stream. *PNAS*. DOI: 10.1073/pnas.1011284108.
61. Burgess, C. P., Lak, A., Steinmetz, N. A. . . . & Carandini, M. (2017). High-yield methods for accurate two-alternative visual psychophysics in head-fixed mice. *Cell Reports*, 20. DOI: 10.1016/j.celrep.2017.08.047.
62. Burns, S. P., Xing, D. & Shapley, R. M. (2010). Comparisons of the dynamics of local field potential and multiunit activity signals in macaque visual cortex. *Journal of Neuroscience*, 30. DOI: 10.1523/JNEUROSCI.0743-10.2010.
63. Busch, N. A., Dubois, J. & VanRullen, R. (2009). The phase of ongoing EEG oscillations predicts visual perception. *Journal of Neuroscience*, 29. DOI: 10.1523/JNEUROSCI.0113-09.2009.
64. Busch, N. A. & Van Rullen, R. (2010). Spontaneous EEG oscillations reveal periodic sampling of visual attention. *PNAS*, 107. DOI: 10.1073/pnas.1004801107.
65. Busse, L., Ayaz, A., Dhruv, N. T. . . . & Carandini, M. (2011). The detection of visual contrast in the behaving mouse. *Journal of Neuroscience*, 31. DOI: 10.1523/JNEUROSCI.6689-10.2011.
66. Buzsaki, G., Anastassiou, C. A. & Koch, C. (2012). The origin of extracellular fields and currents – EEG, ECoG, LFP and spikes. *Nature Reviews Neuroscience*, 13. DOI: 10.1038/nrn3241.
67. Buzsaki, G. (2002). Theta oscillations in the hippocampus. *Neuron*, 33. DOI: 10.1016/S0896-6273(02)00586-X.
68. Buzsaki, G., Leung, L. W. & Vanderwolf, C. H. (1983). Cellular bases of hippocampal EEG in the behaving rat. *Brain Research*, 287. DOI: 10.1016/0165-0173(83)90037-1.
69. Buzsaki, G. (1986). Hippocampal sharp waves: Their origin and significance. *Brain Research*, 398. DOI: 10.1016/0006-8993(86)91483-6.
70. Buzsaki, G. & Wang, X. J. (2012). Mechanisms of gamma oscillations. *Annual Review of Neuroscience*, 35. DOI: 10.1146/annurev-neuro-062111-150444.
71. Canolty, R. T., Edwards, E., Dalal, S. S. . . . & Knight, R. T. (2006). High gamma power is phase-locked to theta oscillations in human neocortex. *Science*, 313. DOI: 10.1126/science.1128115.
72. Cardin, J. A., Carlen, M., Meletis, K. . . . & Moore, C. I. (2009). Driving fast-spiking cells induces gamma rhythm and controls sensory responses. *Nature*, 459. DOI:

10.1038/nature08002.

73. Chang, E. F., Rieger, J. W., Johnson, K. ... & Knight, R. T. (2010). Categorical speech representation in human superior temporal gyrus. *Nature Neuroscience*, 13. DOI: 10.1038/nn.2641.

74. Cho, K., Kemp, N., Noel, J. ... & Bashir, Z.I. (2000). A new form of long-term depression in the perirhinal cortex. *Nature Neuroscience*, 3. DOI: 10.1038/72093.

75. Connors, B. W. & Gutnick, M. J. (1990). Intrinsic firing patterns of diverse neocortical neurons. *Trends in Neurosciences*, 13. DOI: 10.1016/0166-2236(90)90185-D.

76. Cook, P. B. & McReynolds, J. S. (1998). Lateral inhibition in the inner retina is important for spatial tuning of ganglion cells. *Nature Neuroscience*, 1. DOI: doi.org/10.1038/3714.

77. Corbit, L. H., Muir, J. L. & Balleine, B. W. (2001). The role of the nucleus accumbens in instrumental conditioning: Evidence of a functional dissociation between accumbens core and shell. *Journal of Neuroscience*, 21. DOI:10.1523/JNEUROSCI.21-09-03251.2001.

78. Coultrip, R., Granger, R. & Lynch, G. (1992). A cortical model of winner-take-all competition via lateral inhibition. *Neural Networks*, 5. DOI: 10.1016/S0893-6080(05)80006-1.

79. Cui, Y., Liu, L. D., Khawaja, F. A. ... & Butts, D. A. (2013). Diverse suppressive influences in area MT and selectivity to complex motion features. *Journal of Neuroscience*, 33. DOI: 10.1523/JNEUROSCI.0203-13.2013.

80. Das, S., Sadanandappa, M. K., Dervan, A. ... & Ramaswami, M. (2011). Plasticity of local GABAergic interneurons drives olfactory habituation. *PNAS*, 108. DOI: 10.1073/pnas.1106411108.

81. Daw, N., Rao, Y., Wang, X-F. ... & Yang, Y. (2004). LTP and LTD vary with layer in rodent visual cortex. *Vision Research*, 44. DOI: 10.1016/j.visres.2004.09.004.

82. De Baene, W. & Vogels, R. (2010). Effects of adaptation on the stimulus selectivity of macaque inferior temporal spiking activity and local field potentials. *Cerebral Cortex*, 20. DOI: 10.1093/cercor/bhp277.

83. De Gardelle, V., Stokes, M., Johnen, V. M. ... & Summerfield, C. (2013). Overlapping multivoxel patterns for two levels of visual expectation. *Frontiers in Human Neuroscience*, 7. DOI: 10.3389/fnhum.2013.00158.

84. Dehghani, N., Bedard, C., Cash, S. S. ... & Destexhe, A. (2010). Comparative power spectral analysis of simultaneous electroencephalographic and magnetoencephalographic recordings in humans suggests non-resistive extracellular media. *Journal of Computational*

Neuroscience, 29. DOI: 10.1007/s10827-010-0263-2.

85. Desimone, R. & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, 18. DOI: 10.1146/annurev.ne.18.030195.001205.

86. Desimone, R. & Schein, S. J. (1987). Visual properties of neurons in area V4 of the macaque: sensitivity to stimulus form. *Journal of Neurophysiology*, 57. DOI: 10.1152/jn.1987.57.3.835.

87. Dombek, D. A., Khabbaz, A. N., Collman, F. ...& Tank, D. W. (2007). Imaging large-scale neural activity with cellular resolution in awake, mobile mice. *Neuron*, 56. DOI: 10.1016/j.neuron.2007.08.003.

88. Dougherty, K., Cox, M. A., Ninomiya, T. ...& Maier, A. (2017). Ongoing alpha activity in V1 regulates visually driven spiking responses. *Cerebral Cortex*, 27. DOI: 10.1093/cercor/bhv304.

89. Douglas, R. J. & Martin, K. A. (1991). A functional microcircuit for cat visual cortex. *Journal of Physiology*, 440. DOI: 10.1113/jphysiol.1991.sp018733

90. Douglas, R. M., Neve, A., Quittenbaum, J. P. ...& Prusky, G. T. (2006). Perception of visual motion coherence by rats and mice. *Vision Research*, 46. DOI: 10.1016/j.visres.2006.02.025.

91. Drager, U. C. & Olsen, J. F. (1981). Ganglion cell distribution in the retina of the mouse. *Investigative Ophthalmology & Visual Science*. 20. DOI: NA.

92. Dubey, A. & Ray, S. (2016). Spatial spread of local field potential is band-pass in the primary visual cortex. *Journal of Neurophysiology*, 116. DOI: 10.1152/jn.00443.2016.

93. Duhamel, J. R. R., Colby, C. L. & Goldberg, M. E. (1992). The updating of the representation of visual space in parietal cortex by intended eye movements. *Science*, 255. DOI: 10.1126/science.1553535.

94. Dylla, K. V., Raiser, G., Galizia, C. G. & Szyszka, P. (2017). Trace conditioning in *Drosophila* induces associative plasticity in mushroom body Kenyon cells and dopaminergic neurons. *Frontiers in Neural Circuits*, 11. DOI: 10.3389/fncir.2017.00042.

95. Eaton, R. W., Libey, T. & Fetz, E. E. (2017). Operant conditioning of neural activity in freely behaving monkeys with intracranial reinforcement. *Journal of Neurophysiology*, 117. DOI:10.1152/jn.00423.2016.

96. Edden, R. A. E., Muthukumaraswamy, S. D., Freeman, T. C. A. & Singh, K. D. (2009). Orientation discrimination performance is predicted by GABA concentration and gamma oscillation frequency in human primary visual cortex. *Journal of Neuroscience*, 50. DOI: 10.1523/JNEUROSCI.4426-09.2009.

97. El Karoui, I., Kung, J-R., Sitt, J. ... & Naccache, L. (2015). Event-related potential, time-frequency, and functional connectivity facets of local and global auditory novelty processing: an intracranial study in humans. *Cerebral Cortex*, 25. DOI: 10.1093/cercor/bhu143.
98. Elul, R. (1971). The genesis of the EEG. *International Review of Neurobiology*, 15. DOI: 10.1016/S0074-7742(08)60333-5.
99. Engel, A. K., Konig, P., Gray, C. M. & Singer, W. (1990). Stimulus-dependent neuronal oscillations in cat visual cortex: inter-columnar interaction as determined by cross-correlation analysis. *European Journal of Neuroscience*, 2. DOI: 10.1111/j.1460-9568.1990.tb00449.x.
100. Esghaei, M., Reza Daliri, M. & Treue, S. (2017). Local field potentials are induced by visually evoked spiking activity in macaque cortical area MT. *Scientific Reports*, 7: 17110. DOI: 10.1038/s41598-017-17372-4.
101. Faes, L., Stramaglia, S. & Marinazzo, D. (2017). On the interpretability and computational reliability of frequency-domain Granger causality. *F1000Research*. doi: 10.12688/f1000research.1269eCollection 2017.
102. Faivre, N., Mudrik, L., Schwartz, N. & Koch, C. (2014). Multisensory integration in complete unawareness: evidence from audiovisual congruency priming. *Psychological Science*, 25. DOI: 10.1177/0956797614547916.
103. Feldmeyer, D., Lubke, J. & Sakmann, B. (2006). Efficacy and connectivity of intracolumnar pairs of layer 2/3 pyramidal cells in the barrel cortex of juvenile rats. *Journal of Physiology*, 575. DOI: 10.1113/jphysiol.2006.105106.
104. Felleman, D. J. & Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex*, 1. DOI: 10.1093/cercor/1.1.1.
105. Fellous, J. M. & Sejnowski, T. J. (2000). Cholinergic induction of oscillations in the hippocampal slice in the slow (0.5-2 Hz), theta (5-12 Hz), and gamma (35-70 Hz) bands. *Hippocampus*, 10. DOI: 10.1002/(SICI)1098-1063(2000)10:2<187::AID-HIPO8>3.0.CO;2-M.
106. Fischer, T. M. & Carew, T. J. (1993). Activity-dependent potentiation of recurrent inhibition: a mechanism for dynamic gain control in the siphon withdrawal reflex of *Aplysia*. *Journal of Neuroscience*, 13. DOI: 10.1523/JNEUROSCI.13-03-01302.1993.
107. Fiser, A., Mahringer, D., Oyibo, H. K. ... & Keller, G. B. (2016). Experience-dependent spatial expectations in mouse visual cortex. *Nature Neuroscience*, 19. DOI: 10.1038/nn.4385.

108. Fishman, Y. I. (2014). The mechanisms and meaning of the mismatch negativity. *Brain Topography*, 27. DOI: 10.1007/s10548-013-0337-3.
109. Flint, R. D., Lindberg, E. W., Jordan, L. R. ... & Slutzky, M. W. (2012). Accurate decoding of reaching movements from field potentials in the absence of spikes. *Journal of Neural Engineering* 9, 046006. DOI: 10.1088/1741-2560/9/4/046006.
110. Fontolan, L., Morillon, B., Liegeois-Chauvel, C. & Giraud, A-L. (2014). The contribution of frequency-specific activity to hierarchical information processing in the human auditory cortex. *Nature Communications*, 5. DOI: 10.1038/ncomms5694.
111. Ford, J. M. & Hillyard, S. A. (1981). Event-related potentials (ERPs) to interruptions of a steady rhythm. *Psychophysiology*, 18. DOI: 10.1111/j.1469-8986.1981.tb03043.x.
112. Foster, J. J., Sutterer, D., Serences, J. T. ... & Awh, E. (2016). The topography of alpha-band activity tracks the content of spatial working memory. *Journal of Neurophysiology*. DOI: 10.1152/jn.00860.2015.
113. Freund, T. F. & Katona, I. (2007). Perisomatic inhibition. *Neuron*, 56. DOI: 10.1016/j.neuron.2007.09.012.
114. Fries, P. (2005). A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends in Cognitive Sciences*, 9. DOI: 10.1016/j.tics.2005.08.011.
115. Fries, P. (2005). Neuronal gamma-band synchronization as a fundamental process in cortical computation. *Annual Review of Neuroscience*, 32. DOI: 10.1146/annurev.neuro.051508.135603.
116. Fries, P. (2016). Rhythms for cognition: communication through coherence. *Neuron*, 88. DOI: 10.1016/j.neuron.2015.09.034.
117. Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society, Biological Sciences*, 360. DOI: 10.1098/rstb.2005.1622.
118. Friston, K. and Kiebel, S. (2009). Predictive coding under the free-energy principle. *Philosophical Transactions of the Royal Society*, 364. DOI: 10.1098/rstb.2008.0300.
119. Fujisawa, S., Amarasingham, A., Harrison, M. T. & Buzsaki, G. (2008). Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nature Neuroscience*, 11. DOI: 10.1038/nn.2134.
120. Gao, E., DeAngelis, G. C. & Burkhalter, A. (2010). Parallel input channels to mouse

primary visual cortex. *Journal of Neuroscience*, 30. DOI: 10.1523/JNEUROSCI.6456-09.2010.

121. Gardner, M. P. H. & Fontanini, A. (2014). Encoding and tracking of outcome-specific expectancy in the gustatory cortex of alert rats. *Journal of Neuroscience*, 34. DOI: 10.1523/JNEUROSCI.1820-14.2014.

122. Garrido, M. I., Friston, K. J., Kiebel, S. J. ... & Kilner, J. M. (2008). The functional anatomy of the MMN: a DCM study of the roving paradigm. *Neuroimage*, 42. DOI: 10.1016/j.neuroimage.2008.05.018.

123. Glanzman, D. L. (2009). Habituation in Aplysia: The Cheshire Cat of neurobiology. *Neurobiology of Learning and Memory*, 92. DOI: 10.1016/j.nlm.2009.03.005.

124. Glickfeld, L. L., Histed, M. H. & Maunsell, J. H. R. (2013). Mouse primary visual cortex is used to detect both orientation and contrast changes. *Journal of Neuroscience*, 33. DOI: 10.1523/JNEUROSCI.3560-13.2013.

125. Golomb, D., Yue, C. & Yaari, Y. (2006). Contribution of persistent Na⁺ current and M-type K⁺ current to somatic bursting in CA1 pyramidal cells: combined experimental and modelling study. *Journal of Neurophysiology*, 96. DOI: 10.1152/jn.00205.2006.

126. Gonzalez, J., Morales, I. S., Villarreal, D. M. & Derrick, B. E. (2014). Low-frequency stimulation induces long-term depression and slow onset long-term potentiation at perforant path-dentate gyrus synapses in vivo. *Journal of Neurophysiology*, 111. DOI: 10.1152/jn.00941.2012.

127. Gotts, S. J., Chow, C. C. & Martin, A (2012). Repetition priming and repetition suppression: A case for enhanced efficiency through neural synchronization. *Cognitive Neuroscience*, 3. DOI: 10.1080/17588928.2012.670617.

128. Gover, T. D. & Abrams, T. W. (2009). Insights into a molecular switch that gates sensory neuron synapses during habituation in Aplysia. *Neurobiology of Learning and Memory*, 92. DOI: 10.1016/j.nlm.2009.03.006.

129. Govindaiah, G., Wang, Y. & Cox, C. L. (2010). Dopamine enhances the excitability of somatosensory thalamocortical neurons. *Neuroscience*, 170. DOI:10.1016/j.neuroscience.2010.08.043.

130. Graf, R., Longo, J. L. & Hughes, Z. A. (2018). The location discrimination reversal task in mice is sensitive to deficits in performance caused by aging, pharmacological and other challenges. *Journal of Psychopharmacology*, 32. DOI: 10.1177/0269881118779383.

131. Grant, K., Bell, C. & Han, V. (1996). Sensory expectations and anti-Hebbian synaptic plasticity in cerebellum-like structures. *Journal of Physiology-Paris*, 90. DOI:

10.1016/S0928-4257(97)81430-1.

132. Gray, C. M., Engel, A. K., Konig, P. & Singer, W. (1990). Stimulus-dependent neuronal oscillations in cat visual cortex: receptive field properties and feature dependence. *European Journal of Neuroscience*, 2. DOI: 10.1111/j.1460-9568.1990.tb00450.x.

133. Gregory, R. L. (1968). Perceptual illusions and brain models. *Proceedings in the Royal Society of London*, 171. DOI: 10.1098/rspb.1968.0071.

134. Gregory, R. L. (1970). *The intelligent eye*. Weidenfeld & Nicolson: London, UK.

135. Grigorov, A., Moskalyuk, A., Kravchenko, M. ...& Fedulova, S. (2014). Kv7 potassium channel subunits and M currents in cultured hippocampal interneurons. *European Journal of Physiology*, 466. DOI: 10.1007/s00424-013-1406-x.

136. Grill-Spector, K., Henson, R. & Martin, A. (2006). Repetition and the brain: neural models of stimulus-specific effects. *Trends in Cognitive Science*, 10. DOI: 10.1016/j.tics.2005.11.006.

137. Grill-Spector, K. & Malach, R. (2001). fMR-adaptation: a tool for studying the functional properties of human cortical neurons. *Acta Psychologica*, 107. DOI: 10.1016/S0001-6918(01)00019-1.

138. Grosz, D. H., Shapley, R. M. & Hawken, M. J. (1993). Macaque V1 neurons can signal 'illusory' contours. *Nature*, 365. DOI: 10.1038/365550a0.

139. Grossman, E. D., Jardine, N. L. & Pyles, J. A. (2010). fMR-adaptation reveals invariant coding of biological motion on the human STS. *Frontiers in Human Neuroscience*, 15. DOI: 10.3389/neuro.09.015.2010.

140. Grothe, I., Neitzel, S. D., Mandon, S. & Kreiter, A. K. (2012). Switching neuronal inputs by differential modulations of gamma-band phase-coherence. *Journal of Neuroscience*, 32. DOI: 10.1523/JNEUROSCI.0890-12.2012.

141. Grover, L. M., Kim, E., Cooke, J. D. & Holmes, W. R. (2009). LTP in hippocampal area CA1 is induced by burst stimulation over a broad frequency range centered around delta. *Learning & Memory*, 16. DOI: 10.1101/lm.1179109.

142. Gu, N., Vervaeke, K., Hu, H. & Storm, J. F. (2005). Kv7/KCNQ/M and HCN/h, but not KCa2/SK channels, contribute to the somatic medium after-hyperpolarization and excitability control in CA1 hippocampal pyramidal cells. *Journal of Physiology*, 566. DOI: 10.1113/jphysiol.2005.086835.

143. Guan, D., Horton, L. R., Armstrong, W. E. & Foehring, R. C. (2011). Postnatal development of A-type and Kv1- and Kv2-mediated potassium channel currents in neocortical pyramidal neurons. *Journal of Neurophysiology*, 105. DOI: 10.1152/jn.00758.2010.

144. Guillery, R. W., Feig, S. L. & Van Lieshout, D. P. (2001). Connections of higher order visual relays in the thalamus: a study of corticothalamic pathways in cats. *Journal of Comparative Neurology*, 438. DOI: 10.1002/cne.1302.
145. Guo, Z. V., Hires, S. A., Li, N. ...& Svoboda, K. (2014). Procedures for behavioral experiments in head-fixed mice. *PLoS ONE*, 9(2): e88678. DOI: 10.1371/journal.pone.0088678.
146. Gutierrez, R., Simon, S. A. & Nicolelis, M. A. L. (2010). Licking-induced synchrony in the taste-reward circuit improves cue discrimination during learning. *Journal of Neuroscience*, 30. DOI: 10.1523/JNEUROSCI.0855-09.2010.
147. Hadjikhani, N., Kveraga, K., Naik, P. & Ahlfors, S. P. (2009). Early (N170) activation of face-specific cortex by face-like objects. *Neuroreport*, 20. DOI: 10.1097/WNR.0b013e328325a8e1.
148. Hagiwara, S., Kusano, K. & Saito, N. (1961). Membrane changes of Onchidium nerve cell in potassium-rich media. *Journal of Physiology*, 155. DOI: 10.1113/jphysiol.1961.sp006640.
149. Haider, B., Schulz, D. P. A., Hausser, M. & Carandini, M. (2016). Millisecond coupling of local field potentials to synaptic currents in the awake visual cortex. *Neuron*, 90. DOI: 10.1016/j.neuron.2016.02.034.
150. Halgren, M., Fabo, D., Ulbert, I. ...& Halgren, E. (2018). Superficial slow rhythms integrate cortical processing in humans. *Scientific Reports*, 8: 2055. DOI: 10.1038/s41598-018-20662-0.
151. Harms, L., Fulham, W. R., Todd, J. ...& Michie, R. T. (2014). Mismatch negativity (MMN) in freely-moving rats with several experimental controls. *PLoS ONE*. DOI: 10.1371/journal.pone.0110892.
152. Harvey, M. A., Saal, H. P., Damman, J. F. & Bensmaia, S. J. (2013). Multiplexing stimulus information through rate and temporal codes in primate somatosensory cortex. *PLoS Biology*. DOI: 10.1371/journal.pbio.1001558.
153. Hasenstaub, A., Shu, Y., Haider, B. ...& McCormick, D.A. (2005). Inhibitory post-synaptic potentials carry synchronized frequency information in active cortical networks. *Neuron*, 47. DOI: 10.1016/j.neuron.2005.06.016.
154. Hebb DO (1949) The organization of behavior: a neuropsychological theory. Mahwah, N.J.: L. Erlbaum Associates.
155. Heinrichs-Graham, E., Hoburg, J. M. & Wilson, T. W. (2018). The peak frequency of

- motor-related gamma oscillations is modulated by response competition. *Neuroimage*, 165. DOI: 10.1016/j.neuroimage.2017.09.059.
156. Helmholtz, H. (1860). *Handbuch der physiologischenoptik*. Vol. 3. New York, NY: Dover.
157. Henrie, J. A. & Shapley, R. (2005). LFP power spectra in V1 cortex: the graded effect of stimulus contrast. *Journal of Neurophysiology*, 94. DOI: 10.1152/jn.00919.2004.
158. Henze, D. A. & Buzsaki, G. (2001). Action potential threshold of hippocampal pyramidal cells in vivo is increased by recent spiking activity. *Neuroscience*, 105. DOI: 10.1016/s0306-4522(01)00167-1.
159. Hill, N. J., Gupta, D., Brunner, P. ... & Schalk, G. (2012). Recording human electrocorticographic (ECoG) signals for neuroscientific research and real-time functional cortical mapping. *Journal of Visualized Experiments*, 26. DOI: 10.3791/3993.
160. Hillyard, S. A., Hink, R. F., Schwent, V. L. & Picton, T. W. (1973). Electrical signs of selective attention in the human brain. *Science*, 182. DOI: 10.1126/science.182.4108.177.
161. Hodgkin, A. L. & Huxley, A. F. (1952). Currents carried by sodium and potassium ions through the membrane of the giant axon of *Loligo*. *Journal of Physiology*, 116. DOI: 10.1113/jphysiol.1952.sp004717.
162. Horton, J. C. & Adams, D. L. (2005). The cortical column: a structure without a function. *Philosophical Transactions of the Royal Society B*, 360. DOI: 10.1098/rstb.2005.1623.
163. Horvath, J. & Winkler, I. (2004). How the human auditory system treats repetition amongst change. *Neuroscience Letters*, 368. DOI: 10.1016/j.neulet.2004.07.004.
164. Howard, M. W. (2018). Temporal and spatial context in the mind and brain. *Current Opinion in Behavioral Sciences*, 17. DOI: 10.1016/j.cobeha.2017.05.022.
165. Hu, H., Vervaeke, K. & Storm, J. F. (2002). Two forms of electrical resonance at theta frequencies, generated by M-current, h-current and persistent Na⁺ current in rat hippocampal pyramidal cells. *Journal of Physiology* 545. DOI: 10.1113/jphysiol.2002.029249.
166. Hubel, D. H. & Wiesel, T. N. (1969). Anatomical demonstration of columns in the monkey striate cortex. *Nature*, 221. DOI: 10.1038/221747a0.
167. Humphries, C., Liebenthal, E. & Binder, J. R. (2011). Tonotopic organization of human auditory cortex. *NeuroImage*, 50. DOI: 10.1016/j.neuroimage.2010.01.046.
168. Hupe, J. M., James, A. C., Payne, B. R. ... & Bullier, J. (1998). Cortical feedback improves discrimination between figure and background by V1, V2 and V3 neurons. *Nature*,

394. DOI: 10.1038/29537.

169. Iemi, L., Chaumon, M., Crouzet, S. M. & Busch, N. A. (2017). Spontaneous neural oscillations bias perception by modulating baseline excitability. *Journal of Neuroscience*, 37. DOI: 10.1523/JNEUROSCI.1432-16.2016.

170. Ikkai, A., Dandekar, S. & Curtis, C. E. (2016). Lateralization in alpha-band oscillations predicts the locus and spatial distribution of attention. *PLoS ONE*. DOI: 10.1371/journal.pone.0154796.

171. Ito, M. & Komatsu, H. (2004). Representation of angles embedded within contour stimuli in area V2 of macaque monkeys. *Journal of Neuroscience*, 24. DOI: 10.1523/JNEUROSCI.4364-03.2004.

172. Jacob, S. N., Ott, T. & Nieder, A. (2013). Dopamine regulates two classes of primate prefrontal neurons that represent sensory signals. *Journal of Neuroscience*, 33. DOI:10.1523/JNEUROSCI.0210-13.2013.

173. Jacobsen, T. & Schroger, E. (2003). Measuring duration mismatch negativity. *Clinical Neurophysiology*, 114. DOI: 10.1016/S1388-2457(03)00043-9.

174. Jakovcevski, M., Schachner, M. & Morellini, F. (2008). Individual variability in the stress response of C57BL/6J male mice correlates with trait anxiety. *Genes, Brain and Behavior*, 7. DOI: 10.1111/j.1601-183X.2007.00345.x.

175. Javadi, A-H., Glen, J. C., Halkiopoulos, S. ...& Spiers, H. J. (2017). Oscillatory reinstatement enhances declarative memory. *Journal of Neuroscience*, 37. DOI: 10.1523/JNEUROSCI.0265-17.2017.

176. Jia, X., Xing, D. & Kohn, A. (2013). No consistent relationship between gamma power and peak frequency in macaque primary visual cortex. *Journal of Neuroscience*, 33. DOI: 10.1523/JNEUROSCI.1687-12.2013.

177. Jin, M., Beck, J. M. & Glickfeld, L. L. (2019). Neuronal adaptation reveals a sub-optimal decoding of orientation tuned populations in the mouse visual cortex. *Journal of Neuroscience*, 10.1523/JNEUROSCI.3172-18.2019.

178. Jones, S., Paul, E. S., Dayan, P. ...& Mendl, M. (2017). Pavlovian influences on learning differ between rats and mice in a counter-balanced Go/NoGo judgement bias task. *Behavioural Brain Research*, 331. DOI: 10.1016/j.bbr.2017.05.044.

179. Joundi, R. A., Jenkinson, N., Brittain, J-S. ...& Brown, P. (2012). Driving oscillatory activity in the human cortex enhances motor performance. *Current Biology*, 22. DOI: 10.1016/j.cub.2012.01.024.

180. Jurjut, O., Georgieva, P., Busse, L. & Katzner, S. (2017). Learning enhances sensory processing in mouse V1 before improving behaviour. *Journal of Neuroscience*, 37. DOI: 10.1523/JNEUROSCI.3485-16.2017.
181. Kaliukhovich, D. A. & Vogels, R. (2011) Stimulus repetition probability does not affect repetition suppression in macaque inferior temporal cortex. *Cerebral Cortex*, 21. DOI: 10.1093/cercor/bhq207.
182. Kanai, R., Komura, Y., Shipp, S. & Friston, K. (2015). Cerebral hierarchies: predictive processing, precision and the pulvinar. *Philosophical Transactions of the Royal Society* 370: 20140169. DOI: 10.1098/rstb.2014.0169.
183. Kanari, L., Ramaswamy, S., Shi, Y. ...& Markram, H. (2019). Objective morphological classification of neocortical pyramidal cells. *Cerebral Cortex*, 29. DOI: 10.1093/cercor/bhy339.
184. Kandel, E., Schwartz, J., Jessel, T., Siegelbaum, S. & Hudspeth, A. (2013). *Principles of Neural Science* (Fifth Edition). McGraw Hill, New York.
185. Kang, K., Shelley, M., Henrie, J. A. & Shapley, R. (2011). LFP spectral peaks in V1 cortex: network resonance and cortico-cortical feedback. *Journal of Computational Neuroscience*, 29. DOI: 10.1007/s10827-009-0190-2.
186. Kanwisher, N., McDermott, J. & Chun, M. M. (1997). The fusiform face area: a module in human extrastriate cortex specialized for face perception. *Journal of Neuroscience*, 17. DOI: 10.1523/JNEUROSCI.17-11-04302.1997.
187. Katzner, S., Nauhaus, I., Benucci, A. ... & Carandini, M. (2009). Local origin of field potentials in visual cortex. *Neuron*, 61. DOI: 10.1016/j.neuron.2008.11.016.
188. Kekoni, J., Hamalainen, H., Saarinen, M. ... & Naatanen, R. (1997). Rate effect and mismatch responses in the somatosensory system: ERP-recordings in humans. *Biological Psychology*, 46. DOI: 10.1016/S0301-0511(97)05249-6.
189. Keller, G. B. et al. (2012). Sensorimotor mismatch signals in primary visual cortex of the behaving mouse. *Neuron* 74. DOI: 10.1016/j.neuron.2012.03.040.
190. Kilb, W. & Luhmann, H. J. (2003). Carbachol-induced network oscillations in the intact cerebral cortex of the newborn rat. *Cerebral Cortex*, 13. DOI: 10.1093/cercor/13.4.409.
191. Klimesch, W., Sauseng, P. & Hanslmayr, S. (2007). EEG alpha oscillations: the inhibition-timing hypothesis. *Brain Research Reviews*, 53. DOI: 10.1016/j.brainresrev.2006.06.003.

192. Koch, G., Ponzo, V., Di Lorenzo, F. ... & Veniero, D. (2013). Hebbian and anti-Hebbian spike-timing-dependent plasticity of human cortico-cortical connections. *Journal of Neuroscience*, 33. DOI: 10.1523/JNEUROSCI.4988-12.2013.
193. Koch, E., Jin, J., Alonso, J. M. & Zaidi, Q. (2016). Functional implications of orientation maps in primary visual cortex. *Nature Communications*, 7: 13529. DOI: 10.1038/ncomms13529.
194. Kok, P., Bains, L. J., van Mourik, T. ... & de Lange, F. P. (2016). Selective activation of the deep layers of the human primary visual cortex by top-down feedback. *Current Biology*, 26. DOI: 10.1016/j.cub.2015.12.038.
195. Kok, P. & de Lange, F. P. (2014). Shape perception simultaneously up- and downregulates neural activity in the primary visual cortex. *Current Biology*, 24. DOI: 10.1016/j.cub.2014.05.042.
196. Kok, P., Failing, M. F. & de Lange, F. P. (2014). Prior expectations evoke stimulus templates in the primary visual cortex. *Journal of Cognitive Neuroscience*, 26. DOI: 10.1162/jocn.a.00562.
197. Kok, P., Jehee, J. F. M. & de Lange, F. P. (2012). Less is more: expectation sharpens representations in the primary visual cortex. *Neuron*, 75. DOI: 10.1016/j.neuron.2012.04.034.
198. Komatsu, M., Takaura, K. & Fujii, N. (2015). Mismatch negativity in common marmosets: whole-cortical recordings with multi-channel electrocorticograms. *Scientific Reports*, 5. DOI: 10.1038/srep15006.
199. Kovacs, G., Iffland, L., Vidnyanszky, Z. & Greenlee, M. W. (2012). Stimulus repetition probability effects on repetition suppression are position invariant for faces. *Neuroimage*. DOI: 10.1016/j.neuroimage.2012.02.038.
200. Kramer, R. H. & Davenport, C. M. (2015). Lateral inhibition in the vertebrate retina: the case of the missing neurotransmitter. *PLoS Biology*, 13 (12): e1002322. DOI: 10.1371/journal.pbio.1002322.
201. Krauel, K., Schott, P., Sojka, B. ... & Ferstl, R. (1999). Is there a mismatch negativity analogue in the olfactory event-related potential? *Journal of Psychophysiology*, 13. DOI: 10.1027//0269-8803.13.1.49.
202. Krekberg, B., Boynton, G. M. & van Wezel, R. J. (2006). Adaptation: from single cells to BOLD signals. *Trends in Neuroscience*, 29. DOI: 10.1016/j.tins.2006.02.008.
203. Kremlacek, J., Kreegipuu, K., Tales, A. ... & Stefanics, G. (2016). Visual mismatch negativity (vMMN): a review and meta-analysis of studies in psychiatric and neurological disorders. *Cortex*, 80. DOI: 10.1016/j.cortex.2016.03.017.

204. Kronbichler, L., Said-Yurekli, S. & Kronbichler, M. (2018). Perceptual expectations of object stimuli modulate repetition suppression in a delayed repetition design. *Scientific Reports*, 8. DOI: 10.1038/s41598-018-31091-4.
205. Krupa, D. J., Matell, M. S., Brisben, A. J. ...& Nicolelis, M. A. (2001). Behavioral properties of the trigeminal somatosensory system in rats performing whisker-dependent tactile discriminations. *Journal of Neuroscience*, 21. DOI: 10.1523/JNEUROSCI.21-15-05752.2001.
206. Kuravi, P., Caggiano, V., Giese, M. & Vogels, R. (2016). Repetition suppression for visual actions in the macaque superior temporal sulcus. *Journal of Neurophysiology*, 115. DOI: 0.1152/jn.00849.2015.
207. Lante, F., de Jesus Ferreira, M. C., Guiramand, J. ...& Vignes, M. (2006). Low-frequency stimulation induces a new form of LTP, metabotropic glutamate (mGlu5) receptor- and PKA-dependent, in the CA1 area of the rat hippocampus. *Hippocampus*, 16. DOI: 10.1002/hipo.20146.
208. Lawrence, J. J., Saraga, F., Churchill, J. F. ...& McBain, C. J. (2006). Somatodendritic Kv7/KCNQ/M channels control interspike interval in hippocampal interneurons. *Journal of Neuroscience*, 26. DOI: 10.1523/JNEUROSCI.3521-06.2006.
209. Lee, T. S. & Nguyen, M. (2001). Dynamics of subjective contour formation in the early visual cortex. *PNAS*, 98. DOI: 10.1073/ypnas.031579998.
210. Lee, S. H., Kwan, A. C., Zhang, S. ...& Dan, Y. (2012). Activation of specific interneurons improves V1 feature selectivity and visual perception. *Nature*, 488. DOI: 10.1038/nature11312.
211. Lefort, S., Tómm, C., Floyd Sarria, J-C. & Petersen, C. C. H. (2009). The excitatory neuronal network of the C2 barrel column in mouse primary somatosensory cortex. *Neuron*, 61. DOI: 10.1016/j.neuron.2008.12.020.
212. Lerch, J. P., Carroll, J. B., Dorr, A. ...& Henkelman, R. M. (2008). Cortical thickness measured from MRI in the YAC128 mouse model of Huntington's disease. *NeuroImage*, 41. DOI: 10.1016/j.neuroimage.2008.02.019.
213. Letzkus, J. J., Kampa, B. M. & Stuart, G. J. (2006). Learning rules for spike timing-dependent plasticity depend on dendritic synapse location. *Journal of Neuroscience*, 26. DOI: 10.1523/JNEUROSCI.2650-06.2006.
214. Li, Y., Ma, W., Li, L. ...& Tao, H. W. (2012). Broadening of inhibitory tuning underlies contrast-dependent sharpening of orientation selectivity in mouse visual cortex. *Journal of Neuroscience*, 32. DOI: 10.1523/JNEUROSCI.3221-12.2012.

215. Li, T., Tian, C., Scalmani, P. ...& Shu, Y. (2014). Action potential initiation in neocortical inhibitory interneurons. *PLoS Biology*. DOI: 10.1371/journal.pbio.1001944.
216. Lima, B., Singer, W., Chen, N. & Neuenschwander, S. (2009). Synchronization dynamics in response to plaid stimuli in monkey V1. *Cerebral Cortex*, 20. DOI: 10.1093/cercor/bhp218.
217. Linden, H., Tetzlaff, T., Potjans, T. C. ...& Einevoll, G. T. (2011). Modeling the spatial reach of the LFP. *Neuron*, 72. DOI: 10.1016/j.neuron.2011.11.006.
218. Liu, B., Li, P., Li, Y. ...& Tao, H. W. (2009). Visual receptive field structure of cortical inhibitory neurons revealed by two-photon imaging guided recording. *Journal of Neuroscience*, 29. DOI: 10.1523/JNEUROSCI.1915-09.2009.
219. Loewy, D. H., Campbell, K. B. & Bastien, C. (1996). The mismatch negativity to frequency deviant stimuli during natural sleep. *Electroencephalography and Clinical Neurophysiology*, 98. DOI: 10.1016/0013-4694(96)95553-4.
220. Logothetis, N. K., Kayser, C. & Oeltermann, A. (2007). In vivo measurement of cortical impedance spectrum in monkeys: implications for signal propagation. *Neuron*, 55. DOI: 10.1016/j.neuron.2007.07.027.
221. Long, M., Jiang, W., Liu, D. & Yao, H. (2015). Contrast-dependent orientation discrimination in the mouse. *Scientific Reports*, 5. 15830. DOI: 10.1038/srep15830.
222. Lorente de No, R. (1947). A study of nerve physiology. *Studies from the Rockefeller Institute for Medical Research*, 132. DOI: NA.
223. Lovett-Barron, M., Turi, G. F., Kaifosh, P. ...& Losonczy, A. (2012). Regulation of neuronal input transformations by tunable dendritic inhibition. *Nature Neuroscience*, 15. DOI: 10.1038/nn.3024.
224. Lundstrom, B. N. & Fairhall, A. L. (2006). Decoding stimulus variance from a distributional neural code of interspike intervals. *Journal of Neuroscience*, 26. DOI: 10.1523/jneurosci.0225-06-2006.
225. Lytton, W. W. & Sejnowski, T. J. (1991). Simulations of cortical pyramidal neurons synchronized by inhibitory interneurons. *Journal of Neurophysiology*, 66. DOI: 10.1152/jn.1991.66.3.1059.
226. Madison, D. V. & Nicoll, R. A. (1984). Control of the repetitive discharge of rat CA1 pyramidal neurones in vitro. *Journal of Physiology*, 354. DOI: 10.1113/jphysiol.1984.sp015378.
227. Maier, A., Adams, G. K., Aura, C. & Leopold, D. A. (2010). Distinct superficial and

deep laminar domains of activity in the visual cortex during rest and stimulation. *Frontiers in Systems Neuroscience*, 4. DOI: 10.3389/fnsys.2010.00031.

228. Mainen, Z. F. & Sejnowski, T. J. (1995). Reliability of spike timing in neocortical neurons. *Science*, 268. DOI: 10.1126/science.7770778.

229. Mann, E. O. (2005b). Perisomatic feedback inhibition underlies cholinergically induced fast network oscillations in the rat hippocampus in vitro. *Neuron*, 45. DOI: 10.1016/j.neuron.2004.12.016.

230. Manning, J. R., Jacobs, J., Fried, I. & Kahana, M. J. (2009). Broadband shifts in local field potential power spectra are correlated with single-neuron spiking in humans. *Journal of Neuroscience*, 29. DOI: 10.1523/JNEUROSCI.2041-09.2009.

231. Maris, E., Schoffelen, J. M. & Fries, P. (2007). Nonparametric statistical testing of coherence differences. *Journal of Neuroscience Methods*, 163. DOI: 10.1016/j.jneumeth.2007.02.011.

232. Markov, N. T., Vezoli, J., Chameau, P. ... & Kennedy, H. (2013). Anatomy of hierarchy: feedforward and feedback pathways in macaque visual cortex. *Journal of Comparative Neurology*, 522. DOI: 10.1002/cne.23458.

233. McAdams, C. J. & Maunsell, J. H. (1999). Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *Journal of Neuroscience*, 19. DOI: 10.1523/JNEUROSCI.19-01-00431.1999.

234. McGregor, J. E., Yin, L., Yang, Q. ... & Merigan, W. H. (2018). Functional architecture of the foveolar revealed in the living primate. *PLoS ONE*, 13(11): e0207102. DOI: 10.1371/journal.pone.0207102.

235. McCormick, D. A., Connors, B. W., Lighthall, J. W. & Prince, D. A. (1985). Comparative electrophysiology of pyramidal and sparsely spiny stellate neurons of the neocortex. *Journal of Neurophysiology*, 54. DOI: 10.1152/jn.1985.54.4.782.

236. Mehaffey, W. H., Doiron, B., Maler, L. & Turner, R. W. (2005). Deterministic multiplicative gain control with active dendrites. *Journal of Neuroscience*, 25. DOI: 10.1523/JNEUROSCI.2682-05.2005.

237. Meister, M. & Berry, M. J. (1999). The neural code of the retina. *Neuron*, 22. DOI: 10.1016/S0896-6273(00)80700-X.

238. Merigan, W. H. & Katz, L. M. (1990). Spatial resolution across the macaque retina. *Vision Research*, 30. DOI: 10.1016/0042-6989(90)90107-V.

239. Meyer, T. and Olson, C. R. (2011). Statistical learning of visual transitions in monkey

inferotemporal cortex. PNAS, 108. DOI: 10.1073/pnas.1112895108.

240. Meyer, T. et al. (2014). Statistical learning of serial visual transitions by neurons in monkey inferotemporal cortex. *Journal of Neuroscience*, 34. DOI: 10.1523/JNEUROSCI.1215-14.2014.

241. Michalareas, G., Vezoli, J., van Pelt, S. ... & Fries, P. (2016). Alpha-beta and gamma rhythms subserve feedback and feedforward influences among human visual cortical areas. *Neuron*, 89. DOI: 10.1016/j.neuron.2015.12.018.

242. Morgan, S. & Huston, J. P. (1990). The interhemispheric projection from the substantia nigra to the caudate-putamen as depicted by the anterograde transport of [3H]leucine. *Behavioural Brain Research*, 38. DOI: 10.1016/0166-4328(90)90013-5.

243. Mountcastle, V. B., Berman, A. L. & Davies, P. W. (1955). Topographic organization and modality representation in first somatic area of cat's cerebral cortex by method of single unit analysis. *American Journal of Physiology*, 193. DOI: 10.1152/jn.1957.20.4.408.

244. Muckli, L., Kohler, A., Kriegeskorte, N. & Singer, W. (2005). Primary visual cortex activity along the apparent-motion trace reflects illusory perception. *PLoS Biology*, 3:e265. DOI: 10.1371/journal.pbio.0030265.

245. Murakami, S. & Okada, Y. (2006). Contributions of principal neocortical neurons to magnetoencephalography and electroencephalography signals. *Journal of Physiology*, 575. DOI: 10.1113/jphysiol.2006.105379.

246. Myers, N. E., Stokes, M. G., Walther, L. & Nobre, A. C. (2014). Oscillatory brain state predicts variability in working memory. *Journal of Neuroscience*, 34. DOI: 10.1523/JNEUROSCI.4741-13.2014.

247. Naatanen, R., Paavilainen, P., Alho, K. ... & Sams, M. (1987). The mismatch negativity to intensity changes in an auditory stimulus sequence. *Electroencephalography and Clinical Neurophysiology*, 40. DOI: NA.

248. Nassi, J. J., Lomber, S. G. & Born, R. T. (2013). Cortico-cortical feedback contributes to surround suppression in V1 of the alert primate. *Journal of Neuroscience*, 33. DOI: 10.1523/JNEUROSCI.5124-12.2013.

249. Natan, R. G., Briguglio, J. J., Mwilambwe-Tshilobo, L., ... & Geffen, M. N. (2015). Complementary control of sensory adaptation by two types of cortical interneurons. *eLife*. DOI: 10.7554/eLife.09868.

250. Nelken, I. (2014). Stimulus-specific adaptation and deviance detection in the auditory system: experiments and models. *Biological Cybernetics*, 108. DOI: 10.1007/s00422-014-0585-7.

251. Ness, T. V., Remme, M. W. H. & Einevoll, G. T. (2018). H-Type membrane current shapes the local field potential from populations of pyramidal neurons. *Journal of Neuroscience*, 38. DOI: 10.1523/JNEUROSCI.3278-17.2018.
252. Neupane, S., Guitton, D. & Pack, C. C. (2016). Two distinct types of remapping in primate cortical area V4. *Nature Communications*, 7. DOI: 10.1038/ncomms10402.
253. Neupane, S., Guitton, D. & Pack, C. C. (2017). Coherent alpha oscillations link current and future receptive fields during saccades. *PNAS*, 114. DOI: 10.1073/pnas.1701672114.
254. Nguyen, P. V. & Kandel, E. R. (1997). Brief theta-burst stimulation induces a transcription-dependent late phase of LTP requiring cAMP in area CA1 of the mouse hippocampus. *Learning & Memory*, 4. DOI: 10.1101/lm.4.2.230.
255. Niell, C. M. & Stryker, M. P. (2008). Highly selective receptive fields in mouse visual cortex. *Journal of Neuroscience*, 28. DOI: 10.1523/JNEUROSCI.0623-08.2008.
256. Nigro, M. J., Mateos-Aparicio, P. & Storm, J. F. (2014). Expression and functional roles of Kv7/KCNQ/M-channels in rat medial entorhinal cortex layer II stellate cells. *Journal of Neuroscience*, 34. DOI: 10.1523/JNEUROSCI.4153-13.2014.
257. Noudoost, B., Nategh, N., Clark, K. & Esteky, H. (2017). Stimulus context alters neural representations of faces in inferotemporal cortex. *Journal of Neurophysiology*, 117. DOI: 10.1152/jn.00667.2016.
258. Ohki, K., Chung, S., Kara, P. ... & Reid, R. C. (2006). Highly ordered arrangement of single neurons in orientation pinwheels. *Nature*, 442. DOI: 10.1038/nature05019.
259. Oleskiw, T. D., Nowack, A. & Pasupathy, A. (2018). Joint coding of shape and blur in area V4. *Nature Communications*, 466. DOI: 10.1038/s41467-017-02438-8.
260. Onuki, Y., Van Someren, E. J. W., De Zeeuw, C. I. & Van der Werf, Y. D. (2015). Hippocampal-cerebellar interaction during spatio-temporal prediction *Cerebral Cortex*, 25. DOI: 10.1093/cercor/bht221.
261. Orekhova, E. V., Stroganova, T. A., Nygren, G. ... & Elam, M. (2007). Excess of high frequency electroencephalogram oscillations in boys with autism. *Biological Psychiatry*, 62. DOI: 10.1016/j.biopsych.2006.12.029.
262. Oren, I., Mann, E. O., Paulsen, O. & Hajos, N. (2006). Synaptic currents in anatomically identified CA3 neurons during hippocampal gamma oscillations in vitro. *Journal of Neuroscience*, 26. DOI: 10.1523/jneurosci.1580-06.2006.
263. Paavilainen, P. (2013). The mismatch-negativity (MMN) component of the auditory

event-related potential to violations of abstract regularities: a review. *International Journal of Psychophysiology*, 88. DOI: 10.1016/j.ijpsycho.2013.03.015.

264. Paavilainen, P., Kaukinen, C., Koskinen, O. ... & Rehn, L. (2018). Mismatch negativity (MMN) elicited by abstract regularity violations in two concurrent auditory streams. *Heliyon*, 4, DOI: 10.1016/j.heliyon.2018.e00608.

265. Park, P., Volianskis, A., Sanderson, T. M. ... & Collingridge, G. L. (2014). NMDA receptor-dependent long-term potentiation comprises a family of temporally overlapping forms of synaptic plasticity that are induced by different patterns of stimulation. *Philosophical Transactions of the Royal Society, Biological Sciences*, 369. DOI: 10.1098/rstb.2013.0131.

266. Peng, H., Bian, X., Ma, F. & Wang, K-W. (2017). Pharmacological modulation of the voltage-gated neuronal Kv7/KCNQ/M-channel alters the intrinsic excitability and synaptic responses of pyramidal neurons in rat prefrontal cortex slices. *Acta Pharmacologica Sinica*, 38. DOI: 10.1038/aps.2017.72.

267. Penttonen, M., Kamondi, A., Acsady, L. & Buzsaki, G. (1998). Gamma frequency oscillation in the hippocampus of the rat: intracellular analysis in vivo. *European Journal of Neuroscience*, 10. DOI: 10.1046/j.1460-9568.1998.00096.x.

268. Perkel, D. J., Bullier, J. & Kennedy, H. (1986). Topography of the afferent connectivity of area 17 in the macaque monkey: a double-labelling study. *Journal of Comparative Neurology*, 253. DOI: 10.1002/cne.902530307.

269. Peron, S. P., Freeman, J., Iyer, V. ... & Svoboda, K. (2015). A cellular resolution map of barrel cortex activity during tactile behavior. *Neuron*, 86. DOI: 10.1016/j.neuron.2015.03.027.

270. Petersen, C. C. H. (2007). The functional organization of the barrel cortex. *Neuron*, 56. DOI: 10.1016/j.neuron.2007.09.017.

271. Pinto, L., Goard, M. J., Estandian, D. ... & Dan, Y. (2013). Fast modulation of visual perception by basal forebrain cholinergic neurons. *Nature Neuroscience*, 16. DOI: 10.1038/nn.3552.

272. Pitcher, D., Walsh, V. & Duchaine, B. (2011). The role of the occipital face area in the cortical face perception network. *Experimental Brain Research*, 209. DOI: 10.1007/s00221-011-2579-1.

273. Polania, R., Krajbich, I., Grueschow, M. & Ruff, C. C. (2014). Neural oscillations and synchronization differentially support evidence accumulation in perceptual and value-based decision making. *Neuron*, 82. DOI: 10.1016/j.neuron.2014.03.014.

274. Polosecki, P., Moeller, S., Schweers, N. ... & Freiwald, W. A. (2013). Faces in motion: selectivity of macaque and human face processing areas for dynamic stimuli. *Journal of*

Neuroscience, 33. DOI: 10.1523/JNEUROSCI.5402-11.2013.

275. Popov, T., Kastner, S. & Jensen, O. (2017). FEF-controlled alpha delay activity precedes stimulus-induced gamma-band activity in visual cortex. *Journal of Neuroscience*, 37. DOI: 10.1523/JNEUROSCI.3015-16.2017.

276. Porciatti, V., Pizzorusso, T. & Maffei, L. (1999). The visual physiology of the wild type mouse determined with pattern VEPs. *Vision Research*, 39. DOI: 10.1016/S0042-6989(99)00022-X.

277. Price, D. J. & Morgan, J. E. (1987). Spatial properties of neurones in the lateral geniculate nucleus of the pigmented ferret. *Experimental Brain Research*, 68. DOI: 10.1007/BF00255231.

278. Prusky, G. T., West, P. W. & Douglas, R. M. (2000). Behavioral assessment of visual acuity in mice and rats. *Vision Research*, 30. DOI: 10.1016/S0042-6989(00)00081-X.

279. Prusky, G. T. & Douglas, R. M. (2004). Characterization of mouse cortical spatial vision. *Vision Research*, 44. DOI: 10.1016/j.visres.2004.09.001.

280. Puri, A. M., Wojciulik, E. & Ranganath, C. (2009). Category expectation modulates baseline and stimulus-evoked activity in human inferotemporal cortex. *Brain Research*, 1301. DOI: 10.1016/j.brainres.2009.08.085.

281. Ramaswami, M. (2014). Network plasticity in adaptive filtering and behavioral habituation. *Neuron*, 82. DOI: 10.1016/j.neuron.2014.04.035.

282. Randlett, O., Haesemeyer, M., Forkin, G. ... & Granato, M. (2019). Distributed plasticity drives visual habituation learning in larval zebrafish. *Current Biology*, 29. DOI: 10.1016/j.cub.2019.02.039.

283. Rao, R. P. N. & Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2. DOI: 10.1038/4580.

284. Ray, S. & Maunsell, J. H. R. (2011). Different origins of gamma rhythm and high-gamma activity in macaque visual cortex. *PLoS Biology*. DOI: 10.1371/journal.pbio.1000610.

285. Ray, S., Niebur, E., Hsiao, S. S. ... & Crone, N. E. (2008). High-frequency gamma activity (80-150Hz) is increased in human cortex during selective attention. *Clinical Neurophysiology*, 119. DOI: 10.1016/j.clinph.2007.09.136.

286. Reich, D. S., Mechler, F., Purpura, K. P. & Victor, J. D. (2000). Interspike intervals, receptive fields, and information encoding in primary visual cortex. *Journal of Neuroscience*, 20. DOI: 10.1523/jneurosci.20-05-01964.

287. Richter, C. G., Thompson, W. H., Bosman, C. A. & Fries, P. (2017). Top-down beta enhances bottom-up gamma. *Journal of Neuroscience*, 37. DOI: 10.1523/JNEUROSCI.3771-16.2017.
288. Rihs, T. A., Michel, C. M. & Thut, G. (2007). Mechanisms of selective inhibition in visual spatial attention are indexed by alpha-band EEG synchronization. *European Journal of Neuroscience*, 25. DOI: 10.1111/j.1460-9568.2007.05278.x.
289. Ringach, D. L., Mineault, P. J., Tring, E. ... & Trachtenberg, J. T. (2016). Spatial clustering of tuning in mouse primary visual cortex. *Nature Communications*, 7: 12270. DOI: 10.1038/ncomms12270.
290. Roach, B. J. & Mathalon, D. H. (2008). Event-related EEG time-frequency analysis: an overview of measures and an analysis of early gamma band phase locking in schizophrenia. *Schizophrenia Bulletin*, 34. DOI: 10.1093/schbul/sbn093.
291. Roberts, P. D. & Leen, T. K. (2010). Anti-Hebbian spike-timing-dependent plasticity and adaptive sensory processing. *Frontiers in Computational Neuroscience*, 4. DOI: 10.3389/fncom.2010.00156.
292. Roberts, M. J., Lowet, E., Brunet, N. M. ... & De Weerd, P. (2013). Robust gamma coherence between macaque V1 and V2 by dynamic frequency matching. *Neuron*, 78. DOI: 10.1016/j.neuron.2013.03.003.
293. Rockland, K. S. & Pandya, D. N. (1979). Laminar origins and terminations of cortical connections of the occipital lobe in the rhesus monkey. *Brain Research*, 178, DOI: 10.1016/0006-8993(79)90485-2.
294. Rohenkohl, G., Bosman, C. A. & Fries, P. (2018). Gamma synchronization between V1 and V4 improves behavioral performance. *Neuron*, 100. DOI: 10.1016/j.neuron.2018.09.019.
295. Ruyter Van Steveninck, R., Bialek, W. (1988). Real-time performance of a movement-sensitive neuron in the blowfly visual system: coding and information transfer in short spike sequences. *Proceedings of the Royal Society, Biological Sciences*, 234. DOI: 10.1098/rspb.1988.0055.
296. Sabri, M., De Lugt, D. R. & Campbell, K. B. (2000). The mismatch negativity to frequency deviants during the transition from wakefulness to sleep. *Canadian Journal of Experimental Psychology*, 54. DOI: 10.1037/h0087343.
297. Sakata, S. & Harris, K. D. (2009). Laminar structure of spontaneous and sensory-evoked population activity in auditory cortex. *Neuron*, 64. DOI: 10.1016/j.neuron.2009.09.020.7.
298. Salin, P. A. & Bullier, J. (1995). Corticocortical connections in the visual system:

structure and function. *Physiological Reviews*, 75. DOI: 10.1152/physrev.1995.75.1.107.

299. Salinas-Navarro, M., Mayor-Torroglosa, S., Jimenez-Lopez, M. ... & Vidal-Sanz, M. (2009). A computerized analysis of the entire retinal ganglion cell population and its spatial distribution in adult rats. *Vision Research*, 49. DOI: 10.1016/j.visres.2008.09.029.

300. Salisbury, D. F. (2012). Finding the missing stimulus mismatch negativity (MMN): Emitted MMN to violations of an auditory gestalt. *Psychophysiology*, 49. DOI: 10.1111/j.1469-8986.2011.01336.x.

301. Salkoff, D. B., Zagha, E., Yuzgec, O. & McCormick, D. A. (2015). Synaptic mechanisms of tight spike synchrony at gamma frequency in cerebral cortex. *Journal of Neuroscience*, 35. DOI: 10.1523/JNEUROSCI.0828-15.2015.

302. Sarihi, A., Mirnajafi-Zadeh, J., Jiang, B. ... & Tsumoto, T. (2012). Cell type-specific, presynaptic LTP of inhibitory synapses on fast-spiking GABAergic neurons in the mouse visual cortex. *Journal of Neuroscience*, 32. DOI: 10.1523/JNEUROSCI.1386-12.2012.

303. Scheeringa, R., Koopmans, P. J., van Mourik, T. ... & Norris, D. G. (2016). The relationship between oscillatory EEG activity and the laminar-specific BOLD signal. *PNAS*, 113. DOI: doi.org/10.1073/pnas.1522577113.

304. Scheyltjens, I. & Lutgarde, A. (2016). The current status of somatostatin-interneurons in inhibitory control of brain function and plasticity. *Neural Plasticity*, DOI: 10.1155/2016/8723623.

305. Schlagbauer, B., Mink, M., Muller, H. J. & Geyer, T. (2017). Independence of long-term contextual memory and short-term perceptual hypotheses: Evidence from contextual cueing of interrupted search. *Attention, Perception, & Psychophysics*, 79. DOI: 10.3758/s13414-016-1246-9.

306. Schmucker, C., Seeliger, M., Humphries, P. ... & Schaeffel, F. (2005). Grating acuity at different luminances in wild-type mice and in mice lacking rod or cone function. *Investigative Ophthalmology & Visual Science*, 46. DOI: 10.1167/iovs.04-0959.

307. Schnabel, U. H., Kirchberger, L., van Best, E. H. ... & Roelfsema, P. R. (2018). Feedforward and feedback processing during figure-ground perception in mice. *BioRxiv*. DOI: 10.1101/456459.

308. Schroger, E. (1998). Measurement and interpretation of the mismatch negativity. *Behaviour Research Methods, Instruments, & Computers*, 30. DOI: 10.3758/BF03209423.

309. Schuett, S., Bonhoeffer, T. & Hubener, M. (2002). Mapping retinotopic structure in mouse visual cortex using optical imaging. *Journal of Neuroscience*, 22. DOI:

10.1523/JNEUROSCI.22-15-06549.2002.

310. Schwarz, C., Hentschke, H., Butovas, S. ...& Waiblinger, C. (2010). The head-fixed behaving rat – procedures and pitfalls. *Somatosensory & Motor Research*, 27. DOI: 10.3109/08990220.2010.513111.

311. Sculthorpe, L. D., Ouellet, D. R. & Campbell, K. B. (2009). MMN elicitation during natural sleep to violations of an auditory pattern. *Brain Research*, 1290. DOI: 10.1016/j.brainres.2009.06.013.

312. Seabrook, T. A., Burbridge, T. J., Crair, M. C. & Huberman, A. D. (2017). Architecture, function, and assembly of the mouse visual system. *Annual Review of Neuroscience*, 40. DOI: 10.1146/annurev-neuro-071714-033842.

313. Sedley, W., Gander, P. E., Kumar, S. ...& Griffiths, T. D. (2016) Neural signatures of perceptual inference. *eLife* 5:e11476. DOI: 10.7554/eLife.11476.

314. Self, M. W., Kooijmans, R. N., Super, H. ...& Roelfsema, P. R. (2012). Different glutamate receptors convey feedforward and recurrent processing in macaque V1. *PNAS*, 109. DOI: 10.1073/pnas.1119527109.

315. Senzai, Y., Fernandez-Ruiz, A. & Buzsaki, G. (2019). Layer-specific physiological features and interlaminar interactions in the primary visual cortex of the mouse. *Neuron*, 101. DOI: 10.1016/j.neuron.2018.12.009.

316. Seymour, R. A., Rippon, G. & Kessler, K. (2017). The detection of phase amplitude coupling during sensory processing. *Frontiers in Neuroscience*. DOI: 10.3389/fnins.2017.00487.

317. Sham, P. C. & Curtis, D. (1995). Monte Carlo tests for associations between disease and alleles at highly polymorphic loci. *Annals of Human Genetics*, 59. DOI: 10.1111/j.1469-1809.1995.tb01608.x.

318. Sherlock, D. A., Field, P. M. & Raisman, G. (1975). Retrograde transport of horseradish peroxidase in the magnocellular neurosecretory system of the rat. *Brain Research*, 88. DOI: 10.1016/0006-8993(75)90653-8.

319. Sherman, S. M. (2016). Thalamus plays a central role in ongoing cortical functioning. *Nature Neuroscience*, 19. DOI: 10.1038/nn.4269.

320. Sherman, S. M. & Guillery, R. W. (1998). On the actions that one nerve cell can have on another: Distinguishing ‘drivers’ from ‘modulators’. *PNAS*, 95. DOI: 10.1073/pnas.95.12.7121.

321. Shigeno, S. (2017). Effects of auditory and visual priming on the identification of spoken words. *Perceptual and Motor Skills*, 124. DOI: 10.1177/0031512516684459.

322. Shirvalkar, P. R., Rapp, P. R. & Shapiro, M. L. (2010). Bidirectional changes to hippocampal theta-gamma comodulation predict memory for recent spatial episodes. *PNAS*, 107. DOI: 10.1073/pnas.0911184107.
323. Shmuel, A., Korman, M., Sterkin, A. ... & Grinvald, A. (2005). Retinotopic axis specificity and selective clustering of feedback projections from V2 to V1 in the owl monkey. *Journal of Neuroscience*, 25. DOI: 10.1523/JNEUROSCI.4137-04.2005.
324. Silva, L. R., Amitai, Y. & Connors, B. W. (1991). Intrinsic oscillations of neocortex generated by layer 5 pyramidal neurons. *Science*, 251. DOI: 10.1126/science.1824881.
325. Sinex, D. G., Burdette, L. J. & Pearlman, A. L. (1979). A psychophysical investigation of spatial vision in the normal and reeler mutant mouse. *Vision Research*, 19. DOI: 10.1016/0042-6989(79)90018-X.
326. Singer, W. & Gray, C. M. (1995). Visual feature integration and the temporal correlation hypothesis. *Annual Review of Neuroscience*, 18. DOI: 10.1146/annurev.ne.18.030195.003011.
327. Slotnick, B. (2009). A simple 2-transistor touch or lick detector circuit. *Journal of Experimental Analysis of Behaviour*, 91. DOI: 10.1901/jeab.2009.91-253.
328. Smith, A. T. (2001). Estimating receptive field size from fMRI data in human striate and extrastriate visual cortex. *Cerebral Cortex*, 11. DOI: 10.1093/cercor/11.12.1182.
329. Sofroniew, N. J., Vlasov, Y. A., Hires, S. A. ... & Svoboda, K. (2015). Neural coding in barrel cortex during whisker-guided locomotion. *eLife*. DOI: 10.7554/eLife.12559.
330. Sohal, V. S., Zhang, F., Yizhar, O. & Deisseroth, K. (2009). Parvalbumin neurons and gamma rhythms enhance cortical circuit performance. *Nature*, 459. DOI: 10.1038/nature07991.
331. Sohya, K., Kameyama, K., Yanagawa, Y. ... & Tsumoto, T. (2007). GABAergic neurons are less selective to stimulus orientation than excitatory neurons in Layer II/III of visual cortex, as revealed by in vivo functional Ca²⁺ imaging in transgenic mice. *Journal of Neuroscience*, 27. DOI: 10.1523/JNEUROSCI.4641-06.2007.
332. Spaak, E., Bonnefond, M., Maier, A. ... & Jensen, O. (2012). Layer-specific entrainment of gamma-band neural activity by the alpha rhythm in monkey visual cortex. *Current Biology*, 22. DOI: 10.1016/j.cub.2012.10.020.
333. Stark, E. & Abeles, M. (2007). Predicting movement from multiunit activity. *Journal of Neuroscience*, 27. DOI: 10.1523/JNEUROSCI.1321-07.2007.

334. Stefanics, G., Kremlacek, J. & Czigler, I. (2014). Visual mismatch negativity: a predictive coding view. *Frontiers in Human Neuroscience*. DOI: doi.org/10.3389/fnhum.2014.00666.
335. Stefanics, G., Heinzle, J., Horvath, A. A. & Stephan, K. E. (2018). Visual mismatch and predictive coding: a computational single-trial ERP study. *Journal of Neuroscience*, 38. DOI: 10.1523/JNEUROSCI.3365-17.2018.
336. Stokes, M., Thompson, R., Nobre, A. C. & Duncan, J. (2009). Shape-specific preparatory activity mediates attention to targets in human visual cortex. *PNAS USA*, 106. DOI: 10.1073/pnas.0905306106.
337. Storm, J. F. (1987). Action potential repolarization and a fast after-hyperpolarization in rat hippocampal pyramidal cells. *Journal of Physiology*, 385. DOI:10.1113/jphysiol.1987.sp016517.
338. Straus, A., Wostmann, M. & Obleser, J. (2014). Cortical alpha oscillations as a tool for auditory selective inhibition. *Frontiers in Human Neuroscience*. DOI: 10.3389/fnhum.2014.00350.
339. Stroh, A., Adelsberger, H., Groh, A. ...& Konnerth, A. (2013). Making waves: initiation and propagation of corticothalamic Ca^{2+} waves in vivo. *Neuron*, 77. DOI: 10.1016/j.neuron.2013.01.031.
340. Sturgill, J. F. & Isaacson, J. S. (2015). Somatostatin-expressing interneurons provide subtractive inhibition and regulate sensory response fidelity in olfactory cortex. *Nature Neuroscience*, 18. DOI: 10.1038/nn.3971.
341. Summerfield, C., Monti, J. M. P., Trittschuh, E. H. ...& Egner, T. (2008). Neural repetition suppression reflects fulfilled perceptual expectations. *Nature Neuroscience*, 11. DOI: 10.1038/nn.2163.
342. Sumser, A., Mease, R. A., Sakmann, B. & Groh, A. (2017). Organization and somatotopy of corticothalamic projections from L5B in mouse barrel cortex. *PNAS*, 114. DOI: 10.1073/pnas.1704302114.
343. Sun, W. & Dan, Y. (2009). Layer-specific network oscillation and spatiotemporal receptive field in the visual cortex. *PNAS*, 106. DOI: doi: 10.1073/pnas.0903962106.
344. Taaseh, N., Yaron, A. & Nelken, I. (2011). Stimulus-specific adaptation and deviance detection in the rat auditory cortex. *PLoS ONE*. DOI: 10.1371/journal.pone.0023369.
345. Takaura, K. & Fujii, N. (2016). Facilitative effect of repetitive presentation of one stimulus on cortical responses to other stimuli in macaque monkeys – a possible neu-

ral mechanism for mismatch negativity. *European Journal of Neuroscience*, 43. DOI: 10.1111/ejn.13136.

346. Theyel, B. B., Llano, D. A. & Sherman, M. S. (2010). The corticothalamocortical circuit drives higher-order cortex in the mouse. *Nature Neuroscience*, 13. DOI: 10.1038/nn.2449.

347. Thomson, A. M., West, D. C., Wang, Y. & Bannister, A. P. (2002). Synaptic connections and small circuits involving excitatory and inhibitory neurons in Layers 2-5 of adult rat and cat neocortex: triple intracellular recordings and biocytin labelling in vitro. *Cerebral Cortex*, 12. DOI:10.1093/cercor/12.9.936.

348. Thomson, A. M. & Bannister, A. P. (2003). Interlaminar connections in the neocortex. *Cerebral Cortex*, 13. DOI: 10.1093/cercor/13.1.5.

349. Thut, G., Nietzel, A., Brandt, S. A. & Pascual-Leone, A. (2006). α -Band electroencephalographic activity over occipital cortex indexes visuospatial attention bias and predicts visual target detection. *Journal of Neuroscience*, 26. DOI: 10.1523/JNEUROSCI.0875-06.

350. Tierney, A. J. & Harris-Warrick, R. M. (1992). Physiological role of the transient potassium current in the pyloric circuit of the lobster stomatogastric ganglion. *Journal of Neurophysiology*, 67. DOI: 10.1152/jn.1992.67.3.599.

351. Timney, B., Mitchell, D. E. & Giffin, F. (1978). The development of vision in cats after extended periods of dark-rearing. *Experimental Brain Research*, 31. DOI: 10.1007/BF00239811.

352. Todorovic, A., van Ede, F., Maris, E. & de Lange, F. P. (2011). Prior expectation mediates neural adaptation to repeated sounds in the auditory cortex: an MEG study. *Journal of Neuroscience*, 31. DOI: 10.1523/JNEUROSCI.

353. Torrence, C. & Compo, G. P. (1998). A practical guide to wavelet analysis. *Bulletin of the American Meteorological Society*. DOI: 10.1175/1520-0477(1998)079<0061.

354. Traub, R. D., Whittington, M. A., Colling, S. B. ... & Jefferys, J. G. (1996). Analysis of gamma rhythms in the rat hippocampus in vitro and in vivo. *Journal of Physiology*, 493. DOI: 10.1113/jphysiol.1996.sp021397.

355. Traub, R. D., Bibbig, A., LeBeau, F. E. N. ... & Whittington, M. A. (2005). Persistent gamma oscillations in superficial layers of rat auditory neocortex: experiment and model. *Journal of Physiology*, 562. DOI: 10.1113/jphysiol.2004.074641.

356. Treue, S. & Martinez Trujillo, J. C. (1999). Feature-based attention influences motion processing in macaque visual cortex. *Nature*, 399. DOI: 10.1038/21176.

357. Tsao, D. Y., Freiwald, W. A., Tootell, R. B. H. & Livingstone, M. S. (2006). A cortical region consisting entirely of face-selective cells. *Science*, 311. DOI: 10.1126/science.1119983.
358. Tsytsarev, V., Pope, D., Pumbo, E. & Garver, W. (2010). Intrinsic optical imaging of directional selectivity in rat barrel cortex: application of a multidirectional magnetic whisker stimulator. *Journal of Neuroscience Methods*, 189. DOI: 10.1016/j.jneumeth.2010.03.010.
359. Turner, K. M., Simpson, C. G. & Burne, T. H. J. (2017). BALB/c mice can learn touchscreen visual discrimination and reversal tasks faster than C57BL/6 mice. *Frontiers in Behavioral Neuroscience*, 11: 16. DOI: 10.3389/fnbeh.2017.00016.
360. Ulanovsky, N., Las, L. & Nelken, I. (2003). Processing of low-probability sounds by cortical neurons. *Nature Neuroscience*, 4. DOI: 10.1038/nn1032.
361. Van der Meer, M. A., Kalenscher, T., Lansink, C. S. ... & Redish, A. D. (2010). Integrating early results on ventral striatal gamma oscillations in the rat. *Frontiers in Neuroscience*, 4:300. DOI: 10.3389/fnins.2010.00300.
362. Van der Meer, M. A. & Redish, A. D. (2009). Low and high gamma oscillations in rat ventral striatum have distinct relationships to behaviour, reward, and spiking activity on a learned spatial decision task. *Frontiers in Integrative Neuroscience*, 3:9. DOI: 10.3389/neuro.07.009.2009.
363. Van Ede, F., Jensen, O. & Maris, E. (2010). Tactile expectation modulates pre-stimulus beta-band oscillations in human sensorimotor cortex. *Neuroimage*, 51. DOI: 10.1016/j.neuroimage.2010.02.053.
364. Van Kerkoerle, T., Self, M. W., Dagnino, B. ... & Roelfsema, P. R. (2014). Alpha and gamma oscillations characterize feedback and feedforward processing in monkey visual cortex. *PNAS*, 111. DOI: 10.1073/pnas.1402773111.
365. van Pelt, S., Heil, L., Kwisthout, J. ... & Bekkering, H. (2016) Beta- and gamma-band activity reflect predictive coding in the processing of causal events. *Social, Cognitive and Affective Neuroscience*. 11. DOI: 10.1093/scan/nsw017.
366. Vedell, P. T., Svenson, K. L. & Churchill, G. A. (2011). Stochastic variation of transcript abundance in C57BL/6J mice. *BMC Genomics*, 12. DOI: 10.1186/1471-2164-12-167.
367. Veit, J., Hakim, R., Jadi, M. P. ... & Adesnik, H. (2017). Cortical gamma band synchronization through somatostatin interneurons. *Nature Neuroscience*, 20. DOI: 10.1038/nn.4562.
368. Voss, J. L., Federmeier, K. D. & Paller, K. A. (2012). The potato chip really does look

like Elvis! Neural hallmarks of conceptual processing associated with finding novel shapes subjectively meaningful. *Cerebral Cortex*, 22. DOI: 10.1093/cercor/bhr315.

369. Vreysen, S., Zhang, B., Chino, Y. M. ... & Van den Bergh, G. (2012). Dynamics of spatial frequency tuning in mouse visual cortex. *Journal of Neurophysiology*, 107. DOI: 10.1152/jn.00022.2012.

370. Wacogne, C., Labyt, E., van Wassenhove, V. ... & Dehaene, S. (2011). Evidence for a hierarchy of predictions and prediction errors in human cortex. *PNAS*, 108. DOI: 10.1073/pnas.1117807108.

371. Wang, L., Rangarajan, K. V., Gerfen, C. R. & Krauzlis, R. J. (2018). Activation of striatal neurons causes a perceptual decision bias during visual change detection in mice. *Neuron*, 97. DOI: 10.1016/j.neuron.2018.01.049.

372. Wang, Y., Wei, W., Zhang, X. ... & Yang, Y. (2016). A mouse model of visual perceptual learning reveals alterations in neuronal coding and dendritic spine density in the visual cortex. *Frontiers in Behavioral Neuroscience*, 10: 42. DOI: 10.3389/fnbeh.2016.00042.

373. Wassle, H., Grunert, U., Rohrenbeck, J. & Boycott, B. B. (1990). Retinal ganglion cell density and cortical magnification factor in the primate. *Vision Research*, 30. DOI: 10.1016/0042-6989(90)90166-I.

374. Weiler, N., Wood, L., Yu, J. ... & Shepherd, G. M. G. (2008). Top-down laminar organization of the excitatory network in motor cortex. *Nature Neuroscience*, 11. DOI:10.1038/nn2049.

375. White, J. A., Banks, M. I., Pearce, R. A. & Kopell, N. J. (2000). Networks of interneurons with fast and slow γ -aminobutyric acid type A (GABAA) kinetics provide substrate for mixed gamma-theta rhythm. *PNAS*, 97. DOI: 10.1073/pnas.100124097.

376. Whittington, M. A., Traub, R. D. & Jefferys, J. G. (1995). Synchronized oscillations in interneuron networks driven by metabotropic glutamate receptor activation, 373. DOI: 10.1038/373612a0.

377. Winkler, I. (2007) Interpreting the mismatch negativity (MMN). *Journal of Psychophysiology*, 21. DOI: 10.1027/0269-8803.21.34.147.

378. Womelsdorf, T., Schoffelen, J. M., Oostenveld, R. ... & Fries, P. (2007). Modulation of neuronal interactions through neuronal synchronization. *Science*, 316. DOI: 10.1126/science.1139597.

379. Woolsey, T. A. & Van der Loos, H. (1970). The structural organization of layer IV in the somatosensory region (S1) of mouse cerebral cortex. The description of a cortical field composed of discrete cytoarchitectonic units. *Brain Research*, 17. DOI: 10.1016/0006-

8993(70)90079-x.

380. Worden, M. S., Foxe, J. J., Wang, N. & Simpson, G. V. (2000). Anticipatory biasing of visuospatial attention indexed by retinotopically specific α -band electroencephalography increases over occipital cortex. *Journal of Neuroscience*, 20. DOI:10.1523/JNEUROSCI.20-06-j0002.2000.

381. Xing, D., Yeh, C-I. & Shapley, R. M. (2010). Spatial spread of the local field potential and its laminar variation in visual cortex. *Journal of Neuroscience*, 29. DOI: 10.1523/JNEUROSCI.2573-09.2009.

382. Xygonakis, I., Athanasiou, A., Pandria, N. ...& Bamidis, P. D. (2018). Decoding motor imagery through common spatial pattern filters at the EEG source space. *Computational Intelligence and Neuroscience*, 7957408. DOI: 10.1155/2018/7957408.

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

384. Young, A. W., Hellawall, D. & De Haan, E. H. F. (1988). Cross-domain semantic priming in normal subjects and a prosopagnosic patient. *Quarterly Journal of Experimental Psychology*, 40. DOI: 10.1080/02724988843000087.

385. Yue, C. & Yaari, Y. (2004). KCNQ/M channels control spike afterdepolarization and burst generation in hippocampal neurons. *Journal of Neuroscience*, 24. DOI: 10.1523/JNEUROSCI.0765-04.2004.

386. Zeki, S. (2015). Area V5 – a microcosm of the visual brain. *Frontiers in Integrative Neuroscience*, 9. DOI: 10.3389/fnint.2015.00021.

387. Zeki, S. & Shipp, S. (1988). The functional logic of cortical connections. *Nature*, 335. DOI: 10.1038/335311a0.

388. Zhang, W. & Bruno, R. M. (2019). Higher-order thalamic inputs to primary somatosensory cortex are stronger and longer lasting than cortical inputs. *eLIFE*: e44158 DOI: 10.7554/eLife.44158.

389. Zhang, Q., Li, X., Gold, B. T. & Jiang, Y. (2010). Neural correlates of cross-domain affective priming. *Brain Research*, 1329. DOI: 10.1016/j.brainres.2010.03.021.

390. Zhang, Z., Russell, L. E., Packer, A. M. ...& Hausser, M. (2018). Closed-loop all-optical interrogation of neural circuits in vivo. *Nature Methods*, 15. DOI: 10.1038/s41592-018-0183-z.